

A Series of Hamiltonian Cycle Based Solutions to Provide Simple and Scalable Mesh Optical Network Resilience

Hong Huang and John A. Copeland, Georgia Institute of Technology

Abstract

Mesh optical resilience solution has many requirements, some of which are conflicting with others. This article highlights the issues relevant to mesh optical resilience and the challenges of meeting the myriad requirements. A series of Hamiltonian cycle based solutions with different efficiency, complexity, and scalability tradeoffs are introduced. Solutions for moderate-sized and large networks are differentiated with the former emphasizing simplicity and efficiency and the latter stressing resilience from multiple simultaneous failures and isolation of failure events.

Index Terms—Optical networks, network resilience and survivability, Hamiltonian cycle.

I. INTRODUCTION

Resilience from network element failures is an important concern for optical networks because of their carrier class requirements and the possibility of losing large amount of data when failure events occur. Traditional optical networks are ring networks that have simple protection mechanisms, such as SONET self-healing rings. Future optical networks are envisioned to be mesh networks because of their flexibility and efficiency. However, designing a mesh optical network resilience scheme is no easy task because of the complexity of mesh topology and the drive to mine the most possible efficiency from it.

There are many considerations in designing a viable mesh optical network resilience scheme. Efficiency is important, which can be measured by the spare to primary capacity ratio. Recovery time constraint is a requirement, with the current accepted value around 50 ms [1]. Simplicity is desired, since complexity not only causes difficulties in management and control but also compromises robustness. Scalability is another virtue, since an optical network ranges from the metro, the regional, the national, to the global scope. In a large network, the ability to recover from multiple failures and to prevent failure events having a global influence scope is desired or even required. With the convergence of optical and IP networks, the requirement that an optical network resilience scheme should be transparent to the IP layer becomes important, which means that it should be possible to run IP-related protocols such as IGP, BGP, and MPLS over optical networks without major modification to account for

the resilience issues. Transparency of network resilience suggests decoupling the management and control of primary and spare resources, so that IP protocols can run in the space of the primary resources and be oblivious to spare resources provided by the optical layer.

In the following, we review the previously proposed mesh optical network resilience schemes and highlight the difficulties of meeting the requirements outlined above. We introduce our solutions by first discussing the concept of Hamiltonian cycle based resilience schemes and then presenting a solution for moderate-sized networks and three solutions for large networks.

II. CHALLENGES OF MESH OPTICAL NETWORK RESILIENCE

The recovery of a network from element failures can be performed dynamically or statically. The IP resilience mechanism is a familiar example of dynamical resilience schemes. Although conceptually simple, IP resilience has limited use for a carrier class network because of its indeterministic nature (spare resources are not guaranteed) and slow recovery time (on the order of seconds).

Dynamic resilience schemes can be made deterministic by using connection oriented routing and ensuring each primary connection is backed by a secondary connection [2]. Such schemes are made efficient by sharing resources on the secondary connections whose primary connections do not pass the same network elements. However, those schemes make the already difficult problem of constrained routing, which seeks to ensure primary connections have guaranteed resources, much more harder because of the need to route both the primary and the secondary connections and the nontrivial burden to keep track which network element belongs to which secondary connection with respect to which primary connection. In addition to complexity, such schemes are not transparent to IP routing protocols, which have no explicit provision for backup routes.

Static resilience schemes, which are the focus of this article, preconfigure spare resources according to traffic profile or primary network capacity. They can be broadly categorized as ring-based and optimization-based schemes. Ring-based schemes [3][4] are natural extensions of SONET self-healing rings in mesh networks. The concept is to cover a mesh network with a series of self-healing rings with some variations. As such, those schemes are simple and fast. However, they inherit the same inefficiency of SONET self-

healing rings and require the same amount of spare capacity as the primary one.

Optimization-based schemes typically make use of traffic profiles and formulate the provision of the primary and spare resources as an optimization problem [5][6]. There are a couple of variants depending on the constraints imposed. Recover path can be constrained to be between the two endpoints of the connection (end-to-end recovery), between the two points adjacent to the failure (link recovery), or between any two points upstream and downstream along the primary path (general path recovery). A noteworthy scheme was proposed in [7], which has the constraint that its recovery paths are cycles. A distinctive feature of the scheme is that it makes its spare resources available not only to the primary links on the protection cycles but also to the ones straddled on the cycles. Such characteristic makes it efficient and is shared by the Hamiltonian cycle based schemes.

Optimization-based schemes are efficient in network resource usage. However, they have complicated recovery configurations, incur significant signaling delay, and require solving a NP-hard optimization problem, which is time consuming for moderate-sized networks and impractical for large networks. The solution obtained is predicated on the traffic profile that is used as input. Inaccuracy in traffic predication makes the solution questionable and changes in traffic profile entail reoptimization and reconfiguration.

The challenge of mesh optical network resilience is that the many requirements are often in conflict. Efficient schemes, such as optimization-based ones, often involve complicated recovery configurations, are typically slow, and require joint-optimization of primary and spare routes that is not transparent to IP protocols. Fast and simple schemes, such as ring-based ones, are usually inefficient. Clearly, it is desirable to have some solutions that obtain appropriate balance among the myriad requirements of mesh optical network resilience. Our work is an attempt in this direction. In addition, little attention in the research community has been directed to the issues of resilience from multiple failures and isolation of failure events, which are particularly relevant to large networks and are addressed in this article.

III. THE CONCEPT OF HAMILTONIAN CYCLE RESILIENCE

The Hamiltonian cycle based resilience solutions are based on the following observation. We can conceptually partition a resilient network into a primary and a spare network. Working connections are routed in the primary network, and the spare network provides protection from failures. The spare network must span all nodes in the network; otherwise some nodes are excluded to access spare resources. In addition, the spare network must be resilient itself if primary and spare resources are not route-disjoint, since a failure can disrupt both primary and spare resources and make the spare network disconnected. In the case of single link failure, the

resilience requirement dictates that the topology of the spare network must be two-link-connected, which means that removal of one link from the network does not cause the network to be disconnected. We do not consider node failures here, since a node failure can be viewed as failures of the multiple links that are incident on the node. In addition, node failure can be best protected by redundancy and high availability modules inside the node instead of providing redundant link capacity in the network.

A viable spare network should span all nodes and be two-link connected. However, such specification still leaves us countless possible choices for spare network topology. We posit that a Hamiltonian cycle is a good choice. A Hamiltonian cycle is a cycle that traverses each node in the network exactly once. In addition to being a viable spare network topology (spanning all nodes and being two-link connected), it has the characteristic of having minimal number of links among all survivable topologies. Having minimal number of spare links has two consequences. First it reduces spare link capacity. Second, maybe more important, it collapses spare resources into a minimal set of links, which facilitates aggregate protection switching upon network element failures. Aggregate switching means switching in a larger granularity such as fiber or waveband, as opposed to a smaller granularity such as wavelength or digital circuits below the wavelength level. The cost of protection switching is a significant factor in the total cost to provide network resilience. Switching in aggregates (fiber or waveband) can significantly reduce the number of spare switch ports. While aggregate switching in the primary network, which manifests itself as the traffic grooming problem, has been addressed in the literature, its importance in the spare resource dimension seems not be adequately recognized. Lastly, Hamiltonian cycle based resilience solutions segregates network resources into a primary network and a spare network, which are managed and controlled separately. IP routing protocols can run in the primary network with network resilience transparently provided by the spare network. What the spare network protects is the primary network and is not tied with any particular traffic profile. In fact any traffic profile that is admissible to the primary network is protected by the spare network.

To apply Hamiltonian cycle based resilience to practical optical networks, we bifurcate our discussion into two settings: moderate-sized and large optical networks. The two types of networks have very different characteristics and warrant differentiated approaches. A moderate-sized optical network, typically a metro or regional network, is limited in geographical range, has homogeneous network elements, e.g. a pure 32 wavelength system with no 16 or 64 wavelength elements mixed in, is in the scope of a single control domain, and probably consists of network elements coming from a single vendor or a tight vendor alliance. In moderate-sized networks, resilience from single failure is usually adequate because of its limited scope and the high cost of providing resilience from multiple failures. On the other hand, a large

optical network, typically a national or global network, is inhomogeneous, i.e. network elements with different number of wavelengths and spectrum allocations are intermixed, has multiple control domains, and may consist of network elements from a wide range of vendors that do not talk to each other easily. Further, the scenario that multiple failures occur at the same time becomes plausible

IV. HAMILTONIAN CYCLE RESILIENCE IN MODERATE-SIZED NETWORKS

We start by formally stating two theorems that form the basis of Hamiltonian cycle based resilience, the proof of which can be found in [8]:

Theorem 1: A spanning tree of a homogeneous working network that can survive a failure is a minimal set of sufficient spare links to provide protection from single link failure.

Theorem 2: If primary and spare links are not disjoint, a Hamiltonian cycle of a homogeneous Hamiltonian primary network is the minimal set of sufficient spare links to provide resilience from single link failure.

The above two theorems implied two scenarios depending on whether the primary and spare links take disjoint routes. If the two kinds of links are disjoint, we can provision a spanning tree of the primary network as the spare network according to Theorem 1. If the primary and spare links are routed together (in the same cable or conduit), as in most practical cases, a link failure can both disrupt the primary network and disconnect the spanning tree. In such case, we proposed a Hamiltonian cycled based resilience scheme (HCP) [8] based on Theorem 2. With HCP, the primary network is a mesh network and the spare network is a Hamiltonian cycle of the primary network. When a network element fails, the traffic is diverted from the primary network to the spare network at the upstream node adjacent to the failure and switched back to the primary network at the downstream node adjacent to the failure. Because of the homogeneity, primary and spare capacities always match. Because of the spanning nature of Hamiltonian cycle, any two disconnected nodes can be rejoined using the spare network.

HCP essentially uses a spare ring (the Hamiltonian cycle) to protect a mesh primary network. The mechanism of HCP is similar to that of SONET self-healing rings and only two nodes adjacent to the failure are involved in the recovery, so HCP is simple and robust. To measure the spare capacity efficiency of a resilience solution, we introduce a metric: Spare to Primary capacity Ratio (SPR) defined as $SPR = \text{Spare Capacity} / \text{Primary Capacity}$. In a homogeneous network, $SPR = \text{Number of Spare links} / \text{Number of Primary Links}$. In the case of HCP, given a primary network of N nodes and E links, the number of spare links is N , so $SPR = N / E = 2 / d$, where d is the average degree of the network

(average number of links incident on a node). In a typical mesh telecommunications network, d typically is in the range [2.5, 4.5], therefore the SPR of HCP is in the range [0.44, 0.80]. The SPR of HCP is significantly lower than that of ring-based schemes, which is 1.0, and is close to those obtained by optimization-based schemes [5][6]. In fact, the more densely connected and larger the network (in terms of number of nodes), the more efficient HCP becomes. For a fully meshed primary network, HCP has $SPR = 2 / (N-1)$, which decreases drastically with increasing value of N . For instance, the SPR of HCP takes the values of 0.67, 0.29, 0.13, 0.06, when N takes the values of 4, 8, 16, 32, respectively. A small value of SPR means a small portion of capacity is used for spares.

A feature of HCP is that the recovery target is the failed link, not the individual connections. When failure occurs, the traffic in the failed link is switched over as a whole. As such, protection switching can be carried out in the largest granularity possible, which greatly reduces the number of spare switch ports and simplifies failure recovery. This is important, since in metro or regional networks, switching cost play a dominating role in the whole cost structure.

An interesting application of HCP is in the migration of ring networks to mesh networks in a metro or regional setting. A legacy ring network has a primary ring and a spare ring. Using HCP, links can be added between nodes in the primary ring to construct a mesh topology with increased throughput; while network resilience is guaranteed by the original spare ring with no extra cost.

However, HCP has some limitations. It requires that the mesh primary network has a Hamiltonian topology, i.e. a Hamiltonian cycle exists. While many survivable topologies are Hamiltonian, e.g., the much-studied NSF backbone and pan-Europe prototype networks, some are not. Although topology can be chosen by design, HCP has more serious defects: it does not apply to inhomogeneous networks, cannot recover from multiple failures, and has no failure isolation, which limit its use in large networks.

V. HAMILTONIAN CYCLE RESILIENCE IN LARGE NETWORKS

In this section, we introduce three solutions to deal with problems specific to large networks, which are network inhomogeneity, multiple simultaneous failures, and isolation of failure events.

A. Hamiltonian Cycle for Inhomogeneous Networks

A crucial issue for preconfigured resilience schemes is the dimension of spare capacity. For a homogeneous network, the link capacity of the spare Hamiltonian Cycle (HC) is trivially determined, which is equal to that of the primary links. For a large inhomogeneous network, link capacities can be different, e.g. some links use 32 wavelengths and others use 64 wavelengths. The question is how to provision the

capacity of the HC. The answer is that it should be large enough to accommodate worst-case scenario, i.e. the failure of the link with the largest primary capacity. The capacity of the HC (bp_s) is provisioned according to the formula: $bp_s = \max [bw_j, 0.5 bw_i]$, over all i, j , where bw_j indicates the capacity of primary link j that is on the HC and bw_i indicates the capacity for primary link i that is not on the HC, the so-called straddled links [7]. The reason is the following: for a primary link on the HC, there is only one restoration path, i.e. looping back on the other direction of HC. So a HC can accommodate a primary link failure with a capacity no larger than that of HC. On the other hand, for a primary link that is not on HC, there are two restoration paths, each looping back on opposite directions on the HC. Because of this restoration path diversity, a HC can accommodate a primary link failure with twice the capacity of HC. The above situation is illustrated in Figure 1, where solid lines indicate primary links and dashed lines indicate the HC. The solution outlined above is abbreviated as SHI (Single Hamiltonian resilience for Inhomogeneous networks). The numbers on the links indicate primary link capacity in number of wavelengths. The capacity of HC is set to be $3 = \max [3, 2, 6/2, 4/2]$. When link a-d, which is on the HC, fails, there is only one restoration path on the HC: a-b-c-d. However, when link a-c, which is not on the HC, fails, there are two restoration paths: a-d-c and a-b-c. It is easy to check that the capacity of HC is adequate to restore failed links in both cases.

Except for spare capacity provision, SHI operates in the same way as HCP. SHI works when network inhomogeneity is not pronounced. However, when link capacities are widely different, SHI becomes inefficient since the capacity of HC is provisioned according worst-case scenario. In addition, SHI still uses a single HC to protect a mesh network, does not deal with multiple failures, and has no failure isolation.

B. Hamiltonian Cycle Cover

In a large-scale optical network, the possibility of multiple simultaneous failures becomes distinct and it is desirable to isolate failures within a limited influence scope. To confront such situation, we propose a multi-domain resilience strategy, where each domain responds to and recovers from failures within its border. Using a multi-domain solution, failures do not propagate to the global scope, network management is simpler, and multiple failures are dealt with as long as only one failure occurs in each domain.

Hamiltonian Cycle Cover (HCC) is a multi-domain solution based on Hamiltonian cycles. The basic idea is to partition a large mesh network into a set of protection domains, and use one HC to protect each domain. The partition of protection domains could be based by a number of criteria, such as geographical scope, carrier boundary, or same wavelength systems, etc. Care should be taken to ensure relative capacity homogeneity within a domain so that the inefficiency of worst-scenario provision does not arise.

The capacity of a HC in each domain is provisioned in the same way as before, except for links that intersect multiple

domains. In fact, for planner topology, which is our main interest here, a link can only intersect two domains. We provision the capacity of an intersecting link j according to the formula: $bp_j = \max [bp_s]$, where the max operation is taken between the neighboring HCs. Figure 2 is an illustration of HCC. The network is partitioned into two domains, A and B. Each domain is protected by a HC with capacity of 3 and 2, respectively. The spare capacity of the link intersecting domain A and B is provisioned with a capacity of $3 = \max [3, 2]$.

Spare resource configuration in HCC involves the following steps:

- A. Partition the primary network into a set of protection domains.
- B. Find a HC for each domain.
- C. Provision the HC for each domain s with capacity $bp_s = \max [bw_j, 0.5 bw_i]$, where bw_j denotes the primary link capacity for link j that is on the HC, and bw_i denotes the primary link capacity for link i that is not on the HC.
- D. For those links that intersect HCs, their capacity is modified according to the formula: $bp_j = \max [bp_s]$, where the max operation is taken among neighboring HCs intersecting on the link j .
- E. The result is a protection capacity provision, which consists of a set of HCs, each responsible for protecting its local domain.

Muti-domain resilience incurs a cost: the partitioning or intersecting links. While SHI needs N spare links, HCC requires $N + M - 1$ spare links, where N is the number of nodes in the network and M is the number of domains. In addition, protection switching in HCC is more complex than that of SHI at nodes that sit on the borders of domains, because of the need to distinguish failures in different domains.

Aside from the overhead that is necessary to implement multiple protection domains, HCC shares many characteristics of SHI (or HCP). It is simple and robust, employs the minimal number of spare links within each domain, facilitates aggregate protection switching, decouples the primary and the spare networks, and provides network resilience that is transparent to the routing of primary connections. What HCC excels over SHI or HCP is its ability to scale to large networks, to isolate failures, and to protect from multiple simultaneous failures. In addition, the capacity of HC for each domain is provisioned to be adequate for its local domain only, which avoids the inefficiency of global worst-case provision as in the case of SHI. Further, since they are entirely contained within the local domain, recovery paths of HCC are much shorter than those of SHI, which translates into faster restoration.

C. Hamiltonian Cycle Neighbor Capacity Closure

Hamiltonian cycle Neighbor capacity Closure (HNC) is motivated by on the following observation. The partitioning or intersecting links serve two purposes: domain isolation and capacity closure. The first purpose is obvious. The second

one is to bridge the capacity differences between domains. While domain isolation is a value add; capacity difference bridging is of necessity. HNC is an attempt to further reduce spare resource on the intersection links by relaxing strict domain isolation. It is very similar to HCC, the only difference being in step D. Under HNC, the capacity of intersecting links is calculated in the following fashion: $bp_j = \text{abs}(bp_{s1} - bp_{s2})$, where $s1$ and $s2$ are the two neighboring domains and $\text{abs}(\)$ is the absolute value function. The justification is the following. If we remove the intersecting link, the two HCs in the neighboring domains are incomplete, each missing a link with capacity bp_{s1} or bp_{s2} , respectively. Each domain can complete its HC by adding back a path connecting the two nodes disconnected by the removal of the intersecting link. In fact, we can add back multiple paths as long as their aggregate capacity equals that of the HC. We actually have two paths: the intersection link and the path through the neighboring HC. For the HC with the smaller capacity, the path in the neighboring HC is more than adequate. For the HC with the larger capacity, we only need to use the intersecting link to account for the capacity difference between the two neighboring HCs. Figure 3 provides an illustrating example. Domain A, which has spare capacity 3, can use domain B, which has spare capacity 2, to provide one part (2) of its spare capacity requirement. The remaining part of domain A's spare capacity (3-2=1) can be satisfied by that of the intersection link.

The advantage of HNC is that it further reduces the amount of spare resource at intersecting links by an amount that is the difference between $\max[bp_{s1}, bp_{s2}]$ and $\text{abs}(bp_{s1} - bp_{s2})$. The downside is that restoration is no longer local to the domain where failure occurs, i.e., part of spare capacity for a local failure resides in other domains. This reduces network resilience from multiple failures, violates failure isolation, and lengthens the restoration path. In addition, protection switching is more complex, since each node needs to deal with not only the protection switching within its local domain, but also that propagated from other domains.

D. An Example Simulation Study

To study the performance of various resilience solutions, we provide some numerical results for a national scale network with 21 nodes, 29 links, and partitioned into four domains. The performance metric used is SPR, the spare to primary capacity ratio. The solutions compared are a ring-based scheme using ring cover (RC), a optimization-based scheme using integer linear programming (ILP), and three HC based schemes: SHI, HCC and HNC. The results are shown in Figure 4.

Using the result obtained with ILP, which defines the minimum spare capacity achievable, as a benchmark, HNC, HCC, SHI and RC require approximately 12%, 22%, 42%, and 54% more spare capacity, respectively. Clearly ILP is the most efficient scheme in terms of spare capacity requirement but suffers from complexity, slowness in recover, no failure isolation, and incapability to recover from multiple

simultaneous failures. RC is the least efficient scheme but is fast, simple and robust. In between lies our proposed three schemes: SHI, HCC and HNC. SHI seems not be a viable solution, since it is not efficient and does not off other value-adds, such as failure isolation and recovery from multiple failures. HCC and HNC could be contenders, since both are fast, simple, and reasonably efficient, with the tradeoff between strict failure isolation and simplicity of HCC and efficiency of HNC.

VI. CONCLUSION

Mesh optical network resilience exemplifies a complex engineering problem, where many requirements, often conflicting with each other, are in the play. A viable solution needs to achieve a sensible balance among many requirements, be specialized for the type of target network (a moderate-sized or a large network), and, especially for large networks, have the ability to recover from multiple failures and isolate failure events. Our work presented in this paper is an attempt in this direction, and we hope it can evoke further thinking and probably better solutions.

REFERENCES

- [1] S. Ramaswami, and K. N. Sivarajan, *Optical networks: a practical perspective*, Morgan Kaufmann, 1998
- [2] R. Doverspike and J. Yates, "Challenges for MPLS in optical network restoration," *IEEE Commun. Mag.*, vol. 39, no. 2, pp. 59-96, 2001.
- [3] G. Ellinas and T. E. Stern, "Automatic protection switching for link failures in optical networks with bidirectional links," in *Proc. IEEE GLOBECOM*, vol.1, pp. 152-156, 1996.
- [4] M. Medard, S. G. Finn, and R. A. Barry, "WDM loop back recovery in mesh networks" in *Proc. IEEE INFOCOM*, vol. 2, pp. 752-759, 1999.
- [5] B. Van Caenegem, W. Van Parys, F. De Turck, and P.M.Demeester, "Dimensioning of survivable WDM networks," *IEEE J. Select. Area Comm.*, vol. 16, no.7, pp. 1146-1157, 1998.
- [6] S. Ramamurthy and B. Mukherjee, "Survivable WDM mesh networks, part I - protection," in *Proc. IEEE INFOCOM*, vol.2, pp. 744 -751, 1999.
- [7] W. D. Grover and D. Stamatelakis, "Cycle-oriented distributed preconfiguration: ring-like speed with mesh-like capacity for self-planning network restoration," in *Proc. IEEE ICC*, vol.1, pp. 537-543, 1998
- [8] H. Huang and J. Copeland, "Hamiltonian cycle protection: a novel approach to mesh WDM optical network protection," in *Proc. IEEE HPSR*, pp. 31-35, 2001

BIOGRAPHIES

Hong Huang (hong.huang@csc.gatech.edu) received his B. Engr. degree from Tsinghua University, Beijing, China, and M.S. degree in electrical and computer engineering from Georgia Institute of Technology. Currently, he is a Ph. D. candidate in the School of Electrical and Computer

Engineering at Georgia Institute of Technology. His research interests include optical networks, MPLS, and IP quality of service.

John A. Copeland is currently John H. Weitnauer, Jr. Technology Transfer chair, GRA eminent scholar, director of the Communication Systems Center, and professor in the School of Electrical and Computer Engineering at Georgia

Institute of Technology. He is a fellow of IEEE, a recipient of IEEE Morris N. Liebmann Award, served as past editor of the IEEE Transactions on Electron Devices, and was on the Board of Trustees of the Georgia Tech Research Corporation. He received his Ph.D. degree from Georgia Institute of Technology. His research interests include computer networks, digital CATV networks, computer architecture, operating systems, and optical networks.




Figure 1. An illustration of SHI.




Figure 2. An illustration of HCC.




Figure 3. An illustration of HNC.




Figure 4. A PSR comparison of various resilience solutions