Reducing Spare Capacity Through Traffic Splitting

Andrew Zalesky, Hai Le Vu, *Member, IEEE*, and Moshe Zukerman, *Senior Member, IEEE*

***Abstract*—By splitting traffic across an optimal number of disjoint paths, it is shown a capacity saving of up to 11% is possible in randomly generated networks relative to the most capacity efficient mesh-based restoration scheme available to date. An algorithm yielding the optimal number of disjoint paths is developed.**

***Index Terms*—Mesh restoration, survivable network, traffic splitting.**

I. Introduction

SAVING transmission capacity is a key force driving the design of survivable all-optical lightpath networks. By survivable, it is meant resiliency to a single cable-cut is provided through some form of restoration. Switch failures and multiple cable-cuts are not considered. Mesh-based restoration [3] can offer up to a 50%–60% capacity saving relative to ring-based [2] and dedicated one-plus-one restoration [4]. For this reason, it is envisaged that mesh-based restoration will underpin future survivable networks.

Mesh restoration of lightpaths (MRL) [2] is the most capacity efficient mesh-based restoration scheme available to date, which functions as follows. Traffic is transmitted from an ingress point to an egress point by way of a *working lightpath*. When a working lightpath is severed by a cable-cut, its traffic is split across a set of disjoint *spare lightpaths*. Spare lightpaths are configured only when a working lightpath is severed. The capacity allocated to a spare lightpath is referred to as *spare capacity*, and similarly for *working capacity*. Spare capacity can be shared between spare lightpaths so long as sufficient spare capacity is available to restore all working lightpaths severed by any single cable-cut. Working capacity is dedicated to each working lightpath. The capacity efficiency of mesh-based restoration is attributable to sharing of spare capacity. Working and spare lightpaths are often routed to minimize total capacity, which is the sum of working and spare capacity [2], [3].

Minimizing total capacity demanded by a *dynamically* provisioned survivable lightpath network is the optimization problem considered herein. Three properties distinguishing dynamic provisioning are as follows.

1) lightpath requests arrive in real time;
2) reconfiguration of existing working and spare lightpaths to enforce optimality is forbidden;
3) statistics of future lightpath requests are unknown.

Manuscript received January 26, 2004. The associate editor coordinating the review of this letter and approving it for publication was Prof. M. Ma. This work was supported by the Australian Research Council.

The authors are with the ARC Special Research Centre for Ultra-Broadband Information Networks, an affiliated program of national ICT Australia, Department of Electrical and Electronic Engineering, The University of Melbourne, Melbourne, Vic. 3010, Australia (e-mail: a.zalesky@ee.mu.oz.au).

Digital Object Identifier 10.1109/LCOMM.2004.833800

Upon arrival of a request, a working lightpath and a set of disjoint spare lightpaths must be optimally routed and provisioned within a few seconds. As a consequence of the above three properties, routing lightpaths for a new request must be performed independently of existing and future requests.

A previous study [4] considered provisioning a new request by routing a working lightpath and single disjoint spare lightpath that is optimal at the time of provisioning. Optimality is attained by minimizing the sum of capacity on the working lightpath and *unshared capacity*[1] on the spare lightpath. Optimality may however be compromised if existing lightpaths are dismantled or new requests provisioned subsequent to the time of provisioning. A signalling protocol is developed to update each link state upon provisioning of a lightpath request, which is compatible to the concept later proposed herein.

An integer linear program has been formulated [3] to solve the allied problem of minimizing total capacity demanded by a statically provisioned survivable lightpath network. Requests are known and optimally provisioned a priori. Solving an integer linear program upon arrival of a dynamic request within a few seconds is computationally intractable and may require reconfiguration of existing lightpaths to enforce optimality.

Herein a request is dynamically provisioned by uniformly splitting traffic across an optimal number of disjoint lightpaths. When a lightpath is severed, its portion of traffic is uniformly split across the surviving lightpaths, which are provisioned spare capacity in addition to working capacity. It is shown such provisioning offers a capacity saving of up to 11% in randomly generated networks relative to MRL. The capacity saving is not solely attributable to sharing of spare capacity, but rather to a concept to be exposed in Section II.

Finding the optimal number of disjoint lightpaths to minimize total capacity is formulated as an optimization problem and solved in Section III. Cost savings are empirically quantified in Section IV.

II. Underlying Concept

Static provisioning allows spare lightpaths to be optimally routed to maximize sharing of spare capacity. The following intrinsic limitation precludes dynamic provisioning from fully exploiting the capacity saving attributable to sharing of spare capacity. Optimally routing a set of disjoint spare lightpaths for a new request to maximize sharing of spare capacity requires knowledge of future requests at the time of provisioning and may require reconfiguration of existing lightpaths. At best, it

[1]By unshared capacity, it is meant capacity that is not sharable with existing lightpaths. For example, if s_1 and s_2 are spare lightpaths for working lightpaths w_1 and w_2, respectively, s_1 and s_2 can share spare capacity if and only if w_1 and w_2 cannot be concurrently severed by any single cable-cut.

can only be hoped that spare lightpaths are routed such that a near-optimal level of sharing often prevails.

Unlike MRL, the capacity saving offered by the concept considered herein is not solely attributable to sharing of spare capacity, and is thus suited to dynamic provisioning. The intent of the following example is to expose the underlying concept.

Suppose λ units of traffic request transmission from an ingress point to an egress point for which $n(k), k=1,2,\ldots,K$, disjoint k-hop lightpaths can be provisioned. By k-hop lightpath, it is meant a lightpath traversing k fiber links. An arbitrary yet practical modeling choice that reflects a typical network is $n(k) = \theta \cdot k$, $\theta = 1,2,\ldots$, and K a small integer, say two or three to preclude excessively long lightpaths. In words, θ one-hop lightpaths, $2 \cdot \theta$ two-hop lightpaths, etc., up to $K \cdot \theta$ K-hop lightpaths can be provisioned. The modeling parameter θ represents network density. Increasing θ allows a greater number of disjoint lightpaths to be provisioned and models an increase in network density. Let $M = \sum_{k=1}^{K} n(k) = \theta \cdot K \cdot (K+1)/2$ denote the number of disjoint lightpaths that can be provisioned and let m, $m = 2,3,\ldots,M$, denote the actual number of *shortest hop* disjoint lightpaths provisioned upon arrival of a request.

To reflect the cost of provisioning a lightpath, define capacity such that $\lambda \cdot k$ units of capacity are required to transmit λ units of traffic across a k-hop lightpath. Less than $\lambda \cdot k$ units of capacity are required to transmit λ units of traffic across a k-hop spare lightpath if its capacity is sharable with existing spare lightpaths. Defining capacity as such provides an upper bound and represents the worst dynamic provisioning case, in which no spare capacity is sharable. Without loss of generality, assume $\lambda = 1$.

Upon arrival of a request, MRL provisions a one-hop working lightpath and s, $s = 1,2,\ldots,M-1$, shortest hop spare lightpaths requiring total capacity

$$\Lambda_{\theta,m} = 1 + \frac{(k^*+1)\cdot\rho(k^*) - 1}{m-1} + \sum_{k=1}^{k^*} \frac{k \cdot n(k)}{m-1} \quad (1)$$

where $m = s+1$ and $k^* = \max_k \rho(k)$ such that $\rho(k) \triangleq m - \sum_{j=1}^{k} n(j) \geq 0$. Rearrangement of (1) gives

$$\Lambda_{\theta,m} = \frac{6\cdot m\cdot(2+k^*) - \theta\cdot k^*\cdot(k^{*2}+3\cdot k^*+2) - 12}{6\cdot(m-1)}.$$

By provisioning *all* $M-1$ spare lightpaths, spare capacity is maximally dispersed throughout the network, thus maximizing the probability of sharing with existing spare lightpaths.

Consider now the concept of uniformly splitting traffic across m, $m = 1,2,\ldots,M$, shortest hop disjoint lightpaths. When a lightpath is severed, its portion of traffic is uniformly split across the surviving $m-1$ lightpaths, which are provisioned spare capacity in addition to working capacity requiring total capacity

$$\overline{\Lambda}_{\theta,m} = (k^*+1)\cdot\left(\frac{\rho(k^*)}{m} + \frac{\rho(k^*)}{m^2}\right) + \sum_{k=1}^{k^*}\left(\frac{k\cdot n(k)}{m} + \frac{k\cdot n(k)}{m^2}\right). \quad (2)$$

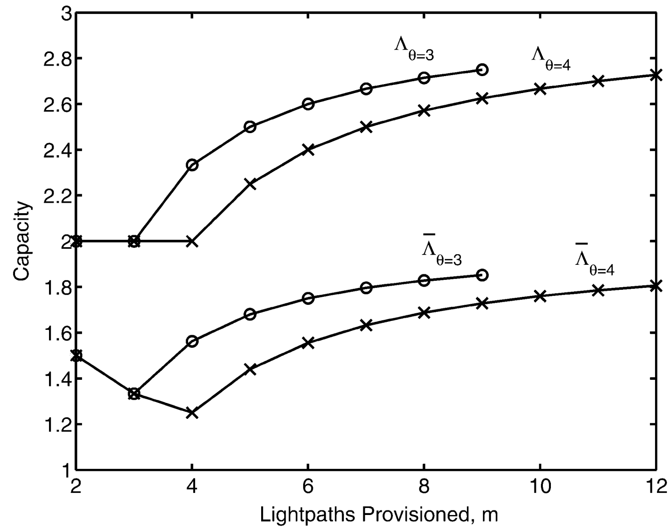

Fig. 1. An example to quantify the capacity saving offered by the underlying concept considered herein relative to MRL.

Rearrangement of (2) gives

$$\overline{\Lambda}_{\theta,m} = \left(\frac{\theta\cdot(m+1)\cdot(k^*+1)}{6\cdot m^2}\right)\left(\frac{6\cdot m}{\theta} - k^{*2} - 2\cdot k\right).$$

As conjectured earlier, it is straightforward to prove $\overline{\Lambda}_{\theta,m} \leq \Lambda_{\theta,m}$ for all $\theta = 1,2,\ldots$ and $m = 2,3,\ldots$. To quantify the capacity saving, a plot of $\Lambda_{\theta,m}$ and $\overline{\Lambda}_{\theta,m}$ for $K=2$ and the cases $\theta = 3,4$, reflecting low and high density networks, respectively, is shown in Fig. 1.

A potential limitation of the concept considered herein is that a capacity saving may require a tradeoff with propagation delay as a consequence of some data packets traversing longer working lightpaths. Increased propagation delay arises in MRL only during the time of a cable-cut because it is only during such a time that traffic is split across many lightpaths.

To motivate the ensuing section, observe in Fig. 1 that capacity can be *optimally* saved by splitting traffic across an optimal number of shortest hop disjoint lightpaths; for example, $\overline{\Lambda}_{4,4} \leq \overline{\Lambda}_{4,m}$ and $\overline{\Lambda}_{3,3} \leq \overline{\Lambda}_{3,m}$ for all $m = 2,3,\ldots$. Therefore, given the worst dynamic provisioning case, in which no spare capacity is sharable, it is optimal to split traffic across four and three shortest hop disjoint lightpaths if $\theta = 4$ and $\theta = 3$, respectively. For the example considered, it is straightforward to prove that splitting traffic across $n(1) = \theta$ shortest hop disjoint lightpaths is optimal for $\theta = 2,3,\ldots$. The ensuing section formulates the problem of finding the optimal number of shortest length disjoint lightpaths for specific networks, in which the length of a lightpath is more practically defined in terms of geographical length rather than a hop count.

III. Problem Formulation

Represent a network with the directed graph, or multi-graph, $G = (S,F)$, where S is the set of switches and $F \subseteq S \times S$ is the set of fiber links. Let l_f denote the geographical length of fiber link $f \in F$ and define capacity such that $\lambda \cdot \sum_{f\in L} l_f$ units of capacity are required to transmit λ units of traffic across lightpath L. Let $\mathbf{l} = (l_1, l_2, \ldots, l_{|F|})$ and $\mathbf{X} = (x_{f,m})_{|F|\times M}$, where

$\overline{M}$=MNF$(G, \mathbf{u})$;
$\Lambda = \infty$;
For $m = 2, 3, \ldots, \overline{M}$
  $\overline{\mathbf{X}}_m$ =MCNF$(G, \mathbf{u}, \mathbf{l}, m)$;
  $\overline{\Lambda}_m = \frac{\lambda \cdot (M+1) \cdot |\mathbf{l}\overline{\mathbf{X}}_m|_1}{M^2}$
  If $\overline{\Lambda}_m < \Lambda$
    $\Lambda = \overline{\Lambda}_m$; $\mathbf{X} = \overline{\mathbf{X}}_m$; $M = m$;
  End
End
Return $(M, \mathbf{X})$

Fig. 2. Algorithm finding the optimal number of shortest length disjoint lightpaths to minimize total capacity. For the MNF and MCNF algorithms, each link $f \in F$ of $G = (S, F)$ supports unity flow.

$x_{f,m} = 1$ if lightpath m, $m = 1, 2, \ldots, M$, traverses fiber link $f \in F$, otherwise $x_{f,m} = 0$. Let $\mathbf{A} = (a_{s,f})_{|S| \times |F|}$ denote the incidence matrix, where $a_{s,f} = 1$ if fiber link f is incident from switch s, $a_{s,f} = -1$ if fiber link f is incident to switch s, otherwise $a_{s,f} = 0$. Finally, let $\mathbf{u} = (u_1, u_2, \ldots, u_{|S|})^T$, where $u_s = 1$ if switch s is the ingress point, $u_s = -1$ if switch s is the egress point, otherwise $u_s = 0$.

Given graph, G, the length of each fiber link, $\mathbf{l}$, and an ingress and egress point, the problem of finding the optimal number of shortest length disjoint lightpaths, M, to minimize total capacity can be formulated as follows:

$$\min_{M=2,3,\ldots} \left(\min_{\mathbf{X}} \left(\frac{\lambda \cdot (M+1) \cdot |\mathbf{l}\mathbf{X}|_1}{M^2} \right) \right)$$

such that

$$\mathbf{A} \cdot \mathbf{x}_m = \mathbf{u}, \qquad m = 1, 2, \ldots, M \tag{3}$$

$$\sum_{f \in F} x_{f,m} \leq 1 \qquad m = 1, 2, \ldots, M \tag{4}$$

where $x \in \{0, 1\}$ and $|\mathbf{y}|_1$ denotes the L_1 norm of vector $\mathbf{y}$. The M shortest length disjoint lightpaths are deducible from $\mathbf{X}$. Constraint (3) ensures a lightpath consists of a contiguous sequence of fiber links from the ingress point to the egress point and constraint (4) ensures all lightpaths are disjoint.

Fig. 2 specifies an exhaustive algorithm that optimally solves the above problem formulation. The algorithm makes use of the maximum network flow (MNF) and minimum cost network flow (MCNF) algorithms [1], in which each link supports unity flow. The MNF algorithm is applied to find the maximum number of disjoint lightpaths, $\overline{M}$, and the MCNF algorithm is successively applied to find the m, $m = 2, 3, \ldots, \overline{M}$, shortest length disjoint lighpaths, $\overline{\mathbf{X}}_m$.

Experimentation shows the optimal number of lightpaths can be found within a few seconds of computation for a fully meshed network comprising of less than 20 switches.

IV. Empirical Quantification of Capacity Saving

Let $\Delta(G) = \sum_{s \in S} \delta(s)/|S|$ be the mean switch degree, where $\delta(s)$ is the number of fiber links incident from switch $s \in S$ of graph $G = (S, F)$. Consider the following experiment. Geographical lengths are randomly assigned to each fiber link of a fully-meshed network denoted $G_1 = (S_1, F_1)$. A set of progressively sparser test networks are iteratively constructed such that $S_t = S_{t-1}$ and $F_t = F_{t-1} \backslash \mathcal{F}_{t-1}$, $t = 2, 3, \ldots$, where $\mathcal{F}_{t-1} \subset F_{t-1}$ is randomly chosen such that G_t is connective and $\Delta(G_t) = \Delta(G_{t-1}) - 1$. Given a fixed ingress and egress point, the algorithm specified in Fig. 2 is invoked to find the optimal number of shortest length disjoint lightpaths for each test network. It is then straightforward to quantify the capacity saving attainable by provisioning the optimal number of shortest length disjoint lightpaths relative to MRL, in which a working and a single spare lightpath are provisioned. The experiment is repeated several times and the average capacity saving is shown in Fig. 3.

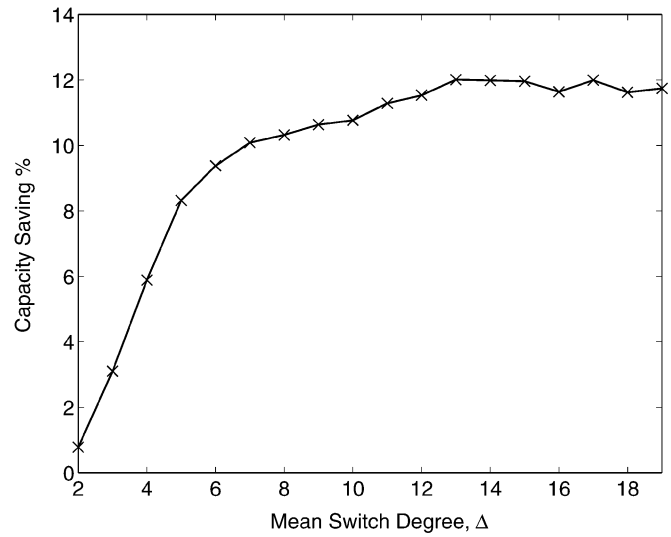

Fig. 3. Capacity saving relative to MRL in randomly generated networks of 20 switches and varying mean switch degree.

As shown in Fig. 3, the capacity saving asymptotes at approximately 11% as the mean switch degree is increased. It can be proven the capacity saving is $(|S| - 3)/3(|S| - 2) \times 100\%$ for equidistant fully meshed networks with $|S|$ switches, that is, a capacity saving that asymptotes at 33%.

V. Conclusion

Although additional propagation delay may be incurred by splitting traffic across an optimal number of disjoint lightpaths, significant capacity savings are attainable, the probability of sharing spare capacity is increased and less traffic requires rerouting at the time of a cable-cut.

References

[1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms and Applications*. New York: Wiley, 1998.

[2] W. D. Grover, J. Doucette, M. Clouqueur, D. Leung, and D. Stamatelakis, "New options and insights for survivable transport networks," *IEEE Commun. Mag.*, vol. 40, pp. 34–41, Jan. 2002.

[3] R. Iraschko, M. H. MacGregor, and W. D. Grover, "Optimal capacity placement for path restoration in STM or ATM mesh-survivable networks," *IEEE Trans. Networking*, vol. 6, no. 3, pp. 325–336, June 1998.

[4] R. Ramamurthy, Z. Bogdanowicz, S. Samieian, D. Saha, B. Rajagopalan, S. Sengupta, S. Chaudhuri, and K. Bala, "Capacity performance of dynamic provisioning in optical networks," *IEEE J. Lightwave Technol.*, vol. 19, pp. 40–48, Jan. 2001.