Effectiveness of Shared Risk Link Group Auto-Discovery in Optical Networks

Panagiotis Sebos, Jennifer Yates, Dan Rubenstein* and Albert Greenberg
AT&T Labs-Research, 180, Park Ave, Florham Park, NJ 07932, {psebos,jyates,albert}@research.att.com
**EE Dept, Colbumbia Univ., 500W 120th Str.,Room 1312, NY, NY 10027; danr@ee.columbia.edu*

Abstract: We evaluate location-based methods for the auto-discovery of *Shared Risk Link Groups (SRLGs)*, guided by the topology and design of AT&T's next generation transport network. Under realistic scenarios, we find the methods extremely effective in identifying SRLGs and associated link diversity relationships.

OCIS codes: (060.0060) Fiber optics and optical communications

1. Introduction

Shared risk link groups (SRLGs), introduced in [1], provide inputs necessary to plan for reliability in transport networks. An SRLG is associated with an entity at risk, typically a fiber span [2], though more general risks might be modeled. SRLGs are fundamental inputs to resource management in the Generalized Multi-Protocol Label Switching (GMPLS) control plane, now being developed in the IETF [3].

SRLGs can provide the technical foundation for Service Level Agreements (SLAs) with customers. For example, an SLA might provide shared mesh restoration [1] for a set of connections, where primary paths for the connections are provisioned (allocated) and shared resources are reserved (but not allocated until failure) for backup paths. If the SLA provides resilience to single SRLG failures, then for each connection, links along the associated primary path must have no SRLG in common with links along the associated backup path. To support such SLAs, auto-discovery of the state of the network is needed to keep service intent synchronized with network reality. Without auto-discovery, databases tracking service intent could become increasingly out of date and erroneous. Auto-discovery is particularly critical for parameters related to resource allocation, such as link bandwidth, encapsulation type, and SRLGs.

In this paper, we evaluate the effectiveness of location-based methods for SRLG auto-discovery as introduced in [4]. Specifically, we consider realistic transport networks overlaid on the AT&T fiber span topology, and SRLGs defined in terms of fiber spans. Location-based SRLG auto-discovery relies on location information collected at active components, from optical amplifiers and Optical Transport Units (OTUs). Varying the amplifier spacing over a realistic range, we find SRLG auto-discovery to be effective. In particular, we are able to identify all SRLGs that exceed one mile, with amplifiers spaced at 50 miles along links. In addition, we are able to correctly identify diversity relationships between all pairs of links.

Fig. 1 illustrates a simple example of a 3 XC network with its associated SRLGs defined in terms of fiber spans. Fiber span failure (e.g., backhoe cut) in general leads to failure of all links (at the XC level) routed within the fiber span. The XC level connectivity graph of Fig. 1a, is embedded in the fiber span level connectivity graph of Fig. 1b. XCs are mapped to fiber span end points, and links between XCs are mapped to fiber span paths. In this example, the link A-C between XCs A and C in Fig 1a is routed over the two-hop fiber span path from A to C in Fig 1b. The three SRLGs in Fig. 1b correspond to the three fiber spans. Accordingly, link A-C belongs to $SRLG_1$ and $SRLG_2$.

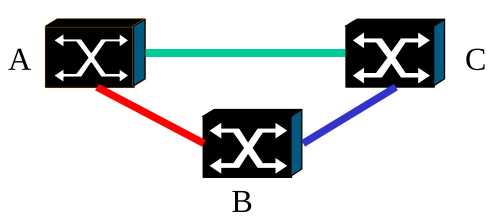

Fig 1a. XC level connectivity. Nodes represent XCs and links represent a set of communication channels.

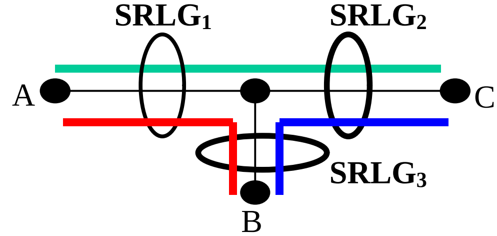

Fig 1b. Corresponding fiber span level connectivity. Links represent fiber spans (collection of fibers in a conduit) and nodes the fiber span end points.

2. SRLG Auto-discovery

Location-based methods for auto-discovery of SRLGs were introduced in [4]. SRLGs are inferred from location information associated with active components (e.g., amplifiers or OTUs) placed along the links. Co-located components are identified and each SRLG corresponds to a set of links that are all routed through the same set of locations.

Consider a fiber span level topology (as in Fig. 1b), and a corresponding embedding of a XC level topology (as in Fig. 1a). In addition, assume active components are situated at locations along each embedded link. Under location-based auto-discovery, SRLGs are defined formally as follows:

- Consider the locations of active components, and let L_i denote the location of the i^{th} active component.
- Let $\sigma(L_i)$ denote the set of links that route through location L_i.
- The maximal list of distinct sets $\sigma(L_i)$ are the SRLGs.

3. Evaluation

We base our investigation on network topologies consistent with AT&T's emerging optical network architecture. Specifically, we start with the AT&T fiber span topology, and then embed a XC level topology of approximately 50 mesh connected XCs, representing a network that provides national coverage. Next, from a fiber span topology database we obtain locations for OTUs situated at fiber span end points. Finally, we place amplifiers along each embedded link between XCs by the following design rule, taking into account that each XC and OTU implies OEO conversion. Let α denote the amplifier spacing. For each link, place amplifiers every α miles along the fiber span path associated with the link. On crossing from a fiber span to the next fiber span, if the spacing rule calls for amplifier placement in the next span instead place it at the end of the current span. On crossing an OTU, reset the amplifier placement (the next amplifier is needed at a distance of α from the OTU location). Note that for this analytic study we have varied the amplifier spacing α beyond the range of typical deployment to test the limits of our proposed SRLG identification scheme.

SRLGs are defined as maximal sets of fiber spans that are used by common sets of links. Thus, if two links have an SRLG in common then the two are not diverse. In general, risk is likely to be proportional to fiber length. In fact, some SLAs may exclude small SRLGs (i.e., the sum of the span lengths determining the SRLG is small; e.g., less than 1 mile), such as those associated with lateral connections to buildings.

Fig. 2a plots the percentage of correctly identified SRLGs as a function of amplifier spacing. If only location information associated with OTUs is used to identify the SRLGs (lowest curve), we are able to identify 50% of the full set of SRLGs (independent of the amplifier spacing). However, if we exclude all SRLGs of length less than one mile, we find that we are able to identify 62% of the SRLGs using OTU location information alone. Fig. 2a additionally depicts the percentage of SRLGs correctly identified using both amplifier and OTU locations. Results are plotted for all SRLGs with non-zero length (middle curve), and for SRLGs with length exceeding 1 mile (upper curve). As we increase the amplifier spacing, the number of locations identified decreases. The effectiveness of our scheme correspondingly decreases.

The most important conclusion to draw from these curves is that we are able to successfully identify all of the SRLGs that exceed 1 mile within our example transport network if we have amplifier spacings of 50 miles or smaller. Even if we increase our amplifier spacing to, say, 70 miles, we are still able to automatically identify 91% of the SRLGs. These results may well be better than that achievable through manual database management.

Fig. 2b plots the percentage of correctly identified SRLGs as a function of SRLG length (in 50 mile increments) with 50 mile amplifier spacing. We observe that we have correctly identified all of the SRLGs of length greater than 50 miles when we utilize location information for both amplifiers and OTUs. However, we have some difficultly identifying very small SRLGs, with only 22% of the SRLGs less than 1 mile being identified and 60% of the SRLGs less than 50 miles identified. In general, longer SRLGs are easier to identify than shorter SRLGs, because these SRLGs are more likely to have active components from which we can obtain location information. This leads us to the highly desirable property that the higher the risk of failure (the longer the SRLG), the easier it is for us to identify the SRLG.

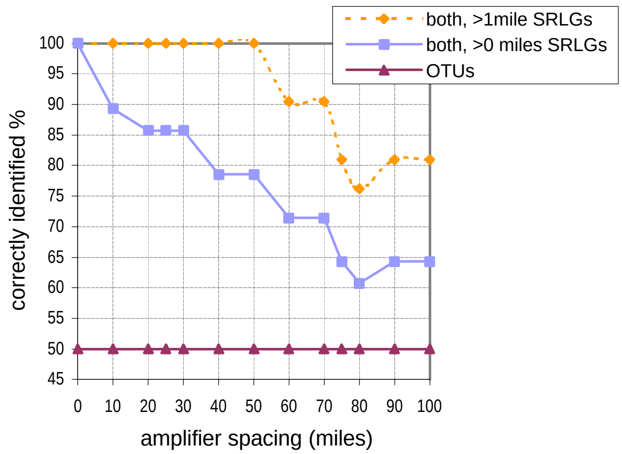

Fig. 2a. percentage correct SRLGs as a function of amplifier spacing.

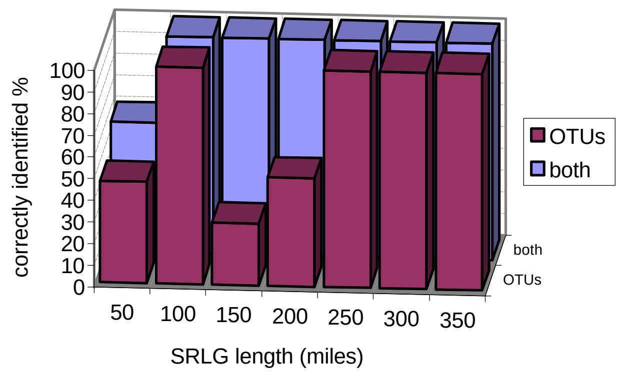

Fig 2b. percentage correct SRLGs as a function of SRLG length.

If we are using only OTU locations, we find that there are some SRLGs of significant length that we are unable to identify. For example, we were only able to identify 29% of the SRLGs that are between 100 and 150 miles long. It is thus highly desirable to utilize both OTU and amplifier location information.

The fundamental purpose of SRLGs is to determine whether two or more links are physically diverse. A pair of links is diverse if the two links share no SRLG. Even if we fail to identify all SRLGs, we may still have adequate information for determining link diversity. We found that if we used all SRLGs, independent of length, that we were able to correctly determine the diversity in *every* case (i.e., 100% correct). This validates the notion that maintaining accurate information of SRLGs enables the deployment of a highly reliable transport network.

Our automated procedure is not able to accurately estimate the length of an SRLG (at least to within distance shorter than the α). This means that we are unable to identify pairs of links that share only very short SRLGs. This may or may not be a limitation, depending on the extent to which short SRLGs are included in different SLAs.

The above results indicate that the most important parameter in terms of the effectiveness of our proposed scheme is the ratio of the amplifier spacing to the SRLG lengths. SRLGs that exceed the amplifier spacing α are identified (assuming our amplifier placement algorithm). Thus, in general, we expect our scheme to be most effective in large networks with relatively long SRLGs, or in networks with closely spaced location identifiers (many active components).

4. Conclusions

We evaluated location-based methods for the auto-discovery of SRLGs, guided by the topology and design of AT&T's next generation transport network. The methods rely on location identifiers on active components. We found these techniques extremely effective in identifying SRLGs and associated link diversity relationships. SRLGs provide inputs necessary to plan for reliability in transport networks, but do not in themselves reveal connection diversity violations. Yet by combining SRLG inputs with SRLG-aware methods for routing links over the fiber topology, a transport service provider can provide diversity assurances that support strong SLAs.

5. References

[1] S. Chaudhuri, G. Hjálmtýsson and J. Yates, "Control of Lightpaths in an Optical Network," Optical Internetworking Forum OIF2000.04, IETF Internet Draft, Jan. 2000.

[2] J. Strand, A. Chiu and R. Tkach, "Issues for Routing in the Optical Layer," *IEEE Communications Magazine*, Feb. 2001.

[3] E. Mannie *et al.*, "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," IETF draft, work in progress, June 2001.

[4] P. Sebos, J. Yates, G. Hjalmtysson and A. Greenberg, "Auto-discovery of Shared Risk Link Groups," *Optical Fiber Commun. Conf.*, March 2001.