Achieving Diversity in Optical Networks Using Shared Risk Groups

Sudheer Dharanikota, Raj Jain

Raj Jain is now at
Washington University in Saint Louis
Jain@cse.wustl.edu
http://www.cse.wustl.edu/~jain/

Yong Xue
WorldCom, Inc.
22001 Loudoun County Parkway
Ashburn, VA - 20147
yong.xue@wcom.com

Abstract

Diverse working and protection paths are used in telecommunication networks to increase survivability and availability in case of faults. In selecting these diverse paths, it is important that the two paths do not share any links that share the same risk. This is done by assigning each link to a Shared Risk Link Group (SRLG). In this paper, we extend this concept to nodes and domains. Diverse paths should not pass through nodes that can both fail together, for example, by being located in the same building. Similarly, for a higher level of survivability, diverse paths may not pass through network domains that share the same risk. Domains are defined as subnetworks with common administrative control, technology, or geography. Shared Risk Group is an extension of the SRLG concept that avoids not only link level risks but also node and domain level risks. We demonstrate how to use the transport network diversity using this concept while computing application level diversity in a multi-layered optical network.

Key words: Diversity, Failures, Optical Networks, Risk, Shared Risk Groups (SRG), SRLG, and Traffic Engineering (TE).

1. Introduction

In connectionless networks, such as in IP networks, a failure in the network is bypassed by rerouting the traffic around the failure. However, the time required to detect the failure and update routing tables throughout the network may be too long for most telecom applications. To alleviate such problems telecom networks have traditionally used connection-oriented architectures with a high-level of link and equipment redundancy. Bi-directional Link Switched Rings (BLSRs) used in SONET/SDH networks are one example of link redundancy mechanisms used in telecom networks. With these ring-based topologies, failure of a single link or node results in a minimal disruption. The operation can be resumed in typically less than 50ms under certain conditions.

Another method for reducing the impact of failures is to have two or more diverse paths such that upon the failure of the working (or primary) path, the traffic can be quickly switched to one of the protection (or secondary) paths. This method applies to networks with both ring and mesh topologies. Also, it allows carriers to use the protection resources for lower priority traffic. Therefore, this is expected to be used commonly in future.

Standard bodies, including Internet Engineering Task Force (IETF), Optical Internetworking Forum (OIF), and International Telecommunication Union (ITU) are all working to formalize mechanisms that will allow automated determination of diverse paths. This work is being done partly under the umbrella of traffic

engineering, which deals with moving traffic to those parts of the networks that will provide the highest utilization of network resources while also providing the best user performance including availability and survivability.

Although the concepts discussed in this paper are relevant to both packet-switched and circuit-switched networks, we concentrate only on circuit-switched (or connection-oriented) networks especially optical networks.

In this paper we introduce the concept of shared risk groups (SRG) and demonstrate how this concept can be used to achieve diversity in optical networks and to assess risks associated with a path. In this process we extend the concept of diversity from a link level, node level and SRLG (Shared Risk Link Group) level diversity to an additional domain level diversity. We demonstrate how to use the transport network diversity using this concept while computing application level diversity in a multi-layered optical network.

In Section 2 we present the background relevant for diversity and risk assessment in telecommunication networks. In Section 3 we provide the relevant definitions whose applicability is then elaborated in Section 4. Achieving diversity and assessing risk with a given path are presented in Sections 5 and 6. In Section 7 we present implementation guidelines of the SRG concept for a specific application. Conclusions, acknowledgements and references follow in Sections 8, 9 and 10, respectively.

2. Background

In data networks, a path consists of an alternating sequence of nodes and links. The nodes originate and terminate data traffic and the links transmit the traffic. One important concept in data path selection is diversity. There are two types of topological path diversities: link disjoint and node disjoint [1]. Two data paths are said to be link disjoint if they do not share a common link to be a single point of failure along the path. Similarly, two data paths are said to be node disjoint if they do not share a common node as a single point of failure along the path except the originating and the terminating nodes. It is obvious that the node disjointness implies the link disjointness. In this paper we use terms link (node)-disjoint interchangeably with term link (node)-diverse.

Consider a 6-node mesh network as shown in Figure 1. Assume that there is a request to make a connection between nodes A and G. Suppose the path specified by set of nodes (A, C, F, G) or to be more specific the (node, link) pairs {{A, 2}, {C, 1}, {F, 2}, G) is the shortest path that satisfies this request. If this path is not protected from failures, a node failure, such as the failure (a) shown in Figure 1.A, can disrupt the traffic flow on the path. In telecommunication networks, to avoid data loss due to a node failure two node diverse paths (i.e., the two paths that do not have any nodes in common excepting the source and destination nodes) are computed and one of them is made the primary path while the other serves as the secondary path. For example, in Figure 1.A, path (A, B, F, G) has been designated as a primary path and (A, C, D, E, G) is the secondary path. In case of the primary path failure, the traffic is switched to the secondary path at node A. Similarly, to protect a path against link failures (such as failure (b) in Figure 1.B), one can compute link-diverse primary and secondary paths. Note that the primary path ({A, 1}, {B, 2}, {F, 1}, G) and the secondary path ({A, 2}, {C, 1}, {F, 2}, G) in Figure 1.B do not share any links in common. If both the paths have links that share the same conduit, as shown in Figure 1.B for links ({B, 2} and {C, 1}), then a failure of the conduit (such as failure (c)) will affect both the primary and the secondary paths. One way to solve this problem is

to assign an identifier to this conduit (the common resource or risk) and avoid including resources with this identifier in both the primary and secondary path during the path computation. This common identifier is called “Shared Risk Link Group (SRLG)”. Hence in addition to node-diverse and link-diverse paths, in telecommunication terms, it is a common practice to compute SRLG-diverse paths [1, 3, 4]. Reference [3] provides enhancements to the concept of SRLGs thereby streamlining its meaning. In this paper, the constraints discussed till now (such as node, link or SRLG diversity) are called “exclusive constraints” meaning they request for excluding risks (such as nodes, links and SRLGs). Diverse path computation can be performed by in-line mechanisms such as the constraint-based routing protocols [5] or by external provisioning tools. The concepts that we introduce in this paper are applicable to both the cases.

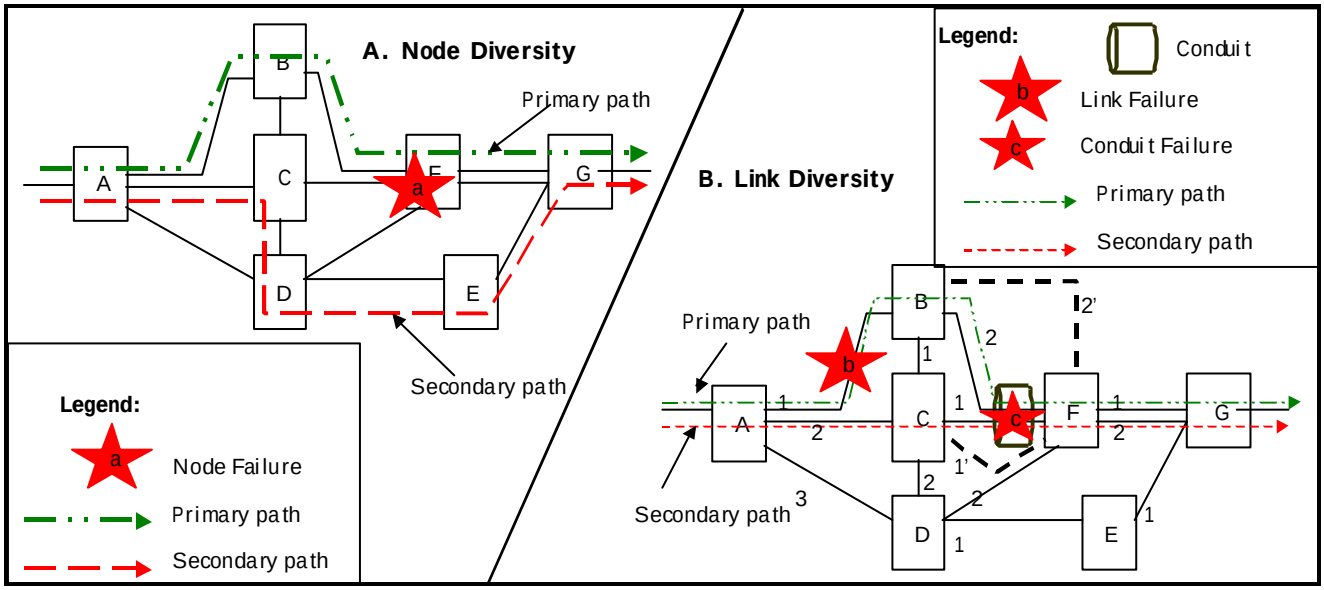

Figure 1: Usage of node and link diversity under different failure conditions

In addition to the SRLG mechanisms to reduce failure time, networks can also be proactively planned to avoid risks. This alternative is called "capabilities." For example, in Figure 1.B, the SRLG failure (c) can be averted by providing a 1:1 shared link protection (link capability) to {B, 2} by {B, 2’} and {C, 1} by {C, 1’} links to F. Of course these new links should not share the risk taken by the original links. These are mainly to localize the fault and hence the recovery, and to reduce time taken to restore the connections from a failure. Such fault or risk localization is done in the telecom networks in the form of shared link (or span) protection, ring protection, and dedicated (such as 1+1) protection. Here, each link has a built-in protection. Hence when a diverse path computation is performed one may want to (specify a constraint to) consider the protection capable links over the unprotected links. Such constraints are called “Inclusive Constraints” in this document.

A risk sharing, however, is not limited to links between adjacent nodes, which has been the case until now. To contain the affect of such failures many precautions are taken. For failures in a node, many hardware and software redundancy mechanisms are employed to contain the failure to the node. Similarly for failures in a region (e.g., central office etc.) redundant equipment and other techniques are employed to contain the failure local to the region and in the same note, a failure in a group of nodes and links operating together is contained to the same

group by employing group restoration mechanisms (such being the case in ring topologies). This leads to the concept of the risk domain (as formally defined in the next section), which defines the perimeter of the components that are involved in solving the failure.

The “shared risk concept” can also be viewed as a mechanism to hide or reduce the amount of topology information propagated in a multi-layered network. Consider a multi-layer network with fiber, optical (for instance G.709 OTN), SONET/SDH and router layer topology in the ascending order of encompassing the previous layer topology. As shown in Figure 2 (which depicts only two layers), the upper layer is called the client layer and the lower layer is called the server layer. In such a topology a link at the client layer (for example, logical link between nodes B and D) can mean many nodes and links in the server layer (for example, it may map to <{B, B’}, {B’, X}, {X, Y}, {Y, D’}, {D’, D}>).

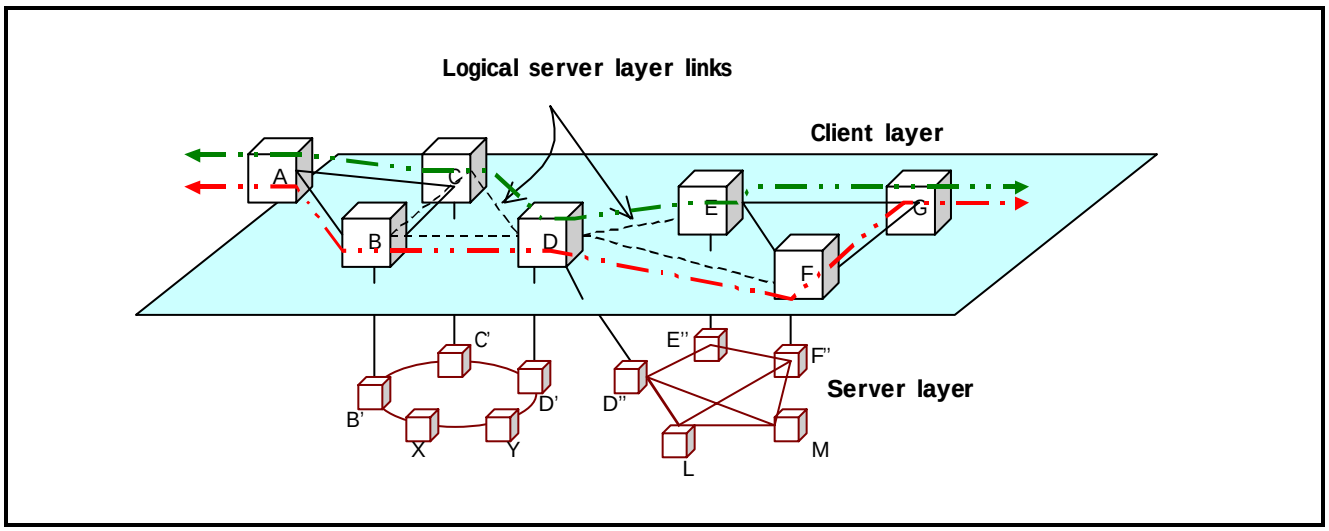

Figure 2 Application of SRG to a multi-layer GMPLS based transport network

To provide diversity at the client layer one should consider failures in the server layer topology. For example, Let us assuming that we have computed two link-diverse paths between (A, G), which are <A, C, D, E, G> and <A, B, D, F, G> as shown in Figure 2. Suppose the client link <C, D> maps to server links <C, C’, D’, D>, and <B, D> maps to <B, B’, C’, D’, D>. A failure on link <C’, D’> can potentially affect both the connections. But since the server topology is a ring it may recover the failure, without knowing that these two connections are part of link diverse paths in the client layer. Now suppose links <D, E> and <D, F> are mapped to <D, D”, M, F”, E”, E> and <D, D”, L, M, F”, F>, respectively. A failure on segment <M, F”> can affect both the connections. If there are no local recovery mechanisms, then this failure may mean that client link layer diversity has failed its purpose. Thus, to provide diverse links or paths (sequences of links) at the client layer requires some means of abstracting the restoration capabilities at the server layer, so that this abstraction can be used by path computation at the client layer.

With the adoption of GMPLS (Generalized Multi-Protocol Label Switching) [7] control plane in the packet and circuit based networks it is now possible to compute diverse paths in multiple layers. The notion of diversity can be abstracted and dynamically computed at many layers. At present, the only way to

provide this abstraction of the server layer topology in the client layers is to use SRLG, which has a very limited usage in diverse path computation. The capability assignment constructs provided in this paper are useful in achieving the task of diverse path computation across multiple layers.

3. Definition of SRG

A *Shared Risk Link Group (SRLG)* is defined as a group of links that share a common risk component whose failure can potentially cause the failure of all the links in the group. For example, all fiber links that go through a common conduit in the ground belong to the same SRLG group because the conduit is a shared risk component whose failure, such as a cut, will cause all fibers in the conduit to break. It's clear from this definition that SRLG is a relation defined within a group of links based upon a specific risk factor. Note that the risk factor can be defined based on various technical or administrative grounds such as "sharing a conduit", "within 10 miles of distance proximity" etc. A SRLG identifier is often defined for each shared risk link group for computational purposes.

The SRLG concept can be used to define link-disjoint path diversity. Two data paths are link-disjoint if no two links on the two paths belong to the same SRLG under consideration. Similarly, shared risk node group (SRNG) concept can be used to define node-disjoint path diversity.

A *Shared Risk Node Group (SRNG)* is defined as a group of nodes that share a common risk component whose failure can potentially cause the failure of all nodes in the group. For example, all routers in a building can be considered as a SRNG group because the failure (e.g. power failure in the building) can cause all routers to fail. Two data paths are said to be SRNG disjoint if no nodes along the two paths belong to the same SRNG under consideration.

The SRLG and SRNG concepts work well for a flat network topology and can be extended to hierarchical networks. A flat network can be partitioned into domains that consist of a set of nodes and associated links. The partition can be based on topological, technological, or administrative reasons. The links between domains are the links between the edge nodes in different domains. Each domain can be further partitioned into sub-domains. Such partitioning can be repeated until the granularity of the domains reaches a certain level, thereby, generating a network hierarchy. In hierarchical networks, a domain is a logical node that has a set of ports corresponding to the incoming and outgoing links.

Domain is an useful logical structure for the description of the networking function and capability. We can treat a domain as a sub-network cloud composed of a set of ports and sub-network connections (SNC) between the ports. The topology details of the domain can be ignored and its networking functions can be defined by the virtual links across the domains that are terminated at the ports.

We define a *risk domain* to be a network domain associated with a defined risk type for which the occurrence of the risk event can cause the failure of the domain in terms of the link connection. *The domain risk* can be quantitatively indicated by a risk factor that represents the probability of the failure occurrence. Another concept we can associate with a domain is the *capability of the domain, which* defines the type of the connections the domain provides. The capability can be defined based on the quality attributes of its sub-network connections, such as protection type, path quality, etc. For example, a SONET BLSR ring is a domain and the SNC links between all ADM ports of the domain have a 1:1 protection capability.

We define a Shared Risk Group as a risk domain with associated risk factors and domain capabilities. Similar to SRLG, there is an SRG identifier associated with the shared risk group for the risk type under consideration. It is important to note that the SRG concept not only defines a shared risk group of nodes and links, but also defines the preference level in terms of path selection by considering the risk level and path quality. SRG is essential in supporting hierarchical routing in which a domain level path can be calculated first and then expanded within each domain.

4. Applications of SRG

The SRG concept is an evolving notion from SRLG. In the past, SRLG has been used by carriers for diversity planning. However, the original SRLG concept has a limited use and is not able to handle some important risk types and path planning requirements. Some of the application scenarios include:

- In planning diverse paths, link diversity is essential, but it can't handle node failures such as central office break-downs due to fire, loss of power, etc. We need to have node-disjoint path planning capability.
- Some risks are associated with larger geographic areas such as metros, regions, etc. We need to have much coarser granularity of network structure than links and nodes to handle these risk avoidance issues. Examples include earthquake or flood prone areas and areas with certain radius of nuclear power plants, etc.
- In some cases, there is a lack of true diverse fiber routes. For example, a tunnel or a bridge may be used by carriers as the only fiber route across certain rivers or mountains. In other cases, a customer may be willing to tolerate larger risk by considering parallel fiber routes in proximity. In these cases, we need to have the capability to select a path with minimum failure risk.
- Carriers may want to take advantage of the protection or restoration capabilities of certain sub-networks such as BLSR SONET ring as a protected "diverse" fiber sub-network connections as long as different access points are used.

SRG and Path Planning:

The SRLG, SRNG and the generalized SRG concepts can be utilized to support two major categories of applications for path selection and planning:

1. One is *diverse path* computation where the goal is to compute two or more mutually link-disjoint, node-disjoint or risk-disjoint (SRLG, SRNG or SRG) diverse paths between two end points and
2. The other is *preferred path* computation where the goal is to find an explicitly routed path that satisfies one or more of the constraints in the following possible formats:
    - Inclusive resource list: a list of physical or logical network structures (node, link, domain or risk type) part or all of which ought to be included in the selected path.
    - Exclusive resource list: a list of physical or logical network structures (node, link, domain or risk type) that have to be excluded for the consideration for the selection in the path route.

- Path quality list: a set of network capability types or risk levels that have to be met cumulatively end-to-end on the selected path.

Some of possible business applications that can be used to illustrate the usage of the path computation model mentioned above are as follows:

- Protected circuit provisioning: Carriers usually can sell either protected or unprotected circuits to their customers. For a protected circuit, a secondary protection circuit that is at least physically diverse from the primary circuit should be computed and provisioned. The protection scheme can be either 1+1 or 1:1 depending upon the customer's needs and affordability. In this case, SRLG associated with the fiber conduit or right-of-way (ROW) can be used to support the diverse path planning.
- Preferably routed circuit provisioning: Carriers are under increasing pressure to provide more value-added services to their customers. One way to offer such services is to allow the customer some level of control as to how their circuits should be routed. It is absolutely essential that the carriers do not have to disclose their network topologies to their customers. However, they can define risk domains at city level or metro sub-network level and allow customers to have a rough view of the network reach and capability without giving the customer detailed visibility into their network topology, thereby allowing customers to select circuit route at the defined domain level. Note that the carriers can define domains as SRG based on their needs. The domain level network topology may not reflect the true topology of the network.
- Preferred quality circuit provisioning: High availability and reliability is one of the most important attributes of the transport network. Customers, internal or external, often request circuits with carrier grade or 99.999% (five 9's) reliability. To support provisioning of this type of circuits, the carriers can associate SRG (their capability and risk attributes) with each link and node. Since a failure probability can be used to calculate the reliability of each network element (link and node) along the path, an end-to-end path-level reliability could be derived.

SRG Information:

Currently, manual provisioning techniques are employed to ensure path routing diversity. This is slow, tedious, and time consuming. The preferred path routing applications are very similar to the constraint-based routing, which should be directly applied to solve the preferred path selection problems for simple flat network topologies. More advanced applications based on the SRG concepts need a new framework to support them.

There are several gating issues that need to be resolved before we can successfully apply the SRG concept to the diverse path and preferred path planning problems. Among them are two major issues: risk modeling and SRG database creation and management.

Risk modeling is the analysis of business application risk requirements. The objective is two fold. On one hand, we need to determine how the network resources should be partitioned or grouped into a hierarchical domain topology that will meet the application requirements. On the other hand, we need to identify all the shared risk components, capabilities and risk factors that can be associated with each one of the defined domain objects.

Depending upon the application, the defined shared risk component could vary greatly. For example, routing diversity can have different meanings and criteria depending upon the actual application requirements. If the application requires two diverse paths not to share the conduit or right-of-the-way (ROW), then the shared risk component will be the conduit or the ROW attributes to be associated with each fiber link. Obviously the conduit and ROW are different in scope.

The SRG database is the key in successful implementation of the diverse path and preferred path routing applications. However, creating and maintaining the SRLG, SRNG, and SRG databases is a very challenging and difficult task. Some of the difficulties lie in:

- A huge amount of SRG data needs to be created and maintained. As discussed in [8], to create the SRLG data for link diverse path provisioning in a carrier network of hundreds of cross-connect nodes, at least one order of magnitude higher number of SRLG specific “nodes” need to be considered. This is because many places along the fiber path such as fiber cross points, amplifiers, fiber splicing points, etc. become fiber span termination points and each span is the unit of network structure for SRLG consideration purposes.
- There are many different SRG’s corresponding to different type of risk considerations. A shared risk domain can belong to multiple SRG’s and a SRG can be associated with multiple risk domains. This many-to-many relation has multiplying effect in terms of the SRG data that need to be created and maintained.
- The nature of the SRG does not lend itself to a self-discoverable process. On the other hand, formation of risk domains and mapping of them to SRG’s normally requires a manual process. Even though, there have been some recent research efforts that tried to automate the discovery of SRLG as described in [4], general SRG discovery is not yet available. It is worth noting that, once created, the risk domains, SRG’s, and mapping between them tend to be static.

Due to the static nature of the SRG information and the need to incorporate them into the path selection process, one can either configure them as link attributes and let the IGP routing protocol advertise and propagate the SRG link attributes across the network or put them into a centralized SRG database and let the path selection process query the database when needed.

5. Diversity and SRG

A diverse path computation algorithm such as modified Dijkstra’s algorithm [1] has three interfaces as shown in **Figure 3**. It takes inputs from three databases namely, *topology, TE, and existing path databases* to provide responses to new path requests. The *topology* and *TE databases* are managed through routing protocols such as OSPF, ISIS or by querying the network management system. The *topology database* contains the nodes, links, and their interconnection. The *TE database* contains the properties or capabilities of the network resources. The *existing path database* contains the path information, such as the nodes and links traversed by various paths in the network. With the above inputs, a request to find diverse paths between a given source and destination pair with given constraints on path selection is processed by the algorithm. The computed path could be a complete enumeration of all intermediate nodes or a partial list of key intermediate nodes between the source and destination pairs. The preceding two options are called

strict explicit path and loose explicit path, respectively. Once the response is accepted, the computed path is recorded in the existing path database.

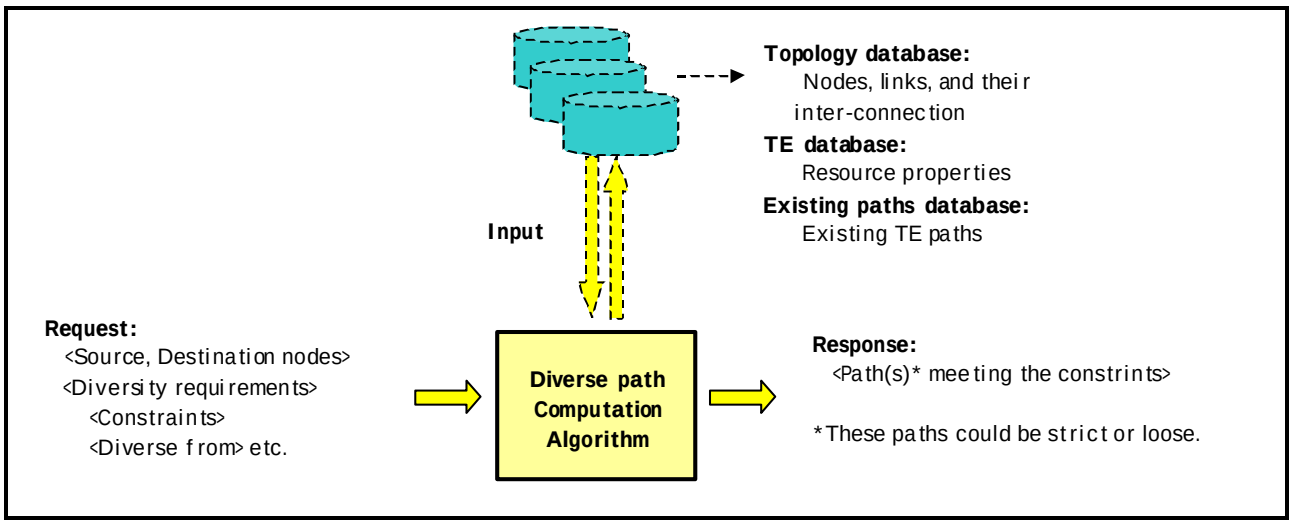

Figure 3 Different interfaces to a diverse path computation algorithm

Traffic Engineering (TE) in IP (Internet Protocol) networks is achieved using MPLS technology as an overlay on the IP networks. The success of the MPLS TE concept has lead to its application in the optical domain as well. By introducing connection-orientedness in the IP technology, we loose the automatic hop-by-hop recovery of data paths in case of failures. Hence, diverse paths are established between the source and destination MPLS nodes for achieving path resiliency.

The diversity requirements of carrier transport networks have some differences from those of packet (Layer 3 and Layer 2 switching) networks, as presented in Table 1. Transport networks inherently provide elaborate protection and restoration mechanisms. These networks are not always structured in mesh topologies as assumed by the packet networks. Transport networks contain multiple sub-layers unlike in packet networks. This leads to different strategies for protection and restoration. At present, in the packet networks only interfaces (or links in some sense) have capability assignment, whereas in the transport networks links, nodes and domains have capabilities. Therefore, the path computation in transport networks allows more elaborate inclusive and exclusive constraints as presented in the table. Also, since the transport networks are grouped differently than packet networks, the path computation mechanisms may not have the complete information about the topology to compute both loose explicit path and strict explicit path. In the following discussion we demonstrate the usefulness of SRG concept to address some of the issues raised by the transport network diversity requirements.

Table 1 Comparison of packet and transport networks for diverse path computation

Category	Packet networks	Transport networks
Inherent protection	Not supported	Supported (using rings etc.)
Topology	Only mesh	Mesh, *Ring*, and *Mesh-Ring interconnects*
Sub-layers of connections	Single layer	*Multiple* layers

Capability assignment	Only to link	Link, *node* and *domain*
Exclusive constraints	None	Link, node, SRLG, and *SRG diversity*
Inclusive constraints	Link, nodes	*Link, node,* and *SRG*
Path computation	Can compute only strict explicit paths	Should compute strict and *loose* explicit paths

The concept of SRG can be used in all parts of diverse path computation module as depicted in Figure 3. SRG can be used to represent any arbitrary subset of the topology as per the definition provided in the previous section. This helps in summarizing and hence in reducing the topology database information. SRG capabilities can be propagated as risk domain capabilities, which will be a part of the TE database. The path computation constraints can be enhanced to include or exclude certain SRGs based on their capabilities. As a final advantage SRG boundaries can be used during the path computation to abstract the risk domain without specifying the explicit nodes in that risk domain. The internal path through that domain is expanded only during the path setup. This is the so called loose path computation.

6. Risk assessment and SRG

Risk assessment is defined as the evaluation of the potential risk associated with the inclusion of a given resource in a given path. For Example, consider the following client requests to the optical network:

- Request a persistent connection with 99.999% (widely known five 9's) availability or equivalently a downtime of less than 5 minutes per year or
- Request a higher protection for a portion of the traffic (at the expense of paying a higher charge) compared to other low-priority traffic.

Such requirements will be translated into constraints in path computation. Such constraints can be grouped into path selection constraints and path characterization constraints. The *path selection constraints* typically dictate which physical path should be taken to achieve the client's availability requirements. These requirements are typically the logical and physical diversity. The *path characterization constraints* typically dictate the protection mechanisms as requested by the client. This can be achieved in the form of optical rings, mesh protection mechanisms, etc. These constraints can be satisfied using the link, node, and domain capabilities as discussed in the previous section on diversity.

The components that need formalization in this example are specifying the user requirements, translating the requirements into path computation constraints, configuring the network in a way that helps in assessing its features (such as the availability), propagating the information, and finally using information in path computation. In the following discussion we address all these components except the specification, which is not in the scope of this paper.

A convenient way to achieve risk assessment is by associating a conditional risk value with each of the SRGs. Also, by associating a weight factor with the SRGs, we can change the probability of selecting specific SRGs. This calls for configuring a risk factor and a weight factor per SRG. In addition to the SRG capabilities, as discussed before, the above values can also be propagated via routing protocols. With the help of these two configuration parameters, the use of

typical CSPF algorithms to compute a path can be extended to assess the risk associated with the computed path. For example, if a path traverses SRGs 1, 3, 5, then one may infer that the risk associated with this path is (Risk 1 x Risk 3 x Risk 5).

7. Implementation guidelines

In this section, we summarize the discussion till now by applying SRG concept to distributed path computation, an example application considered before. These protocol extensions are considered at length in [9]. As shown in Figure 4, a typical transport network can be configured into different domains based on the capabilities. The topology information of all the domains is not known to the end-points, which determine the constraint-based path for a connection request. In the following paragraphs, we discuss the tasks involved in applying SRG concept during configuration, routing, path computation, and path setup (or signaling).

To achieve diversity in path computation and risk assessment, one has to configure relevant parameters on the network elements. A summary of these parameters is presented in Figure 4. As a first step, one has to determine the domain boundary based on any of the criteria discussed in Section 3. This is the configuration of domain to provide a boundary for summarization or hiding the capability information. This is similar to the concept of an area boundary in the existing IGPs. An SRG value is allocated to this domain. This SRG can be a flat 32-bit number or can have an elaborate encoding mechanism as proposed for SRLGs in [3]. For each SRG, a set of capabilities of the domain is assigned. These capabilities are dependent on the reasons behind the domain creation. For example, these capabilities could be the protection capability of the domain. Reliability parameters of a domain such as the risk associated with the domain depends on the protection and restoration mechanisms inherent to the domain. Since the domain belongs to a single operator, we can assign these parameters per mechanism per domain. This can be assigned for each type of failure per SRG. To influence a diverse path computation algorithm, one can configure preferred path allocation parameters, such as weights to each SRG. This weight can be statically configured once at the initial stage or dynamically determined since it can be defined as additive metric whose individual values are the link TE metrics.

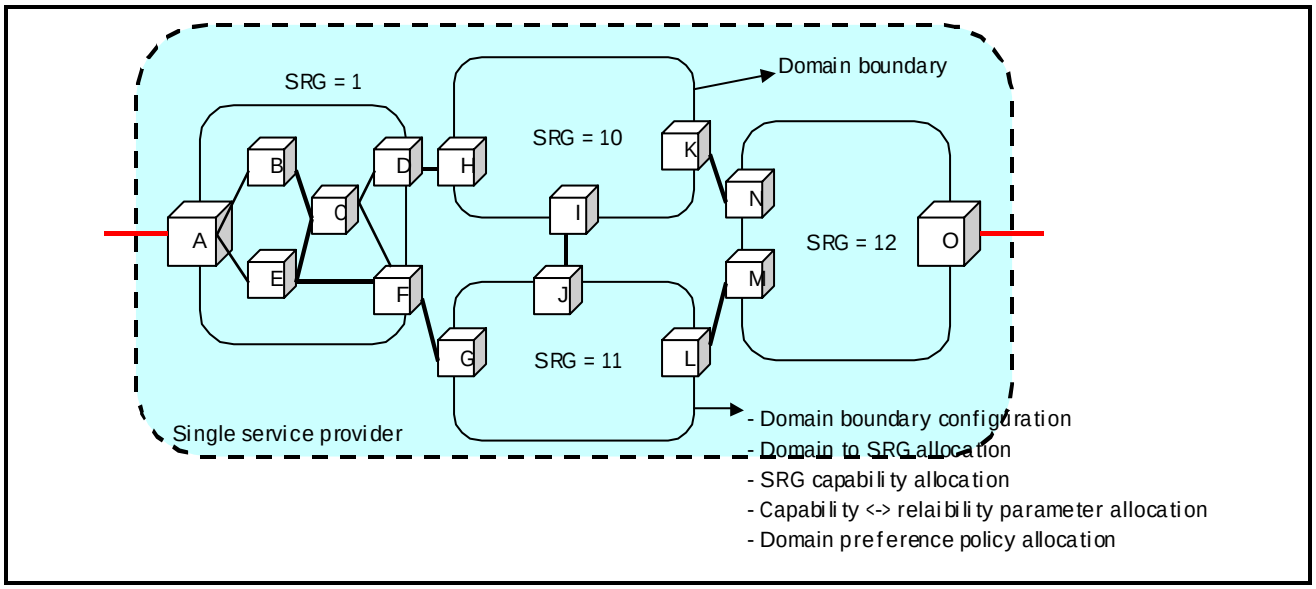

Figure 4 Configuration in a distributed path computation across multiple domains using SRG

Once the network is operational, the topological and domain information is disseminated into the network using the routing protocols, as shown in Figure 5A. As shown in the figure, topology and TE databases host the information about the interconnection of the network elements and their capabilities, respectively. Note that abstraction can be achieved in either of the databases. A domain is represented as an abstract point-to-multipoint link for the sake of topology representation and for path computation in the routing protocols. This is exemplified in Figure 5B for the risk domain represented by SRG 11. The exit points of the domain are interconnected by abstract links called “risk links”. This notation of abstract links helps in loose path computation. The capabilities of the risk domain and their current values for each of the risk links are captured in the TE database as shown in Figure 5C. Whenever there is a change in the topology or in the capability information, it is propagated throughout the network by the routing protocols.

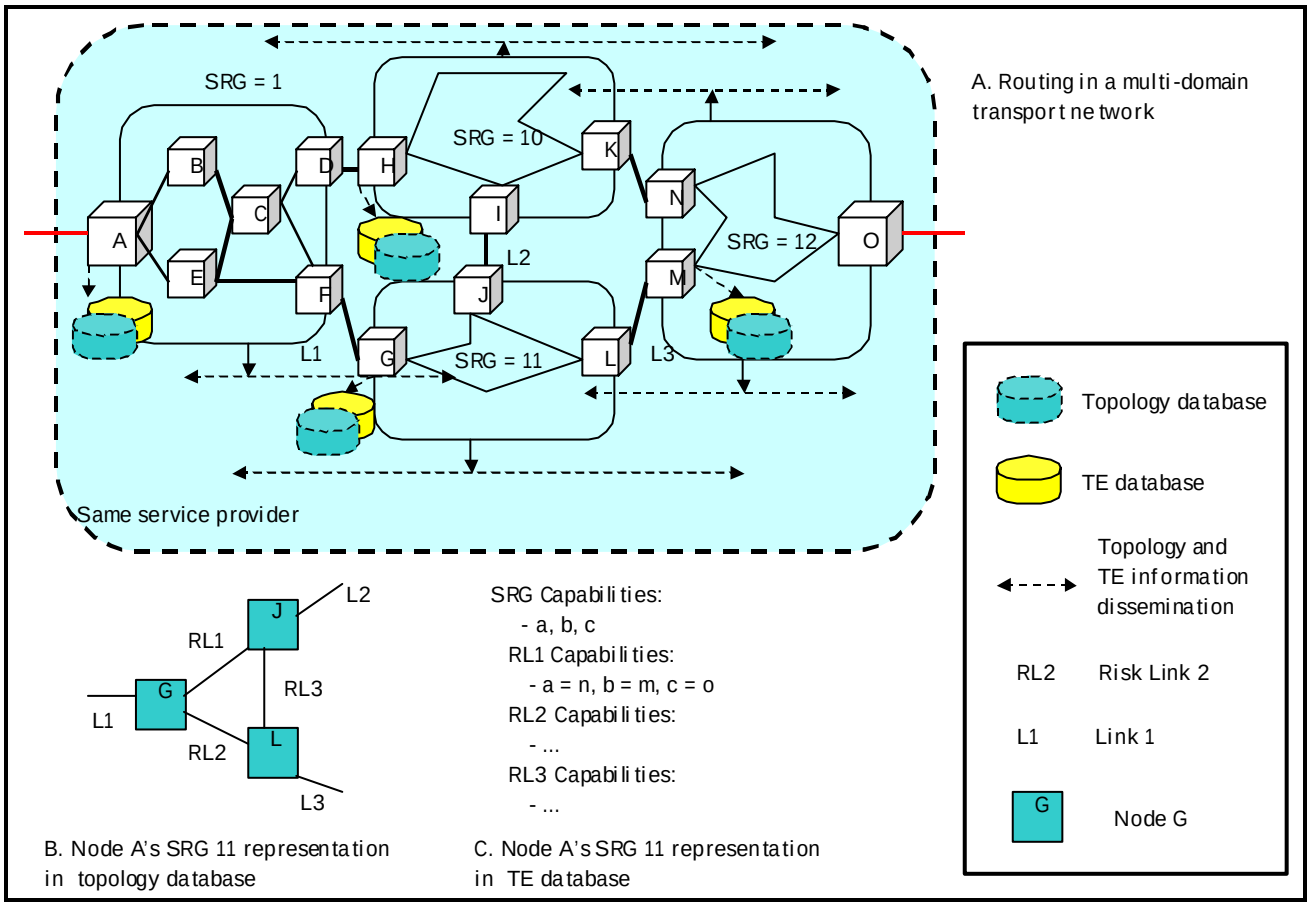

Figure 5 Routing in a distributed path computation across multiple domains using SRG

Once the topology and TE databases are created, the diversity computation module can attend to any constraint-based path request. Assume that a request is made at node A to compute a constraint-based path between A and O, as shown in Figure 6. Since the complete topology between A and O is not known at A, A can compute strict explicit path in SRG 9 but only loose explicit paths through SRGs 11 and 12 (or SRGs 10 and 12, or SRGs 10, 11, 12). The selection of intermediate domains such as (11, 12) or (10, 12) or (10, 11, 12) is based on the compatibility of the domains with the constraints posed by the connection request. For the sake of argument, assume that the domain capabilities are not passed across domains. In such a case, the path computation mechanisms has to make an arbitrary selection of SRGs. During the path setup, if the selected domain does not satisfy the constraints then it has to crankback to the originating node (in this case node A). The computed path, for example, <A, E, F, G, L, M, O> is then sent as a response to be setup either by the network management system or using the signaling protocols. In case of a signaling protocol, the strict path is setup from A to G via E, F, but the loose segments G, L and M, O cannot be expanded. Hence, when the signaling request reaches G, it is expanded by node G for the SRG 11 domain by referring to its topology and TE databases. Similarly, node M expands for the loose segment M, O and then forwards the signaling request.

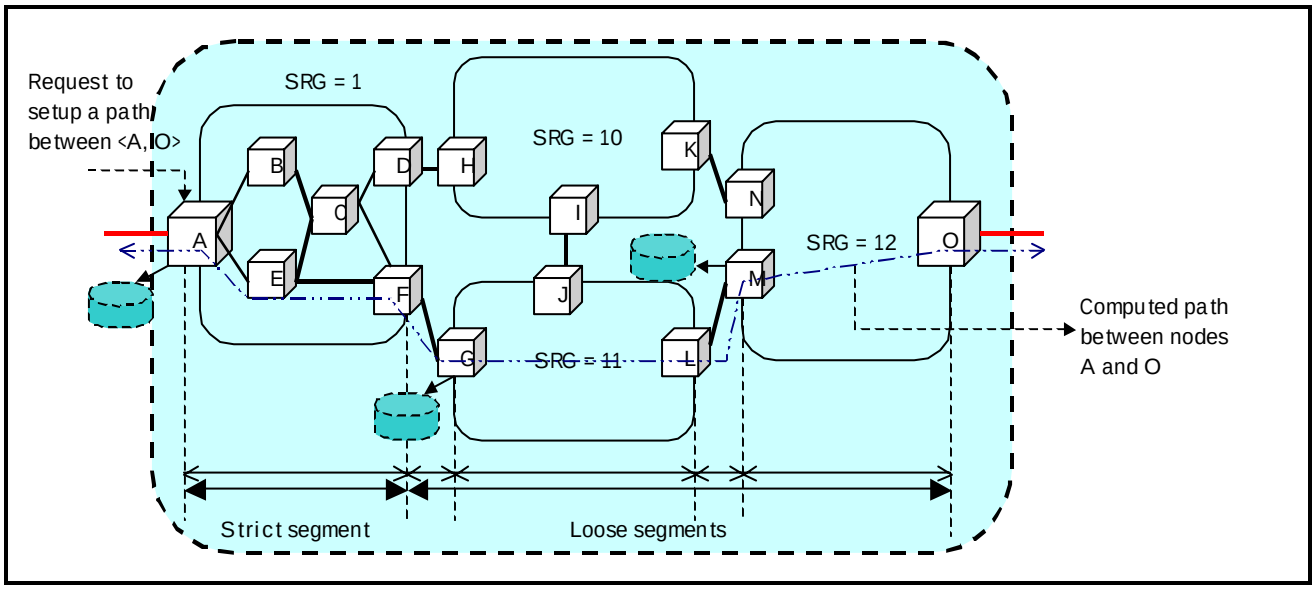

Figure 6 Path computation and signaling in a distributed path computation across multiple domains using SRG

8. Conclusions

In this paper the diversity principles of link, node and SRLG are extended to domain diversity using SRG concept. Many real-world applications are presented to motivate the concept of SRG. We introduced the assignment of capabilities to domains and demonstrated how this can be used in achieving diversity with both inclusive and exclusive constraints. We then argued how transport level diversity can be used in diverse path computation using the SRG concept. Risk associated with a path can also be assessed using the reliability parameters that are assigned to an SRG. At the end, we glued all these concepts together for a distributed path computation application.

9. Acknowledgements

The authors acknowledge Dimitri Papadimitriou of Alcatel, Riad Hartani of Caspian Networks, Curtis Brownmiller of Worldcom, John Strand of AT&T, Vishal Sharma of Metanoia, and Greg Bernstein of Ciena for their valuable comments as co-authors of some of these ideas at IETF and OIF standard body meetings.

10. References

1 Daniel O. Awduche et al., “A Framework for Internet Traffic Engineering,” draft-ietf-tewg-framework-04.txt, work in progress.
2 Ramesh Bhandari, “Survivable networks – Algorithms for diverse routing,” Kulwer Academic Publishers.
3 D. Papadimitriou et al., “Inference of Shared Risk Link Groups,” draft-many-inference-srlg-00.txt, work in progress.
4 Anagiotis Sebos, Jennifer Yates, Gisli Hjalmtysson and Albert Greenberg, “Auto Discovery of Shared Risk Link Group”, OFC 2001.
5 K. Kompella et al., “OSPF Extensions in Support of Generalized MPLS,” draft-kompella-ospf-gmpls-extensions-01.txt, work in progress.

6 Kireeti Kompella et al., “IS-IS Extensions in Support of Generalized MPLS,” draft-ietf-isis-gmpls-extensions-02.txt, work in progress.
7 E. Mannie et al., “Generalized Multi-Protocol Label Switching (GMPLS) Architecture,” draft-many-gmpls-architecture-00.txt, work in progress.
8 John Strand and Angela Chiu, “Issues for Routing in the Optical Layer”. IEEE Communication Magazine, February, 2001.
9 S. Dharanikota et al., “Inter domain routing with SRGs – Protocol extensions,” draft-many-ccamp-srg-01.txt, work in progress.

11. Author’s biography

Sudheer Dharanikota: Sudheer Dharanikota obtained his Master of Engineering from Indian Institute of Science (IISc) in 1990 and Ph.D. from Old Dominion University in 1997. He worked as a scientific officer in ERNET at IISc for two years on many networking technologies. After his Ph.D, he worked at Racal Datacom as a Manager of Routing, Bridging and Frame Relay compression groups from 1996-1997. Then he worked at Alcatel USA as a manager in many data products, including RCP 7770 - a 640 Gbps core router, from 1997-2000. He is also a research associate professor at Old Dominion University. He is currently working at the capacity of a systems architect at Nayna Networks, addressing the data over optical related issues. He has 5 patents pending in the networking area and has many research papers to his credit. He is a member of IEEE, ACM and Phi Kappa Phi.

Raj Jain: Raj Jain is a Co-founder and Chief Technology Officer of Nayna Networks, Inc. He is currently on a leave of absence from the Ohio State University, Columbus, Ohio, where is a professor of Computer and Information Science. He is a Fellow of IEEE and a Fellow of ACM. He is on the Editorial Boards of Computer Networks: The International Journal of Computer and Telecommunications Networking, Computer Communications (UK), Journal of High Speed Networks (USA), and Mobile Networks and Applications. He is currently a Distinguished Lecturer for the IEEE Communications Society and is on Technical Advisory Boards of several companies. He is the author of ``Art of Computer Systems Performance Analysis,'' published by Wiley and winner of the 1991 ``Best Advanced How-to Book, Systems'' award from Computer Press Association. His second book ``FDDI Handbook: High-Speed Networking with Fiber and Other Media'' was published in 1994 by Addison Wesley. His papers and publications can be found at http://www.cis.ohio-state.edu/~jain/.

Yong Xue: <>