Diversity Requirements for Selecting Candidate Paths for Alternative-Path Routing

Jane M. Simmons
Monarch Network Architects, Holmdel, NJ 07733

Abstract: In alternative-path routing, it is shown that requiring candidate paths to be diverse only with respect to the most congested links in the network performs better than requiring the candidate paths to be completely diverse.

OCIS codes: (060.4265) Networks, wavelength routing; (060.4510) Optical communications

1. Introduction

Selecting the path on which to route a network demand is an important aspect of network design, as it affects network blocking, cost, and reliability. This is especially true as networks become more dynamic. The simplest routing strategy is to pre-calculate one path for each source/destination pair (typically the shortest path), and always follow that path when a demand request arrives for that pair. This strategy, known as fixed-path routing, is well known to produce premature blocking due to the unbalanced link loading that typically results. A preferable strategy is alternative-path routing, where *K* paths are pre-calculated for each source/destination pair [1]. When a demand request arrives, one of the *K* paths is selected based on the current state of the network, thus leading to better load balancing and lower blocking.

The choice of the *K* candidate paths plays an important role in network performance. In much of the previous research employing alternative-path routing (including recent research), the candidate paths are selected to be link-disjoint. However, requiring that the candidate routes be completely disjoint (or maximally-disjoint) can be overly restrictive. We investigate a methodology where the candidate paths are required to be diverse only with respect to the most heavily congested links in the network. We show that the added flexibility of this strategy results in a lower-cost design, due to fewer regenerations, in addition to producing a lower probability of blocking.

The next section provides more details on generating candidate paths with 'bottleneck' diversity rather than complete diversity. We then quantify the improvement that is provided by the bottleneck-diversity strategy, where Section 3 considers unprotected traffic and Section 4 considers protected traffic. In the examples provided, it is assumed that the network employs optical-bypass technology, where traffic that is transiting a node can remain in the optical domain, thereby eliminating much of the need for signal regeneration. However, the methodology is equally applicable to networks with optical-electrical-optical (OEO) signal regeneration at each node.

2. Bottleneck-Link Diversity

Generating diverse paths relative to the most congested links is a relatively simple process [2]. One can take the forecasted traffic demands and perform a preliminary design, where each demand is routed over its shortest path. This yields the links that are likely to be the most heavily loaded. (Alternatively, one can combine the traffic forecast with the max-flow technique of [3] to determine the most critical links.) Of course, the traffic forecast is unlikely to be completely accurate, but the areas of congestion in a network tend to remain consistent, even with some change in traffic pattern. If the traffic pattern changes significantly over time, the list of bottleneck links can be updated based on the actual network loading.

For each required source/destination pair, the bottleneck links are eliminated one at a time, and a shortest-path algorithm is run on the reduced topology. This process produces a set of paths, each of which avoids a particular bottleneck link (the path set may have duplicates). If some of the bottleneck links are adjacent, it is desirable to eliminate the sequence of links prior to running the shortest-path algorithm, to avoid the whole pocket of congestion.

The paths produced from this procedure are then narrowed down to *K* candidate paths, where *K* is typically small. The shortest path is usually selected (i.e., the shortest path in the full topology), plus *K*-1 of the paths that avoid one or more bottlenecks. The selection process can be based on cost, where the cost of a path is largely determined by its required number of regenerations. Note that it is not desirable to select only those paths that avoid all expected areas of congestion, as this would likely lead to excessively long paths, and would likely have the effect of shifting congestion to other areas of the network.

This process is illustrated with the simple topology of Fig. 1. It is assumed that the optical reach of the system is 2,500 km; i.e., after a signal travels 2,500 km, it has degraded to a point where it needs to undergo regeneration. Regeneration is usually accomplished by passing the signal through back-to-back transponders, and thus incurs cost.

Assume that we are interested in generating two candidate paths (i.e., $K = 2$) to be used for unprotected demands between Nodes A and Z. Additionally, assume that links CD and DE are likely to be congested; the remaining links are expected to be less loaded. There are three possible paths between Nodes A and Z:

Path 1: A-B-C-D-E-Z, with a distance of 1,300 km, and therefore no required regeneration

Path 2: A-B-C-F-G-H-Z, with a distance of 3,000 km, requiring one regeneration

Path 3: A-I-J-K-L-M-N-O-Z, with a distance of 8,000 km, requiring three regenerations

Using the bottleneck-diversity strategy, Path 1 would be selected as one candidate path as it is the shortest/cheapest path. The second candidate path is obtained by running the shortest-path algorithm with bottleneck links CD and DE eliminated, producing Path 2. This set of paths requires a total of one regeneration. Compare this to a strategy where completely diverse paths are selected as the candidate paths. Taking the shortest 2-diverse paths results in Paths 1 and 3. This set of candidate paths requires a total of three regenerations. Alternatively, Paths 2 and 3 could be chosen as the two diverse candidate paths, however, this path set requires a total of four regenerations. As this small example illustrates, requiring complete diversity can result in longer, more costly candidate paths.

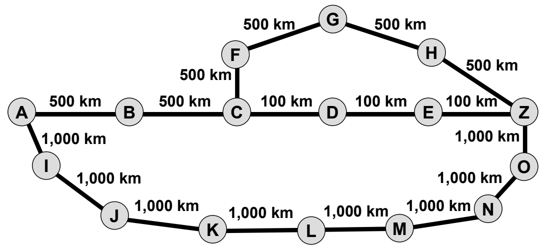

Fig. 1. Three possible paths exist between Nodes A and Z. The link distances are indicated.

3. Network Study Results – Unprotected Traffic

To investigate the impact of candidate-path selection, a study was performed using a network topology representative of a United States national network, with 60 nodes and an average nodal degree of 2.6. The optical reach was assumed to be 2,500 km. A realistic traffic pattern forecast was generated for the network; in this section, we consider only unprotected demands. The traffic pattern was used as input to the bottleneck-diversity methodology described in Section 2, where up to three candidate paths were generated for each relevant source/destination pair. (Three candidate paths typically provide adequate load balancing. Having more alternatives leads to longer paths, and can also lead to more problems with wavelength contention.) A second set of candidate paths was generated using a shortest K-diverse algorithm, where K was taken to be three. In either strategy, paths were not considered if they required more than two extra regenerations as compared with the shortest path. It is assumed that routing traffic over such circuitous paths is undesirable (it should be noted, however, that this eliminated only a very small number of paths that might otherwise have been used).

Demands were added to the network one at a time, where the demands had Poisson arrivals and exponential holding times. All demands were assumed to be at the line-rate. The average offered load was 550 demands (representing 5.5 Tb/s of traffic if the line-rate is 10 Gb/s).

As demand requests arrived, one of the candidate paths was selected based on the least-loaded methodology. With this strategy, the path with the most residual path capacity is selected. (Note that this is different from Fixed Alternative-Path Routing, where the paths are considered in a prescribed order, and the first one with available capacity is selected.) Initially, the selection process did not consider the cost of the paths.

First-fit wavelength assignment was used. If the wavelength assignment process failed, adding a small number of regenerations for purposes of wavelength conversion was permitted.

Two experiments were run, one with bottleneck-diverse candidate paths and one with completely diverse candidate paths. Several runs were performed for each experiment, with 100,000 demand requests per run; the resulting averages and 95% confidence intervals across the runs are shown in the second and third columns of Table 1. First, somewhat more candidate paths were generated with the bottleneck-diversity strategy: on average, 2.6 vs. 2.2 candidate paths per source/destination. Searching for fully diverse paths limits the alternatives, as if either endpoint has degree-two, only two diverse paths can be found. Second, the average length and the average number of hops of the routed paths were about 5% larger when completely diverse candidate paths were used. The combination of fewer alternatives and longer paths resulted in a blocking probability that was three to four times higher with completely diverse paths than with the bottleneck-diversity strategy. Furthermore, the completely diverse strategy is more costly, as the average number of regenerations required per routed demand was roughly 15% higher. Additionally, while not explicitly calculated, the availability of the demands routed via the completely diverse strategy would be lower, due to the paths being longer, with more nodes traversed and more regeneration equipment along the paths.

One of the limitations with the selection methodology described above is that the cost of the path (i.e., the number of regenerations) was not considered. A second set of experiments was run using the same candidate paths but where the least-loaded scheme was applied only to the lowest-cost paths. If none of the lowest-cost paths were available, then candidate paths with one extra regeneration were considered; if none of these paths were free, candidate paths with two extra regenerations were considered. The results are shown in the fourth and fifth columns of Table 1. The average path length, number of path hops, and number of required regenerations per routed demand decreased for both candidate-path strategies, with the bottleneck-diversity strategy maintaining its advantage. The blocking probability also decreased, with the decrease more pronounced in the complete-diversity strategy.

Table 1: Network Study Results with 95% Confidence Intervals

	Unprotected; Regens not Considered		**Unprotected; Regens Considered**		**1+1 Protected; Regens Considered**	
	Completely Diverse	Bottleneck Diverse	Completely Diverse	Bottleneck Diverse	Completely Diverse	Bottleneck Diverse
Avg. # of Candidate Paths per s/d pair	2.2	2.6	2.2	2.6	1.4	2.7
Avg. Routed Path Length	2,161 ± 11 km	2,020 ± 9 km	1,895 ± 8 km	1,814 ± 7 km	4,748 ± 13 km	4,863 ± 20 km
Avg. Routed Hops	5.2 ± .03	4.9 ± .02	4.8 ± .02	4.5 ± .02	11.6 ± .04	11.8 ± .07
Avg. # of Regens per Routed Demand	0.50± .01	0.43± .01	0.35± .01	0.29± .01	1.02± .01	1.00± .01
Avg. Blocking Probability	$8.2x10^{-4} \pm 0.9x10^{-4}$	$2.4x10^{-4} \pm 0.6x10^{-4}$	$3.7x10^{-4} \pm 0.7x10^{-4}$	$1.6x10^{-4} \pm 0.3x10^{-4}$	$107x10^{-4} \pm 5.0x10^{-4}$	$3.7x10^{-4} \pm 1.4x10^{-4}$

4. Network Study Results – Protected Traffic

The bottleneck-diversity strategy can be applied to protected traffic as well. Here, we consider 1+1 traffic, where two paths are established for each demand. For purposes of protection, it is important to have link/node diversity; thus, we required the path-pair to be completely diverse. However, the candidate path-pairs do not need to be completely diverse from one another; again, diversity with respect to the congested links is sufficient. The path-pairs are found by eliminating a bottleneck link and running a shortest dual-path algorithm on the reduced topology. In the example of Fig. 1, one candidate path-pair would be Path 1 and Path 3, and the other candidate path-pair would be Path 2 and Path 3. In the experiment, up to three alternative path-pairs per source/destination were generated.

With the complete-diversity methodology, the number of alternative path-pairs may be restricted due to the topology. If either endpoint of a demand request has a degree of two, then at most two diverse paths exist, leading to a single path-pair for routing. For source/destination pairs with at least three diverse paths (say A, B, and C), then three alternative path-pairs can be generated (A/B, A/C, and B/C).

A network study was run similar to that of Section 3, except that all demands were required to be 1+1 protected. The average offered load was 200 demands. The least-loaded scheme was applied to the lowest-cost path-pairs. Only if these path-pairs were blocked were more costly path-pairs considered. (The path cost is determined by the sum of the regenerations on the two paths; path-pairs with more than three extra regenerations as compared to the shortest dual-path were excluded from the alternative-path-pair set.)

The results are shown in the final two columns of Table 1. The complete-diversity strategy has many fewer alternative paths from which to choose. This resulted in 30 times higher blocking as compared with the bottleneck-diversity strategy. There was little difference in the average path lengths, path hops, and number of regenerations for the routed paths (these statistics are shown in the table for the combination of the work and protect paths).

This same technique can be used with shared restoration, although it can be advantageous to have a larger K so that the sharing of restoration capacity can be better optimized.

5. Conclusion

The network studies presented here illustrate the improved performance that can be achieved by selecting candidate paths with diversity relative to congested links rather than complete diversity. The scheme can be enhanced if dynamic routing is performed when all candidate paths fail, assuming the setup time requirements allow for this.

6. References

[1] A. Birman and A. Kershenbaum, "Routing and wavelength assignment methods in single-hop all-optical networks with blocking," in *Proc. IEEE INFOCOM*, Apr. 1995, vol. 2, pp. 431–438.

[2] J. M. Simmons, *Optical Network Design and Planning*, Springer, New York, 2008.

[3] K. Kar, M. Kodialam, and T. V. Lakshman, "Minimum interference routing of bandwidth guaranteed tunnels with MPLS traffic engineering applications," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 12, pp. 2566 – 2579, Dec. 2000.