Towards a Deeper Understanding of Managing Dynamic Optical Networks Under Link Failures

Sun-il Kim, Xiaolan J. Zhang and Steven S. Lumetta
CS and ECE Dept., Univ. of Illinois at Urbana-Champaign
Coordinated Science Lab., 1308 W. Main St., Urbana, IL 61801
email: {sunilkim,xzhang29,lumetta}@uiuc.edu

Abstract: We quantify the impact of link failures after initial rerouting and show that some nodes continue to drop a significant portion of its traffic. We discuss the practical tradeoffs in rebalancing the network using redimensioning.

OCIS codes: (060.4265) Fiber optics and optical communications/Networks, wavelength routing; (060.4257) Fiber optics and optical communications/Networks, survivability

1 Introduction

Failures in optical networks can result in loss of enormous amounts of data. For critical applications that require rapid service recovery, many protection schemes have been proposed. However, many applications can tolerate short term outages and can be exempt from the added expense of providing protection. These applications can then be classified under best-effort traffic, which is rerouted when there is a failure. This process can take seconds to minutes in contrast to protected traffic with rapid recovery (50 to 100ms). Differentiating traffic based on the level of survivability on IP/GMPLS networks has been studied in the past to take advantage of the cost saving opportunity the different types of traffic presents [1, 2].

In this paper, we show the impact of link failures on a dynamically routed network, beyond the initial failure management stage (where broken connections are rerouted). First, we start with the reroutability ratio of failed connections on a well dimensioned (low blocking with balanced capacity) network. We then quantify the longer term impact of link failures, which we show is critical to network performance. Our results show that the performance can be severely limited for traffic on certain nodes. Under failures, some nodes experience very high blocking at load ratios that allow below 1% to 5% blocking under normal operation.

Finally, we quantify the costs associated with restoring the network. We apply a redimensioning technique, which is a logical step in effectively re-balancing the network and addressing the issue of fairness among nodes. The tradeoff in utilizing such measures is clearly shown in terms of the typical as well the worst case additional capacity required at each link and node. The results clearly outline the limitations of going to a best-effort mode for non-critical applications and motivates further study in both traffic engineering techniques as well as measures for systematically over-provisioning the network to allow graceful degradation. The rest of the paper is organized as follows. In the next section, we discuss our simulation model and the technique used for network provisioning. The results from studying the impact of link failures and the costs associated with network redimensioning are provided in Section 3. Finally we conclude the paper with a short summary.

2 Network Model and Balanced Provisioning

Proper resource dimensioning improves network performance and offer a more stable service to systems and applications utilizing the network [3]. In [4], we showed that balanced provisioning is important and showed that we can significantly improve blocking performance if capacity is properly balanced rather than uniformly assigned to each link given the same amount of total capacity. We use the same technique to provision the network with an average of 1000 randomly chosen, fully loaded, projected traffic matrices from the same uniform distribution. The blocking probability is obtained by averaging 100 runs of randomly chosen traffic.

We focus on wavelength routed networks with full wavelength conversion capability at every node. Poisson processes are used to model call arrivals and hold times given that the burstiness of data traffic is usually well suppressed by huge amount of aggregate data on optical backbones. Each connection requests a uniform 1 unit capacity and has the same departure rate. The departure rate is determined by the network load (product of the projected departure rate and the load ratio). The arrival rates are uniformly distributed between 1 to 10 and the load ratio is varied at some fraction of the projected traffic load. The projected load is computed by using Equation 1. Each request demands the same capacity and hold time, and shortest path first routing is used. We use ARPANET for illustrate the ideas and the performance results covered in this

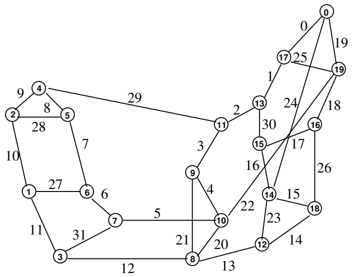

Fig. 1. ARPANET w/ link & node IDs.

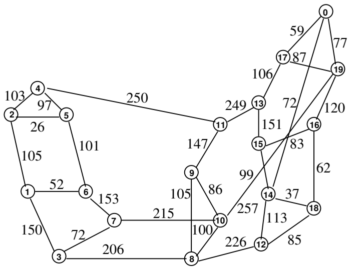

Fig. 2. Balanced capacity

SPF primary with	capacity	% overhead
no protection (reroute)	3851	0%
SPP N-hop backup	5924	53%
SPP shortest backup	6212	61%
DPP	9601	149%

Fig. 3. Provisioned capacity (for same load)

C_l	Current capacity of link $l \in E$.
R	Set of all connection request pairs. $R \subseteq N \times N$.
r_i	Requested capacity of a node request pair $i \in R$.
λ_i	Poisson arrival rate per request pair i.
μ_i	Poisson departure rate per request pair i. The average holding time is $1/\mu_i$.
len_i^{SPF}	topological shortest path length for pair i.

$$projected\ load = \frac{\mathsf{E}(\lambda)\mathsf{E}(r)\sum_i len_i^{SPF}}{\mathsf{E}(\mu)\sum_{l\in E} C_l} \qquad (1)$$

```
1: Set infinite available capacity on each link $C_l \leftarrow \infty\ \forall l \in E$
2: Set capacity to 0 for the link assumed to have failed.
3: while System has not reached steady state do
4:     Generate a new request $j$ from traffic matrix $T^{prj}$.
5:     Route $j$ by shortest path first algorithm (SPF).
6: end while
7: $\widetilde{C}_l$ is the actually used network capacity of each link.
8: Repeat from Line 1 to get the distribution of $\widetilde{C}_l$, mean $\mathsf{E}_l(\widetilde{C}_l)$
```

Algorithm 1: Redimensioning algorithm for link failure.

paper, but other well-known networks such as NJLATA and NSFNET were also evaluated (these results are left out for brevity given that they exhibit similar trends).

Figure 1 shows the link and node ID's for ARPANET. Figure 2 illustrates capacity provisioned on each link using our dimensioning algorithm. In Figure 3, we report the total capacity (computed by counting the total number of wavelength channels provisioned) for ARPANET with and without protection for the same projected load. Under the dynamic model, we computed the cost both arbitrarily long backup paths and shortest only backup paths for shared path protection (SPP) as well as the cost for dedicated path protection (DPP). The amount of overhead clearly supports the motivation for utilizing a reroutable class of traffic.

3 On Impact of Link Failures

Even though the class of traffic in question is considered non-critical, and thus has no protection assigned, it is important to guarantee stability in performance. For this traffic, rerouting is used to restore the broken connections. As no extra capacity is reserved, one can reasonably expect to see less than 100% recovery. The metric used to measure the immediate impact of a link failure is reroutability ratio. This measure is obtained by taking the ratio of the number of connections that can be successfully rerouted to the total number of connections affected by the failure. This idea is well understood and has been studied in the past. Figure 4 shows this reroutability under a well balanced network. Starting with half load (where all connections can be rerouted), most of the connections at load ratios below 0.75 can be rerouted. However, we next show that the real impact on performance occurs after the rerouting stage, before the failure can be physically repaired.

3.1 *Intermediate Term Impact*

Until the failure is physically repaired, the network is left operating under suboptimal conditions. Due to the failed link, many of the connections are also forced to use (topologically) non-shortest paths as well as limited number of path choices, which exacerbates the impact of the failures. Blocking probabilities of the network before and after a link failure at different load ratios are shown in Figure 5. The numbers for after failure results are averaged over all single link failure scenarios. The average blocking performance under a link failure may seem reasonable up to relatively high load ratios (for example, 0.77 where blocking is around 5%). However, when the blocking probability for each node is computed, we can see that there is a detrimental impact in terms of fairness. Figure 6(a) shows each network node's blocking probability under normal operation and after failure along with the worst case blocking for each node. At load ratio of 0.75, requests may be blocked over 30% of the time (load ratio of 0.75 was chosen for demonstration purposes because the post-failure blocking remains reasonable at slightly below 5%). At load ratio of 0.85 (around 3% blocking without failures as shown in Figure 5), a link failure causes some nodes to drop about a third of their traffic, with the rest of the network dropping around 13% on average.

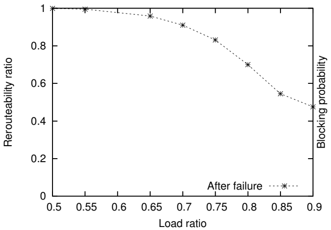

Fig. 4. Reroutability

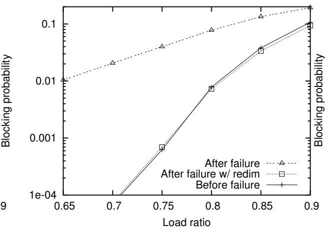

Fig. 5. Overall blocking

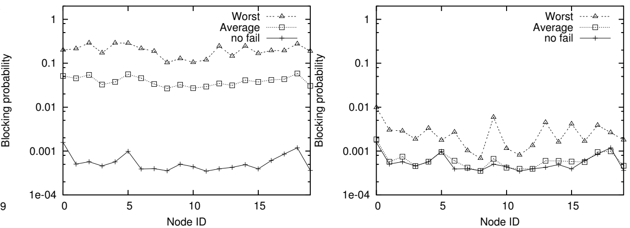

Fig. 6. (a)Before redim. (b)After redim. (load ratio=0.75)

	capacity
original	3851
avg. redim (% overhead)	3966.7 (+3%)
max. redim (% overhead)	4278 (+11%)

Fig. 7. Redimensioning overhead

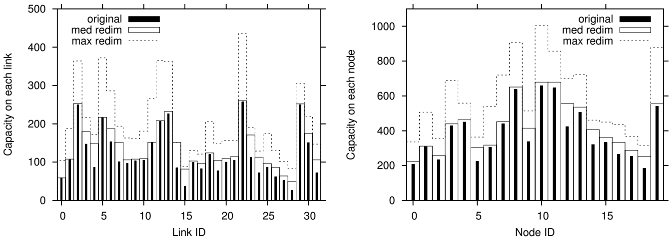

Fig. 8. (a)Link capacity (b)Node capacity

3.2 *Cost and Limitations of Redimensioning*

We quantify the cost and limitations of using redimensioning, which is needed to bring the network back to the normal state without pre-allocating extra capacity. Using a method similar to our network resource dimensioning technique [4], we can preplan for each type of failure and redimension the network upon failure (Algorithm 1). Figure 6(b) shows that, as expected, the average performance of the redimensioned network is very close to the original network, and the worst case impact is reduced by two orders of magnitude.

Figure 7 shows both the average and the worst case overhead for using the redimensioning technique to restore network balance. The overhead is a result of an increase in many topological shortest paths with a link missing. Overall, the cost overhead may seem low enough, especially compared to protection, but the overall capacity can be misleading as far as practical tradeoffs are concerned. We present more detailed cost models in terms of per-link capacity as well as per-node capacity. Figures 8(a) and 8(b) show the original capacity, typical capacity (median) after redimensioning and the worst case capacity over all link failures. For links, this cost is subject to available number of wavelength channels (and fibers). The node capacity costs hint at possible equipment costs such as maximum switch capacity or the number of linecards needed. For example, the typical node capacity suggests that, often, linecards may be reused by means of manual reconfiguration since it is close to the original node capacity. However, in both measures, the maximum capacity required for redimensioning clearly outlines the potential limitations of relying on redimensioning to manage the network. Our results motivate a more careful study of classification of applications and traffic engineering as well as detailed post-failure management plans to to better address the needs of reliable communication. To this end, we are currently investigating techniques to maximize post-failure performance given a set of constraints as discussed above, but the discussion is outside the scope of this work.

4 Summary

We measured the impact of link failures on dynamically routed networks beyond the recovery stage, and showed significant degradation in longer term performance. We then illustrated that a network redimensioning algorithm can be used to restore the network balance, and quantified the practical costs for redimensioning to show that its applicability can be limited. At the same time, the impact of failures on un-redimensioned networks can be detrimental to many applications. Therefore, a network must be carefully designed in order to reliably support a best-effort class(es) of traffic. Finally, our work motivates a more careful classification and management of different applications as well as designing limited redimensioning techniques.

References

1. A. Autenrieth and A. Kirstädter, "Engineering End-to-End IP Resilience Using Resilience-Differentiated QoS," *IEEE Communications Magazine*, January 2002.
2. D. Choi and D. Taylor and R. Seibel, "Priority-based Optical Network Protection and Restoration with Application to DOD Networks," *IEEE MILCOM*, 2003.
3. T. K. Nayak and K. N. Sivarajan, "Routing and Dimensioning of Optical Networks Under Traffic Growth Models: An Asymptotic Analysis," *IEEE Journal of Selected Areas in Communications*, 21(8), Oct. 2003.
4. X. J. Zhang, S. Kim, and S. S. Lumetta, "Resource Dimensioning in WDM Networks under State-Based Routing Schemes," *IEEE BROADNETS*, 2007.