

Service Availability: A New Approach to Characterize IP Backbone Topologies

Ram Keralapura and Chen-Nee Chuah
Univ. of California, Davis
{rkerlapura,chuah}@ucdavis.edu

Gianluca Iannaccone
Intel Labs., Cambridge, UK
gianluca.iannaccone@intel.com

Supratik Bhattacharyya
Sprint ATL, CA
supratik@sprintlabs.com

Abstract—Traditional SLAs, defined by average delay or packet loss, often camouflage the instantaneous performance perceived by end-users. We define a set of metrics for *service availability* to quantify the performance of IP backbone networks and capture the impact of routing dynamics on packet forwarding. Given a network topology and its link weights, we propose a novel technique to compute the associated service availability by taking into account transient routing dynamics and operational conditions, such as BGP table size and traffic distributions.

Even though there are numerous models for characterizing topologies, none of them provide insights on the expected performance perceived by end customers. Our simulations show that the amount of service disruption experienced by *similar* networks (i.e., with similar intrinsic properties such as average out-degree or network diameter) could be significantly different, making it imperative to use new metrics for characterizing networks. In the second part of the paper, we derive *goodness factors* based on service availability viewed from three perspectives: ingress node (from one node to many destinations), link (traffic traversing a link), and network-wide (across all source-destination pairs). We show how goodness factors can be used in various applications and describe our numerical results.

I. INTRODUCTION

Service-Level Agreements (SLAs) offered by today's Internet Service Providers (ISPs) are based on four metrics: end-to-end delay, packet loss, data delivery rate and port availability. The first three metrics are usually computed network-wide and averaged over a relatively long period of time. For the fourth 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.

The main problem with the existing SLA specifications is that they do not capture the effect of instantaneous network conditions like failures and congestions. A recent study [1] shows that failures occur on a daily basis due to a variety of reasons (e.g., fiber cut, router hardware/software failures, and human errors) and can impact the quality-of-service (QoS) delivered to customers. When a link/node fails, all routers will independently compute a new path around the failure. At that time, routers may lack or have inconsistent forwarding

information, resulting in packet drops or transient routing loops [2][3]. However, not all failures impact the network equally. The failure of a critical backbone link carrying heavy traffic may be more detrimental than the failure of an access link connecting a single customer. Yet, these service degradations are camouflaged by the *average* parameters reported in current SLAs. Therefore, to *measure* IP network performance, it is essential to consider the network routing configuration and traffic pattern during link or node failures.

Reports from various tier-1 ISPs suggest that IP backbone networks are usually over-provisioned where the link utilization of backbone links is less than 50% of their total capacity [4][5]. The reports also confirm that congestion due to link or router overload is a very rare event in backbone networks. During a link failure event traffic on the failed link is rerouted and may congest links along alternate paths. However, such congestions are usually not significant for a single failure event. Heavy congestions may occur when there are multiple failures, but such events are relatively rare. Hence in our current work we ignore the effect of congestions on network performance and only consider failures.

In this paper, we define a set of metrics for *service availability* of IP backbone networks that capture the impact of routing dynamics on packet forwarding. Instead of relying on active or passive measurements, we propose a methodology to *estimate* the service availability of a network in the presence of independent link failures. Specifically, given a topology (nodes, links and connectivity) and routing information (link weights, link delays and BGP peering points), we are able to compute the potential impact on service due to link failures. To achieve this, we carefully model the factors identified in the measurement based study by Iannaccone et al [6] that contribute to routing convergence. Convergence refers to the amount of time it takes for traffic forwarding to resume correctly on the backup path after a link failure. We wish to point out here that we focus mainly on single link failures for IP backbone networks since these are the dominant class of failures (i.e. over 70% of all failures in IP backbone networks) as observed by Markopoulou et al [1].

We use the novel concept of service availability to evaluate known topologies, such as full-mesh, ring and Tier-1 ISP backbones. Our simulations show that the performance of a network not only depends on routing dynamics, but also on various other factors like IGP link weight assignment and BGP prefix distribution. This brings out a necessity to identify new

⁰This research project was supported by National Science Foundation CAREER grant 238348 and University of California Micro with matching funds from Fujitsu and Sprint

metrics that customers can use to differentiate networks. There have been many attempts to characterize the Internet topologies or to model their graph-theoretic properties [7][8][9]. The resulting models are useful for re-generating topologies that best model real networks, such as GT-ITM [10] and BRITE [11], but they do not provide any insights on the QoS that a particular network can provide.

Using the concept of service availability, we derive *goodness factors* based on three different perspectives: ingress node (from one node to many destinations), link (traffic traversing a link), and network-wide (across all source-destination pairs). The goodness factors reveal how topologies with similar intrinsic graph properties, such as average out-degree or network diameter, do not necessarily offer the same level of service availability.

Finally, we describe several applications for the goodness factors in network planning and provisioning. For example, goodness from an ingress node perspective allows customers to choose the best place to connect to a network (or to choose among different providers), while link-based goodness helps an ISP to identify the set of critical links to be upgraded.

The rest of the paper is organized as follows. Section II identifies the importance of routing dynamics in estimating the end-to-end performance of a network and motivates this work. In Section III, we describe the proposed metrics and introduce the concept of service availability to characterize network topologies. We also define a set of network goodness factors. We discuss our numerical results in Section IV. We show the various applications of the goodness factors in Section V and conclude our paper in Section VI.

II. ANALYZING IMPACT OF ROUTING DYNAMICS

Intra-domain routing protocols, such as IS-IS [12] and OSPF [13], define how each node in the network responds to changes in the topology. Such protocols are also known as *link state protocols* where each node has complete knowledge of the network topology including all the links present in the network.

Upon detection of a link/node failure or a configuration change in the network, each node is responsible for disseminating the new topology description to all its neighboring nodes and recomputing the forwarding information in its own routing table. From the time of the failure or configuration change to the time all nodes have been informed of the change and have updated their routing tables, traffic disruptions (like packet drops and routing loops) are possible as the nodes may have an inconsistent view of the network.

We define the “convergence time” (CT) of a node due to a failure event in the network as the time taken by the node to update its routing and forwarding information in response to the failure. A node that does not have to update its routing or forwarding information has a convergence time of zero. The convergence time for any node n (that has to update its routing or forwarding information) due to a failure can be summarized as a combination of 3 components:

- *Detection time*: This is the time taken by the adjacent nodes to detect the failure. Today’s IP routers provide several mechanisms to perform this function [14], but all of them are based only on local information exchanged between neighboring nodes. For example, *hello* messages at the IP layer or alarms at the optical layer. Hence detection time represents a fixed price that is independent of the network topology or configuration.
- *Notification time*: This represents the time taken by the routing update to propagate through the network to reach n . In link state protocols, messages are flooded throughout the network. Therefore, the notification time strongly depends on the hop distance between n and the adjacent nodes. Each node along the forwarding route needs to process the message update before forwarding it, thus introducing a delay in the propagation of information.
- *Route Computation and Update time*: This is the time spent by node n to compute the new shortest path routing tree (that incorporates the failure information) and then update its forwarding information based on the new routing tree. The update procedure involves applying the changes to all the network prefixes that have been learnt via the BGP inter-domain protocol. The result of this computation is a forwarding table where each prefix is associated with a neighboring node as the *next hop*. The route computation and update time at any node heavily depends on the number of prefixes for which the next hop information needs to be changed. In turn, the number of prefixes to be updated not only depends on the location of the failure, but also on the distribution of prefixes to the next hops in the forwarding table. Indeed, the closer the failure occurs to a node, the larger will be the number of prefixes affected. Similarly, if a large number of prefixes share the same next hop node, a change in the topology close to that node will result in long route update time for all nodes in the network.

Based on the per-node convergence time (CT) we define the “network convergence time” as the maximum value of CT among all the nodes in the network. This indicates the time at which all the nodes in the network have learned about the failure, updated their routing tables and have a consistent view of the network.

Service availability of a network depends on the convergence time. Providers attempt to increase the overall availability of their networks by reducing convergence time in order to speed up the recovery after a failure. As we described above convergence time depends on (i) router technology used for failure detection, (ii) network topology, (iii) routing protocol configuration such as IGP weights and timers, and (iv) location of peering points with other networks that determines the distribution of network prefixes among egress nodes. Clearly, it is not possible to look at a subset of the above mentioned factors to derive the service availability of a network.

Consider the network illustrated in Figure 1. The number on each link indicates the IGP link weight. Let node A be the traffic source and node D be the traffic sink. Consider the




Fig. 1. Example of convergence time and service disruption

failure of link E – D. We assume that the nodes and the links have similar characteristics with a *detection time* of 500ms, a *notification time* between neighboring nodes of 100ms, and a *node route computation and update time* of 400ms (except for node C for which we assume this to be 100ms due to the fact that it does not have to change its route to reach D after the failure of link E – D). Consider the disruption observed for the traffic sent from node A to node D following the failure event at time $t_0 = 0$. Table I shows the routing events, changes in the traffic forwarding path and service availability from node A to node D. In this example, we assume that a node notifies neighboring nodes only after it has completed the update of its own routing table. However, it is easy to verify that if the updates are sent before updating the routing table, the example yields similar results.

Interestingly, as the message update propagates across the network, the forwarding path from node A to node D changes four times. Some of these intermediate paths are valid thus restoring service between A and D, while some are not, causing packet drops (or *traffic black-hole*) and routing loops. For example, packets are dropped until node E has computed a new forwarding path to reach D and routing loops occur when nodes B and F have conflicting forwarding information.

Based on the example we can derive some initial observations:

- The network convergence time provides a rough upper bound on service disruption time (i.e. the time for which service is not available). The service is not available when the packets cannot reach their destinations due to the lack of forwarding information. Given that traffic forwarding may resume even if all the nodes in the network have not updated their routing table, there may be a significant lack of correlation between network convergence and service availability. For example, in our illustration above, the cumulative time for which the service is not available amounts to 1.1s while the network takes 1.9s to converge. Therefore, network convergence time does not capture the entire routing dynamics.
- The network topology by itself does not help in understanding the service availability of a network. Additional information such as the location of peering points and size of routing tables needs to be considered in characterizing service availability. The time taken to update the routing table depends on the number of prefixes that

are affected. Table II shows the impact of varying the number of prefixes that have D as the next hop node on convergence time and service availability. Even with the same topology and identical failure scenario, increasing route update time may lead to significant differences in service availability.

In the following section, we introduce a new topology metric that exploits the knowledge of routing dynamics to define and compare the *goodness* of topologies. It is based on the algorithm to compute the cumulative time for which service is not available (i.e. the service disruption time as presented in Table I). The details of this algorithm are presented in Table III.

III. SERVICE AVAILABILITY TO CAPTURE NETWORK ROUTING DYNAMICS

A. Service Availability

We define service availability in three perspectives:

- *Ingress node perspective*: This is a measure of the network performance as seen by a particular ingress node where the traffic enters the network. It provides an insight about the level of service to expect when a customer connects to different ingress nodes of a network. It also helps an ISP to ensure that it can meet SLA specifications for customers connecting to different ingress nodes.
- *Link perspective*: The performance of a network should not heavily depend on the reliability of a few links in the network. In other words, the network should not have critical links whose failure results in serious performance degradation. Service availability from a link perspective evaluates the importance of various links for network performance.
- *Network perspective*: This measures the performance of the entire network from the perspective of an ISP. This is useful in designing a network to achieve high end-to-end performance.

We propose three metrics that capture routing dynamics in a network due to single link failures for analyzing service availability:

- *Service Disruption Time* (SD time): This represents the time for which service between a particular source-destination (O-D) pair is disrupted. From the point of view of an ingress node, it indicates the loss in connectivity with all/some parts of the Internet due to a link failure.
- *Traffic Disruption* (TD): This metric captures the total traffic disrupted between a particular source and destination node due to single link failures. It is the product of traffic rate between the O-D pair along the failed link and the service disruption time. For an ISP, TD is more important than SD time because of the fact that customers are compensated for the amount of traffic lost, irrespective of the duration for which the service is disrupted.
- *Number of Delay Parameter Violations* (DV): In a well-designed network, the end-to-end delay along alternate

Time	Event	Forwarding Path (A – D)	Service from A to D	Notes
0s	Failure of link E – D	A-F-E-D	NO	
0.5s	D,E: failure detection	A-F-E-D	NO	
0.9s	D,E: route update	A-F-E-C-D	YES	Forwarding is restored
1.0s	C,F: notified of failure	A-F-E-C-D	YES	
1.1s	C: route update	A-F-E-C-D	YES	Path C to D is not affected
1.2s	B: notified of failure	A-F-E-C-D	YES	
1.4s	F: route update	A-F-B-F-...	NO	Routing loop B – F
1.5s	A: notified of failure	A-F-B-F-...	NO	Routing loop B – F
1.6s	B: route update	A-F-B-C-D	YES	
1.9s	A: route update	A-B-C-D	YES	Network Convergence

TABLE I

SUMMARY OF ROUTING EVENTS (WITH A ROUTE UPDATE TIME OF 400MS). NETWORK CONVERGENCE TIME = 1.9S; SERVICE DISRUPTION TIME = 1.1S

Routing update time	100ms	200ms	300ms	400ms	500ms	1000ms
Convergence time	1.0s	1.3s	1.6s	1.9s	2.2s	3.7s
Time - service not available	0.8s	0.9s	1.0s	1.1s	1.2s	1.7s

TABLE II

CONVERGENCE TIME DEPENDS ON ROUTING UPDATE TIME IN NODES

STEP 1: Initialize the *Service Disruption Time*, $\phi_l(x, y)$, for the path from x to y due to the failure of the link l to 0 i.e. $\phi_l(x, y) = 0$. If the original path from x to y does not contain link l then QUIT.

STEP 2: Find the convergence time (CT) for each node and list the nodes in the increasing order of convergence time. Let the convergence time of the first node in the list be CT^1 , second node be CT^2 and so on. In general, the convergence time for the n^{th} node in the list is CT^n . Set the *current node*, k , (which is the node number on the sorted list) to 0.

STEP 3: Increment k to 1. Set $\phi_l(x, y) = CT^1$. At the time instant CT^1 after the failure event, find the path that a packet from the source node, x , follows to reach the destination node, y , taking into account that the first node in the list has converged and others have not. If the path has a routing loop or black-hole, then set $previousDisruption = true$ else set $previousDisruption = false$

STEP 4: Increment k by 1. At the time instant CT^k after the failure event, find the path that a packet from the source node, x , follows to reach the destination node, y , taking into account that the intermediate nodes might have converged or not.

STEP 5: If the path that the packet follows does not contain the failed link l and has no routing loop, then set $previousDisruption = false$. Go to Step 7. Else go to Step 6.

STEP 6: If the path that the packet follows contains the failed link l or has a routing loop, then the path from x to y is still disrupted. If $previousDisruption = false$, then do not update the *Service Disruption Time* but set $previousDisruption = true$ else if $previousDisruption = true$, then update the *Service Disruption Time* in the k^{th} iteration as, $\phi_l(x, y) = \phi_l(x, y) + CT^k - CT^{k-1}$

STEP 7: If there are more nodes in the list then go to STEP 4. Else QUIT.

TABLE III

ALGORITHM TO CALCULATE SERVICE DISRUPTION TIME FROM NODE x TO NODE y DUE TO A SINGLE LINK FAILURE

paths is generally higher than the original path. We assume here that the delay parameter is representative of the maximum end-to-end delay that can be tolerated by delay sensitive traffic in the network. If the end-to-end delay along the alternate path exceeds the delay parameter specification, then there is a delay parameter violation. DV measures the number of such delay parameter violations from the perspective of an ingress node due to single link failures.

Note that SD time and TD capture the effect of a link failure during service disruption, while DV is a post-convergence effect. However, given complete network topology specifications

(i.e, nodes, links, connectivity, link weights, link delays and delay parameter), BGP prefix distribution and traffic rate in the network, all the metrics can be pre-computed and used to characterize the end-to-end performance of a network.

B. Goodness Factors

To use the notion of service availability in characterizing network topologies, it is necessary to capture it using a quantitatively measure. Intuitively, this measure should yield a numerical value that can estimate the end-to-end performance of a network and help in differentiating various topologies. In order to accomplish this, we define a set of *goodness factors*

based on different perspectives of service availability. These factors can be defined differently depending on the specific scenario (for example, it could be driven by the cost involved, SLA specifications etc.).

We first introduce the notations that we use in the rest of this section. Consider a network with N nodes and M links. Let Γ and Λ represent the set of nodes and links respectively. For any node, i , there are $(N - 1)$ different destinations in the network and hence $(N - 1)$ different paths with node i as the ingress node. For the failure of link j , all/some/none of $(N - 1)$ paths could be affected. Let Q_{ihj} and T_{ihj} denote the SD time and TD for the source-destination pair $i - h$, due to the failure of link j . Similarly, let S_{ij} denote the number of delay parameter violations from ingress node i along all $(N - 1)$ paths due to the failure of link j .

1) *Goodness from ingress node perspective:* Typically, many customers are connected to a network at an ingress node. TD represents the total traffic affected for all the customers connected to the ingress node due to a link failure and does not provide valuable information to individual customers. Instead we use SD time and DV to measure the goodness of a network from an ingress node perspective. We define,

$$GI_i = f(Q_i, S_i) \quad (1)$$

where GI_i is the goodness factor of the network from the perspective of ingress node i . Q_i is the average SD time for node i across all $(N - 1)$ paths and M possible single link failures:

$$Q_i = \frac{1}{M(N-1)} \sum_{\forall h \in \Gamma, h \neq i, \forall j \in \Lambda} Q_{ihj} \quad (2)$$

Similarly, S_i is the average of the number of delay parameter violations from node i , to all other nodes in the network due to various single link failures:

$$S_i = \frac{1}{M} \sum_{\forall j \in \Lambda} S_{ij} \quad (3)$$

The function f depends on SLA specifications between ISP and customers. For simplicity, we assume that goodness is inversely proportional to various metrics. Hence,

$$GI_i = \frac{C}{(Q_i)^q (S_i)^p} \quad (4)$$

where C is a constant. The exponents q and p are SLA dependent. In our simulations we assume the exponent values to be 1.

2) *Goodness from link perspective:* The impact of a link failure on the network performance directly depends on total TD and total DV (i.e. sum of DV for all nodes) due to the failure. High values of these metrics for a link implies that the link is critical to network performance. We define the goodness factor from a link perspective based on similar assumptions on f as in Equation 4:

$$GL_j = \frac{C}{(T_j)^t (S_j)^p} \quad (5)$$

where GL_j is the goodness from the perspective of link j . T_j is the total TD due to the failure of link j :

$$T_j = \sum_{\forall i, h \in \Gamma, i \neq h} T_{ihj} \quad (6)$$

Similarly, S_j is the total number of delay parameter violations due to the failure of link j i.e.,

$$S_j = \sum_{\forall i \in \Gamma} S_{ij} \quad (7)$$

To find critical links in a network it is more useful to compute the *badness* of a link rather than its goodness. Hence we define the badness of link j in the network as the inverse of it's goodness:

$$BL_j = C(T_j)^t (S_j)^p \quad (8)$$

3) *Goodness from network perspective:* We define the goodness of the entire network as the sum of link goodness factors for various links in the network:

$$GN = \sum_{\forall j \in \Lambda} \frac{C}{(T_j)^t (S_j)^p} \quad (9)$$

IV. RESULTS AND DISCUSSION

In this section, we use the metrics proposed in Section III to quantify the impact of link failures on service availability of a network. First, we illustrate how the algorithm presented in Section II is used to compute SD time, TD time and DV for different classes of network topologies. Then, we examine how traditional graph properties such as out-degree, network diameter, increase in tree depth, and disconnecting sets correlate to QoS offered by the network. We will show the effectiveness of goodness factors in differentiating various network topologies by capturing the routing dynamics that affect traffic forwarding performance. Lastly, we will discuss the implication of the graphs and how goodness factors can be used in evaluating "quality" of connectivity and network design applications.

A. Simulation Setup

We built a java-based simulator to emulate intra-domain routing dynamics in the presence of link failures and to implement the algorithm presented in Table III. The inputs to the simulator are complete network topology specifications, BGP prefix distribution, and traffic load along different links in the network. In our simulations, each node in the network topology is mapped to a geographic location. The delays for individual links are then calculated based on the geographical distance between the nodes that the link connects. We categorize the nodes in a network as large, medium and small, depending on the amount of traffic they generate. We consider 20% of the nodes as large nodes, 30% as medium nodes and the rest as small nodes. The classification of nodes into large, medium and small nodes in a network is based on PoP (Point of Presence) level network topologies of various tier-1 ISPs. The traffic matrix for all the networks are generated using this model. We distribute BGP prefixes proportional to the

traffic between nodes. Large traffic flow from a source node to destination node implies that the source node reaches a large number of prefixes in the Internet through the destination node. Our results with different assumptions for prefix distributions yielded interesting results. We will discuss this further in Section V.

Based on these inputs, the simulator runs Dijkstra's SPF algorithm to find the shortest path from every node to all other nodes in the network. It then simulates single link failures and executes Dijkstra's SPF algorithm again, to find new paths in the network. The SD time for various O-D pairs are calculated based on the algorithm in Section II. TD is then calculated using SD time and network traffic distributions between different source-destination pairs. DV is determined using the link delay values and delay parameter specification. We assume that the maximum tolerable delay to be 100 ms in our simulations.




Fig. 2. ISP-A topology




Fig. 3. ISP-B topology

We consider the following two sets of topologies.

Set A: The first set of topologies represent standard topologies which includes ring, full-mesh and two PoP-level tier-1 ISP topologies (ISP-A and ISP-B in Figure 2 and Figure 3). We use these topologies to show that the metric values calculated from our simulations are intuitively correct. All the topologies considered in the first set have 20 nodes in their networks. ISP-A and ISP-B have 44 links each while the ring and mesh topologies have 20 and 190 links respectively.

Set B: For the second set of topologies, we consider 10 topologies which are "similar" in terms of their intrinsic properties like average out-degree and network diameter. The diameter refers to the maximum depth of all routing trees in the network. One of the topologies considered here is ISP-A (Figure 2) from Set A. The other 9 topologies were generated by changing link connectivity in ISP-A. Each topology was generated using the following procedure:

- Randomly delete x links from ISP-A topology. We consider $10 \leq x \leq 15$ in our simulations.
- Randomly add x links back into the topology while honoring the following rules: (i) The minimum out-degree of any node in the network is 2. (ii) The resulting topology is connected i.e., a spanning tree for the resulting topology has 19 links.

All the topologies considered in Set B have 20 nodes and 44 links with an average out-degree of 4.4 per node. We found that the network diameter for the topologies lie in the range of 4-6.

In all cases, we assign equal weights to all the links in a network, thus making it a minimum-hop routing scheme. In reality, different IGP link weight assignment schemes yield different metric values and we will further explore this in Section V. To achieve a fair comparison, we consider the same traffic distributions in all the topologies. Traffic rate to/from large, medium and small nodes remain the same in all the topologies. Finally, we want to point out that the values of goodness factors are normalized to 1 in all the results presented.

B. Service Availability for Known Topologies

In this section, we use the metrics proposed for service availability to study the performance of four known topologies (Set A). Intuitively, a mesh topology should perform the best among all the networks with the same number of nodes while a ring topology should perform the worst. Figure 4 shows the cumulative distribution (cdf) of the total TD in each of the four networks for various single link failure scenarios. Every link in the mesh topology carries traffic between a single O-D pair, thus resulting in small values of total TD for single link failures while every link in a ring topology carries traffic between multiple O-D pairs resulting in very high values of total TD. These two topologies are the extreme cases for any topology with the same number of nodes. TD values for other topologies, like ISP-A and ISP-B, lie in-between these extreme values.

Figure 5 shows the cdf of total DV for single link failures. There are no delay parameter violations for the mesh topology after single link failures while ring topology is significantly worse compared to other topologies for obvious reasons.

C. Goodness Factors vs. Static Graph Properties

The following case studies illustrate the limitations of static graph properties (like out-degree distributions or network diameter) to evaluate the performance of a network. This provides the motivation to characterize network graphs using




Fig. 4. CDF of Traffic Disruption due to single link failures




Fig. 5. CDF of Number of Delay Parameter Violations due to single link failures

goodness factors, which are directly derived based on estimated performance, i.e., service availability. All the following results are for topologies in *Set B*.

In the past, out-degree of a node has been used as a metric for characterizing a source node [15]. Figure 6 shows the maximum, minimum and average goodness factor values for various source nodes as a function of their out-degree. The goodness factors for source nodes with the same out-degree are considerably different. Also, higher out-degree does not always result in higher goodness factor values. The main problem in using out-degree to estimate the performance of a source node is that it does not capture the effect of various network characteristics such as BGP prefix distribution and link weight assignment scheme which have a significant impact on source node performance. To show this, we repeat the simulations by distributing BGP prefixes such that all the

prefixes have a single exit point in the network (we refer to this as *extreme prefix distribution*). None of the other topology properties were altered. The graph in Figure 7 shows that changing BGP prefix distribution in the network changes the source goodness factor for various nodes. This implies that out-degree is not an appropriate metric to measure the performance of source nodes.




Fig. 6. Goodness factor from an ingress node perspective Vs out-degree in various topologies with normal prefix distribution




Fig. 7. Goodness factor from an ingress node perspective Vs out-degree in various topologies with extreme prefix distribution

Network diameter gives an estimate of the maximum convergence time in the network. One would expect that a network with a small diameter would exhibit small convergence time and hence offer better service availability. The top graph of Figure 8 shows the network goodness factor for various topologies against network diameter. However our results show that topologies with smaller diameter do not always result in higher goodness factors. Also, topologies with same

diameter exhibit different network goodness factor values. Like out-degree, network diameter does not account for several characteristics of a network and hence is not a good metric for predicting its performance.

Another metric related to network diameter is the increase in tree depth due to single link failures. In a well-designed network, the end-to-end delay depends directly on the depth of the routing tree from source to destination. Hence DV depends on the increase in tree depth after a failure. Typically, lower values of increase in tree depth due to single link failures should result in better topologies. From our simulations we find that this is not true. The bottom graph in Figure 8 shows the average increase in tree depth for various topologies against the network goodness factors. Even though increase in tree depth can be used to roughly estimate DV, we find that it is not a good metric for predicting network goodness.

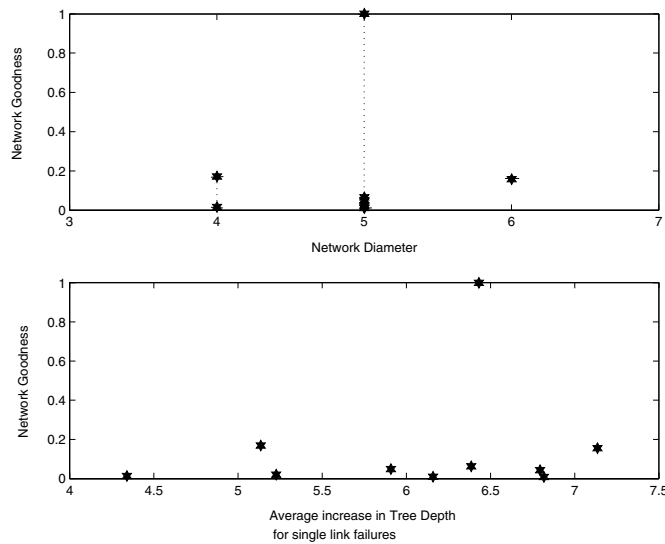


Fig. 8. Top graph shows the variation of network goodness factor with network diameter while the bottom graph shows the variation of network goodness with average increase in tree depth due to single link failures

Another important traditional metric used to estimate network performance is the number of *disconnecting sets* in a network [7]. Disconnecting set of a topology is defined as a set of links, whose cardinality is less than the minimum node out-degree in the network and whose removal disconnects the topology into two or more smaller topologies [16]. The topologies that we consider here, have nodes with out-degree of two. Hence, by the definition above, the cardinality of the disconnecting set should not be more than 1. Since this definition yields a null set for all the topologies, we define a disconnecting set, D , with the following modifications:

- D partitions a topology into 2 parts with each part having at least 2 nodes.
- Exclude Supersets: If D_1, D_2, \dots, D_n are n disconnecting sets for a network, then $D_i \not\subseteq D_j$ when $i \neq j$.

There are numerous solutions to the above definition of D , but we consider only those sets for which the cardinality

of D is less than 3, i.e., sets with at most 3 links removed. Disconnecting sets determine links that have a big impact on network performance when they fail. The failure of links in the disconnecting set could potentially result in high TD and DV. As the cardinality of the disconnecting set increases, the criticality of the links in the set decrease. For example, a single link connecting two sub-topologies is more critical than two links connecting two sub-topologies. The bottom graph of Figure 9 shows the number of disconnecting sets with different cardinalities for various topologies. The top graph in the same figure shows the network goodness factor for various topologies. Even though topology-4 and topology-8 have no disconnecting sets, we see that they do not result in the best network goodness factor values. Hence, similar to the other traditional metrics, disconnecting sets are not good metrics to capture service availability of a network.

The top graph in Figure 9 shows that even though the topologies are similar, the end-to-end performance based on service availability is significantly different. ISPs can take advantage of this network differentiation to design their networks to provide high service availability to customers. It is also helpful in estimating the cost of compensating the customers for SLA violations of the network.

V. APPLICATIONS OF GOODNESS FACTORS

This section explores the various applications of the proposed *goodness factors*, and present our initial numerical results using topologies in *Set B* as case studies.

A. Goodness Factors from Ingress-node and Link perspectives

Given an ISP network, the top graph in Figure 10 shows the performance that a customer can expect when connected to different ingress nodes. This helps a customer to choose an ideal location to get connected to the network. It also helps the ISP to ensure that it can meet the SLA specifications for customers at a given source node.

The bottom graph in Figure 10 shows badness factors of the links (inverse of goodness, Eq. 9) for one specific network topology (Topology-6). The graph shows that there are a handful of critical links (e.g., Link-4 and Link-14) that have a big impact on network performance when they fail. The rest of the link failures only result in minor service degradation. The ability of the badness factors to clearly distinguish the critical links in the network is extremely useful for capacity planning and traffic engineering purposes of an ISP. For example, this information allows the ISP to make an informed decision about bringing down a link for maintenance to minimize the impact on network performance. ISPs can also re-negotiate peering relationships to divert traffic away from critical links. With a better picture of how the failure of individual links can impact performance, the ISP can better estimate whether certain SLAs can be met.

B. Network-wide Goodness Factors

Our results from the previous sections have established the importance of characterizing topologies based on meaningful




Fig. 9. Top graph shows the network goodness factors while the bottom graph shows the disconnecting sets for different topologies




Fig. 10. Top graph shows the goodness factors from the perspective of various ingress nodes in topology-6 while the bottom graph shows the link badness factor for various links in topology-6

performance metrics such as service availability. The proposed *network-wide goodness* aims to capture the impact of routing dynamics on service degradation measured in terms of TD and DV across all source-destination pairs. This goodness metric forms the basis for optimizing network design decisions. We explore three such applications in the following discussion.

IGP Link Weight Assignment: IGP link weights determine the routing trees and hence can significantly influence the service availability of a network. We explore how goodness factors can be used to “evaluate” different link weight assignment schemes in Figure 11. We consider the same ten topologies in *Set B* but with three different link weight assignments: (i) weights proportional to link delays, (ii) equal weights, and (iii) random weights. Figure 11 shows that the

same topology can behave quite differently when different link weights are used. As part of our future work, we will explore how goodness factors can be used as an optimization metric for selecting link weights to provide the best service availability.

BGP peering points: As discussed in Section IV-C, the location of BGP peering points determine the prefix distribution size at different nodes, and hence also influences service availability. In Figure 12, we compare network goodness for topologies with three ways of distributing BGP prefixes across different nodes: (i) equal distribution, (ii) extreme distribution where all the BGP prefixes are located at one exit point, and (iii) typical prefix distribution observed in a tier-1 ISP. Depending on the network topology, different locations of BGP peering points result in different network goodness in terms of overall service availability. This illustration shows that goodness factors can be a useful metric in determining “ideal” BGP peering points that result in the most desired network performance.

Network upgrade: To cope with customer demands and meet SLAs, an ISP may have to schedule network upgrade to introduce a new node or link into its network. Deciding *where* in the existing network to connect this new node/link to, becomes a design challenge. In this case study, we consider adding a new node with 3 links into ISP-A network (Figure 2). Figure 13 shows five possible solutions and compares the resulting network-wide goodness of the new network. Solution-2 clearly shows the best network goodness in terms of the offered service availability in the presence of failures.




Fig. 11. Network Goodness for topologies with different link weight assignment scheme

VI. CONCLUSIONS AND FUTURE DIRECTIONS

In this paper, we examined the importance of incorporating network dynamics in characterizing topologies. Our simulations show that traditional metrics like out-degree, network diameter and disconnecting sets that disregard network dynamics do not effectively capture the performance of a network.

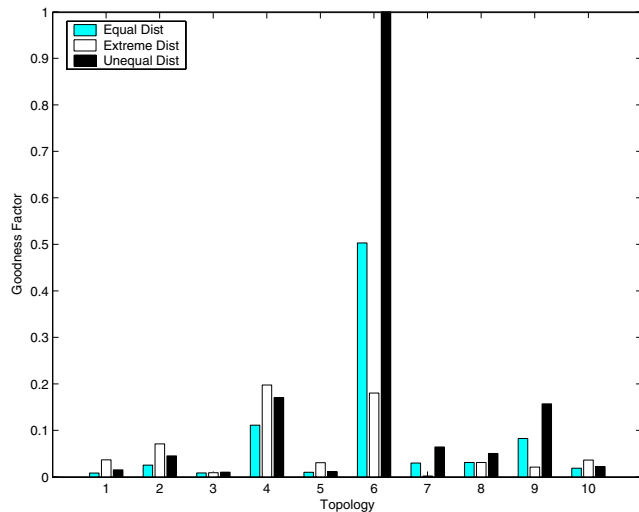


Fig. 12. Network Goodness for topologies with different prefix distribution schemes




Fig. 13. Adding a new node into the network - Network Goodness for various solutions

Hence it calls for a new approach to characterize topologies. To fill this void, we proposed a novel methodology based on the concept of service availability and demonstrated its effectiveness using simulations.

To the best of our knowledge, this is the first work to consider network dynamics in characterizing topologies. The approach is the first step in the right direction and is appealing to both ISPs and customers alike. We have identified numerous applications for goodness factors in capacity planning and network design, but their detailed analysis is a part of our future work. In addition, much work remains in extending this work to incorporate realistic failure patterns that includes multiple simultaneous link failures. As a complement to simulations we plan to validate our results through experiments on a test-bed or measurements in real world networks.

REFERENCES

- [1] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. N. Chuah, and C. Diot, "Characterization of Failures in an IP Backbone Network," in *Proceedings of IEEE Infocom*, 2004.
- [2] C. Boutermans, G. Iannaccone, and C. Diot, "Impact of Link Failures on VoIP Performance," in *Proceedings of NOSSDAV*, May 2002.
- [3] A. Sridharan, S. Moon, and C. Diot, "On the Causes of Routing Loops," in *Proceedings of ACM Sigcomm Internet Measurement Conference*, Oct. 2003.
- [4] C. Filsfils, "Deploying Tight-SLA services on an IP Backbone," June 2002. [Online]. Available: <http://www.nanog.org/mtg-0206/ppt/filsfils>
- [5] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and C. Diot, "Packet-Level Traffic Measurements from the Sprint IP Backbone," *IEEE Network*, Nov. 2003.
- [6] G. Iannaccone, C. 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.
- [7] P. Radoslavov, H. Tangmunarunkit, H. Yu, R. Govindan, S. Shenker, and D. Estrin, "On Characterizing Network Topologies and Analyzing Their Impact on Protocol Design," Tech. Rep. USC-CS-TR-00-731, Mar. 2000.
- [8] C. Faloutsos, P. Faloutsos, and M. Faloutsos, "On Power-Law Relationships of the Internet Topology," in *Proceedings of ACM Sigcomm*, Sept. 1999.
- [9] R. Govindan and H. Tangmunarunkit, "Heuristics for Internet Map Discovery," in *Proceedings of IEEE Infocom*, 2000.
- [10] K. Calvert, M. Doar, and E. W. Zegura, "Modelling Internet Topology," in *IEEE Communications Magazine*, 1997.
- [11] A. Medina, A. Lakhina, I. Matta, and J. Byers, "BRIT: An Approach to Universal Topology Generation," in *Proceedings of the International Workshop on Modeling, Analysis and Simulation of Computer and Telecommunications Systems (MASCOTS)*, Aug. 2001.
- [12] D. Oran, "OSI IS-IS Intra-domain Routing Protocol," RFC 1142, Feb. 1990.
- [13] J. Moy, "OSPF Version 2," RFC 2328, Apr. 1998.
- [14] G. Iannaccone, C.-N. Chuah, S. Bhattacharyya, and C. Diot, "Feasibility of IP Restoration in a Tier-1 Backbone," *IEEE Network*, Mar. 2004.
- [15] B. Waxman, "Routing of Multipoint Connections," *IEEE Journal on Selected Areas in Communication*, vol. 6, no. 9, pp. 1617-1622, Dec. 1988.
- [16] R. Ahuja, T. Magnanti, and J. Orlin, *Network Flows : Theory, Algorithms and Applications*. Prentice Hall, 1993.