Benefits of *p*-Cycles in a Mixed Protection and Restoration Approach

François J. Blouin[1], Anthony Sack[2,3], *Student Member, IEEE,*
Wayne D. Grover[2,3], *Fellow, IEEE*, Hadi Nasrallah[1]
[1] Nortel Networks, PO Box 3511 Station C, Ottawa, Ontario, Canada K1Y 4H7
[2] TRLabs, 7th Floor, 9107 116 Street NW, Edmonton, Alberta, Canada T6G 2V4
[3] Department of Electrical and Computer Engineering, University of Alberta
E-mail: {blouin, hadinasr}@nortelnetworks.com, {asack, grover}@trlabs.ca

Abstract—**There are reasons to expect that the protection switching times of *p*-cycles will allow them to be considered for protecting services where previously only 1+1 automatic protection switching would have been specified, for the highest speed. Here we introduce *p*-cycles as an option in a mixed-method environment to protect demands that do not strictly require dedicated 1+1 APS, but need a faster switching time than provided by shared backup path protection. Results show that in a mixed services environment, realizing such demands with *p*-cycles rather than 1+1 APS can produce a significant capacity saving. To further explore the potential of mixed-method design, we also test an approach where individual demands are not strictly assigned to be realized by a particular method. Instead, each path may be protected with any method that gives the minimum required service guarantee, but where possible or advantageous in the overall design, paths can be protected by an even faster method than required. For instance, a requirement for 200 ms, usually served by an SBPP realization, may (selectively) be realized instead with *p*-cycle protection at 80 ms. This optimization of the protection method selection produces a small further capacity savings, but provides a remarkable increase in the number of service paths that enjoy faster protection than strictly required by their class of service specification. When modularity is included, this new design approach provides even greater improvement in protection quality.**

Index Terms—**Optical networks, protection and restoration, automatic protection switching, *p*-cycles, shared backup path protection, network planning, network optimization.**

I. INTRODUCTION

A variety of protection and restoration methods exist for optical transport networks [1],[2]. These include automatic protection switching (APS), *p*-cycles, shared backup path protection (SBPP), mesh span restoration, and mesh path restoration, among others. These methods vary in their spare capacity requirements, restoration speed, and other aspects. There can also be more than one restoration technique employed within the same network providing different qualities of protection service to different connections.

One of the commonly found combinations of schemes today is the use of 1+1 diverse dedicated APS arrangements for premium or "high availability" services and SBPP for services with less stringent restoration requirements and where sharing of protection capacity allows greater effi-ciency. There may also be unprotected or best-efforts services for which no protection capacity (either dedicated or shared) is provided at all, as well as legacy domains of SONET or optical ring-based protection. Restorability may also be provided by different layers, for example, by the optical layer or the MPLS layer [2],[3]. In addition to these schemes, the recent technique of *p*-cycles also seems to offer an attractive combination of features, with both ring-like speed and mesh-like efficiency [4]. In this work we consider the possible benefits of introducing *p*-cycles as an additional option for protected service realizations in an optical-layer transport environment that would otherwise use only 1+1 APS or SBPP methods to realize customer service requirements.

A. Background

Two key properties of protection and restoration methods are capacity efficiency and restoration speed. In general, methods that provide high efficiency have longer restoration times, while the faster methods usually require more redundant capacity. Each method mentioned above has its own compromise between these characteristics.

In 1+1 APS, the origin node for a demand will send the signal simultaneously on two paths, a working path and a protection path. The paths are span-disjoint and preferably node-disjoint. The destination node monitors the working path, and when a failure is detected, the destination node switches to the protection path. 1+1 APS thus provides very fast protection, but the dedicated capacity on the protection path implies greater than 100% redundancy.

In SBPP, each working path is assigned a pre-calculated backup path, and a number of backup paths can share spare capacity on some of their spans if their corresponding primary paths are mutually disjoint [5]-[8]. Enough protection capacity is allocated to accommodate all the backup paths that are activated simultaneously by each single span failure. When a failure occurs, the nodes along a backup route must re-configure to establish the backup path corresponding to the working path that failed. This introduces additional delays to signal to the intermediate nodes and configure them, thus increasing restoration time compared to methods where restoration paths are pre-configured.

Among the methods used here, SBPP is the slowest but most capacity efficient [9].

A *p*-cycle is a cyclic pre-connected closed path of spare capacity [4],[10]. It provides protection for any span that has both end nodes on the cycle (as either an on-cycle span or a straddling span). *p*-Cycles provide a unique combination of properties. It is a span protection method, and therefore has a lower fault detection and rerouting time than a path-oriented method. The *p*-cycle protection path is also pre-connected, and is therefore faster than methods which must signal and configure switches along a restoration path at the time of failure. As a fully pre-connected span-switched protection method, *p*-cycles should be at least as fast[1] as BLSR rings under appropriate implementations, and may be almost as fast as 1+1 APS, especially in long-haul applications. Unlike a BLSR, however, spare capacity in a *p*-cycle is shared among multiple protected on-cycle and straddling spans, providing better capacity efficiency than methods where protection capacity is dedicated, or where protected spans must be on-cycle (as in rings). The capacity efficiency of *p*-cycles is generally much higher than that of 1+1 APS, but not quite as high as that of SBPP [9]. *p*-Cycles also do not constrain the working paths in any way, allowing them to take the best (shortest) routes through the network.

Characteristic restoration times for the methods we consider are summarized in Table I. The actual values for all methods will depend on product implementations and network topologies, but most would agree with these characteristic values. In what follows, each method (*m*) constitutes a Class of Protection (CoP) to which a service application can be mapped.

In practice, a fourth method which attempts to dynamically re-establish connections upon failure based on GMPLS shortest-path routing may also be available in WDM networks. This is, however, inherently only a best-efforts scheme where restorability is not guaranteed. It is difficult to make a fair comparison of the spare capacity requirements of this dynamic redial scheme with other methods, since the restorability it provides will vary with the conditions in the network at the time of failure. It is also difficult to quantify the restoration times for the same reason. Thus, we do not include it in the set of methods used here.

B. Motivation and Objectives

Having a wider range of protection methods available should enable the service provider to choose a method for each service that meets its restoration time requirements with an overall minimum in redundant capacity. More

---

[1] One reason we say "at least as fast" is that when a *p*-cycle protects a straddling span failure, the average protection path length is half that of the corresponding BLSR ring and for an on-cycle failure it is identical to a BLSR.

TABLE I: CHARACTERISTIC RESTORATION TIMES OF VARIOUS METHODS

Class of Protection (CoP)	Restoration time	*m*
1+1 APS	< 60 ms	1
p-Cycles	< 80 ms	2
Shared Backup Path Protection	< 200 ms	3

specifically we are motivated by the idea that *p*-cycles may be essentially as fast as 1+1 APS, and considerably faster than end-to-end SBPP. Thus, *p*-cycles may be able to protect a number of service paths which would otherwise require 1+1 APS realizations only because SBPP is too slow for them. But this is where the unique combination of efficiency and relatively high speed of *p*-cycles provides a potentially attractive third option—customers that need faster-than-SBPP restoration may see 1+1 APS and *p*-cycles as "essentially just as fast" alternatives. Compared to 200 ms, the difference between APS and *p*-cycles is not important (especially if the greater efficiency of *p*-cycles relative to 1+1 APS is reflected in the service cost). Thus, the question arises: to what extent could a set of *p*-cycles be established to replace some of the "brute force" 1+1 APS service realizations? The motivation from the standpoint of a network operator is that a *p*-cycle solution may require much less spare capacity than an equivalent set of 1+1 APS paths. The acceptability from the customers' standpoint is that some of them may really only need "faster than SBPP" restoration, but not a dedicated 1+1 realization.

Given this background, we examine the benefits of adding *p*-cycles to a set of methods that network operators can use to meet customer requirements (which otherwise contain only 1+1 APS and SBPP) and quantify the capacity savings and, in some cases, the service quality improvements generated by doing so.

C. Outline and Approach

To assess the potential benefits, we start with test-case designs that serve a mixed service environment with 1+1 APS or SBPP methods only, as a benchmark. The same mixed-service cases are then served with designs in which some or all of the 1+1 APS service class is realized by *p*-cycles (instead of 1+1 APS). Once *p*-cycles are introduced into our planning framework, we proceed a step further to examine the benefits of using a self-selecting approach where the choice of exactly which method to employ for each lightpath is made by the optimization solver. This has the potential to give paths a higher class of protection for an equivalent or lower cost than would otherwise be the case.

Given a set of demands with different restoration time requirements, a basic CoP assignment strategy to minimize cost is to map each individual demand unit to the CoP with the highest restoration time that can still meet its requirement (assuming slower methods are more capacity efficient).

This is the strategy we use in the first part of the study where individual demand units are assigned to a specific method that must be used for their service realization. We subsequently allow "self-selection" of the methods employed for each demand. Here it may be possible to serve a number of lightpaths with a better method than necessary. We use a generic design model for both scenarios, with slight changes for each case. This model is defined in the next section, along with a description of the test networks and demand patterns that were used. The two sections following that contain results showing the benefits for both the assigned and self-selecting mixed-method designs. A final section gives results based on a real-life planning test case and includes modularity to see what extra benefit can possibly be gained in that context. Concluding remarks follow.

II. GENERIC OPTIMIZATION MODEL AND TEST NETWORKS

In this section, we define the generic Integer Linear Programming optimization model that will be used to obtain our results. The formulation finds a minimal-cost allocation of capacity for 100% single-span restorability using either assigned or self-selected protection methods for each demand. Working routes and capacity, as well as spare capacity, are optimized together. This is necessary because the methods being compared vary in their requirements for both working and spare capacity. APS working capacity is placed on both sides of the single shortest cycle between end-nodes so there is no APS "spare" capacity per se. For p-cycles, working demands are shortest-path routed and working capacity placed accordingly. Spare capacity is then placed to support an optimal p-cycle set as in [4]. For SBPP, primary and backup paths are placed on the short and long sides, respectively, of whatever disjoint path pair is chosen for the SBPP realization. Here, there is no constraint on the extent of spare capacity sharing in SBPP [9]. The model assumes adequate wavelength conversion abilities so that wavelength mismatch blocking is not significant. The three protection methods are denoted by their method numbers as given in Table I.

SETS

M — Set of restoration methods, indexed by m.

S — Set of spans, indexed by i (failed) or j (surviving).

D — Set of demand relations, indexed by r.

P — Set of candidate p-cycles, indexed by p.

R_r — Set of candidate SBPP path pairs for r, indexed by b.

PARAMETERS

ε — A small positive constant (0.0001).

Δ — A large positive constant (100000).

c_j — Cost of one unit of capacity on span j.

$d^{m,r}$ — Number of demand units for relation r that require method m.

t_j^r — Equal to 1 if span j is on the shortest cycle between the end nodes of relation r, 0 otherwise.

e_j^r — Equal to 1 if span j is on the shortest route between the end nodes of relation r, 0 otherwise.

$x_{i,p}$ — Equal to 2 if span i straddles p-cycle p, 1 if span i is on p-cycle p, 0 otherwise.

$\theta_{j,p}$ — Equal to 1 if span j is on p-cycle p, 0 otherwise.

$\delta_{j,b}^r$ — Equal to 1 if span j is on the primary (shorter) side of SBPP path pair b for relation r, 0 otherwise.

$\phi_{j,b}^r$ — Equal to 1 if span j is on the backup (longer) side of SBPP path pair b for relation r, 0 otherwise.

VARIABLES

w_j — Working capacity on span j.

w_j^m — Working capacity on span j for method m.

s_j — Spare capacity on span j.

s_j^m — Spare capacity on span j for method m.

$\alpha^{m,r}$ — Number of units for relation r that use method m (equal to $d^{m,r}$ for assigned case).

n_p — Number of unit-capacity copies of p-cycle p.

q_b^r — Number of units for relation r that use SBPP path pair b.

v_b^r — Equal to 1 if relation r uses SBPP path pair b, 0 otherwise.

OBJECTIVE

Minimize: $\sum_{\forall j \in S} c_j \cdot \left(w_j + s_j \right) + \varepsilon \cdot \sum_{\forall m \in M} \sum_{\forall r \in D} m \cdot \alpha^{m,r}$ (1)

(Minimize capacity cost and use the best methods.)

CONSTRAINTS

$\sum_{\forall m \in M} w_j^m = w_j$ $\forall j \in S$ (2)

(For every span, the working capacity must equal the sum of working capacity for all methods.)

$\sum_{\forall m \in M} s_j^m = s_j$ $\forall j \in S$ (3)

(For every span, the spare capacity must equal the sum of spare capacity for all methods.)

$\sum_{\forall m \in M} \alpha^{m,r} = \sum_{\forall m \in M} d^{m,r}$ $\forall r \in D$ (4)

(For every relation, the total quantity of demand assigned to all methods must equal the total quantity required.)

$$\sum_{\forall o \in \{1..m\}} \alpha^{o,r} \geq \sum_{\forall o \in \{1..m\}} d^{o,r} \qquad \forall m \in \boldsymbol{M}, \forall r \in \boldsymbol{D} \quad (5)$$

(For every relation, the total demand is protected with an assigned method or better.)

$$\sum_{\forall b \in \boldsymbol{R}_r} q_b^r = \alpha^{3,r} \qquad \forall r \in \boldsymbol{D} \quad (6)$$

(For SBPP working demands, use method 3 protection.)

$$\sum_{\forall b \in \boldsymbol{R}_r} v_b^r = 1 \qquad \forall r \in \boldsymbol{D} \quad (7)$$

(For SBPP demands, only one path pair can be used.)

$$\Delta \cdot v_b^r \geq q_b^r \qquad \forall r \in \boldsymbol{D}, \forall b \in \boldsymbol{R}_r \quad (8)$$

(For SBPP demands, only the backup path of the path pair selected in (7) can be used for restoration.)

$$\sum_{\forall r \in \boldsymbol{D}} t_j^r \cdot \alpha^{1,r} \leq w_j^1 \qquad \forall j \in \boldsymbol{S} \quad (9)$$

(For APS demands, working capacity is asserted on all spans of the shortest cycle between O-D nodes.)

$$\sum_{\forall r \in \boldsymbol{D}} e_j^r \cdot \alpha^{2,r} \leq w_j^2 \qquad \forall j \in \boldsymbol{S} \quad (10)$$

(For p-cycle demands, working capacity is asserted on all spans of the shortest route between O-D nodes.)

$$\sum_{\forall r \in \boldsymbol{D}} \sum_{\forall b \in \boldsymbol{R}_r} \delta_{j,b}^r \cdot q_b^r \leq w_j^3 \qquad \forall j \in \boldsymbol{S} \quad (11)$$

(For SBPP demands, working capacity is asserted on the primary side of the chosen path pair between O-D nodes.)

$$\sum_{\forall p \in \boldsymbol{P}} x_{i,p} \cdot n_p \geq w_i^2 \qquad \forall i \in \boldsymbol{S} \quad (12)$$

(For each span failure, the set of p-cycles must be adequate to protect all p-cycle working capacity.)

$$\sum_{\forall p \in \boldsymbol{P}} \theta_{j,p} \cdot n_p \leq s_j^2 \qquad \forall j \in \boldsymbol{S} \quad (13)$$

(For every span, there must be sufficient spare capacity to support the set of p-cycles used in (12).)

$$\sum_{\forall r \in \boldsymbol{D}} \sum_{\forall b \in \boldsymbol{R}_r} \delta_{i,b}^r \cdot \phi_{j,b}^r \cdot q_b^r \leq s_j^3 \qquad \forall (i,j) \in \boldsymbol{S}^2 \mid i \neq j \quad (14)$$

(For each span failure, there is enough spare capacity on surviving spans to support all activated SBPP backups.)

The objective function is bi-criteria (as in [11]) to minimize total capacity cost while using the overall best (i.e., lowest) set of methods when possible. ε is set low enough not to interfere with capacity placement decisions, but can still influence method choices where all else is equal in the self-selected case. In the assigned case, $\alpha^{m,r}$ becomes a parameter equal to the value of $d^{m,r}$, making (4) and (5) strictly redundant. Values for c_j are the real distances (in km) for each span, as given in the topology definition. This formulation was implemented in AMPL and solved with

CPLEX 7.5 on a Sun Fire V480 server with four 900 MHz UltraSparc III processors and 16 GB RAM.

A. Test Networks

The COST 239 European network [12] (with span distances from [13]) and a generic U.S. network planning model were used. Both are shown in Figure 1. A realistic demand pattern for the Generic U.S. network was also available and used in Section V. For the two sections to follow, however, a "flat" demand pattern with 20 lightpaths between each node pair was used, and the demand mix was systematically varied on each O-D pair based on test scenarios of 15%, 30%, and 55% CoP requirement mixes. Using these percentages with the 20-unit flat demand pattern, each CoP can be variously assigned exactly 3, 6, or 11 units of the 20 demands. This method for studying effects of multi-service demand mixes is similar to that in [14].

For each test network the model was provided with at least ten candidate SBPP path pairs for each node pair. The Generic U.S. network has 341 distinct cycles in its graph and COST 239 has 3531. In both cases the entire cycle set was provided as eligible cycles for p-cycle realizations.

III. Mixed Protection Design with Assigned Methods

To obtain the first set of results, we used the generic model with $\alpha^{m,r}$ as parameters. In this case $\alpha^{m,r}$ stipulates exactly how many of the demand units for each relation, r, are *assigned* to be protected by 1+1 APS ($m = 1$), p-cycles ($m = 2$), and SBPP ($m = 3$), respectively. This has the effect of forcing the demand units to be assigned to their minimum required method. To obtain benchmark results we eliminate p-cycles from the set of methods considered by setting the $d^{2,r}$ parameters to 0.

Generic U.S.:
$N = 19$
$S = 31$
$\bar{d} = 3.26$

COST 239:
$N = 11$
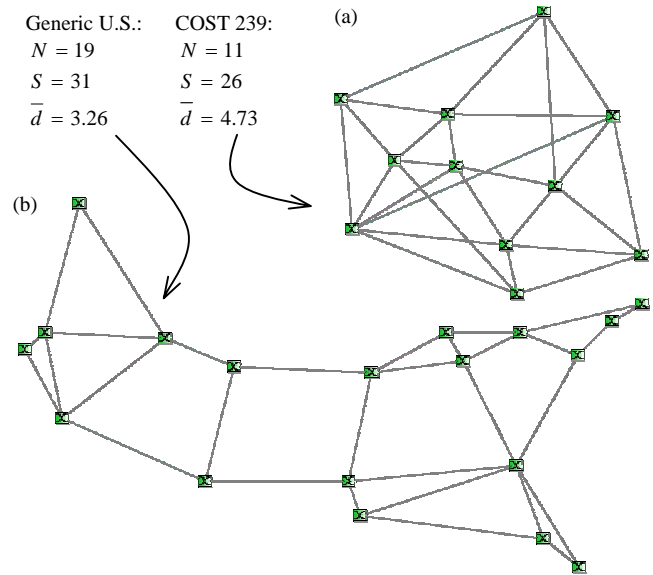$S = 26$
$\bar{d} = 4.73$

(a)

(b)

Fig. 1. COST 239 (a) and Generic U.S. (b) test networks.

All designs for the Generic U.S. network were solved to within 0.01% of optimality; those for COST 239 were within 0.1% of optimality. In Figures 2 and 3, each pair of bars (after the first three) is a comparison in total capacity required with and without *p*-cycles for one of the systematically varying demand mixes. For instance "15-30-55 w/o" is a capacity design where 15% of demands on each O-D pair require 60 ms, 30% require 80 ms, and 55% require 200 ms, but only APS and SBPP methods are available. The bar then shows the total and breakdown of the capacity used by APS and SBPP to serve this mix of requirements. Without *p*-cycles, the middle 30% has to be served by APS. Accordingly, the "15-30-55 with" bar is for the solution where the middle 30% of demands are protected by *p*-cycles. Each subsequent pair of bars is the result for other combinations of 15%, 30% and 55% of demands in each requirement class. At the left hand side, capacities for pure (i.e. 100%) APS, *p*-cycle, and SBPP solutions are included for reference. For Generic U.S., *p*-cycles and SBPP capacities were 9.8% and 22.7% lower than APS, respectively. In COST 239, they were 30.0% and 41.3% lower.

Figure 2 is for the Generic U.S. network. By adding *p*-cycles, the improvement is (on average) 1.6%, 3.2%, or 5.7%, depending on the amount of APS demand that is shifted to *p*-cycle protection (15%, 30%, or 55% respectively). These rather small changes were initially somewhat puzzling when we consider how many opportunities for spare capacity sharing *p*-cycles provide over 1+1 APS. However, the issue is whether those opportunities for capacity sharing can actually be taken advantage of in the
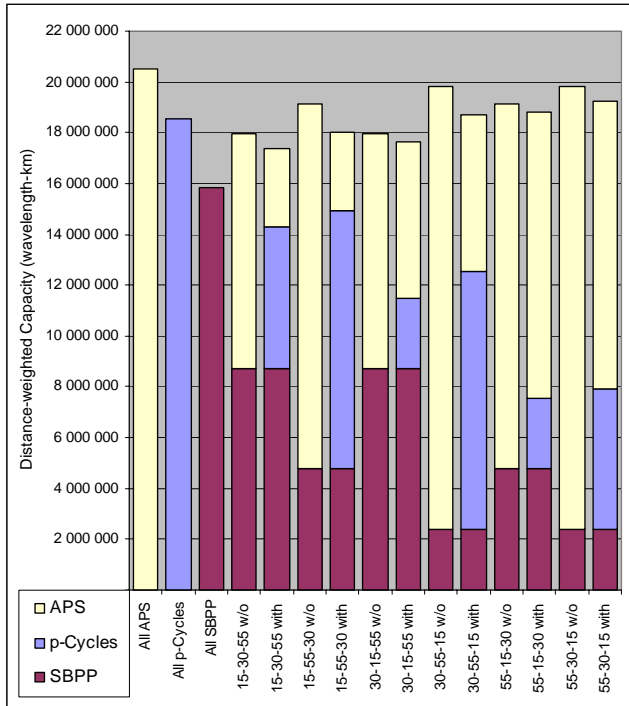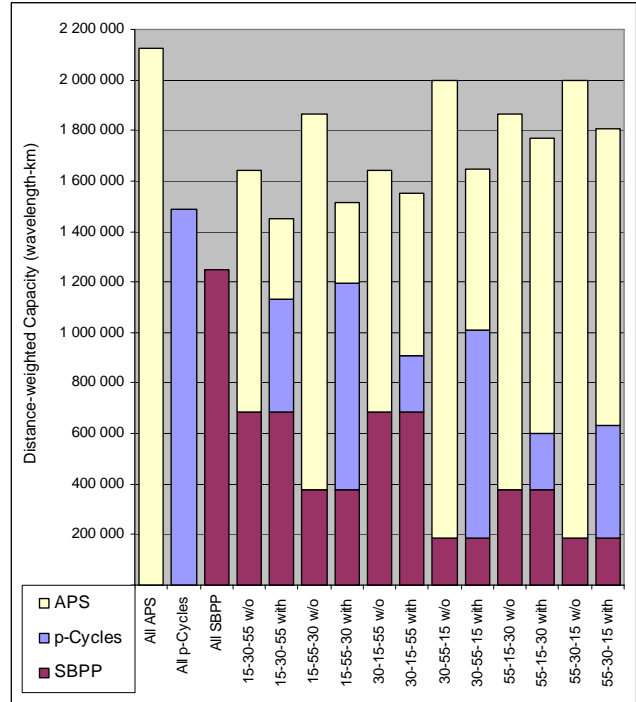

Fig. 3. Capacity effects with assigned model in COST 239 network.

given topology. The key to taking advantage of *p*-cycles is to have enough working flows which can share the protection capacity, *especially working flows which cross the straddling spans*. Without working capacity on the straddlers, a *p*-cycle will be limited in its efficiency. In Figure 1(b), we can see that this network contains a "ladder" sub-network which will tend to result in the routing of working paths on spans other than the "rungs" of the ladder. In some cases, making use of the rungs will be impossible due to the presence of a routing infeasibility as described in [9], where routing the primary path over a certain span would disconnect the network so no disjoint backup path existed. The rungs, however, will frequently appear as straddling spans in candidate cycles, reducing the possible savings in a *p*-cycle solution.

Figure 3 shows the corresponding results for COST 239. Here we observe a 5.4%, 10.6% and 18.1% improvement in total capacity cost as 15%, 30%, and 55%, respectively, of the demand is shifted from 1+1 APS to *p*-cycles. These are much better results, especially when we consider that we are not migrating *all* demand to *p*-cycles, but only a portion of it. COST 239 is a richly-connected network where demand routing will more frequently cross spans that are straddlers to an efficient candidate *p*-cycle. Thus, in general, we can see that adding *p*-cycles to the set of methods can produce significant savings in capacity, but the effect depends on the topology. We expect that the benefits will be greatest in networks with higher average nodal degree.

Fig. 2. Capacity effects with assigned model in Generic U.S. network.

IV. SELF-SELECTED MIXED-METHOD PROTECTION DESIGN

We now look at an extension to the design problem involving multiple protection methods in which we will let the solver decide on the assignment of the method to be used for each demand, given that its minimal requirement is at least met. For instance a requirement for 60 ms will still always be realized by 1+1 APS in the solution but a requirement for 200 ms may be realized by SBPP *or* either of the two faster methods (*p*-cycles or APS). Individual demands may receive protection from a better method than they strictly require. The rationale for permitting this is that in an environment of shared spare capacity, it may be more efficient to do so. For example, a better CoP assignment strategy might be to provide *p*-cycle protection for a path requiring a 200 ms restoration time, which may be servable within existing spare capacity for *p*-cycles while an SBPP setup might not currently have the same sharing efficiency opportunities depending on the number and routing of other SBPP-class paths. In addition, there is no reason not to serve each customer with the best CoP that is possible—under this design scheme we will be interested to see how many demands get "upgraded" at no additional cost.

To effect this design approach, the generic formulation is again used but now $\alpha^{m,r}$ are decision variables (not assigned input parameters). While minimizing cost, this allows demands to move up to a better grade of protection (i.e., a lower *m*-number in the right hand term of the objective function) if the solver sees that it is advantageous (or at least without penalty) from a capacity point-of-view to do so. We again use both test networks with the 20-unit flat demand pattern and 15%-30%-55% mixes to see the effects of the self-selecting framework. We are interested in quantifying the further capacity improvements, but now also the number and percentage of upgraded demand units.

The capacity comparisons are presented in Figures 4 and 5. We observe that the total capacity decreases only slightly more—relative to the decreases already obtained in the assigned approach—on average 4.8% for Generic U.S., and 2.2% for COST 239. Relative to the baseline capacity in the assigned approach where *p*-cycles are not present, the reductions are actually 8.1% for Generic U.S., and 13.3% for COST 239. Interestingly, however, considerable capacity once assigned to SBPP is reallocated to 1+1 APS or to *p*-cycles and 1+1 APS capacity sometimes increases at the expense of *p*-cycles as well, meaning that *p*-cycle demands are also being upgraded opportunistically. The shifts are especially dramatic for those demand mixes with a high percentage of demand units requiring SBPP.

For each demand mix, Table II gives the percentage of individual lightpaths (out of the total demand quantity) that enjoy an upgraded protection speed under the self-selecting best-method model. The demand mixes listed (in both Table II and in various figures) are given in order of increasing method number, i.e. % APS - % *p*-cycles - % SBPP. It
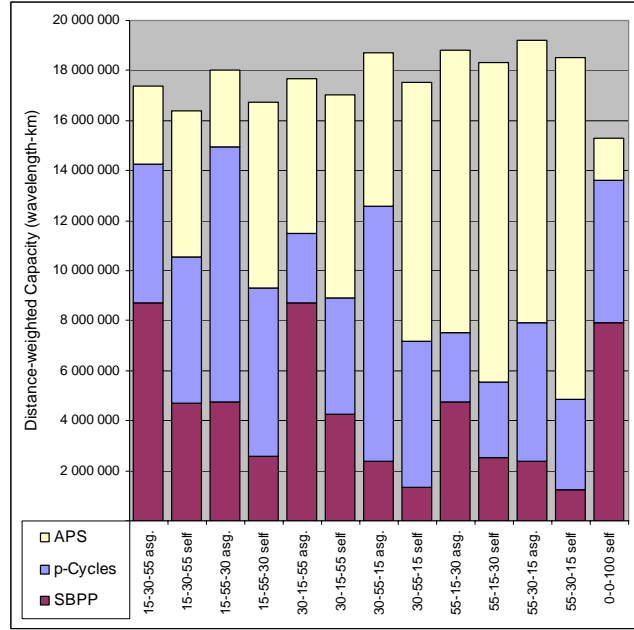

Fig. 4. Capacity effects with self-selected model in Generic U.S. network.

is also interesting to look at the before and after aggregate demand profiles (in Figures 6 and 7), although these do not highlight the benefits quite as well. When viewing the aggregate quantity of demand assigned to each method, we are only able to see the totals and not the number of individual units moving up. This is why the percentages in Table II may seem higher than the improvement apparent in Figures 6 and 7—they are higher, and this is correct. In COST 239, we see that the number of 1+1 APS units stays
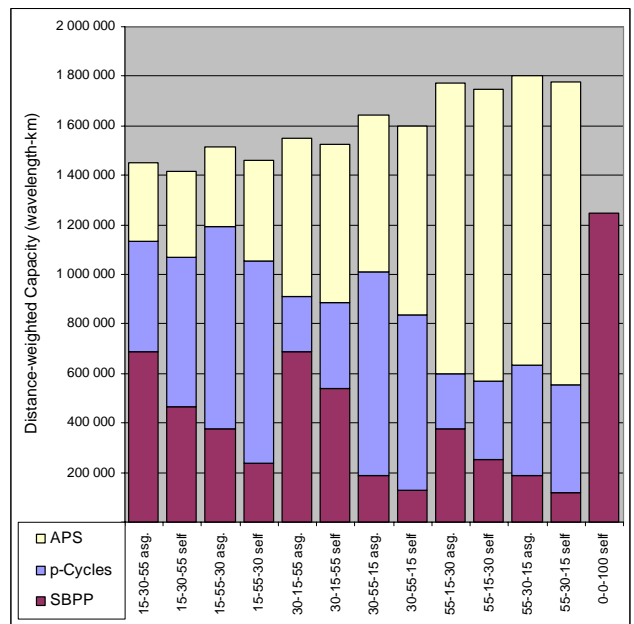

Fig. 5. Capacity effects with self-selected model in COST239 network.

TABLE II: PERCENTAGE OF DEMANDS WITH IMPROVED PROTECTION

Demand mix	Generic U.S.	COST 239
15-30-55	31.9%	24.4%
15-55-30	28.6%	18.0%
30-15-55	29.9%	16.8%
30-55-15	21.9%	11.7%
55-15-30	17.4%	13.4%
55-30-15	15.1%	9.5%
0-0-100	45.4%	0.0%

Fig. 7. Demand upgrading effects in COST 239 network.

relatively constant, and most of the improvement is with *p*-cycles. In Generic U.S., there is improvement in both. This network likely has more cases where APS will be cheaper than *p*-cycles, as noted earlier.

Results for the 0%-0%-100% demand mix are especially interesting. This is the case where all demands only *require* SBPP protection, but may be served by SBPP, *p*-cycles or 1+1 APS. In COST 239, the solver determines that a pure SBPP solution is optimal. But for Generic U.S., a significant number of demands (45.4%) are allocated to 1+1 APS and *p*-cycles! Comparing against the pure SBPP cost for this network in Figure 2, we see that *the self-selecting solution including all three methods is actually cheaper than pure SBPP*. This is an indication that a mix of protection and restoration methods, in some cases, can provide superior performance *over any single method alone*, in terms of both capacity and average restoration speed. From the point-of-view of a network operator, revenue enhancements could come not only from the (slight) additional savings in capacity, but also from the ability to offer a greater percent-age of premium services to customers.
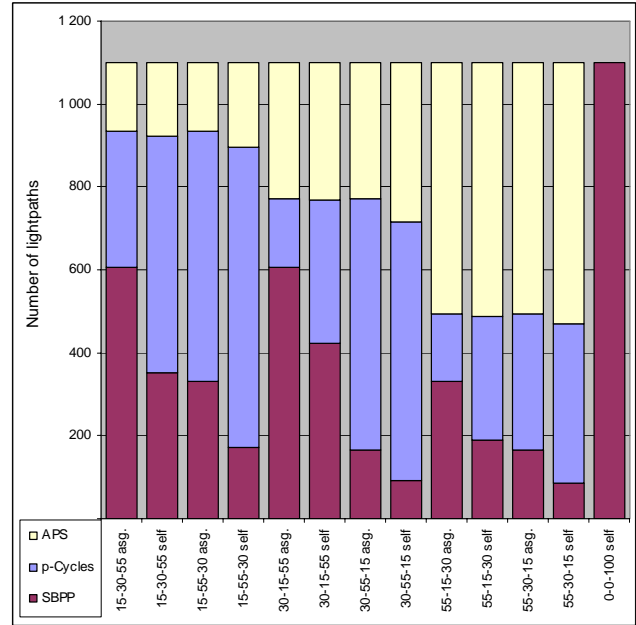
Based on only these two networks, we can postulate that a pure SBPP solution may indeed be the most economical for richly-connected cases like COST 239, but for sparser topologies like Generic U.S., a mixed solution may be superior. Note, however, that the Generic U.S. has an average nodal degree of 3.26 which is still reasonably high—many networks will be even lower. The effects of network topology on these optimal solutions from the self-selecting design model are a promising area for future work.

V. GENERIC U.S. "REAL-WORLD" TEST CASE

In this section we add modularity to the design model and consider a specific mix of CoP requirements based on realistic customer assessments. The aggregate "real-world" mix of service requirements used was 20.5% (60 ms), 23.7% (80ms) and 55.8% (200 ms). Without modularity the solution shows a 2.3% capacity reduction with the assigned mixed design model (relative to the case where 44.2% of demand is all APS-served and the remainder handled by SBPP). With the self-selecting mixed model, a 4.8% reduction in capacity is achieved relative to the assigned approach with *p*-cycles (compared to the assigned approach without *p*-cycles, the decrease in capacity is 7.1%). But remarkably, within this reduced capacity, 24.3% of lightpaths (318 out of 1310) are able to upgrade to a better method!

Throughout this study, the optimization model has made use of the bi-criteria objective function. It would be reasonable to wonder if it really has any effect. As a quick check, this realistic test case was solved again, with and without the bi-criteria term in the objective function, to the very
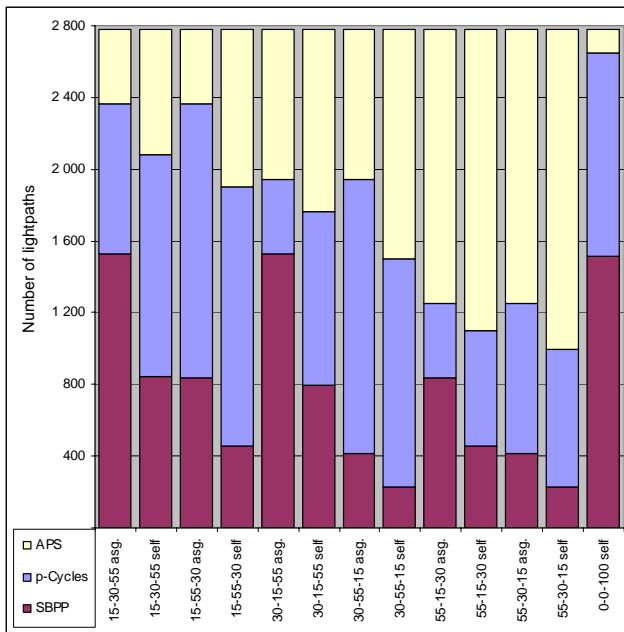

Fig. 6. Demand upgrading effects in Generic U.S. network.

tight tolerance of within 0.001% of optimal to see the effect more precisely. With the bi-criteria term removed, 271 demand units moved "up" to a better (faster) method of protection. When present, however, with $\varepsilon = 0.0001$, 318 demand units (also indicated above) took advantage of a better method, showing that the bi-criteria objective function does have the intended effect of biasing the design towards the best method choices for each demand, as long as total capacity does not suffer.

Modularity was next introduced using the methods in [15]. Here we used a set of modules with sizes of 1, 4, 8, 16, 32, and 160 wavelength channels, with relative costs of 1, 2, 4, 6, 9, and 25 respectively, approximating a "3x2x" effect economy of scale model. The following changes are made to the generic model for modularity:

ADDITIONAL SET

Y Set of modules, indexed by y.

NEW PARAMETERS

c^y Cost of one module of type y.

τ^y Size of one module of type y.

NEW VARIABLE

η_j^y Number of modules of type y placed on span j.

REVISED OBJECTIVE FUNCTION

Minimize:
$$\sum_{\forall y \in Y} \sum_{\forall j \in S} c_j \cdot c^y \cdot \eta_j^y + \varepsilon \cdot \sum_{\forall m \in M} \sum_{\forall r \in D} m \cdot \alpha^{m,r} \qquad (15)$$

(Minimize total *module* costs and use best methods.)

ADDITIONAL CONSTRAINT

$$\sum_{\forall y \in Y} \tau^y \cdot \eta_j^y \geq w_j + s_j \qquad\qquad \forall j \in S \;\; (16)$$

(Total modular capacity must support all logical channel requirements on each span.)

Solved on the Generic U.S. network (to within 0.3% of the optimum) with the realistic demand pattern and mix, the distance-weighted capacity requirements (not including unused module "overhead") were higher than in the non-modular self-selecting approach, but were still lower than they were for the basic assigned model with p-cycles. The total module cost is of course much lower than anything seen so far because of the economy-of-scale effect (an 81.9% decrease from the self-selecting solution without modularity). But the most interesting effect of modularity is that now 37.6% of all demands are up-served (compared to ~24% without modularity).

VI. CONCLUDING DISCUSSION

In this paper we have explored the benefits of allowing p-cycle protected service realizations to displace a certain number of 1+1 APS service implementations. More generally we have taken a first step towards design of survivable networks that serve mixed requirements with multiple different protection methods and optimize the choice of method for each service in the overall network context. Results showed that two main benefits can be obtained:

1) In a relatively well-connected network, significant capacity savings can arise (~20% in total for COST 239 under the self-selected design compared to the assigned approach without p-cycles where 55% of the demand would require APS instead).

2) In any network, significant up-grading of the service realization (speed of restoration) is possible beyond customers' stated CoP requirements (over 45% in the Generic U.S. model if all units require only SBPP).

Not only was the self-selecting model able to produce designs with even somewhat lower overall capacity, but also allowed a remarkable number of demand units to get "bumped up." The bi-criteria model contributes to this improvement. We also observed an interesting phenomenon where a demand mix requiring only SBPP for all lightpaths in a sparser network was optimally served with a *variety* of methods, with a large number of units moving up to p-cycles and even a few to 1+1 APS. This points a direction for further work toward gaining fundamental insight about the interaction between topology and the capacity efficiency of various survivability schemes. Philosophically, we note that this is also consistent with the generalized understanding from [16] that a hybrid combination of restoration schemes is in general always part of a strictly minimum cost network design.

REFERENCES

[1] S. Ramamurthy, B. Mukherjee, "Survivable WDM mesh networks, Part I – Protection," *Proc. 18th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '99)*, New York, USA, 21-25 Mar. 1999, vol. 2, pp. 744-751.

[2] D. Zhou, S. Subramaniam, "Survivability in optical networks," *IEEE Network*, vol. 14, no. 6, Nov./Dec. 2000, pp. 16-23.

[3] R. Aubin, H. Nasrallah, "MPLS and ASON restoration synergy," presented at *MPLS World Congress*, Paris, France, 4-7 Feb. 2003.

[4] W. D. Grover, D. Stamatelakis, "Cycle-oriented distributed preconfiguration: Ring-like speed with mesh-like capacity for self-planning network restoration," *Proc. IEEE International Conference on Communications (ICC) 1998*, Atlanta, USA, 7-11 Jun. 1998, pp. 537-543.

[5] R. Kawamura, K. Sato, I. Tokizawa, "Self-healing ATM networks based on virtual path concept," *IEEE Journal on Selected Areas in Communications*, vol. 12, no. 1, Jan. 1994, pp. 120-127.

[6] Y. Miyao, H. Saito, "Optimal design and evaluation of survivable WDM transport networks," *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 7, Sept. 1998, pp. 1190-1198.

[7] S. Sengupta, R. Ramamurthy, "Capacity efficient distributed routing of mesh-restored lightpaths in optical networks," *Proc. IEEE Global Telecommunications Conference (GLOBECOM) 2001*, San Antonio, Texas, USA, Nov. 2001, vol. 4, pp. 2129-2133.

[8] E. Bouillet, J.-F. Labourdette, G. Ellinas, R. Ramamurthy, S. Chaudhuri, "Stochastic approaches to compute shared mesh restored lightpaths in optical network architectures," *Proc. 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2002)*, New York, USA, 23-27 Jun. 2002, vol. 2, pp. 801-807.

[9] J. Doucette, W. D. Grover, "Comparison of mesh protection and restoration schemes and the dependency on graph connectivity," *Proc. Third International Workshop on the Design of Reliable Communication Networks (DRCN 2001)*, Budapest, Hungary, 7-10 Oct. 2001, pp. 121-128.

[10] W. D. Grover, "Understanding *p*-cycles, enhanced rings, and oriented cycle covers," *Proc. First International Conference on Optical Communications and Networks (ICOCN 2002)*, Singapore, 11-14 Nov. 2002, pp. 305-308.

[11] J. Doucette, W. D. Grover, T. Bach, "Bi-criteria studies of mesh network restoration path-length versus capacity tradeoffs," *Proc. Optical Fiber Communication Conference (OFC) 2001*, Anaheim, California, USA, 17-22 Mar. 2001, pp. TuG2-1-TuG2-3.

[12] P. Batchelor et al., "Ultra high capacity optical transmission networks: Final report of action COST 239," Faculty of Electrical Engineering and Computing, University of Zagreb, 1999.

[13] C. G. Gruber, "Deployment of *p*-cycles in WDM networks," Diploma thesis, Technical University of Munich, Munich, Germany, 2001.

[14] W. D. Grover, M. Clouqueur, "Span-restorable mesh network design to support multiple quality of protection (QoP) service-classes," *Proc. First International Conference on Optical Communications and Networks (ICOCN 2002)*, Singapore, 11-14 Nov. 2002, pp. 321-323.

[15] J. Doucette, W. D. Grover, "Influence of modularity and economy-of-scale effects on design of mesh-restorable DWDM networks," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 10, Oct. 2000, pp. 1912-1923.

[16] W. D. Grover, R. G. Martens, "Optimized design of ring-mesh hybrid networks," *Proc. Second International Workshop on the Design of Reliable Communication Networks (DRCN 2000)*, Munich, Germany, 9-12 Apr. 2000, pp. 291-297.