

Impact of multi-port card diversity constraints in mesh optical networks¹

Eric Bouillet

*IBM Watson Research, 19 Skyline Drive, Hawthorne, New York, 10532
email:ebouillet@ieee.org*

Jean-François Labourdette

*Verizon, 1095 Avenue of the Americas, New York, New York 10036
email:labourdette@ieee.org*

Subir Biswas

*Electrical and Computer Engineering, Michigan State University, East Lansing, Michigan 48824
email:sbiswas@egr.msu.edu*

Abstract: We model and measure the impact of multi-port cards and circuit-pack diversity constraints on the routing of protected lightpaths in optical networks and the resulting network capacity. We describe an offline routing and channel assignment algorithm combined with a multi-port card configuration algorithm, and experiments with these algorithms on various network designs. We show in particular that the impact is negligible, even in the case of shared-mesh protection architectures protecting against single multi-port card failures.

©2005 Optical Society of America

OCIS codes: (060.1810) Couplers, switches, and multiplexers, (060.4250) Networks

1. Introduction

Wavelength Division Multiplexed (WDM) networks that route optical connections using intelligent optical cross-connects are firmly established as the core constituent of next generation networks. It is a well-known fact that the demand for bandwidth has been growing at a steady pace, requiring denser optical switches capable of packing more bandwidth into congested central office spaces. Higher density is in part achieved with the use of circuit-packs, with multiple pluggable modules, which minimize hardware redundancy, footprint and power consumption by sharing functionalities among multiple optical interfaces. The optimum number of pluggable optical interfaces sharing a circuit pack depends on two competing factors. In one hand the port density of the optical switch is proportional to the density of the circuit-pack itself, and it is thus desirable to increase the number of interfaces per circuit pack in order to maximize the port density of the switch. On the other hand, as we migrate functionality from the interfaces into the circuit-packs, the interfaces sharing the circuit-pack must have some level of commonality, such as bit-rates, or signal format (e.g., TDM, SONET, Ethernet, storage networking protocols such as FC/FICON). As a consequence it is often necessary to equip a switch with multiple types of circuit-packs to manage the larger variety of interfaces. This is particularly true on the access side of the network, where the traffic emerges in various forms from the customer premises and is aggregated into a format suitable for transport through the core network. Under those circumstances, the use of circuit packs leads to capacity fragmentation, known as de-rating, and a less efficient utilization of the switch capacity if all the ports of the modules are not occupied. This is further exacerbated as illustrated below, by the fact that each circuit-pack constitutes a single point of failure, whereby all the interfaces it contains could be simultaneously affected if the module happens to fail, or is decommissioned for maintenance. With connection rates reaching tens of Gigabits/s, preventing and repairing failures is increasingly becoming an integral part of the network design process. They are two categories of end-to-end path restoration[1][2]. As illustrated in Fig 1, in end-to-end dedicated (1+1) mesh protection, the ingress and egress optical switches of the failed connection attempt to restore the signal on a predefined backup path that is disjoint, or diverse, from the primary path. Path diversity guarantees that primary and backup paths will not simultaneously succumb to the same failure. This approach requires large amount of backup capacity that is more than the working capacity since backup paths are usually longer than working paths. However the backup path remains “live” in permanence, thus saving crucial path-setup latency when recovery takes place. As shown in Fig 2, in shared mesh restoration, backup paths can share capacity if the corresponding primary paths are mutually diverse. Compared to dedicated (1+1) mesh protection, this scheme allows considerable saving in terms of capacity required. In addition, the backup resources can be utilized for lower priority pre-emptible traffic in normal network operating mode. However recovery is

¹ This work was performed while the authors were with Tellium.

slower than dedicated (1+1) mesh protections, essentially because it involves signaling and path-setup procedures to establish the backup path.

There are two policies to assign channels to shared-mesh restored backup paths. A failure-independent strategy assigns the protection channel at the time of provisioning before failures occur. A failure-dependent strategy assigns the protection channels along a pre-computed route after failure occurrence and relies on the restoration signaling mechanism to select the channels on each link along the route from a pool of reserved channels [9]. In the following we consider the case of the failure-independent strategy.

Both dedicated protection and shared mesh restoration architectures support a variety of protection level requirements, ranging from the most affordable protection against single optical span failures to the most expensive protection against multiple optical switches failures. In the following we consider the case of multiple optical span failures, and use the term Shared Risk Optical Group (SROG) to denote a group of optical spans that share a common risk of failures [2]. The concept of SROG is used for instance to indicate that a conduit cut affects simultaneously all the fibers bundled in it. Protection against single SROG failure hence requires SROG diversity between all primary paths and their corresponding backup paths, and between all primary paths of backup paths that share an optical channel. Furthermore, if protection against node failures is desired then (1) the primary and backup paths must also be node diverse, with the exception of the ingress and egress nodes, for which a failure cannot be remedied if the primary and backup paths access the network through the same optical switches; and (2) end-to-end diversity is required among all primary paths protected by backup paths that share a channel. End-to-end diversity is necessary in this case because even though a demand is not restorable if one of its end nodes fails, the surviving end will nevertheless attempt to restore the demand and contend for the shared channel.




Fig 1. Dedicated mesh (1+1) protection




Fig 2. Shared mesh restoration

In this context, the diversity requirements described above apply to the circuit-packs, as well. As illustrated in Fig 3, a protected circuit can (I) have all its network port plugged on the same circuit pack, or (II) use a separate module for each of its access and network ports. Although this is not shown in the figure for the sake of simplicity, the network operator can further have the choice to protect the access link using 1+1 or 1:N protection on the same module or on redundant modules. Configuration (I) ignores the circuit-pack diversity constraint, and the circuit is susceptible to the failure of a single module. Configuration (II) is the most robust, since the circuit is restored upon the failure of a single circuit-pack, however it also requires the largest number of circuit-packs, and is hence likely to be penalized by a lower overall switch utilization than other configurations. The impact of the diversity and sharing constraints is further illustrated in Fig 4. The figure depicts two shared mesh restorable demands routed on an optical network between a pair of nodes A and Z. Circles are used in the figure to indicate groups of ports that can be collocated on a same circuit-pack without violating the diversity and sharing constraints on protection against circuit-pack failures.

In the remainder of this paper, we explore these issues. We present a circuit-pack diversity routing algorithm, and experiment with different sizes of circuit packs assuming a shared-mesh restorable architecture. In shared mesh restoration, two circuits must be link and circuit-pack diverse to be able to share a backup channel. We then measure under various demand forecast scenarios the impact of de-rating on the cost-efficiency of the network.

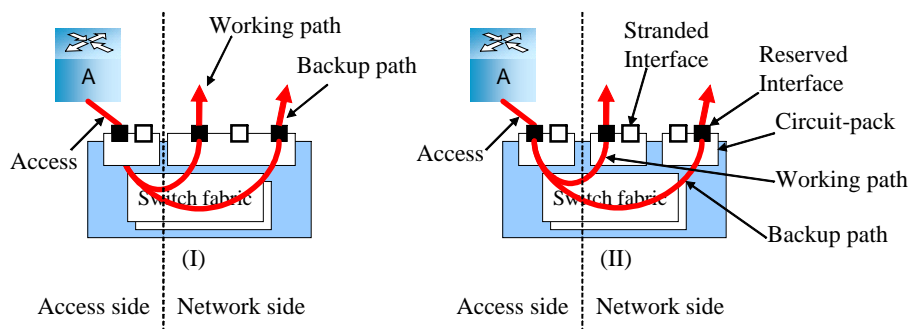


Fig 3. Ingress of an end-to-end protected circuit using various circuit-pack configurations with increasing level of protection.

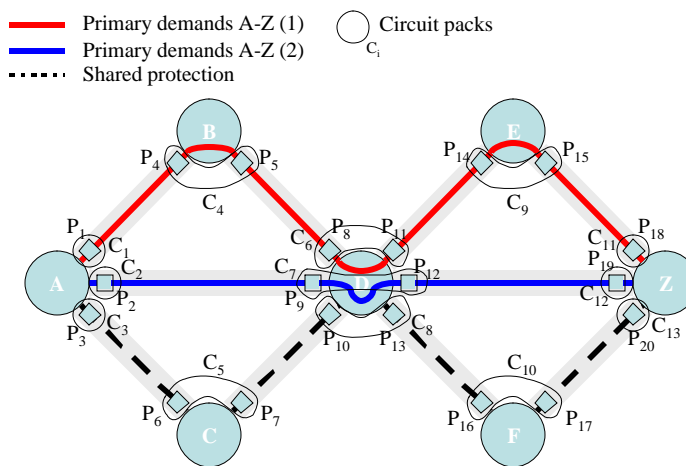


Fig 4. Primary and backup paths of a shared-mesh restorable demand, and circuit pack configuration.

2. Routing Algorithm

The routing of a protected circuit arises principally under two contexts. In online-routing, the network physical topology, the circuit-pack configuration, and the available capacity are prescribed inputs. Under this context, circuits are routed sequentially in the prescribed order, and strictly in accordance to the constraints imposed by the inputs. In particular having a given circuit-pack configuration allows us to express and solve the circuit-pack diversity constraint in terms of node diversity constraints. In the context of an offline design, some of the constraints are relaxed, such as the configuration of the circuit-packs, or the sizing of the network capacity, both of which can be optimized in the process in order to obtain the most cost-efficient solution. For our experiments we investigate the case of the offline design. We thus solve the routing problem and circuit pack assignment problem by first finding a cost-efficient routing of the demands, assuming shared mesh restoration, and without considering the circuit-pack diversity constraints. We achieve this by iteratively applying a k-shortest path-based algorithm on the demand set in conjunction with small link-weight adjustments, until the solution converge to a desirable solution. We then group the ports required by the resulting routes into circuit-packs using a set-cover heuristic, so that the number of circuit-packs is minimized while at the same time respecting the circuit pack diversity required for the shared-mesh protection.

The demand provisioning algorithm (without considering circuit pack diversity constraints) operates as illustrated in insets Algorithm 1 and Algorithm 2. In the algorithm we define the length of a path p as the sum of the weights of the links that constitute the path; the weight of the link being as defined in step 4.a. of Algorithm 2. The constant ε in the ROUTE_PATH algorithm determine the weight of “shareable” protection channels and is used to adjust the tradeoff between capacity requirements and backup path lengths ($0 \leq \varepsilon \leq 1$) [4,5]. Lower values of ε results in solutions that maximize sharing and minimize capacity requirements, but that have longer backup paths and therefore longer restoration times in case of failure. Higher values of ε limits the importance of sharing in the routing objective, resulting in solutions that have a faster restoration time but have higher capacity requirements. In our experiments, the value of ε is set to $\varepsilon=0.3$, as suggested in [4].

ROUTE_DEMANDS(T)

1. Set $min_total_length = \text{infinity}$
2. For each demand (A,Z) of demand set **T**, let $r(A,Z) = (\{p,q\}, length)$ be the primary and backup path pair of demand (A,Z), and their combined $length = length(p) + length(q)$.
Initially $r(A,Z) = (\{\emptyset, \emptyset\}, 0)$.
3. Let **R** denote the set $\{r(A,Z)\}$ for all (A,Z) of **T**.
4. Repeat
 - a. Set $current_total_length = 0$
 - b. For every demand (A,Z) of demand set **T**
 - i. If $r(A,Z) = (\{p, q\}, length) \neq (\{\emptyset, \emptyset\}, 0)$ then
Remove $\{p, q\}$ from network: free corresponding capacity, and update sharing information.
 - ii. Set $(\{p^*, q^*\}, length) \leftarrow ROUTE_PATH(A,Z)$.
 - iii. $current_total_length \leftarrow current_total_length + length$
 - c. If $current_total_length = min_total_length$, then continue to step 5.
5. Return **R**.

Algorithm 1. ROUTE_DEMANDS(T) assign paths to a set T of demands that minimize a total capacity utilization.

ROUTE_PATH(A, Z)

1. For every edge e set weight to cost c_e of one channel in edge (cost of transponders and OAs.)
2. Compute set \mathbf{P} of k minimum-weight paths connecting node-pair A-Z, or all feasible paths if they are less than k of them.
3. Set $min_length < \infty$, and $\{p^*, q^*\} < \{\emptyset, \emptyset\}$.
4. For each path p in \mathbf{P} , do
 - a. Assign weight to every edge e
 - i. If e intersects SROG of primary path p , set weight to infinity.
 - ii. If e has at least one channel that is shareable with p , set weight to $s_{e=\varepsilon} \times c_e$.
 - iii. Otherwise, set weight to $w_e = c_e$.
 - b. Using metric defined in a., compute minimum-weight path q connecting node pair A-Z.
 - c. If q does not exist, continue at step 4. with next path p in \mathbf{P} .
 - d. If $min_length > length(p) + length(q)$, then
 $\{p^*, q^*\} < \{p, q\}$ and $min_length < length(p) + length(q)$.
5. Return $(\{p^*, q^*\}, min_length)$

Algorithm 2. ROUTE_PATH(A,Z) returns the optimum primary and backup paths for a demand (A,Z).

The output of Algorithm 1 is a routing configuration and channel assignment that accommodates the prescribed demand set with protection against single SROG failure. Using this routing configuration we identify for each switch all pairs of ports that are in conflict with the diversity and sharing requirements. According to these requirements, we cannot collocate on the same circuit-pack: (1) two ports that are respectively traversed by the primary and the backup of a demand (diversity constraint), and (2) two ports that are traversed by the primaries of two demands sharing a same backup channel (sharing constraint). We model these constraints by way of a conflict graph for each optical switch. The vertices of the conflict graph represent the ports of the optical switch, and the links indicates pairs of conflicting ports. The inverse of the conflict graph is the association graph, in which links indicate pairs of ports that can be collocated on the same circuit pack. Theoretically, if we enumerate all the cliques $\{K_i\}$ of the association graph, eliminate the cliques that exceed the size of a circuit pack, and group the vertices (i.e. ports) of the remaining cliques into sets $\{C_i\}$, then finding the optimum allocation of ports to circuit packs requires finding the minimum set number that covers all the ports. Note that a routing and channel assignment algorithm designed to protect against node failures always results in a fully connected association graph, since under these diversity and sharing constraints, ports on the same switch cannot be in conflict with each other [2]. The routing and channel assignment, the clique enumeration, and the set cover are known to be NP-complete[3]; and all of them must be solved simultaneously if the optimal solution is sought. In practice we are solving these steps independently, and use the greedy algorithm described in Algorithm 3 for the port to circuit-pack allocation after computing the routing and channel assignment by way of Algorithm 1 and Algorithm 2.

ALLOCATE_CIRCUIT_PACKS(R)

1. Set $min_total_length = \infty$
2. For each optical switch S of the network, do:
 - a. Let Q define the set of conflicting port pairs on S . Pairs of ports in that set cannot be collocated on a same circuit pack. Set $Q \leftarrow \emptyset$.
 - b. For each pair of ports p_1 and p_2 of S , do:
 - i. Using routing configuration R , if p_1 and p_2 are used by the primary and backup of a same demand (A,Z) , or if p_1 and p_2 are used for the primaries of two demands that share a same backup channel, then set $Q \leftarrow Q + \{p_1, p_2\}$.
 - c. Rank ports of S , such that ports appearing the least in Q have lowest rank.
 - d. Sort ports according to increasing rank order, for each port p_i in that order do:
 - i. If p_i can be located to an existing circuit pack C on S , then locate it to C .
 - ii. Otherwise if there is no circuit-pack on S , or if either all the ports of the circuit packs are used or the circuit-packs already contain ports that are in conflict with p_i , then create new circuit-pack on S and assign port p_i to it.

Algorithm 3. Port to circuit-pack allocation algorithm.

3. Experiments

We experiment the algorithm on a 45-node topology, with various demand loads. For each scenario we assume a case with a circuit-pack size of 8 ports, and a case with a circuit pack size of 16 ports. Each case is then solved once without circuit-pack but link diversity constraints (no Circuit Pack Diversity Constraint), and once with circuit-pack diversity constraints (with CPDC). We compare the two solutions in order to determine the combined effect of circuit-pack sizes, and circuit-pack diversity on the network cost. The results of our experiments are summarized below in figures Fig 5 and Fig 6.

Fig 5 shows the number of circuit-packs that are required to carry the demand as we increase the loading on the network. The load is expressed in units of the total number of ports required to accommodate the demand set. Note that as mentioned earlier, the routes are computed before assigning circuits to circuit packs, and therefore each demand set is routed identically independently of the circuit-pack size whether circuit-pack diversity constraint is required or not. We observe from the figure that as expected the number of circuit-pack increases linearly with the demand load. We also note that the impact of the circuit-pack diversity constraint can be significant for the larger circuit-pack size when the demand load is relatively small, but rapidly becomes negligible as the demand increases or the sizes of the circuit pack is reduced.

Fig 6 shows the average circuit-pack utilization as a function of the total number of ports used by the demand. For instance in the network under consideration, under heavy demand load, if each circuit-pack can accommodate 16 ports, then only 80% of the circuit-pack capacity is used on the average.

4. Conclusions

In this presentation, we have discussed, formalized, and analyzed the impact of multi-port cards and corresponding circuit-pack diversity constraint on routing in the context of optical mesh network design. The results show that beyond a certain amount of traffic, the impact on total network cost should become relatively small despite the lower circuit-pack utilization. Furthermore, with the use of pluggable optical modules², the cost penalty remains even smaller as the optical modules can be added as needed.

² Referred to in the industry as SFP (Small Form Pluggable) or XFP (Extra-small Form Pluggable) modules.

5. References

- [1] J-F. Labourdette, et. al., "Routing Strategies for Capacity-Efficient and Fast-Restorable Mesh Optical Networks", *Photonic Network Communications*, **4:3/4**, 219-235 (2002)
- [2] G. Ellinas, et. al., "Routing and Restoration Architectures in Mesh Optical Networks", *Optical Networks*, **4**, 91-106 (2003)
- [3] M. Garey and D. Johnson. "Strong NP-completeness results: motivation, examples and implications", *Journal of the ACM*, **25**, 499-508 (1978)
- [4] E. Bouillet, et. al., "Enhanced Algorithm Cost Model to Control tradeoffs in Provisioning Shared Mesh Restored Lightpaths", in *OFC'02*, Anaheim, CA (2002)
- [5] C. Qiao, Y. Xiong and D. Xu, "Novel models for efficient shared-path protection", in *OFC'02*, Anaheim, CA (2002)
- [6] D. Xu, Y. Chen, Y. Xiong, C. Qiao and X. He, "On Finding Disjoint Paths in Single and Dual Link Cost Networks", in *IEEE INFOCOM'04*, Hong Kong (2004)
- [7] D. Xu, Y. Xiong and C. Qiao, "Novel Algorithms for Shared Segment Protection", *IEEE Journal on Selected Areas in Communications (JSAC)*, **21**, 1320-1331 (2003)
- [8] D. Xu, Y. Xiong, C. Qiao and G. Li, "Trap Avoidance and Protection Schemes in Networks with Shared Risk Link Groups", *IEEE Journal of Lightwave Technology (JLT)*, **21**, 2683-2693 (2003)
- [9] S. Datta et al. "Efficient channel reservation for backup paths in optical mesh networks", in *IEEE INFOCOM'01*, San Antonio, TX, (2001)

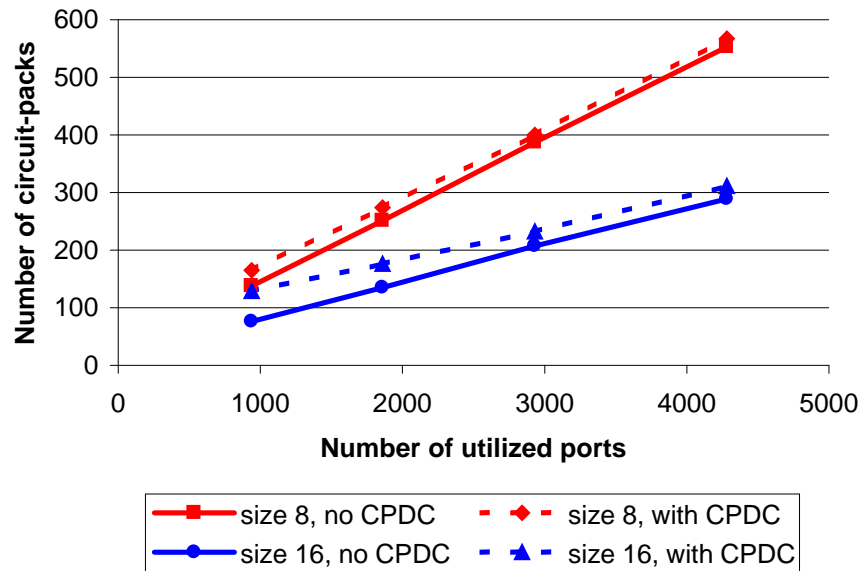


Fig 5. Number of circuit-pack modules under various demand loads, circuit-pack sizes, and diversity constraints (CPDC)

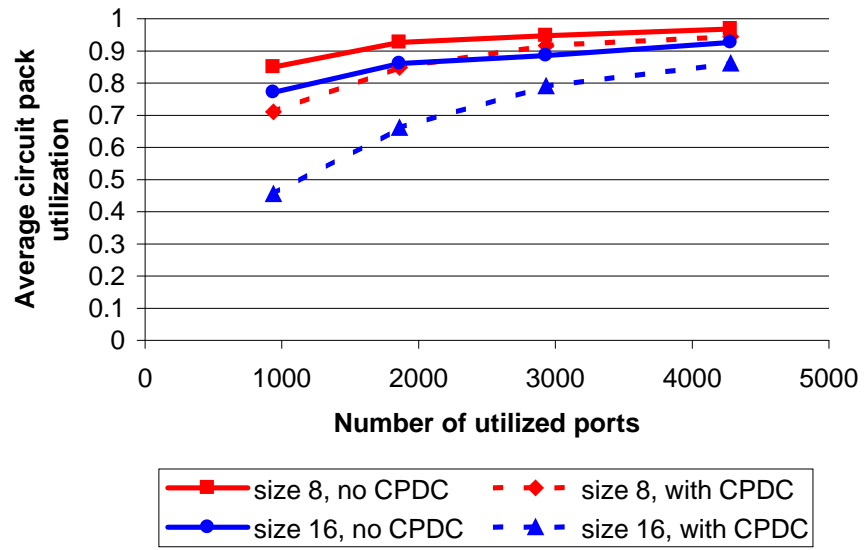


Fig 6. Average circuit-packs utilization under various demand loads, circuit-pack sizes, and diversity constraints (CPDC).