Influence of Modularity and Economy-of-Scale Effects on Design of Mesh-Restorable DWDM Networks

John Doucette, *Student Member, IEEE* and Wayne D. Grover, *Senior Member, IEEE*

***Abstract*—This work is motivated by interest in restorable mesh architectures for WDM optical networking. DWDM technology is expected to create an extremely modular capacity-planning situation and to produce potentially strong nonlinear economy-of-scale effects in capacity. How will this influence the design of cost-optimized mesh-restorable networks? Will it be essential to do true modular design optimization, or will the traditional rounding-up procedure still be adequate? Can a true modular design method exploit these effects for capital cost savings in the network design? What influence would strong modularity and economy-of-scale have on the evolution of the fiber facilities graph topology for these networks? We address these questions with three mathematical programming formulations that allow a comparative study of these issues in terms of the cost and architectural differences between networks designed with different treatments of the modularity issue. Results show that there are worthwhile savings to be had by bringing modularity aspects directly into the basic design formulation, rather than postmodularizing a continuous integer result, as done in most prior practice. The most significant research finding may be the demonstration of topology reduction (or paring down of the facilities graph) arising spontaneously in optimized designs under the combined effects of high modularity and economy of scale. This is the first quantitative indication and explanation of why *less* highly connected graph topologies may be preferred (at least from an economic standpoint) in future WDM networks, even though the spare capacity efficiency for mesh-based restoration is improved by higher connectivity.**

***Index Terms*—Mesh restoration, network design, network fault tolerance, optical fiber communication, optimization methods, protection, wavelength division multiplexing.**

I. Introduction

A. Background and Objective

TODAY'S BUSINESS culture is highly dependent on a reliable and continuously functioning communications infrastructure. Network failure can severely disrupt crucial services such as telephony, web commerce, debit and credit card transactions, banking and travel booking systems, and even emergency calling systems and air traffic control. With tens of thousands of route-kilometers in typical national or regional networks, fiber optic cable cuts are among the most frequent and serious types of disruption. Hermes, a "pan-European carriers' carrier," for instance, estimates an average of one cable cut every four days on their network [1]. The design of restorable networks is therefore an essential aspect of WDM networking. For competitive carriers, the adoption of a network design and operational methods to ensure split-second survivability is virtually a requirement to retain major customers. The costs of the redundancy to ensure restorability can, however, be high and are well worth optimizing. Without careful choices of architecture and design methods, it is easy to find the costs of supporting 100% restorability against any cable cut approaching twice that of nonsurvivable point-to-point transmission network.

Our focus in this paper is on span-restorable mesh networking in the wavelength cross-connection layer of a DWDM transport network. More specifically, our attention is on the logical design of the routing and capacity plan using operations research (OR)-based methods to minimize total cost. Our aim is to take into account three particular aspects or issues that we think DWDM technology raises for the network planning process: 1) the large capacity modularity of anticipated DWDM systems, 2) the effects of economy of scale on the cost of capacity in WDM systems, and 3) the possibility of paring down the topology of existing fiber networks for further design cost reductions enabled by the effects of the just-mentioned issues.

Manuscript received October 15, 1999; revised May 15, 2000.

The authors are with TRLabs, Edmonton, AB, T5K 2P7, Canada, and the Department of Electrical and Computer Engineering, University of Alberta, Edmonton, AB, Canada (e-mail: {john.doucette; wayne.grover}@edm.trlabs.ca).

Publisher Item Identifier S 0733-8716(00)09019-3.

B. Outline

The remainder of Section I briefly explains the role of transport layer restoration relative to higher service-layer schemes, and reviews some prior work on the optimized mesh-restorable network capacity design problem. Section II then introduces the three design formulations that are implemented and compared in this work. This includes an essentially status-quo reference method (integer design followed by modular rounding-up) and two steps of more complex and aggressive optimization approaches to exploit modularity and economy-of-scale effects. Section III gives details of the network topology, demand patterns, and cost and capacity models of the test case trials on which the three approaches were implemented and compared. The results are presented and discussed in Section IV. Our overall conclusions, insights, and observations are offered in Section V.

C. Optical Transport Layer Restoration

The main characteristic of restoration in the optical layer, whether by ring or mesh technologies, is that prefailure transmission capacity is directly replaced with a set of equal capacity transmission paths. Such carrier-signal level restoration can be so fast,[1] complete, and accurate in protecting against cable cuts

[1] 30 to 50 ms typically for optical rings, and a few hundred milliseconds for optical mesh networks using distributed restoration protocols.

or single light-path failures that higher-level services experience no outage at all. The effect on user applications and client networks is more like that of a single error-second, or a few packets to be retransmitted, not the more extensive data loss and recovery delay associated with a service layer reconfiguration such as, for example, a routing table update and reconvergence associated with Internet protocols.

However, the redundant capacity required for restorability in the transport layer can be expensive[2] and cannot by itself address all restoration problems. With WDM rings, or 1+1 diverse routing, there will be an investment of over 100% in transmission capacity redundancy. With restorable mesh alternatives this may be reduced, but 60%–80% physical redundancy levels are still typical [8], [10]. There are also types of failures where lower-layer restoration is not an option. For instance, a router (node) failure in an IP network can only be addressed by reconfiguration among peer-layer routers. In fact, *node* failures in any service layer in general (either physical or involving software bugs) can only be restored by peer-level network element actions. In addition, while reconfiguration at the optical transport layer may be appropriate in response to a cable cut, failure of a wavelength-terminating interface card on a router is best handled by redundancy in the router design itself or by IP layer reconfiguration, not a transport layer response. For these reasons we can expect to see IP layer restoration routing strategies, as well as continued evolution of ring and mesh restoration concepts in the transport layer.

D. Rings and Mesh

The fastest possible scheme for transport restoration has always been to route the payload-bearing signals over two physically disjoint paths (1+1 DP) and perform selection of the surviving signal at the receiver. This requires an investment of over 100% redundancy in terms of the bandwidth-distance product consumed, but can be economical for some of the largest point-to-point demands in a metropolitan network [2], [3]. The next closest relative to 1+1 DP is the unidirectional path-switched ring (UPSR) or, in a WDM context, the optical path protection ring (OPPR) [19]. The OPPR/UPSR consists logically of numerous 1+1 DP relationships between nodes at a tributary or light-path level supported on a single higher-bandwidth ring. Next come schemes that provide some form of protection bandwidth sharing. The Sonet bidirectional line-switched ring (BLSR) is perhaps the most widely used protection-sharing structure today. Its WDM version is the optical shared protection ring (OSPR) [5]. In an OSPR/BLSR, the protection bandwidth is shared over all spans of the ring through a line-level loop-back switching mechanism. BLSRs usually have better economics than 1+1 DP or UPSRs for higher capacity or longer distance applications because more demands can be served for the same installed base of transmission capacity.

When it comes to restoration capacity sharing, however, "mesh-restorable" networks can be another two to three times more efficient than even a BLSR-based network in terms of the total demand served for a given amount of installed transmission capacity [20], [21]. In long-haul networks, the greater distance-related costs make capacity efficiency important, so there has been continuing interest in design and operation of mesh-restorable networks for long-haul applications, whereas ring-based networks have generally dominated in metropolitan applications.

In the Sonet era, however, the capacity advantage of the mesh alternative has usually been coupled to significantly greater restoration times, due to the use of centralized control or, in distributed restoration, due to the more general rerouting mechanism involved and the slow cross-connection times of some DCS. In mesh restoration, the set of all working channels damaged in a failure are rerouted in bundles or individually, following many diverse paths network-wide, using relatively small amounts of spare capacity on each span. The restoration path set is adaptive to the actual spare capacity and demands on each span. Each unit of spare capacity in a mesh network is thus reusable in many more ways than in ring-based networks. This is what makes mesh-restorable networks efficient, but also what makes them strictly slower than schemes which reserve a dedicated restoration path or reduce the sharing to the spans within each ring only. Because of the speed issue, the lack of a standard for mesh restoration, and the relatively high cost of Sonet-era DCS relative to ring ADMs, interest in mesh restoration waned for a number of years.

Several recent factors are, however, causing mesh-restorable networks to be revisited as a prime architectural alternative for next generation WDM networks. First, is the emergence of a new generation of optical cross-connect systems (OCS) [4], [6] which are planned to have very fast (microsecond to millisecond range) restoration switching in mind from the start of their development. Second, it is now understood how even a relatively slow highly adaptive distributed restoration algorithm (DRA) can still be the basis of very fast restoration through the method of distributed preplanning (DPP) [7]. DPP offers a way to decouple the real-time phase of restoration from the slower path set finding phase, without sacrificing the autonomy and network adaptability that is inherent to a DRA. Third, there has been considerable advance and dissemination of methods for the efficient capacity design of mesh-restorable architectures. This is the literature to be reviewed in the next section. These three developments, coupled with a growing industry appreciation of the planning complexity and capacity requirements of multiring networks, have set the stage for a serious reconsideration of the mesh networking architecture for optical WDM networks.

There are two classes of mesh-restorable networks. The first is called a span- (or link-[3]) restorable network wherein

[2]People sometimes tend to dismiss the cost of transport capacity, presuming it to be virtually "infinite and free." However, the incremental budget for transport equipment alone (neglecting other large costs for basic rights-of-way, staff, maintenance, and other transport infrastructure) of any major carrier is measured in hundreds of millions of dollars annually. Even if cost is dismissed, the converse is that more *revenue* could have been earned with an existing transmission base if less capacity has to be set aside for restoration. This is one of the main reasons that mesh restoration stands ultimately to become more cost-effective than rings, especially in a long-haul environment with rapid demand growth.

[3]The terms *span* and *link*—as used here—have their origins in the transmission networking community. As Bhandari [18] explains, the point is to distinguish between the logical links of higher levels, in this case the logical network of light-path connectivity links of the higher service layers and the physical transmission "spans" over which all end to end logical links are established. "...Spans are the set of physical transmission fibers/cables in the physical facility graph. Links (or edges) of the logical connectivity graph are built from spans. A given span can thus be common to a number of links" [18].

demands routed over a failed span are rerouted between the immediate end nodes of the failed span without consideration for the various demands' origin–destination (O–D) node pairs. Span restoration thus deploys a logical detour around the break while demands remain on their previous routes on either side of the failure. In the alternative *path* restoration, demands that traversed the failed span are simultaneously reprovisioned end to end between their original O–D node pairs within the surviving network. Path restoration is the more capacity-efficient technique but it is considerably more complex in terms of capacity design and real-time implementation. For present purposes we work in the context of span-restorable mesh networks.

E. Prior Work on the Mesh Spare Capacity Placement (SCP) Problem

An early heuristic for dimensioning the spare capacity of a mesh network was developed in [8]. This heuristic involved a hill-climbing approach where spare capacity units are added one at a time, seeking a maximal step increase in network restorability for each addition, until the target restorability level (usually 100%) is achieved. An optional "tightening phase" improves the design by seeking removals and rearrangements that reduce total spare capacity but do not reduce the restorability. A very fast heuristic for generating feasible but suboptimal mesh spare capacity plans is called "max-latching." The basic idea with this method is to set the spare capacity of each span according to the largest spare capacity requirement seen on it over a set of mock restoration trials of all spans one at a time [9].

OR methods were introduced by Sakauchi and by Herzberg. Sakauchi *et al.* [10] provided a linear programming (LP) representation of the problem based on min-cut max-flow considerations. This basic approach was further developed and tested with a number of enhancements by Venables *et al.* [11]. In these models, the spare capacity assignments dimension various cuts of the network graph until the minimum cut for each failure has enough capacity for the required restoration level to be feasible. A technical challenge with this approach is that the number of cut-sets in a network is $O(2^S)$, so the computational problem is to find a suitably small set of cut-sets that fully constrains the system while also permitting an optimal capacity design. The approach is therefore to use successive solutions of an LP to detect and add missing cut-set constraints in the problem tableau. The final values are rounded up either at the end, or at each iteration, to obtain an integer and/or modular solution.

Herzberg and Bye [12] used an LP formulation in which the graph topology is first processed to find all the distinct logical routes that are "eligible" for use in the restoration routing for each failure scenario. To reduce the problem size, hop limits restrict the length of eligible restoration routes. Spare capacity placement is then formulated as an LP to assign restoration flows to the eligible restoration route set so that a minimum of spare capacity is required to support the required flow assignments. Rounding and adjustment at the end approximate the integer solution. In a strict sense, the complexity of this approach is as great as the cut-set oriented formulation because the number of distinct routes is also $O(2^S)$. In practice, however, it is easier to reduce the problem size in Herzberg's method by reducing the number of eligible routes without loss of solution quality. The method also gives detailed information about the routing solutions, while the cut-set approach implicitly assumes only that a max-flow equivalent restoration routing is achieved. Herzberg's basic approach of hop limits and flow assignments to eligible routes underlies each of the models in this paper.

All the methods above first route the working demands (usually through shortest path routing), and then optimize the spare capacity to restore the resultant working capacities. A *jointly optimized* working and spare capacity solution was developed by Iraschko *et al.* in [13] in the form of a mixed integer program (MIP) for both span and path restoration. This method allows working paths to be routed in other than a shortest path manner such that, in conjunction with the spare capacity needed for restoration, the total (working plus spare) capacity requirement is minimized.

More recently, Miyao and Saito [14] and Van Caenegan *et al.* [15] provide OR-based problem formulations for WDM networks. In [14], the authors consider restoration which is based on pre-determined restoration paths that have a one-to-one correspondence with working paths and a linear dependence of cost on nonmodular span capacity. In [15], the authors emphasize issues such as wavelength conversion versus tunability at the end points of a demand, and allow determination of facility span choices in a framework where span capacities are integer but nonmodular, as in the sense here.

Importantly, for the present effort, these prior works all solve for capacity assignments that are integer but nonmodular and use linear models for span cost as a function of installed capacity. In practice, transmission capacity is modular as, for example, exemplified in Sonet (commercial systems exist only at 3, 12, 48, 96, 192, and 768 STS-1 units of capacity). Of course, the prior researchers were aware of this, but the community generally agreed that while trying to solve or study other more basic issues, the practice of producing a modular solution (if needed) by rounding up to the nearest module size was accepted. Additionally, in prior studies, the purpose was often to obtain idealized research comparisons of fundamental questions such as limiting efficiency and limiting ring versus mesh-type comparisons, so modularity was not a concern, or was even an unnecessary confounding factor.

For reasons given next, however, we think that the combined effects of DWDM technology in terms of modularity and economy-of-scale effects are such that we should now consider these effects. At the very least, we can obtain a check at this point on whether the working assumption of "just rounding up" at the end of a nonmodular design method is or is not very different from an explicitly modular treatment.

II. Modular-Capacity Design Strategies

A. The Nature of Modularity and Economy-of-Scale Effects in DWDM

ITU-T recommendation G.692 initially defined 43 wavelength channels in the range from 1530 to 1565 nm with a spacing of 100 GHz. More recently (October 1998), this standard defines 81 wavelengths in the C-band starting from 1528.77 nm incrementing in multiples of 50 GHz (0.39 nm). In addition, commercial systems with 16, 40, 80, and 128 wavelengths per fiber have been announced, and systems with many more wavelengths are soon to emerge from the laboratories. Theoretically more than 1000 channels may be multiplexed in a fiber [4]. It is therefore reasonable that we will see a range of commercial transmission systems with a variety of *modularities* ranging from 4 to 1000 wavelengths. It also seems reasonable that, as with Sonet, DWDM vendors and standards organizations will define a discrete set of modular waveband or system capacities. Based on private conversations with industry experts, a plausible view is that a Sonet STS-like (3, 12, 48,$\cdots$) or SDH STM-like (4, 16, 64,$\cdots$) progression of standard modular wavelength numbers is likely to arise. For this work we assume a representative family of system capacities in the set of {12, 24, 48, 96} wavelengths.[4] The methods, insights, and findings that follow are not fundamentally dependent on the assumed modularity values. The methods can accept any eventual set of standard values as simple input constants. The research findings depend primarily on the fact that we use enough distinct modularity values to be characteristic of an eventual standard set in terms of number of members (at least four) and the total dynamic range covered (a factor of at least 96/12 = 8). If anything, this may underestimate the dominance of the modularity effect in the future.

The usual benefits of standardization will apply for network operators; namely the ability to obtain multivendor supply for various subsystems, and to interconnect between independently administered networks. The adoption of a discrete set of modular capacities will also aid system designers because it allows specific technology choices (such as frequency spacing, carrier generation and EDFA noise, gain, bandwidth, etc.) to be combined and optimized to realize a product at discrete target rates. Some vendors may address the market for 16 wavelength systems, while others may use quite different technology to specialize in, say, 768 wavelength systems.

The relevance to this paper is that, from a network design perspective, we think this means that DWDM technology will present a more overtly modular capacity planning and network design problem than before. Moreover, we expect that a more highly nonlinear cost-capacity relationship may also arise. The reason is that, perhaps more so with DWDM systems than in prior technology generations, the total cost of an installed operating system may depend more on the costs of right-of-way and common equipment than on the actual number of wavelengths operating. Consider, for instance, an optical amplifier that can span the entire 1300 or 1500 bands, or both, with adequate total power handling abilities. For this postulated system element, there is no dependency of cost at all on the actual number of wavelengths operated in the system. Other elements such as the per-wavelength electrical interfaces and laser diodes will be more capacity dependent; but, again, techniques like direct optical comb generation tend to make the one-time cost more significant than the capacity-dependent costs. The intense vendor competition and rapid technical advances that are being seen may also contribute to this economy-of-scale effect in favor of higher system capacities available at far less than proportional cost increases.

To be more specific in illustrating what we mean by modularity, and also giving definition to the items that in practice would be considered in the cost coefficients of the models that follow, let us pick a postulated system, say a 48 wavelength transmission system. The cost for this system will reflect all the following physical items (in each direction):

a) one fiber (and associated per-fiber prorated allocation of right-of-way, duct, cable, installation, and repeater/amplifier housing costs),
b) 48 electrical (transmit) channel interfaces,
c) generation and modulation of 48 optical carriers,
d) optical WDM mux,
e) in-line optical amplifiers with bandwidth–power capabilities suitable for 48 wavelengths, every 60 to 100 km typically,
f) an average cost for 48-channel 3R regenerators, every 1000 km, say,
g) optical WDM demux,
h) 48 electrical (receive) channel interfaces,
i) redundant common power, maintenance processor, rack, cabling, and equipment bay installation costs.

Items a), d), e), g), i), and a large part of f), are one-time costs common to the whole system's existence. The electrical interfaces and per-channel optical carrier generation functions could be provisioned on a one-by-one basis as needed. To the extent, however, that the former cost contributors dominate, it illustrates why we think planners will be increasingly faced with planning decisions of the type where an n-wavelength system would serve current needs, but a $4n$ or even a $6n$-wavelength system may only cost twice as much. This is what we mean by a network design environment that is both highly modular and with strong economy-of-scale effects in capacity.

We now introduce the three main design formulations to be studied. The first method is a benchmark design formulation, equivalent to Herzberg's approach where the solution is optimal for sparing placed in a nonmodular integer manner. The resulting working and spare link totals on each span are then postmodularized for a benchmark case representing current practice. The second method places spare capacity in a modular-aware sense such that the already-existing working link quantities are integer, but the formulation attempts to minimize modular totals required once sparing is added to every span. This is considered to be possibly the most practical compromise between the benchmark and the third approach, which is a joint working and spare formulation. In the third approach we solve simultaneously for the routing of all working light-paths, and the placement of spare capacity so that total modular capacity cost is minimized. This approach has the very interesting prospect of spontaneously eliminating (by complete disuse) some spans from the fiber facilities graph.

B. Postmodularized SCP (PMSCP): Benchmark

To provide a benchmark—representing today's practice—for comparison to the two modular design methods that follow,

[4]The work was initially conducted with a 192 wavelength module as well. However, it was found that with the given demand patterns and test networks, a 192 module was only ever used in the solutions for two of the nine test networks, and only under the extreme $6\times 2\times$ economy-of-scale model, under MJCP. To ease the computational effort in obtaining optimal or nearly optimal solutions for the other cases, the 192 module size was therefore dropped from the study.

we use a version of Herzberg's formulation [12], solved as a pure integer program. First, working capacity is assigned by shortest-path mapping of the demand matrix and spare capacity is optimized in an integer but nonmodular manner. The integer span capacity assignments $(w_i + s_i)$ are then modularized by rounding-up to the nearest module size. For this and subsequent formulations, we use the following notation.

Parameters (Inputs):

C_j Cost of each unit of capacity on span j. (N.B.: This is *during* the nonmodular cost optimization only. Postmodularized cost *assessment* uses modular costing—see below.)

L_i Target restoration level for span i (L_i = 1 assumed).

S Number of spans in the network.

P_i Number of eligible routes for restoration of span i.

w_j Number of working links (capacity units) on span j.

$\delta_{i,j}^p$ Equal to one if pth eligible route for span i uses span j, zero otherwise.

Variables:

f_i^p Restoration flow assigned to pth route for span i.

s_j Number of spare capacity units placed on span j.

The PMSCP benchmark formulation is

$$\text{Minimize} \quad \left\{ \sum_{j=1}^{s} C_j \cdot s_j \right\} \tag{1}$$

$$\text{subject to} \quad \sum_{p=1}^{P_i} f_i^p \geq \lceil w_i \cdot L_i \rceil \qquad \forall i = 1, 2, \cdots, S \tag{2}$$

$$s_j \geq \sum_{p=1}^{P_i} \delta_{i,j}^p \cdot f_i^p \qquad \forall (i, j) = 1, 2, \cdots, S \quad i \neq j. \tag{3}$$

The constraint set in (2) ensures that restoration for span failure i meets the target level (henceforth 100% restoration will be assumed). The set of constraints in (3) forces sufficient spare capacity on each span j such that the sum of the restoration paths routed over that span is met for every failure span i. The largest simultaneously imposed set of restoration paths imposed on a span effectively sets the minimum s_i value on each span in the solution. Once the link-by-link spare capacity placement is determined, modularity is implemented by rounding the total working and spare capacities on each span up to the smallest module of sufficient size. If the total on a span exceeds the largest module, a largest module is placed and the remainder is similarly used to select the smallest additional module on top of the first, etc.

To understand the numerical comparisons that follow, regarding the application of economy-of-scale effects, C_j (in the optimization model just given) corresponds to a single wavelength and is *defined as* a relative cost of 1, during the nonmodular optimization above. The *modular* designs that result from postmodularizing these basic solutions are, however, the relevant "design cost" that is meaningful for later comparison against results of the two intrinsically modular formulations that follow.

C. Modular Spare Capacity Placement (MSCP)

In this model, we continue to route demands in a separate step before optimally placing spare capacity; but here, the spare capacity is decided upon in a "modularity-aware" manner. This is a compromise of sorts in that capacity decisions are modular, but any advantages due to potential modularity impacts on *working* path routing are not yet captured. While the formulation is general, the actual values used for the results that follow are noted in brackets. We reuse the prior notation with the addition of the following new parameters and variables.

Additional Parameters:

C_j^m Cost of a module, of the mth size, on span j ($m = 1, \cdots, 4$).

M Number of different module capacities (M = 4).

Z^m Number of capacity units for the mth module size ($Z^m \in \{12, 24, 48, 96\}$).

New Variables:

n_j^m Number of modules of mth size placed on span j.

The objective is now to minimize the total cost of *modules* placed on the network:

$$\text{Minimize} \quad \left\{ \sum_{m=1}^{M} \sum_{j=1}^{s} C_j^m \cdot n_j^m \right\} \tag{4}$$

and we add a new set of constraints to (2) and (3) above:

$$s_j + w_j \leq \sum_{m=1}^{M} n_j^m \cdot Z^m \qquad \forall j = 1, 2, \cdots, S. \tag{5}$$

This new constraint set assigns a sufficient number of modules of each size so that the total of working and spare capacity required will be met. Note that this does not replace constraint systems (2) and (3), but is required in addition to them. We say this formulation is "modularity-aware," but not a completely modular optimization, because the spare capacity decisions are inherently taking into effect that the *sum* of working and spare quantities will be subject to modularization, but the working path routing still makes no concession or consideration of modularity. The economy-of-scale effect is represented in this model in the C_j^m coefficients. For example, to represent a "four times capacity for twice the cost" economy-of-scale model (denoted by $4\times2\times$), applied between the 12- and 48-wavelength systems, we could assign C_j^1 = 12 (cost of a 12-wavelength span is 12) and C_j^3 = 24 (cost of a 48-wavelength span is 24, or twice that of a 12-wavelength span.). The cost coefficients for each other modular size in the family are correspondingly worked out according to the same economy-of-scale model. As used here, the equivalence between the benchmark nonmodular linear-cost model and the modular economy-of-scale models is in the smallest module size, but it is a simple exercise to work out the C_j^m values so that the average linear cost equivalence is moved out to any of the other module sizes. Note that in engineering practice with these design models, one would not "*create*" an economy-of-scale model in the way we do here for research purposes. Rather, one would simply input the *actual cost* estimate for each system option of the given type, length, and number of wavelengths (plus any cost allocation for EFI, maintenance, net present value considerations, etc.).

D. Modular Joint Capacity Placement (MJCP)

The final model is a wholly modular optimization approach where the routes for working light-paths and the spare capacity for restoration are simultaneously decided with respect to the set of modular capacities available. This is the most complete approach to the problem from a standpoint of modularity effects, but also the most computationally complex. The additional complexity comes mainly from having to define an eligible route set for each working demand in the prefailure network, as well as the eligible route set for restoration of each failure scenario. However, this formulation enables the design to reflect the intuitive notion that under strong economy of scale a very large module may, in effect, "attract" working routes to itself, making it economic to detour the routing of some (typically smaller) working flows. This effect and its extreme manifestation of complete disuse of one or more spans (i.e., $w_i + s_i$ = 0) were made hypotheses of the study associated with this model. Again, we build on the previous formulation and introduce new notation as needed.

Further Parameters:

- D Total number of O–D pairs with nonzero demand.
- d^r Number of demand units for O–D pair r.
- Q^r Number of eligible working routes available for demand pair r.
- $\zeta_j^{r,q}$ Equal to 1 if the qth eligible route for demands between O–D pair r uses span j.

New Variables:

- $g^{r,q}$ Working capacity required by the qth eligible route for demand pair r.
- w_j Number of working capacity units on span j.

We continue to use the objective function in (4) but we add two new sets of constraints associated with solving for the working path routing:

$$\sum_{q=1}^{Q^r} g^{r,q} = d^r \qquad \forall r = 1, 2, \cdots, D \tag{6}$$

$$\sum_{r=1}^{D} \sum_{q=1}^{Q^r} \zeta_j^{r,q} \cdot g^{r,q} = w_j \qquad \forall j = 1, 2, \cdots, S. \tag{7}$$

The constraint set in (6) ensures that all working demands are routed. The constraints in (7) generate the logically required working capacity required on each span j to satisfy the sum of all prefailure demands routed over it. The assignment of values to the C_j^m to effect the economy-of-scale models is identical to that above for MSCP.

E. A Modeling Refinement

As presented, the cost model for transmission systems in the two modular formulations is implicitly one where the system is fully equipped for operation of all channels when the system is "placed." In reality, some of the listed cost elements in Section II-A may be equipped on a per-wavelength, as needed, basis. In other words, there is a second level of unit-capacity cost modularity that can be considered in the optimization. For example, if a given span actually required $w_i + s_i = 17$, then the common system cost for a 24-wavelength system would be required, plus strictly only 17 per-channel unit costs (not 24).

This refinement is not difficult to add to the models, however, because each of the modular formulations still implicitly resolves both w_i and s_i values, which represent the actual number of wavelength channels that need to be turned up within each modular system. This can be reflected by changing the objective function for MSCP and MJCP to

$$\left\{ \sum_{m=1}^{M} \sum_{j=1}^{S} C_j^m \cdot n_j^m + \sum_{j=1}^{S} c_j \cdot (w_j + s_j) \right\}$$

where c_j is the per-channel cost of equipping an additional wavelength on an already installed system on span j, and C_j^m is redefined to represent only the common equipment cost for establishing a system of type m on span j. (A further extension to make the per-channel cost dependent on the system type also follows easily.) This is not much more complex to solve but requires a more detailed set of cost assumptions where the per-wavelength dependent relative costs are separated from the complete-system common costs. For present purposes, we have avoided adding such dimensionality to the presentation and results. The fully equipped system model, implicit in our results, is still quite characteristic of the actual economics if any of the following conditions apply. 1) If the rate of growth is fast enough, and the cost of dispatching maintenance crews to populate new cards one-by-one is high enough that it makes sense to simply fully equip the system when installed. 2) If "per-channel" incremental costs are relatively small compared to the get-started investment required to establish the common-equipment parts of the system. 3) If the resulting designs exhibit high system utilization.[5] 4) If the C_j^m coefficients for the "fully equipped" model are actually based on common equipment costs plus a characteristic *average fill* factor for the per-channel costs.

III. Experimental Design

A. Optimization Software

The three formulations were implemented and compiled in the AMPL modeling system and solved using CPLEX linear optimizer 6.0. The eligible route sets for restoration and working path routing for each network model were generated using purpose-specific programs. They are based on a depth-first search of the graph topology up to the hop limits to enumerate all, or a budgeted number, of distinct route options for each restoration scenario or working path routing decision.

B. Network Models: Topologies and Demand Patterns

The different design models were tested on six networks of varying size and other characteristics shown in Table I. WDM makes it possible to transport traffic of many different types and

[5]As they are in this study. This is seen by comparing the "required capacity" columns of the results (Tables IV–XII) to the "total modular capacity" columns. The former is the actual number of channels required. The latter is the fully installed modular capacity provided. Typically the results show very high fill levels in the actual designs, validating the initial "fully turned-up systems" cost model that is implied.

speeds over each wavelength individually. For instance, the optical network may actually bear a mixture of IP/LAN, DS-n, STS-n, ATM, and leased whole-wavelength services. An important concept in the transport network is, however, that all such higher level service requirements, once forecast or otherwise assessed, can be viewed in aggregate and distilled down into a total demand for wavelengths between each O–D pair. Therefore, we do not need to deal with an array of traffic types, only a model for their aggregated totals between nodes. We created such postulated wavelength demand matrices for the test cases from a *gravity-based demand* model. The gravity demand model assumes that all node pairs may exchange some amount of demand. Here, the number of demand units that a node pair exchanges is proportional to the product of the degrees of the two nodes and inversely proportional to the distance between the two nodes. This model tends to reproduce plausible expectations about the real world in that large centers have strong communities of interest and an inherent hubbing tendency is apparent.[6] The demands are calculated as

$$\text{demand}(a,b) = \text{int}\left[\frac{\text{nodal degree}_a \times \text{nodal degree}_b}{\text{distance}} \times \text{constant}\right].$$

The constant is set empirically so that the individual demand quantities are realistic in comparison to other data sets we have obtained from industry sources. Additionally, there is an aspect of experimental design involved in setting the constant so that the demand quantities are meaningful relative to the module sizes in the test cases. The intent for a general interpretation of the results is to ensure that the largest available module size is indeed fairly large relative to the average baseline w_j (working capacity) quantity, while the smallest module size is quite frequently exceeded. Very similar gravity demand models have been used in other work such as [16] and [17].

To keep the problem sizes computationally manageable for the MJCP formulation, eligible routes for working path routing were adjusted so that there were typically fewer than 10 (but at least three) eligible working routes for each demand pair. The average hop-limits and numbers of eligible routes generated for the formulations are shown in Table II.

C. Test Case Economy-of-Scale and Modularity Models

As mentioned, the four module sizes used were {12, 24, 48, and 96} wavelengths. To investigate the effects of economy of scale, three different cost-scaling rules were employed, namely $3\times 2\times$, $4\times 2\times$, and $6\times 2\times$ ($3\times 2\times$ means that a tripling of capacity results in a doubling of cost, etc.). We do not assert that any of these characteristics will be specifically true of DWDM systems. They are meant only to explore a range of possibilities. Anecdotally, however, we are told by industry sources that $3\times 2\times$ is today fairly characteristic of Sonet rings over the OC-24 to OC-96 range. In all three economy-of-scale models, the cost of a 12-unit module was arbitrarily given an absolute cost of 120, and costs for all other module sizes were generated according to the particular economy of scale. The exact modular system costs used are shown in Table III.

[6]In cases where the nodes are real cities with population data or number of businesses available, the gravity model is preferably based on that data. We use nodal degree as a surrogate here for the presumed size of a population center associated with the node.

TABLE I
TEST NETWORK CHARACTERISTICS

Network	Nodes	Spans	Demand Pairs	Total Demand
9n17s1	9	17	36	167
9n17s2	9	17	36	183
10n19s1	10	19	45	251
10n19s2	10	19	45	244
11n21s1	11	21	55	285
11n21s2	11	21	55	282

TABLE II
ELIGIBLE ROUTES PROVIDED TO THE FORMULATIONS

Network	Average Working Hop-Limit	Eligible Working Routes	Average Restoration Hop-Limit	Eligible Restoration Routes
9n17s1	3.3	240	3.8	119
9n17s2	3.3	246	3.8	121
10n19s1	3.3	290	3.9	138
10n19s2	3.5	292	4.1	134
11n21s1	3.5	353	4.0	141
11n21s2	3.5	357	4.0	141

IV. TEST RESULTS

A. General

The main results appear in Tables IV–XII. For each design model there are three tables, for each of the three economy-of-scale models ($3\times 2\times$, $4\times 2\times$, $6\times 2\times$). Tables IV–VI are for the nonmodular benchmark results. There is really one basic solution for each network under PMSCP, but the tables reflect the different cost totals arising after the round-up modularization step under each of the three sets of different module costs. PMSCP itself is a linear, nonmodular, formulation that is invariant under the economy-of-scale models except in the sense that the total cost of the round-up resultant module set varies as we vary the costs for each module under the economy-of-scale models for the true modular designs. Tables VII–IX give MSCP model results, and Tables X–XII follow the same pattern for MJCP.

All designs are for 100% restorability to any single span failure, under a span-restoration model. In each table, column #2 shows the total purely logical design capacity (e.g., the sum of all logical $si + wi$ quantities). By examining this column relative to the total modular capacity placed, we get an indication of the average module fill in each resulting design. The next four columns give the exact numbers of each module size used in each design. The "Cost" column is the total objective function cost of each design. The "Span Elim." column represents the number of spans having zero modules of any size placed on them under the MJCP formulation. This is an effect which can only emerge from modularity and economy-of-scale cost effects combining in the MJCP case to "pull" both working and spare capacity away from spans where it would have been in a purely minimum capacity (as opposed to a minimum modular

TABLE III
MODULE SIZES AND COSTS

Cost Scheme	Module Size 12	Module Size 24	Module Size 48	Module Size 96
3x2x	120	186	288	446
4x2x	120	170	240	339
6x2x	120	157	205	268

cost) design result. Such "span elimination" cannot arise in any of the other models because the independent preliminary step of shortest path mapping of the demands uses every span to some extent. The "Tot. Mod. Capacity" column shows the total modular capacity actually provided, i.e., the sum of the capacities of all modules called for in the design. MSCP and MJCP results tables also have a final column (% cost improvement) where the total cost for each design is compared (in percentage reduction terms) to that of the postmodularized benchmark cost (PMSCP) for that network and economy-of-scale model. All results for the PMSCP benchmark and MSCP formulations were obtained with full CPLEX terminations and therefore happen to represent provably optimal results for the given cases. Results for the MJCP test cases were obtained with a CPLEX "mipgap" of 10%. This means that the MJCP results may be optimal but are only proven to be within 10% of the lower bound set by the LP relaxation of the problem. The actual benefits of the MJCP design method may therefore be strictly greater than reported here.

B. Discussion of Results

By comparing the "required" capacities for the PMSCP designs to the actual modular total capacities arising after rounding up (Table IV), it is apparent that 23% to 34% (average 29.6%) of the total capacity in these networks is excess capacity due to modularity effects.

Comparison shows that the *required* (logical $w+s$) capacity using the MSCP method increases relative to PMSCP, but in fact this is to be expected because the restoration routing is now being influenced (lengthened slightly on average) to take advantage of modular economy-of-scale effects. However, in nearly all cases, total module capacity decreased relative to simple postmodularized designs, and there was a reduction in the proportion of unused module capacity. This is evidence of the restoration routings being redirected to exploit modular capacity placement opportunities. On average, under MSCP, 12.3% of modular capacity placed is in excess of design use. This is about half the excess arising from simple postmodularization of the nonmodular designs. This translated into total modular cost savings averaging 7.9% with MSCP with a high of 12.3%.

The greatest improvements in design efficiency and costs arise with the MJCP design formulation. Cost savings as high as 18.3% were realized (average of 12.2%), and excess modular capacity was decreased to an average of 3.8%. This is all consistent with the notion that in MJCP both working paths and restoration routes are being carefully aggregated and coordinated to fit together in well-filled cost-effectively chosen modules of the largest sizes that can be economically exploited. Fig. 1 graphically summarizes the relative costs of the solutions achieved for each network averaged over all three economy-of-scale models using the three design methods. The bars represent the cost of each solution relative to the corresponding postmodularized benchmark design (PMSCP).

C. Economy-of-Scale Effects on Topology

In the MJCP formulation, one of the main effects observed is the tendency, as expected, for working light-path routes and restoration routes to go "out of their way" and be coordinated so as to take advantage of design opportunities for larger and more cost-efficient modules. A somewhat unexpected outcome, however (illustrated in Fig. 2), is the frequency with which this effect could go to the extent of entirely disusing one or more spans. This occurred even under the weakest economy-of-scale model ($3\times2\times$). The meaning in such cases is that the cost-optimal routing and capacity design is realized fully on a subgraph of all the possible spans present. This topology sparsening effect increases as the economy of scale strengthens. This is a significant experimental demonstration because it shows that topology *reduction* may be cost-effective even in a mesh restorable network, despite a mesh's tendency toward higher connectivity to increase the spare capacity sharing efficiency of a mesh network. An example is shown in Fig. 3 of how the subgraph associated with the cost-optimal solution evolves (for network $11n21s2$) as the economy-of-scale effect in capacity increases.

The explanation is clear enough: the greater the economy-of-scale factor and the modular capacity available, the more incentive there is for a working path or restoration route to detour to help fill a larger module. The cost-optimal solution is characterized by use of the fewest, largest modules, with the routings being subservient to this internal dominance effect. In contrast, a capacity-minimal solution is dominated by shortest path routing considerations. Only when capacity cost is linear (and also nonmodular) is the min-*capacity* solution also the cost-optimal solution. The min-capacity solution always benefits from greater network average nodal degree, but the modular min-cost solution can evidently counteract and overtake this effect under sufficient economy-of-scale effects.

Even for the small test networks, with the not-unrealistic $3\times2\times$ economy-of-scale model, some spans are eliminated in the MJCP model. The greatest impact is obviously with the extreme "what if" case of $6\times2\times$ where an average of 19.8% of the spans are eliminated, and one network ($11n21s2$) saw a 28.6% reduction in spans. Other inspections of the results suggest that the effect is greater in larger networks and when larger hop-limits are used—both of which are also consistent with the explanation above.

V. CONCLUDING DISCUSSION

This work has been motivated by the idea that DWDM technology may produce a network-planning problem that is dominated by very large and nonlinear capacity-cost modularity effects. We have asked how the capacity design of mesh-restorable networks could be influenced by this. The vehicle used was a set of three integer program formulations; one serving as a benchmark for current practice, and two original multi-module-size modular formulations, studied on a suite of test network cases. Results are sufficient to support a recommendation

TABLE IV
POST-MODULARIZED SCP (BENCHMARK) RESULTS FOR $3\times 2\times$ ECONOMY OF SCALE

Network	Req'd Capacity	Size 12 Modules	Size 24 Modules	Size 48 Modules	Size 96 Modules	Cost	Span Elim.	Tot. Mod. Capacity
9n17s1	323	0	14	3	0	3468	N/A	480
9n17s2	368	0	13	4	0	3570	N/A	504
10n19s1	540	1	6	13	0	4980	N/A	780
10n19s2	550	0	9	9	1	4712	N/A	744
11n21s1	653	3	10	10	1	5546	N/A	852
11n21s2	650	4	9	12	0	5610	N/A	840

TABLE V
POST-MODULARIZED SCP (BENCHMARK) RESULTS FOR $4\times 2\times$ ECONOMY OF SCALE

Network	Req'd Capacity	Size 12 Modules	Size 24 Modules	Size 48 Modules	Size 96 Modules	Cost	Span Elim.	Tot. Mod. Capacity
9n17s1	323	0	14	3	0	3100	N/A	480
9n17s2	368	0	13	4	0	3170	N/A	504
10n19s1	540	0	6	12	1	4239	N/A	816
10n19s2	550	0	9	9	1	4029	N/A	744
11n21s1	653	0	10	7	4	4736	N/A	960
11n21s2	650	0	9	8	4	4806	N/A	984

TABLE VI
POST-MODULARIZED SCP (BENCHMARK) RESULTS FOR $6\times 2\times$ ECONOMY OF SCALE

Network	Req'd Capacity	Size 12 Modules	Size 24 Modules	Size 48 Modules	Size 96 Modules	Cost	Span Elim.	Tot. Mod. Capacity
9n17s1	323	0	14	3	0	2813	N/A	480
9n17s2	368	0	13	4	0	2861	N/A	504
10n19s1	540	0	6	12	1	3670	N/A	816
10n19s2	550	0	9	9	1	3526	N/A	744
11n21s1	653	0	10	7	4	4077	N/A	960
11n21s2	650	0	9	8	4	4125	N/A	984

TABLE VII
MSCP DESIGN RESULTS FOR $3\times 2\times$ ECONOMY OF SCALE

Network	Req'd Capacity	Size 12 Modules	Size 24 Modules	Size 48 Modules	Size 96 Modules	Cost	Span Elim.	Tot. Mod. Capacity	% Cost Improvement
9n17s1	370	4	11	2	0	3102	N/A	408	10.6%
9n17s2	404	3	11	3	0	3270	N/A	444	8.4%
10n19s1	565	1	12	7	0	4368	N/A	636	12.3%
10n19s2	569	2	10	6	1	4274	N/A	648	9.3%
11n21s1	704	0	12	7	2	5140	N/A	816	7.3%
11n21s2	655	2	12	9	0	5064	N/A	744	9.7%
Average									9.6%

TABLE VIII
MSCP DESIGN RESULTS FOR $4\times 2\times$ ECONOMY OF SCALE

Network	Req'd Capacity	Size 12 Modules	Size 24 Modules	Size 48 Modules	Size 96 Modules	Cost	Span Elim.	Tot. Mod. Capacity	% Cost Improvement
9n17s1	377	3	11	3	0	2950	N/A	444	4.8%
9n17s2	401	3	11	3	0	2950	N/A	444	6.9%
10n19s1	570	0	12	6	1	3819	N/A	672	9.9%
10n19s2	565	2	10	6	1	3719	N/A	648	7.7%
11n21s1	710	0	12	7	2	4398	N/A	816	7.1%
11n21s2	717	0	15	2	4	4386	N/A	840	8.7%
Average									7.5%

TABLE IX
MSCP DESIGN RESULTS FOR $6\times 2\times$ ECONOMY OF SCALE

Network	Req'd Capacity	Size 12 Modules	Size 24 Modules	Size 48 Modules	Size 96 Modules	Cost	Span Elim.	Tot. Mod. Capacity	% Cost Improvement
9n17s1	374	4	11	2	0	2617	N/A	408	7.0%
9n17s2	393	3	11	3	0	2702	N/A	444	5.6%
10n19s1	567	0	12	6	1	3382	N/A	672	7.8%
10n19s2	575	2	10	6	1	3308	N/A	648	6.2%
11n21s1	712	0	12	7	2	3855	N/A	816	5.4%
11n21s2	717	0	15	2	4	3837	N/A	840	7.0%
Average									6.5%

TABLE X
MJCP DESIGN RESULTS FOR $3\times 2\times$ ECONOMY OF SCALE

Network	Req'd Capacity	Size 12 Modules	Size 24 Modules	Size 48 Modules	Size 96 Modules	Cost	Span Elim.	Tot. Mod. Capacity	% Cost Improvement
9n17s1	344	5	12	0	0	2832	0	348	18.3%
9n17s2	404	6	8	3	0	3072	0	408	13.9%
10n19s1	573	0	13	6	0	4146	0	600	16.7%
10n19s2	568	4	12	5	0	4152	0	576	11.9%
11n21s1	790	0	14	2	4	4964	1	816	10.5%
11n21s2	711	0	10	10	0	4740	1	720	15.5%
Average							0.3		14.5%

TABLE XI
MJCP DESIGN RESULTS FOR $4\times 2\times$ ECONOMY OF SCALE

Network	Req'd Capacity	Size 12 Modules	Size 24 Modules	Size 48 Modules	Size 96 Modules	Cost	Span Elim.	Tot. Mod. Capacity	% Cost Improvement
9n17s1	340	5	12	0	0	2640	0	348	14.8%
9n17s2	438	1	12	3	0	2880	1	444	9.1%
10n19s1	584	0	13	6	0	3650	0	600	13.9%
10n19s2	681	0	4	11	1	3659	3	720	9.2%
11n21s1	851	1	2	13	2	4258	3	876	10.1%
11n21s2	813	2	1	16	0	4250	2	816	11.6%
Average							1.5		11.5%

TABLE XII
MJCP DESIGN RESULTS FOR $6\times 2\times$ ECONOMY OF SCALE

Network	Req'd Capacity	Size 12 Modules	Size 24 Modules	Size 48 Modules	Size 96 Modules	Cost	Span Elim.	Tot. Mod. Capacity	% Cost Improvement
9n17s1	427	1	10	4	0	2510	2	444	10.8%
9n17s2	525	3	2	8	1	2582	3	564	9.8%
10n19s1	805	0	2	9	4	3231	4	864	12.0%
10n19s2	663	0	5	10	1	3103	3	696	12.0%
11n21s1	1027	0	1	8	7	3673	5	1080	9.9%
11n21s2	1031	0	0	5	10	3705	6	1200	10.2%
Average							3.8		10.8%

that mesh capacity design should move to truly modular formulations to be relevant for DWDM. The long practice of postmodularizing continuous integer solutions leaves significant potential cost savings unexploited. The MSCP formulation is perhaps the most practical formulation to adopt as it does not require working path provisioning processes to change from shortest paths, but does exploit modularity in the assignment of restoration routes and spare capacity. However, MSCP cannot generate any spontaneous span eliminations (assuming networks do not contain any spans that already bear no working capacity). The MJCP formulation is the most computationally challenging and may be the most difficult in practical use because of the aspect of deviating working light-paths when provisioned, from their shortest routes. It is not clear at present how the incremental growth provisioning process could effect MJCP considerations which, in the above, emerge only when there is a chance to collectively optimize a group of demands and their restoration routes. MJCP may be of more value in longer range planning studies where, based on forecast demand patterns and projected equipment modularities and costs, the span eliminations that it suggests may be taken as indications of the direction in which network topology should evolve.

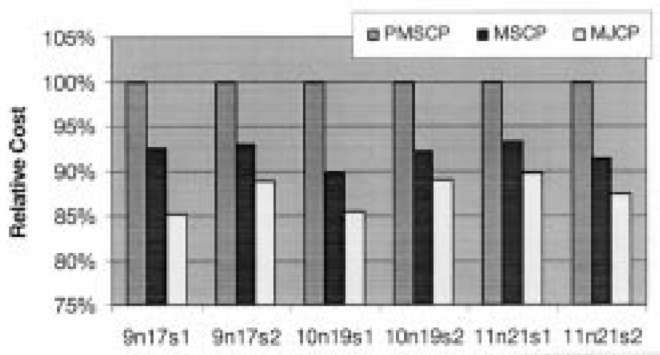

Fig. 1. Relative network average costs.

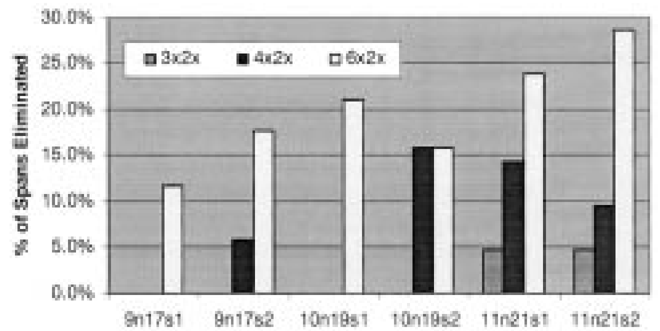

Fig. 2. Percentage span elimination by economy of scale.

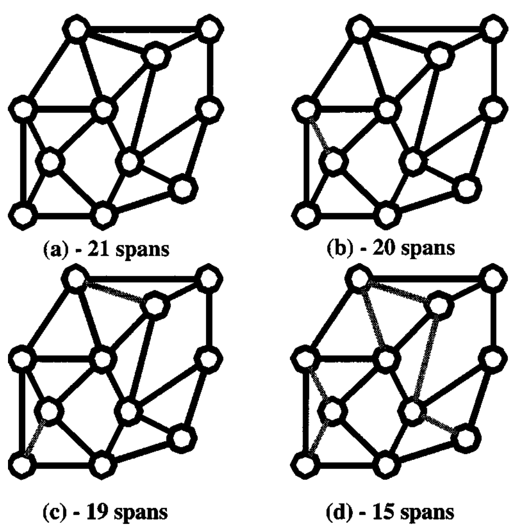

Fig. 3. Effects of economy of scale on optimal network topology: network $11n21s2$ using PMSCP (a), MJCP with $3\times2\times$, $4\times2\times$, and $6\times2\times$ economies of scale [(b), (c), and (d), respectively].

To the extent that future technology advances in WDM may present the increasingly large modular economy-of-scale effects that we postulate, it may be useful to keep the prospect of topology simplification in mind as a future opportunity. Of course, the purely economic indication that an opportunity exists for a reduced topology has to be tempered by counter-acting considerations of dual-failure unavailability and signal propagation lengths, etc. But the demonstration of spontaneous span elimination under MJCP is a deepening of our overall understanding of mesh-restorable networks. While it is still true that the spare capacity reduces in more highly connected networks, we now see that effect is in counterbalance to the aggregation effect of large modularity and strong economy of scale. It also suggests the possibility that future WDM mesh networks may be cost effective when implemented directly on some of today's relatively sparse topologies, which are the legacy of ring-based transport. In addition, the practical benefit of each span elimination is greater than the modular transport system costs captured in the formulations alone, in terms of further reduced right-of-way, leasing, and/or operations and maintenance costs, etc.

REFERENCES

[1] P. F. Fonseca, "Pan-European multi-wavelength transport networks: Network design, architecture, survivability and SDH networking," in *Proc. 1st Int. Workshop Design Reliable Comm. Networks (DRCN'98)*. Brugge, Belgium, May 17–20, 1998, Paper P3.

[2] T. S. Wu, *Network Service Survivability*. Norwood, MA: Artech House, 1992.

[3] T. Flanagan, "Fiber network survivability," *IEEE Commun. Mag.*, June 1990.

[4] S. V. Kartalopoulos, *Introduction to DWDM Technology*. New York: IEEE Press and SPIE Optical Engineering Press, 2000, ch. 14, 15.

[5] P. Neusy and R. Habel, "Availability analysis of optical shared protection rings for long haul networks," in *Proc. Opt. Fiber Commun. Conf. (OFC'99)*, Feb. 1999, pp. TuL5-1–TuL5-4.

[6] M. Koga *et al.*, "Large capcacity optical path cross-connect system for WDM photonic transport network," *IEEE J.Select. Areas Commun.*, vol. 16, pp. 1260–1269, Sept. 1998.

[7] W. D. Grover, "Distributed restoration of the transport network," in *Network Management into the 21st Century*, T. Plevyak and S. Aidarous, Eds. New York: IEEE/IEE Press, Feb. 1994, ch. 11, pp. 337–417.

[8] W. D. Grover, T. D. Bilodeau, and B. D. Venables, "Near optimal spare capacity planning in a mesh restorable network," in *Proc. IEEE GLOBECOM'91*, pp. 2007–2012.

[9] W. D. Grover, V. Rawat, and M. MacGregor, "A fast heuristic principle for spare capcity placement in mesh-restorable sonet/SDH transport networks," *Electron. Lett.*, vol. 33, no. 3, pp. 195–196, Jan. 30, 1997.

[10] H. Sakauchi, Y. Nishimura, and S. Hasegawa, "A self-healing network with an economical spare channel assignment," in *Proc. IEEE GLOBECOM'90*.

[11] B. D. Venables, W. Grover, and M. H. MacGregor, "Two strategies for spare capacity placement (SCP) in mesh restorable networks," in *Proc. IEEE ICC'93*, Geneva, May 1993, pp. 267–271.

[12] M. Herzberg and S. J. Bye, "An optimal spare-capacity assignment model for survivable networks with hop limits," in *Proc. IEEE GLOBECOM'94*.

[13] R. R. Iraschko, M. H. MacGregor, and W. D. Grover, "Optimal capacity placement for path restoration in STM or ATM mesh-survivable networks," *IEEE/ACM Trans. Networking*, vol. 6, pp. 325–336, June 1998.

[14] Y. Miyao and H. Saito, "Optimal design and evaluation of survivable WDM transport networks," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 1190–1198, Sept. 1998.

[15] B. Van Caenegem, W. Van Parys, F. De Turck, and P. M. Demeester, "Dimensioning of survivable WDM networks," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 1146–1157, Sept. 1998.

[16] A. Lardies and A. Aguilar, "Planning methodology for SDH + optical networks," in *Proc. 1st Int. Workshop Design Reliable Comm. Networks, DRCN'98*, Belgium, 1998.

[17] M. Bettin, G. Ferraris, and G. Pignari, "Comparison of protection and restoration schemes for SDH networks," in *Proc. 1st Int. Workshop Design Reliable Comm. Networks, DRCN'98*, Belgium, 1998.

[18] R. Bhandari, *Survivable Networks: Algorithms for Diverse Routing*. Norwell, MA: Kluwer Academic, Nov. 1998.

[19] M. W. Maeda, "Management and control of transparent optical networks," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 1005–1023, 1998.

[20] G. Brown, W. D. Grover, J. B. Slevinsky, and M. H. MacGregor, "Mesh/arc networking: An architecture for efficient survivable self-healing networks," in *Proc. IEEE ICC'94*, New Orleans, May 1994, pp. 471–477.

[21] W. D. Grover, "Case studies of survivable ring, mesh, and mesh-arc hybrid networks," in *Proc. of IEEE GLOBECOM'92*, Dec. 1992, pp. 633–638.

John Doucette received the B.Sc. degree in mathematics from Dalhousie University, Halifax, NS, Canada, in 1992, and the B.Eng. degree in industrial engineering from the Technical University of Nova Scotia (TUNS, now DalTech), Halifax, NS, in 1996. He is currently pursuing the Ph.D. degree in electrical engineering at the University of Alberta, Edmonton, AB, Canada.

From 1996 to 1997, he was with Syncrude Canada Limited, Fort McMurray, Alberta. He joined the Network Systems group at TR*Labs* (Edmonton) in 1997 where he is a Research Engineer investigating fiber route topology evolution, optimal capacity placement, and network design for survivable networks. Since 1998 he has been a Sessional Lecturer in the Department of Electrical and Computer Engineering at the University of Alberta, where he teaches courses in telecommunication systems engineering and in probability and statistics.

Dr. Doucette is a finalist in the Alberta Science and Technology (Astech) Leaders of Tomorrow Award, October 2000.

Wayne Grover (S'74–M'76–SM'90) received the B.Eng degree from Carleton University, Ottawa, ON, Canada, the M.Sc. degree from the University of Essex, U.K., and the Ph.D. degree from the University of Alberta, Edmonton, AB, Canada, all in electrical engineering.

He was with BNR Ottawa and Edmonton in scientific staff and management before joining the senior management of TRLabs, Edmonton, Canada, in 1986. Here, he is currently Chief Scientist of Network Systems and has been instrumental in the development of TRLabs from a start-up to its present level. He is also a Professor in Electrical and Computer Engineering at the University of Alberta. He has authored over 100 journal and conference publications and patents on 25 topics.

Dr. Grover has served as Associate Editor for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS and on the Editorial Board for the *Journal Of Network And Systems Management*. In 1996 he received the TRLabs Technology Commercialization Award for the licensing of restoration-related technologies to industry. He is also the recipient of the 1996/97 McCalla Research Professorship in Engineering. In 1999 he received the IEEE Baker Prize Paper Award for his work on self-organizing broadband networks, the IEEE Canada Outstanding Engineer Award, the Martha Cook-Piper Research Prize from the University of Alberta, and the Alberta Science and Technology (ASTECH) award for technology leadership. He is also an NSERC E.W.R Steacie Memorial Fellowship holder for the years 2001–2002. He is a P.Eng. in the Province of Alberta.