Towards Quantification of IP Network Reliability

Hao Wang* Alex Gerber† Albert Greenberg§
Jia Wang† Yang Richard Yang*
†AT&T Labs §Microsoft Research *Yale University

1. Problem Motivations

There has been intense interest on the level of reliability and the techniques to improve the reliability of IP networks. As more mission-critical applications and network-based computing move to the Internet, the importance of IP network reliability will only increase.

However, so far the discussions on the level of reliability of IP networks are mainly informal or on relative scales. The estimation of IP network reliability varies over a wide range, from 99% to 99.99%. How this reliability level is calculated is unclear. Furthermore, although there is much research on traffic engineering (TE), the focus tends to be on network specific metrics such as maximum link utilization (MLU). The effectiveness of the proposed techniques on the reliability perceived by customers is unknown.

In this paper, we conduct an initial quantification on the level of IP network reliability. We ask two simple but important questions: 1) *what is the reliability level of IP networks?* In particular, how close are they to be five nines (*i.e.*, achieving 99.999% reliability)? 2) *how effective are the IP layer techniques in improving the reliability of IP networks?* In particular, do they make a difference in improving IP network reliability? The evaluations presented in this paper are our initial step in answering the preceding two questions.

2. Model and Methodology

Our methodology is to define IP network reliability using service level agreements and then evaluate the effects of various factors on IP network reliability.

Overview: The reliability level of an IP network depends on many factors. Figure 1 shows the factors we evaluate in this paper. We emphasize that there are other factors affecting reliability, and a comprehensive evaluation is beyond the scope of this poster. Below we briefly go over these factors.

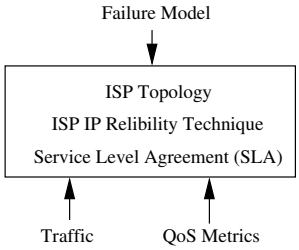

Figure 1. Reliability Factors.

IP network topology: We use the router-level topology of Abilene in our evaluation, which consists of 11 routers and 28 uni-directional links, with link bandwidths of 10 Gbps for all links except one at 2 Gbps.

Connectivity	Delay	Loss
100%	45 ms	0.3%

TABLE I
Summary of Internap SLA QoS targets.

Scenario	MTTF	MTTR	FRL
Low failure	7 days	1 hour	0.5, 1
High failure	1 day	10 minutes	0.2, 0.4, 0.6, 0.8, 1

TABLE II
Summary of failure scenarios.

Traffic demands: We use real traffic demand matrices of Abilene from October 4, 2006 to October 31, 2006, at 5-minute intervals. We assume a random portion (up to 50%) of the traffic is VPN traffic.

QoS metrics: We consider connectivity, end-to-end delay and end-to-end packet loss rate.

Service level agreement (SLA): We use the public SLA of Internap available on its web page (www.internap.com). It is summarized in Table I.

Failure model: The failure model captures failure characteristics in the IP network, and is used to generated synthetic failure events to stress the network. We construct our failure model based on the measurement study of Markopoulou *et al.* [1]. We focus on failures occuring within the backbone (i.e., edge links and routers are not captured in the current model). Specifically, we evaluate a high failure scenario and a low failure scenario. A scenario is characterized by three parameters: mean-time-to-failure (MTTF), mean-time-to-repair (MTTR), and failure ratio level (FRL). Table II summarizes the failure model parameters corresponding to these two failure scenarios. The failure ratio level parameter is a set of discrete failure ratios that indicate the fraction of lost capacity due to a link failure. The low failure scenario models failures that are less frequent, but have larger impact, and take a longer time to repair, such as scheduled maintenance and fiber cuts. The high failure scenario models failures that are more frequent, but have less impact, and take a shorter time to repair, such as aging network equipment with intermittent and recurring faults. For both scenarios, we choose the MTTF so that approximately the median value of the time between two failures of a link is similar to the corresponding measured value in a real IP backbone network [1].

IP network reliability techniques: In the past, most IP networks relied on link layer techniques such as SONET rings to protect against link failures. Later, due to the relatively high cost of SONET protection and the lower

cost and greater flexibility of IP, many IP networks have turned to the IP layer to handle failures. When a failure happens, IP routers detect the failure and trigger routing reconvergence. We refer to such techniques as *restoration* techniques.

A potential drawback of the restoration techniques is their relatively long transient convergence time, which may not be sufficient to meet the requirements of some mission-critical applications such as VPN networks carrying delay sensitive traffic. Thus, the restoration techniques are enhanced via MPLS to bypass the failed equipment before routing convergence. Such techniques are referred to as *protection* techniques. A key research focus of the protection techniques is computing the protection paths. A widely used approach is constrained shortest path (CSPF).

We evaluate the effects of the following representative IP layer restoration and protection techniques.

- Standard IGP re-convergence (IGP): This is shortest-path-first (SPF) IGP routing (such as IS-IS and OSPF) with restoration.
- IGP fast-rerouting (IGP-FRR): This is SPF IGP routing with CSPF-based protection.
- TE-based traffic engineering with fast re-routing (TE-FRR) [2]: This technique determines not only the routing of traffic when there is no failure, but also the backup paths for various failure scenarios.
- TE-based traffic engineering with SPF fast-rerouting (TE-SPF): This technique is the same as TE-FRR except that the backup paths are computed using the CSPF algorithm.

For all shortest path based routing, we use the real link metrics used in Abilenene.

Reliability metric: We define the reliability levels of an IP network as the percentage of time that it can satisfy its SLA. We measure reliability level for each origin-destination (O-D) pair and report aggregated statistics about the reliability level of OD pairs. Note that another metric is the reliability level from on origin node to all other nodes in the network, as typically an SLA is for a customer at a node to other nodes in the network.

3. Major Findings

Figures 2 and 3 show the reliability of Abilene under low and high failure scenarios respectively. The reliability metrics are evaluated on the time scale of a 4-week span from October 4, 2006 to October 31, 2006. For each O-D pair, we classify it according to the highest level of reliability (in terms of nines) that it can achieve. We make the following observations.

Decent yet insufficient reliability: For both failure scenarios, a large number of O-D pairs (over 70%) enjoy a reliability level of at least 99%. However, the network is still far from achieving 99.999% reliability. Under both failure scenarios, no more than 60% of the O-D pairs can reach 99.999% reliability, no matter what reliability technique is in use.

IP network reliability techniques help: IP layer reliability techniques can make a difference. IGP re-covergence

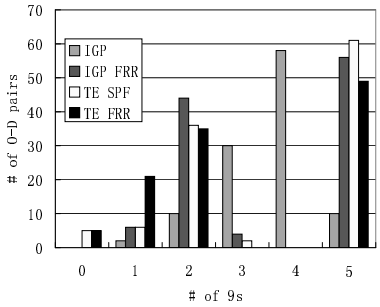

Figure 2. Reliability of Abilene under low failure.

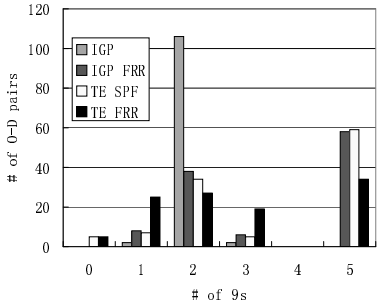

Figure 3. Reliability of Abilene under high failure.

(IGP) in the Abilene network suffers SLA violations for both failure scenarios. Less than 10% of O-D pairs maintain 99.999% reliability even under low failure; while no O-D pair reaches that level under high failure. The three protection techniques are significantly better at maintaining SLAs. Under both failure scenarios, between 30% and 55% of the O-D pairs can enjoy 99.999% reliability.

No single reliability technique prevails: As failures last longer, restoration (SPF) begins to gain performance benefits over the protection techniques. This can be seen by observing that under TE-SPF instead of TE-FRR, the number of O-D pairs achieving five nines is higher under both scenarios. An explanation is that as failures last longer, a better routing, although it takes time to compute, leads to better routing after the computation is done.

4. Conclusions

In this paper, we have conducted an initial, yet important evaluation on the level of reliability of IP networks. Using the Abilene network, we find that although its level of reliability is already pretty high, it still needs significant improvement to achieve higher levels of reliability such as five nines. Also, we find that different reliability techniques have different performance characteristics. An appropriate choice should consider both failure characteristics and the reliability goal. Note that, since the Aliene network is not considered as highly representative, it is not clear if the same observation would hold on commercial ISP networks. A more comprehensive evaluation on other IP networks and incorporating edge router reliability are left for future work.

References

[1] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. Chuah, and C. Diot. Characterization of failures in an IP backbone network. In *Proceedings of IEEE INFOCOM '04*, Hong Kong, China, Apr. 2004.

[2] H. Wang, H. Xie, L. Qiu, Y. R. Yang, Y. Zhang, and A. Greenberg. COPE: Traffic engineering in dynamic networks. In *Proceedings of ACM SIGCOMM*, Pisa, Italy, Sept. 2006.