Service Availability in IP Networks

Supratik Bhattacharyya, Sprint ATL
Christophe Diot, Intel Research
Gianluca Iannaccone, Sprint ATL
Athina Markopoulou, Stanford University
Chen-Nee Chuah, UC Davis

***Abstract*— Recent studies reveal that Internet backbones are capable of providing very low loss rates and end-to-end delays close to the speed of light (in optical fiber). However link failures may lead to packet losses, delay variations and outages due to missing or inconsistent forwarding information in routers. Link failures occur everyday, and may ultimately affect the quality of service experienced by an ISP's customers. On the other hand, not all link failures impact customer service. Therefore we advocate the need for a new metric to characterize IP service availability in the presence of link failures. We analyze the various factors affecting service availability in IP networks, and explore how to define this metric.**

I. Introduction

Service-Level Agreements (SLAs) offered by today's Internet Service Providers (ISPs) are based on three metrics: loss, delay and port availability. The first two metrics are usually computed network-wide and averaged over a relatively long period of time (e.g., a month). For the third metric, the term "port" refers to the point at which a customer's link attaches to the edge of an ISP's network. Port availability therefore refers to the fraction of time this port is operational, and measures a customer's physical connectivity to the ISP's network. None of these SLA metrics capture the ability of the network to carry customer traffic to Internet destinations at any point in time.

Current provisioning and engineering practices in IP backbones makes it relatively easy for ISPs to meet their SLA guarantees for loss and delay. For example, an ISP may be able to guarantee a loss of less than 1%, end-to-end delay of 55 msec within continental USA, and port availability of 99.9%.

However, current SLAs do not offer a metric to quantify how often a network is unable to forward packets from a source to a destination due to failures (although this may be indirectly reflected in the loss rate). IP-level link failures occur as part of everyday operations and result from a variety of causes such as optical fiber cut, optical equipment malfunction, router crashes, router software bugs, protocol implementation errors, etc.[1]. Protection and restoration mechanisms at or below the IP layer are used to guard against the occurrence of failures (e.g., SONET protection switching) or to recover quickly from them (e.g., MPLS fast restoration, IS-IS/OSPF fast convergence, etc.).

Network failures may affect the IP-level performance experienced by backbone customers in several ways. When a link/router fails, traffic is rerouted via alternate paths, and may congest links along those backup paths[2]. Congestion may lead to packet losses and may violate the SLA offered by an ISP. Such congestion may be severe in the event of widespread outages involving several links.

Failures may also affect the packet delays experienced by a customer. The primary and backup paths at the IP layer are usually mapped onto disjoint fiber paths at the optical layer to improve robustness to fiber cuts. Hence end-to-end delays on the primary and backup IP-level paths may differ by tens of milliseconds. It has been shown that this is the major source of jitter in IP backbones[3].

There are other ways in which link failures may impact a network. While an alternate path is being set up around a link failure, routers may lack forwarding information or may have inconsistent forwarding information, resulting in black-holes and transient loops[4, 5]. Backbone link failures can also change the exit point for routes to external networks learned through BGP[6]. This may further prolong routing convergence around failures.

Therefore ISPs need to incorporate a metric in their SLAs that quantifies the disruption in packet forwarding due to failures. We refer to this metric as *service availability*. We believe that an SLA that truly differentiates among the service offered by ISPs must include service availability.

Service availability from a source to a destination in an IP network refers to the ability of the network to deliver IP packets from the source to the destination. We analyze here the factors influencing service availability, and explore the challenges in defining and measuring it on a network-wide basis. While the existence of physical connectivity between two points is key to service availability, we identify and discuss a number of other factors as well, such as restoration time around failures, frequency of failures and likelihood of simultaneous failures on primary and backup paths. We argue that service availability has to be defined in a manner that allows for clear quantitative comparison among ISPs. Furthermore, application-level requirements such as the the maximum duration of a single outage that an application can tolerate, must be considerd. Our discussion is illustrated with observations from Sprint's tier-1 IP backbone network.

This paper describes work in progress. It is intended to draw attention to a topic that has high practical significance for network users, but has not been addressed in a systematic manner. From a research perspective, this paper represents a departure from the traditional approach towards Quality of Service (QoS) based on loss, delay and jitter. The proposed service availability metric holds the promise of influencing future research in topics such as network design, network management, protection and restoration techniques, router

Contact author. Postal addres: Sprint ATL, 1 Adrian Court, Burlingame CA 94010, USA. Phone: +1 650-375-4124. Email: supratik@sprintlabs.com.

architecture and routing protocol design.

The rest of the paper is organized as follows. Section 2 describes related work. Section 3 enunciates the notion of IP service availability and discusses the various factors affecting it. Section 4 uses observations from the Sprint IP backbone network to demonstrate the importance and necessity of defining service availability. Section 5 outlines some challenges in defining and measuring this metric, and section 6 lays out directions for future work.

II. Relevant Work

The issue of availability has been primarily studied in the context of telephone networks, where the gold standard is “five nines” availability, i.e., a network which is available for 99.999% of the time. [7] has addressed the causes of failures in public switched telephone networks, while [8] has qualitatively evaluated an ANSI-developed survivability metric for telecommunications. This metric is referred to as the *outage-index*, and allows an outage episode to be characterized as a single metric by incorporating service, duration and magnitude of the outage. The notion of availability has long been used to measure the quality of telephone networks. IP networks are different from telephone networks in that they provide best-effort packet delivery with no call admission control. Hence the notion of availability and service outage also differs significantly from telephone networks.

The availability of spare capacity together with sound engineering practices in commercial IP backbones makes it easy to achieve traditional quality of service (QoS) objectives such as low loss and delay. Recent results from Sprint’s IP backbone show that the network can support toll-quality voice services [4]. However observations on the same backbone have shown that equipment failure can create significant outages and delay variations [9]. But currently there is no metric to capture the degradation in service due to these events.

Common approaches for ensuring network survivability in the presence of failures include protection and restoration mechanisms at the optical layer or the IP layer[10]. In addition, a number of new approaches have been proposed to account for failures in IP networks. These include the selection of ISIS/OSPF link weights in the presence of failures [11, 12] a deflection routing algorithm to alleviate link overloads due to failures [2], and failure insensitive routing [13].

[14] has performed pioneering work in proposing availability-based service differentiation for IP networks. A precise and standardized definition of service availability will help in realizing such service differentiation.

Recently, the media has become concerned about comparing the performance of various ISPs. [15] has conducted a comparative study of the performance of several IP backbones in the USA. The metrics considered were uptime, jitter and loss. For the computation of uptime, scheduled periods of maintenance (when the network may not be available) was first excluded. Then an outage period was defined as any interval of 10 seconds of more during which no packet was delivered. The availability of a path between two routers was computed on the basis of these outage periods. This definition has several shortcomings. First it is ad-hoc - it is debatable whether transient bursts of packet loss should be accounted for in the availability metric or the loss metric. Second, it provides an unfair advantage to ISPs that have longer maintenance periods. Finally, it does not address the notion of network-wide availability for an ISP.

III. Factors Impacting Service Availability

Service availability between two points in an IP network refers to the ability of the network to deliver IP packets from one point to another. Therefore it captures the ability of an IP network to perform its primary function, i.e., IP packet forwarding. Today’s ISPs offer their customers a guarantee of *port availability* as part of their SLAs. This simply represents the uptime of a single network element, i.e., the hardware by which the customer attaches to the ISP’s network. However, this does not, in any way, provide guarantees about the destinations that a customer can reach at a given point in time.

A key requirement for service availability between two points is the existence of physical connectivity between the points. We refer to this as *path availability*. A link failure that impacts path availability will also impact service availability. For example, if there is a single path between two edge points of an IP network, a link failure anywhere on that path will make the path unavailable and also disrupt packet delivery.

Not all link failures affect path availability. If there are several link-disjoint paths between two points, then path availability is not affected by multiple failures on a subset of these paths. Note however, that in such a case, the traffic will have to be carried on the available paths and may lead to packet loss on these paths due to congestion. In case that there is severe congestion, all (or most) packets may be dropped for short periods of time. This has been factored into network uptime in [15]. However, in this paper, we choose not to include this kind of transient outages in the notion of service availability. The impact of such outages will be captured by the loss and delay SLA metrics. In fact, accounting for congestion losses in the service availability metric would be a loss of information for network users and operators.

Even when a link failure does not affect path availability, it may still impact service availability. In the event of a link failure, every router in an IP network recomputes an alternate path (assuming that such a path is available) around the failed link to each destination. This is usually referred to as routing protocol convergence. While this computation is being performed, different routers may have inconsistent views of how to forward packets towards a given destination. This may result in the creation of a black-hole where packets reach a router and then get dropped due to lack of forwarding information for the packet.s destination. Alternatively, routing loops may be created where packets are sent back and forth among a set of routers until they are discarded due to TTL expiration. In either case, packet forwarding between a source and a destination is unavailable; hence service availability is affected.

Thus a definition of IP service availability must consider the following factors:

- network topology.
- mapping of IP links onto the underlying physical infrastructure.
- inter-dependence of IP network elements.
- failure characteristics of links/routers.
- routing protocol convergence time.

Network topology determines the number of alternate paths between two points, and whether they are link/router-disjoint. The mapping of IP links to the physical infrastructure determines the likelihood of simultaneous failures on these paths. For example, if two IP links may share an optical fiber conduit, they will fail simultaneously every time this conduit is cut. Therefore physical path diversity in IP to physical layer mapping[16] is a very important consideration for service availability. Inter-dependence of IP network elements may cause one failed element (e.g., a link or router interface card) to trigger the failure of another. This may arise from errors in the control protocol messages exchanged between routers, or from router software bugs. The frequency of failures determines how often traffic has to be re-routed, with accompanying forwarding disruptions due to routing protocol convergence. This has been traditionally captured in the notion of mean time-to-failure (MTTF) which is the time interval between the end of a failure and the start of the next. Finally router protocol convergence time determines the forwarding disruption associated with each failure.

Note that our focus is on a measure of service availability for IP networks. Therefore we do not take into consideration the availability of end-systems or services (such as DNS) in our proposed metric. For example, the "availability" of a web-based service, depends not only on IP service availability, but also on whether the web-server is up and has sufficient resources to respond to a client. However, such considerations are beyond the scope of this paper.

IV. Observations from the Sprint Backbone

In this section, we present observations from Sprint's IP backbone network that support of the need for an IP service availability metric, and validate the factors that need to be considered in its definition. In particular, we examine three factors that affect service availability - network design, failure characteristics and routing protocol convergence. We also discuss the effect of failures on VoIP traffic to emphasize the need for a service availability metric.

A. The Sprint Backbone

Sprint's IP backbone consists of a collection of Points-of-Presence connected via high-speed OC-48 and OC-192 links. The "logical" IP network is layered over a DWDM optical network. A PoP consists of a set of IP routers in a single physical location (usually a city or a metropolitan area). Each PoP connects customers ranging from large corporate networks to regional ISPs and data centers to the Sprint backbone.

The Sprint backbone has been provisioned and engineered to sustain packet forwarding even in the event of widespread failures. Multiple parallel links are provisioned between each PoP pair to ensure fault tolerance to the failure of one or more of these links. However, in some cases, parallel links between a given PoP pair share the same optical fiber, segment or optical equipment, and are therefore susceptible to simultaneous failures. The engineering rule for capacity provisioning is to maintain the average utilization of each link under 50%. The routers inside a PoP are connected in a highly meshed topology to guarantee connectivity even when multiple links/routers fail simultaneously. An added benefit of this low link utilization is that network congestion is extremely rare, and Sprint can easily adhere to the loss and delay guarantees that it provides in its SLAs to customers. The number of IP routes available in the Sprint network has been studied in [16, 17].

Sprint's network uses the link state protocol ISIS[18] for intra-domain routing. Every link in the network is assigned a weight and the cost of a path is measured as the sum of the weights of all links along a path. Each node independently computes its minimum cost path to every other node using Dijsktra's Shortest Path Forwarding (SPF) algorithm, and routes packets along this path. Link weights are chosen such that packets traverse paths with low propagation delays (link loads are a less important criterion, given current provisioning and engineering practices).

If there are multiple paths between a given node pair with the same minimum cost, then IS-IS splits traffic evenly among these paths. This capability of IS-IS is referred to as Equal Cost Multipath (ECMP), and is a key component in balancing traffic load among multiple parallel links between each PoP pair.

From the above discussion, we note that the Sprint network is designed to absorb the effect of failures without overloading links. However link overloads may occur when several links fail simultaneously[2]. In addition, service disruption may occur during routing convergence, while IS-IS recomputes a new minimum cost path around a failure.

Note that we use the Sprint network only as an illustration. We do not expect all networks to be designed in the same way. However, our discussion on service availability applies to all types of IP networks including community networks, wireless networks, etc.

B. Failure Characteristics

A detailed understanding of failure characteristics is essential for defining service availibility. However, very little is known about failure characteristics of operational networks, largely due to the lack of data. [1, 9] address this deficiency by analyzing IS-IS logs to derive failure information.

Figure 1 shows the distribution of failure events in the Sprint USA network between April and August 2002 on a daily, weekly and monthly basis. Failures are fairly well-spread out across weeks and months, and even over each day. Clearly, they need to be taken into account as part of everyday operations, and not just as extraordinary events.

The Sprint IP network is mapped onto a DWDM optical infrastructure with SONET framing. The network relies on IP layer restortion (via IS-IS) to recover from failures, with no protection at the optical level. Therefore link failures result from a variety of causes at or below the IP layer. At the optical

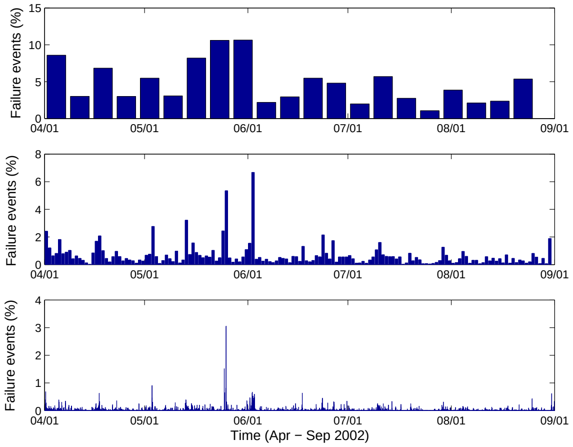

Fig. 1. Failure notifications over three time scales: weekly (top), daily (middle) and hourly (bottom)

layer, fiber cuts or optical equipment failure may led to loss of physical connectivity. At the IP level, problems such as hardware failures (e.g., router interface cards), router processor overloads and protocol implementation bugs can lead to link failures.

Understanding the causes of failures is central to deriving failure characteristics that are needed for defining service availability. In [1], we develop a cause-based classification of failures. The classifiction shows that 20% of the failures occur during scheduled maintenance windows that last only a few hours every week. Of the remaining unplanned failures, almost 30% are shared by multiple links. This can be attributed to the sharing of either routers or optical equipment or fiber by these links. The remaining 70% of unplanned failures affect one link at a time. This classification is used to derive important characteristics of each class such as failure duration, time between successive failure, and spatial distribution of failures by links and routers. These characteristics are needed to derive parameters based on which service availability is computed (as discussed further in Section V).

C. *IS-IS convergence Times*

Routing protocol convergence time as one the factors affecting service avaialibility (Section III), since forwarding state on routers may be inconsistent or unavailable during convergence. In [9] we analyzed IS-IS convergence times based on controlled link-failure experiments in the Sprint network. The convergence process can be broadly divided into three phases: (i) failure detection (ii) failure notification to the control plance and (iii) re-computation and re-establishment of forwarding informtion on routers.

We found that packet forwarding disruption due to IS-IS convergence can be as high as 6-8 seconds. This duration is dominated by a number of protocol timers, introduced to reduce the risk of network instabilities. Fine-tuning these timers allowed Sprint to bring down this convergence time to about 1 second. However, further improvements are needed to router architectures and computation of forwarding information to further lower convergence times.

D. *Effect of Outage on VoIP*

So far, the discussion in this section has focused on some of the factors that affect service availability. Let us now consider why service availability metric is needed for an emerging application - Voice-over-IP (VoIP). [4] conducted measurements on the Sprint network to determine the quality of voice calls that the network can support. Results show that voice packets experience almost no queuing, and the component that dominates end-to-end latencies is the propagation time over optical fiber. In the absence of failures, voice calls had toll quality, measured using subjective quality metrics.

However, there was degradation in call quality during failures when no packets were delivered for several seconds at a time. We contend that this happened because the network does not offer any gurantees on the maximum duration of any single outage period due to failures. The service availability metric should guarantee this, as discussed further in Section V.

V. Towards a Definition of Service Availability

In this section, we discuss some key issues in defining a service availability metric. We begin with a strawman definition of service availability between a single source-destination pair. Let there be P paths between the source-destination pair, numbered $1, 2, \cdots, P$, where $P \geq 1$. Assume that traffic between the source and the destination is split equally among the P paths, i.e., each path carries a fraction $1/P$ of the traffic. However, each path has the capacity to absorb all the traffic. Let E_k denote an event where any k of the P paths fail simultaneously, where $k = 1, 2, \cdots, P$. Let t_k be the mean time between successive events of type E_k, assuming that it is the same for every subset of k paths. A path fails due to the failure of one or more links on that path[1].

When all paths fail simultaneously, i.e., an event of type E_P takes place, no path is available between the source and the destination. So packet forwarding is restored only when the failure event terminates. Let O_P denote the average duration of failure events of type E_P. For all other types of failure events, traffic fails over to the available paths. However packet forwarding is disrupted on the failing paths during protocol convergence. For an event of type E_k, $k = 1, 2, \cdots, P-1$, the fraction of source-destination traffic affected is k/P. Let D be the constant time[2] for which packet forwarding is disrupted during routing protocol convergence.

The expected number of events of type E_k over an interval T is T/t_k. The expected time of packet forwarding disruption over interval T due to all events of type E_k is $D * T/t_k$, for $k = 1, 2, \cdots, P-1$. The expected time of packet forwarding disruption due to all events of type E_P over interval T is $O * T/t_P$.

The total disruption due to all failure events, weighted by the fraction of traffic affected by each event, is then given by

$$\tau = D*T/K*[1/t_1+2/t_2+\cdots+(P-1)/t_{P-1}]+O_P*T/t_P$$

[1] Router failures are not considered separately since the failure of a router leads to the failure of all links attached to the router.

[2] In practice, this time is variable due to interaction between consecutive failures, timer settings on different routers, propagation time of IS-IS updates, etc.

We define service availability A as $A = 1 - \tau/T$, which leads to

$$A = 1 - [D/K * \sum_{k=1}^{P-1} k/t_k + O_P/t_P] \qquad (1)$$

The above scenario contains several simplifying assumptions. Nevertheless it illustrates how an availability metric can be derived based on the factors identified in Section III:

- The parameter D depends on a large number of routing protocol parameters, and is a variable value in practice. Deriving this parameter is a research problem in itself, and requires a deep understanding of routing convergence behaviour [9].
- The parameter P depends on the network topology.
- The mean time between successive events of a given type E_k depends on a number of factors including network topology design, IP to physical layer mapping, and the sharing of routers/links along these paths. Even when the IP to physical layer mapping and the IP-level topology are known, it is challenging to derive these values, since it requires a precise understanding of the causes behind link failures [1]. Moreover, all paths (or sets of paths) may not have the same failure characteristics [1], which may further complicate the derivation of A.
- The parameter O_P depends on network design and failure characteristics. For example, if the paths share an optical conduit (i.e., a single point of failure), then a fiber cut can result in several hours of disruption in packet forwarding.
- Traffic may be split unequally among the P paths, in which case it is not enough to consider the failure of any subset of k paths out of P in the derivation of A. Instead, we would have to consider separately the simultaneous failure of every subset of the P paths, and the cumulative fraction of traffic on those failing paths. Again, only a subset of the paths (primary paths) may be used to carry traffic in the absence of failure, whereas the rest (backup paths) will be used only in the event of primary path failures. This will have to be considered in deriving A.

There are a number of other challenges in the definition of service availability:

- Application requirements need to be taken into account. For example, certain applications such as VoIP or distributed games may not be able to tolerate an outage longer than a few seconds at a time, whereas traditional applications such as ftp may be more tolerant. Therefore the service availability expression derived in Equation (1) may have to be supplemented by guarantees on the maximum duration of any single outage, e.g., 99% of outages will last no longer than 2 seconds.
- Any definition must allow for direct comparison among different ISPs (or IP networks). For example, one can imagine "service availability-based peering" where a customer selects from one among multiple upstream providers for Internet connectivity based on the service availability offered by these providers. This property will also open a set of interesting research issues. For example, different topologies may be compared with respect to the service availability metric in order to determine the best way of inter-connecting a set of nodes.
- The definition of service availability may have to be extended from a single source-destination pair to many source-destination pairs or for an entire networks. For example, an ISP may want to offer a service availability guarantee to a VPN customer with sites in multiple cities. Consider for example a customer that connects to the rest of the Internet via an ISP's network. Service availability of the customer should to be based on how often the customer can send a packet out of the ISP's network towards each Internet destination. Different subsets of Internet destinations will have different exit points from the ISP's network. The service availability of the customer therefore has to be calculated as a function of the service availability to each of these exit points.
- A customer must have the ability to measure service availability in order to verify whether the ISP is meeting its SLA guarantees. The customer is not likely to have access to the information used by the ISP to compute service availability and may not wish to perform very complex computations. Furthermore, the customer may be more interested in measuring service availability based on observed network characteristics such as sudden bursts of consecutive packet losses, a sequence of packets with very low TTL values, etc. Bridging this gap between definition (how an ISP defines service availability and designs a network to guarantee it) and measurement (how a customer verifies the ISP's compliance with offered SLA) is one of the most important and difficult problems in defining and using a service availability metric.

VI. Summary

The goal of this position paper is to advocate the need for a new SLA metric for IP networks - service availability. Loosely speaking, this represents the fraction of time that the network is able to deliver IP packets to destinations. We have identified various factors impacting service availability, and used observations from the Sprint network to illustrate our discussion. We have also outlined the key challenges in defining, using and measuring service availability. In recognition of the difficulty and complexity of defining service availability accurately, we have only provided a strawman definition to illustrate the open problems. As a step towards defining service availability, we are investigating various factors affecting availability such as the frequency and duration of link failures, the occurrence of simultaneous failures and routing protocol convergence times.

From an operations perspective, this paper aims to foster discussions among network operators and users about how to standardize this metric, and how to use it. From a research perspective, this paper proposes a new direction for QoS in IP networks, and has the potential to influence future research on topics such as network design, network management, protection and restoration techniques, router architecture and routing protocol design.

REFERENCES

[1] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C.-N. Chuah, and C. Diot, "Characterization of failures in an IP backbone," Sprint ATL, Tech. Rep. RR03-ATL-070100, July 2003.

[2] S. Iyer, S. Bhattacharyya, N. Taft, and C. Diot, "An approach to alleviate link overload as observed on an IP backbone," in *Proceedings of IEEE Infocom*, Mar. 2003.

[3] "Sprint IP Monitoring Project (IPMON) web-site," http://ipmon.sprint.com.

[4] C. Boutremans, G. Iannaccone, and C. Diot, "Impact of link failures on VoIP performance," in *Proceedings of NOSSDAV*, May 2002.

[5] A. Sridharan, S. Moon, and C. Diot, "On the causes of routing loops," in *Proceedings of ACM Sigcomm Internet Measurement Conference*, 2003.

[6] J. W. Stewart, *BGP4: Inter-domain Routing in the Internet*. Addison-Wesley, 1998.

[7] R. Kuhn, "Sources of failure in the public switched telephone network," *IEEE Computer*, vol. 30, no. 4, Apr. 1997.

[8] A. P. Snow, "A Survivability Metric for Telecommunications: Insights and Shortcomings," http://www.cert.org/research/isw/isw98/all_the_papers/no32.html.

[9] G. Iannaccone, C.-N. Chuah, R. Mortier, S. Bhattacharyya, and C. Diot, "Analysis of link failures in an IP backbone," in *Proceedings of ACM Sigcomm Internet Measurement Workshop*, Nov. 2002.

[10] A. Fumagalli and L. Valcarenghi, "IP restoration versus WDM protection: Is there an optimal choice?" *IEEE Network Magazine*, vol. 14, no. 6, pp. 34–41, Nov. 2000.

[11] A. Nucci, B. Schroeder, S. Bhattacharyya, N. Taft, and C. Diot, "IS-IS link weight assignment for transient link failures," in *Proceedings of the International Teletraffic Conference*, Sept. 2003.

[12] B. Fortz and M. Thorup, "Optimizing OSPF/IS-IS weights in a changing world," *IEEE Journal on Selected Areas in Communications*, Feb. 2002.

[13] S. Nelakuditi, S. Lee, Y. Yu, and Z.-L. Zhang, "Failure insensitive routing for ensuring service availability," in *Proceedings of International Workshop on Quality of Service (QoS)*, June 2003.

[14] M. Duvry, C. Diot, N. Taft, and P. Thiran, "Network availability based service differentiation," in *Proceedings of International Workshop on Quality of Service (QoS)*, June 2003.

[15] D. Newman, "ISP backbones stand up in grueling 30-day performance test," Network World Global Test Alliance, Network World Fusion. http://www.nwfusion.com., Dec. 2002.

[16] F. Giroire, A. Nucci, N. Taft, and C. Diot, "Increasing the robustness of IP backbones in the absence of optical level protection," in *Proceedings of IEEE Infocom*, Apr. 2003.

[17] R. Teixeiera, K. Marzullo, S. Savage, and G. Voelker, "Characterizing and measuring path diversity in internet topologies," in *Proceedings of ACM Sigmetrics*, June 2003.

[18] D. Oran, "OSI IS-IS intra-domain routing protocol," RFC 1142, Feb. 1990.

[19] M. Dahlin, B. Chandra, L. Gao, and A. Nayate, "End-to-end WAN service availability," *IEEE/ACM Transactions on Networking*, vol. 11, no. 2, Apr. 2003.

[20] G. Iannaccone, C.-N. Chuah, S. Bhattacharyya, and C. Diot, "Feasibility of IP restoration in a tier-1 backbone," Sprint ATL, Tech. Rep. TR03-ATL-030666, Mar. 2003.