New Options and Insights for Survivable Transport Networks

*Wayne Grover, John Doucette, Matthieu Clouqueur, Dion Leung, TR*Labs*, University of Alberta*
Demetrios Stamatelakis, Network Photonics Canada

ABSTRACT

This article is devoted to a selection of recent topics in survivable networking. New ideas in capacity design and ring-to-mesh evolution are given, as well as a systematic comparison of the capacity requirements of several mesh-based schemes showing how they perform over a range of network graph connectivity. The work provides new options and insights to address the following questions. How does one evolve from an existing ring-based network to a future mesh network? If the facilities graph is very sparse, how can mesh efficiency be much better than rings? How do the options for mesh protection or restoration rank in capacity requirements? How much is efficiency increased if we enrich our network connectivity? We also outline *p*-cycles, showing this new concept can realize ring-like speed with mesh-like efficiency. The scope is limited to conveying basic ideas with an understanding that they could be further adapted for use in IP or DWDM layers with GMPLS-type protocols or a centralized control plane.

CONTEXT AND OUTLINE

Ring-based survivable networking dominated in the SONET era, and DWDM optical rings are now also being deployed. Although the operation of a *single* ring is very simple, a decade of planning, operating, and growing ring-based networks has shown the *multi*-ring network planning problem to be extremely complex, and that even well designed ring networks are surprisingly inefficient and inflexible. In contrast, there is a growing appreciation today of the inherent capacity efficiency and flexibility of mesh-based survivable architectures. Indeed, many network operators are now convinced of the long-term benefits of establishing a mesh-based DWDM transport architecture, but currently have ring-based transport and may even have just deployed some optical rings. The problem is one of knowing where we would like to go, but not seeing an evolution path from where we are today. Our first topic, ring mining, views this quandary as an opportunity for deferral and reduction of capital expense during a period of ring-to-mesh conversion, through the ring mining strategy [1].

A valid and widely held intuition about mesh networking is that capacity efficiency is best on highly connected graphs. In this regard, a concern of some North American carriers considering the evolution to mesh is that their facilities graph may be very sparse. European networks often have average nodal degrees ($\bar{d}$) as high as 4.5, but some global and North American networks have $\bar{d}$ as low as 2.2. Figure 1 exemplifies these extremes of physical-layer connectivity. With a low $\bar{d}$ it is natural to question the ability of a mesh network to pay-in over rings. The *meta-mesh* idea is a refinement of existing methods for span-restorable capacity design (or for pre-planned "link protection") that specifically targets increases in capacity efficiency on sparse facilities graphs [2].

In our next topic we respond to the growing interest in mesh-based survivability by comparing the capacity requirements of the five most commonly considered mesh schemes, plus the meta-mesh approach. This particular study also provides a previously unseen perspective on how each of these schemes reacts to varying connectivity in the underlying facilities graph [3].

Our final topic, *p-cycles* [4], has its origins in the longstanding issue of "50 ms" restoration. For many years it was always assumed that one could have *either* mesh efficiency *or* ring speed, but not both. The remarkable thing about *p*-cycles is that they offer a third approach, distinct from both ring and mesh concepts, that *can* provide ring-like speed, but without sacrificing mesh-like capacity efficiency and operational flexibilities.

 0163-6804/02/$17.00

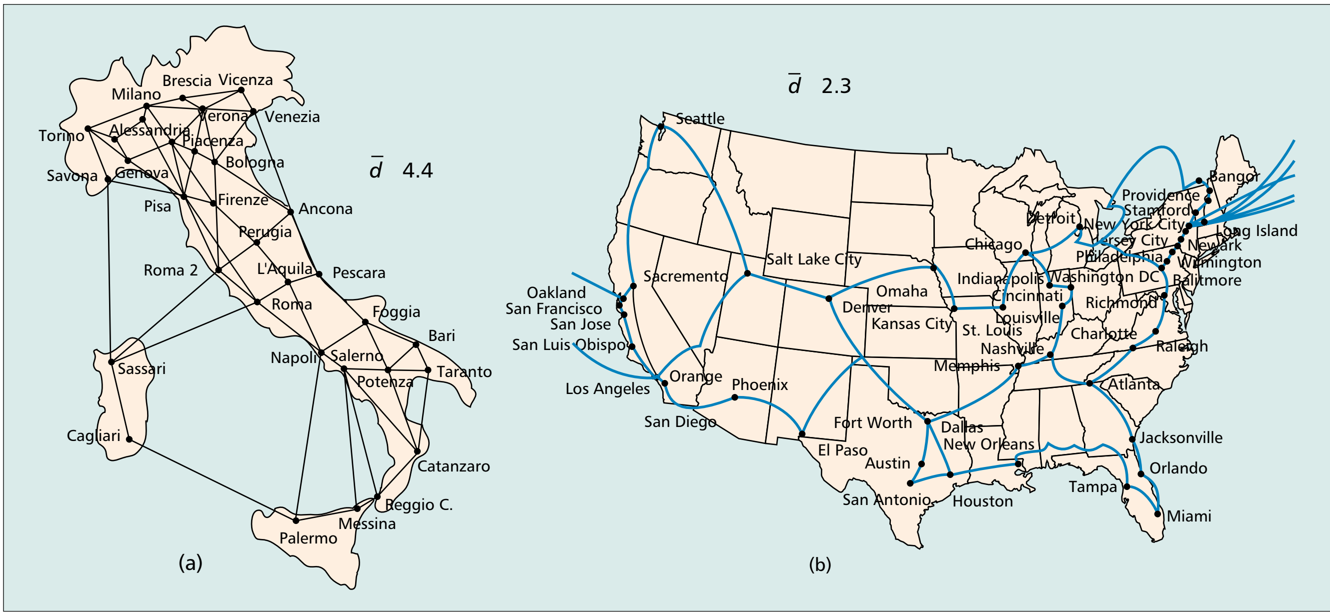

■ **Figure 1.** *Illustrating the range of physical-layer transport network connectivity.*

Ring Mining for Ring-to-Mesh Evolution

Ring mining is a generic name for strategies for ring-to-mesh evolution that involve logical *disassembly and reuse* of ring transmission capacities within a target mesh architecture. A baseline strategy for ring-to-mesh evolution is to "cap the rings" and serve growth with a new mesh overlay. Eventually with growth the network will be almost all mesh-based and demands served in the residual rings can be rolled into the mesh. In [1], however, we have looked at the potential for serving ongoing growth in demand without any additional capacity by mining the inefficiently used protection and working capacity resources in existing rings. Additional service-bearing capacity is extracted through routing and restoration redesign using mesh principles within the *existing* ring capacities. Thus, the installed capacity of rings is viewed as a sunk investment to be "mined." The enticing prospect is that for certain operators ring mining could very significantly defer and partly eliminate new capital expenditures for incremental transport equipment.

In researching the potential benefits of the ring mining idea, a first simple approach is to ask: *To what extent could an existing fully loaded ring network satisfy additional demand simply by providing access to its span capacities for use as a restorable mesh architecture?* We call this pure ring mining in that there would be literally no new capacity added while sustaining growth by conversion to a mesh mode of operation. The pure ring mining potential is easily assessed as an optimization problem. If the potential benefit is high, a next step is to allow for possible costs associated with converting ring nodes for full access by mesh cross-connects and to permit selected additions of new capacity to unlock even more ring capacity as the ring-to-mesh transition proceeds. This leads to a complete minimum cost transitional planning model as detailed in [1].

This planning model was studied on 17 efficient ring network designs from separate research on ring design methods. Results show a dramatic potential for ring mining in some cases. *In more than 35 percent of test cases a complete doubling or more in demand could be supported with no additional capacity investment*. In one case a total growth to 290 percent over the initial demand served by the ring network was sustained simply through conversion to mesh with no added capacity. In the minimum-cost evolution models where selected capacity additions were allowed we saw complete deferral of new expense for up to 50–200 percent service growth followed by an approximately linear cost at a point when all growth is finally being picked up by mesh growth. In comparison, the cap-and-grow strategy shows a large immediate cost for the new mesh overlay and no deferral in the cost curve for ongoing growth. Even at the highest Internet growth rates, the findings suggest that some operators could see a one-time opportunity from ring-to-mesh transition, to *grow their revenue base for a year or more without major capital additions* for transport capacity. Analysis shows that such large growth multipliers arise from:

- Access to the 100 percent ring protection capacity
- Unlocking stranded ring working capacity
- Shortening of the working path routes

Mesh access to the ring capacity is through ring ADM nodes that have been converted for ring mining. There are different technical means of arranging such access, but the important characteristics for assessing the network benefits are the total ring capacity made accessible to a co-located cross-connect, and the cost of arrangements to make this capacity accessible. Ideally, the protection channel is accessible through an extra traffic feature and the ADM has 100 percent add-drop access. An ADM is then converted for ring mining by freezing it in a 100 percent drop configuration and connecting its extra traffic ports to the cross-connect. The converted

Recently deployed long-haul WDM rings may be especially interesting candidates for ring mining because most of their cost is in the extensive line systems (which are kept and more efficiently re-used), not the ring terminals themselves.

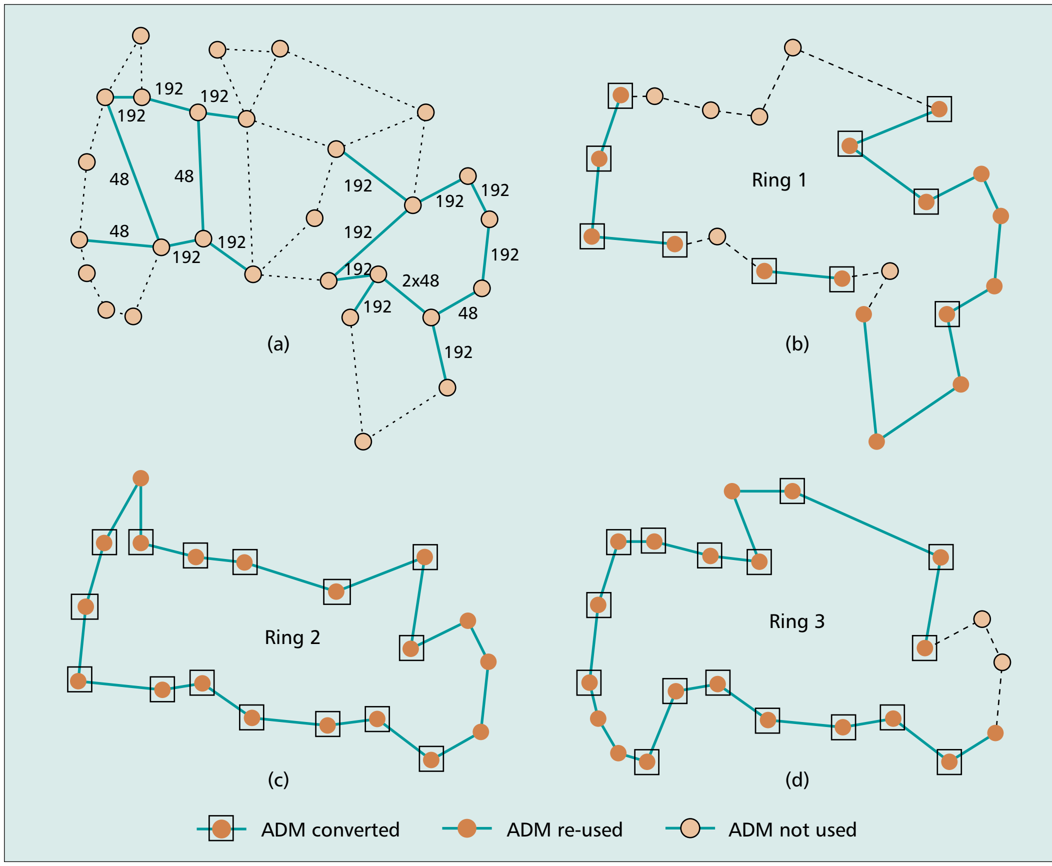

Figure 2. *Example of a detailed ring-mining plan.*

ADM thus emulates four unprotected lightwave terminals at the ring line rate. Alternately there may be salvage or reuse value for permitting outright ADM removal and re-termination of the repeatered span-lines directly on the cross-connects. Minimum-rearrangement planning models are also under study.

Figure 2 is a selected result from the most complete transition model showing the plan for mining the three largest rings of a seven-ring test case. Figure 2a shows the layout of selected additions of new capacity. Figs. 2b–d show which ADMs are converted for mesh access and which stay in place to serve as degree-2 nodal elements within the resulting mesh network. Because the research was based on initial ring designs that were highly optimized to a known demand pattern, the potential for ring-mining of real networks that are only partway to their planning horizon, or less than optimally loaded, may be even greater. Recently deployed long-haul WDM rings may be especially interesting candidates for ring mining because most of their cost is in the extensive line systems (which are kept and more efficiently re-used), not the ring terminals themselves. Each situation can be studied in detail with the methods in [1].

The Meta-Mesh Concept

In the *meta-mesh* concept [2] the spare capacity in a low-degree mesh network using span restoration (or link protection) is reduced by paying special attention to the manner in which chain subnetworks are structured and restored. The approach views a sparse network as a *meta-mesh of chain sub-networks*. To explain this, we begin with the practical observation that a sparse (but still bi-connected) network as shown in Figure 1b tends to contain *chains* of degree-2 nodes. A property of a degree-2 node under span restoration is that the spare capacity on each side of the node must match (or exceed) the working capacity on the other side of the same node. The degree-2 topology simply dictates that the spare capacity must be sufficient to support *loop-back* of the failed working capacity on the other side of the node, just as in a BLSR ring.

In a conventional span-restorable design the spare capacity in a chain is set by the largest working capacity of any span within the chain, as if the chain was part of a BLSR ring, given the degree-2 considerations just mentioned. Under failure of a *chain span*, all the failed working capacity is looped-back in the opposite direction on the chain until it encounters a degree-3+ node, called an *anchor node*, at the end of the chain. Between the anchor nodes a mesh-like diversity of restoration flows then follows within the wider network as a whole. Figure 3a-b illustrates this conventional capacitation of spare capacity, and the ring-like loop-back phase of restoration that normally occurs.

In the meta-mesh, however, we make a distinction between types of demand in the chain spans. Demands that originate and/or terminate at a chain node are referred to as local or *intra-chain*. Any demand flow that is traveling through the chain as a whole having its origin and destina-

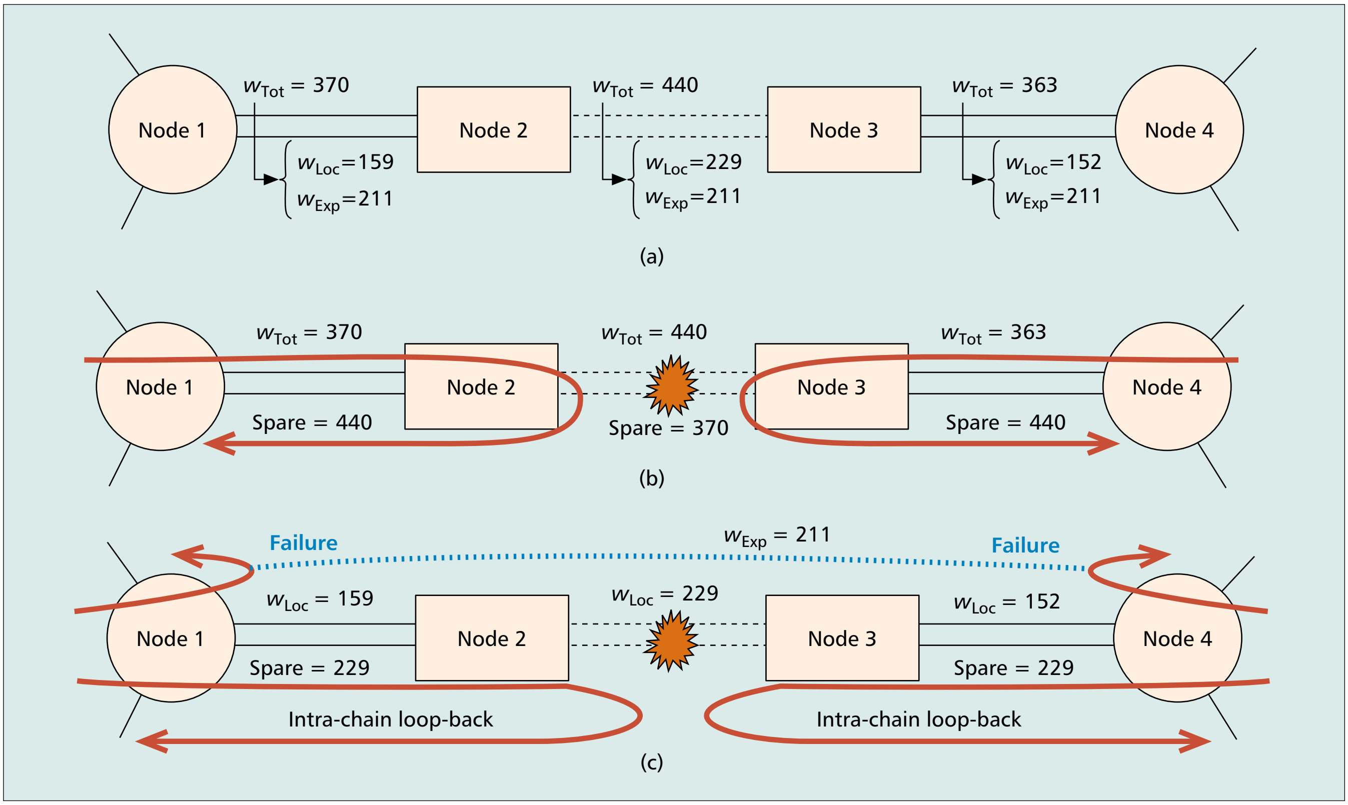

Figure 3. *Reduced spare and altered restoration of chain subnetworks in meta-mesh designs.*

tion outside of the chain, or at one of the anchor nodes, is referred to as *express flow*. The meta-mesh design will provide spare capacity only sufficient to provide for loop-back of the intra-chain demands. No spare capacity is provided at all for the loop-back of express flows to the anchor nodes. Instead, we allow the express demands simply to *fail back* to the anchor nodes, as in Fig. 3c. The rationale is that because the demand composition of the express flow crossing the failure span is not changed by add/drop effects at the chain nodes, it can be returned to the anchor nodes for mesh restoration by simply *letting those demands fail all the way back to the anchor nodes*.

Let us now move up from the view of a single chain to view the chain as a constituent part of a meta-mesh network. The meta-mesh is the topology that arises when all direct spans and chain sub-networks are viewed equivalently as edges of another graph, namely, the meta-mesh graph. The meta-mesh is thus comprised of only degree-3 or higher nodes connected by direct physical spans or chain sub-networks — both are just logical spans of the meta-mesh. The significance of the meta-mesh is that it is only at this level of abstraction that true mesh-type spare capacity sharing efficiencies can arise. As an example, the meta-mesh graph for the network of Fig. 1b has only 15 nodes and 23 "spans" ($\bar{d}$ = 3.08), while the complete network has 55 nodes and 62 spans ($\bar{d}$ = 2.25). By its nature, *the meta-mesh graph is always of at least degree 3*. Using a well known lower bound on capacity redundancy in span-restorable networks, $1/(\bar{d} - 1)$ [3], the potential effect in going from $\bar{d} \sim 2$ to $\bar{d} \sim 3$ is to halve the redundancy. Therefore, to the extent that any part of the restoration problem can be logically shifted to the meta-mesh, there may be capacity savings. This is the main theoretical idea behind the meta-mesh concept.

Thus, by letting express flows *fail* back (rather than *loop*-back) to the anchor nodes, we save spare capacity in the chain and *we refer the express flow entirely out to the meta-mesh graph for restoration*. Express flows through chains are now treated entirely with mesh-based restoration principles and never enter into the spare-capacity sizing of a chain. The same is *not* an option for local flows because the makeup of the local demand flow is altered in each span by add/drop actions on the local flow. Explicit loop-back is therefore required to return the demands from the point of the failure to the anchor nodes in the composition that they had at the failure span.

The changes to the restoration mechanism are minor. If a span cut occurs within a chain, the adjacent OADMs perform their loopback function as before, but now only local flows are even routed through the ADMs. Express flows logically or physically bypass the ADMs on either a separate wavelength or fiber, thereby also saving ADM core bandwidth. At the same time a failure causes the ADMs to loop-back, the express flow failure condition simply propagates out to both anchor nodes. "Loss of Signal" or "Alarm Inhibit Signal" may alert the anchor nodes of the failure. At the anchor nodes, failed-back express channels and looped-back local channels are then unified from a restoration viewpoint and appear as a single logical span failure for restoration among the meta-mesh graph of OXC nodes. Only the nodes of the meta-mesh require optical-cross-connect functionality. Chain nodes go on using simpler ADMs.

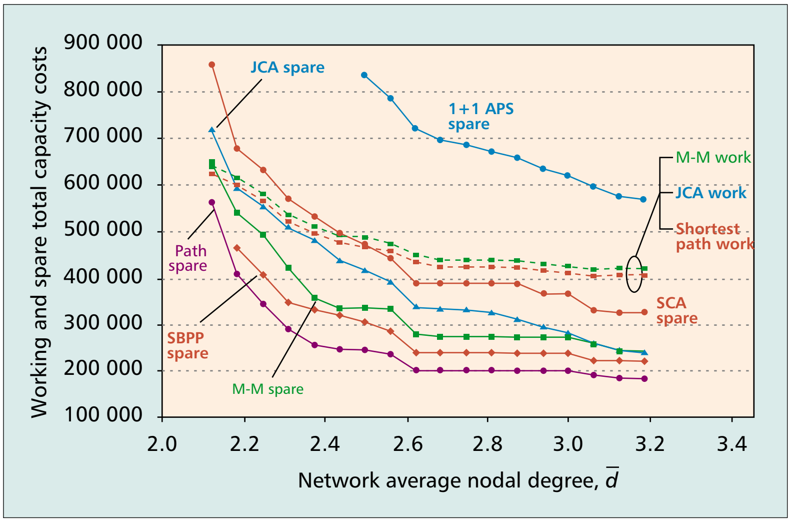

Figure 4. *Breakdown of working and spare capacity versus network average nodal degree.*

The corresponding differences in the overall network capacity design are explained and implemented in [2] and selected results are discussed further below.

Comparative Capacity Requirements and Topology Dependence of Mesh Schemes

Let us now consider how the most widely considered types of mesh survivability schemes compare in their intrinsic capacity requirements, and given the range in topology characteristics in Fig. 1, how each scheme depends on the physical-layer graph connectivity. In [3], these mesh schemes were tested on the basis of comparing optimal designs to serve a common demand matrix with 100 percent restorability for any single complete span cut. Listed in general order of increasing capacity-efficiency, the schemes considered are:

1+1 Path Protection (1+1 APS). Here, demands are shortest path routed and an equal allocation of capacity reserved on the next shortest span-disjoint route. A tail-end switch selects the surviving (or best) signal. In certain cases where no route is disjoint from the shortest path, the shortest cycle containing both end nodes is found with Suurballe's algorithm or an equivalent. 1+1 (or equivalently, UPSR) service is the only scheme other than *p*-cycles (below) that literally assures 50 ms restoration and is typically used for the most critical services only.

Span Restoration — Spare Capacity Assignment (SCA). Here demands are shortest path routed followed by optimization of spare capacity to support 100 percent restorability with minimum total *spare* capacity. Restoration occurs via k-shortest paths, e.g., re-routing between the end nodes of the failure span. A dynamic restoration mechanism can adaptively construct the replacement paths on demand [5], or the switching arrangements can be preplanned within the same spare capacity (in which case this scheme is often called "link protection.")

Span Restoration — Joint Capacity Assignment (JCA). Here the choice of route taken by each working demand is optimally coordinated with the spare capacity assignment decisions so as to minimize *total* working and spare capacity. The options for dynamic or pre-planned restoration are the same as for SCA. The routes of the working paths typically deviate very little from shortest paths but extra capacity savings arise over SCA primarily from working-flow leveling effects.

Meta-Mesh (M-M). This is a variant of JCA where express flows through chain subnetworks are provided with logical express bypasses. The same basic restoration mechanism applies as for SCA, with minor extensions at the anchor nodes to integrate looped-back chain flow with failed-back express flow for span restoration within the logical meta-mesh of higher degree nodes [2].

Shared Backup Path Protection (SBPP). Here demands are shortest path routed and a single fully disjoint backup route pre-selected for each. Spare capacity on backup routes is shared across working path demands that have no spans or other physical-layer dependencies in common, and hence should not need the backup capacity simultaneously. A real-time signaling phase seizes and cross-connects the shared capacity. This scheme is currently favored in IETF deliberations for MPLS-layer protection and MPLS-controlled optical path protection. The scheme is logically identical to *backup VP protection in ATM* [6]. In the MPLS layer, a strategy of controlled over-subscription of restoration capacity is possible to further optimize capacity use [7], whereas in the optical layer the problem is the special case of [7] where there is no over-subscription.

Path Restoration (Path). Here demands are shortest path routed and just enough spare capacity placed to support a multi-commodity maximum-flow type end-to-end re-routing of all simultaneously failed working paths. A centralized or distributed re-routing mechanism [8] deploys a collectively coordinated set of replacement paths in response to the specific failure. Surviving ("stub") portions of failed working paths can be re-used for restoration.

We compare these schemes in a way that also shows their topology dependencies by using a family of 18 test networks derived from a single high-degree *master network* of 32 nodes and 51 spans ($\bar{d}$ = 3.18). Progressively sparser networks were derived from the master by random removal of one span at a time subject to retaining bi-connectivity. This provides a reasonably continuous variation of $\bar{d}$ while keeping nodal positions and the end-to-end demand pattern common over all test networks. The mathematical models for design of each network type, the demand patterns used, and computational aspects are provided in [3, 9–12].

Figure 4 shows the working and spare capacity requirements of each scheme over the range of test networks. Total network capacity required to serve the common demand pattern is the sum of these curves, and their ratio (spare/working) is called the *redundancy*. 1+1 APS is extraordinarily redundant. It is never less than 140 percent redundant and surpasses 200 percent on the sparse graphs. The span-restorable schemes are much more capacity-efficient, with JCA improving considerably over SCA through relatively small changes to the routing of working paths that tend to level out the working capacity quantities. Notice there is hardly any difference between the working capacity requirements of any of the schemes. Path restoration is the most efficient scheme of all. It is remarkable to note that at its lowest, the path spare capacity corresponds to about 40 percent redundancy for 100 percent restorability.

SBPP is generally intermediate in efficiency between the meta-mesh designs and path restoration, consistent with its end-to-end orientation, but its non-failure-specific response which cannot re-use working capacity on the failed paths.

It is interesting that in comparing span-oriented schemes (SCA, JCA) to path oriented schemes (SBPP, Path), meta mesh shows capacity efficiencies and a rate of efficiency increase with connectivity that is more like a member of the path group, although its mechanism remains span-restoration. This is consistent with interpreting the express bypass of the chain subnetworks as a specific type of partial path restoration. The meta-mesh designs provide an attractive option for networks with $2.4 < \bar{d} < 2.8$. In this region they are essentially as efficient as SBPP, although they require only a span restoration (or pre-planned link protection) mechanism. The capacity savings relative to JCA/SCA peaks at $\bar{d} \cong 2.4$ with up to 12 percent reduction in total network capacity, 30 percent reduction in spare channel counts, and 21 percent reduction in working channel counts. Note that at both extremes of high and low $\bar{d}$ the meta-mesh curve essentially merges with JCA. This is because toward the limit of $\bar{d} = 2$, we have long chains dominated by local flows. At the other extreme there are no chains and MM is identical to JCA.

In practice, the steeper rate of drop on the path-oriented curves suggests that if an initially sparse network is undergoing topology additions to increase its connectivity, the path schemes or meta-mesh will give a faster pay-back than span-oriented schemes as the topology is enriched. It is also quite notable that aside from 1+1 APS and JCA, the decrease of spare capacity beyond $\bar{d} \sim 2.6$ is extremely slow compared to the benefit from increases in $\bar{d}$ in the 2.2 to 2.5 range. The exact point of this cutoff effect is somewhat network- and demand-specific, but its existence can provide planners with a valuable guideline for setting targets for topology evolution.

p-Cycles: Ring Speed with Mesh Efficiency

Aside from 1+1 APS, in each of the mesh schemes above, survivability is achieved through the active formation of re-routing paths out of a pool of *normally unconnected* spare capacity. What if protection paths could be obtained on-demand by breaking into some *pre-connected* storage structure of spare capacity? Rather than *assembling* spare capacity into paths as needed, we would form fully-connected structure(s) of spare capacity and break into them to obtain protection paths. This seemingly unusual approach is something we pursued in a line of research on mesh "preconfiguration" from 1995 on, initially with the aim of obtaining an average-case speed-up for mesh networks where cross-connect time was the limiting speed factor. This lead to realization of the *p*-cycle concept, where in fact *all* spare capacity is preconnected and only two ring-like switching actions are needed for any path being protected.

The most striking property is that one gets the *best* characteristics of ring and mesh together. *p*-cycles retain ring-like switching characteristics but can be designed to capacity effiencies that are essentially the same as SCA or JCA mesh designs. This is something that after 10 years of research on either ring or mesh approaches, many had surmised to be fundamentally impossible to have together. Ring-speed arises because only two nodes perform any real-time actions for each restored signal unit, and such actions are fully predetermined before failure. The surprising capacity efficiency is the less obvious property but it is ultimately attributable to the aspect of protecting *straddling* span failures as well as *on-cycle* failures. This seemingly small difference between a ring and a *p*-cycle actually leads to major differences in protection capacity requirements. Our present aim is to convey the basic idea of *p*-cycles, make it intuitive why they do offer such high efficiency, and to point out some other advantageous properties.

The simplest way to think about *p*-cycles is that they are like rings, but with support for the protection of straddling failure spans as well as the usual ring protection of spans of the ring itself ("on-cycle" failures). A straddling span is one that has its end-nodes on the *p*-cycle, but is not itself part of the *p*-cycle, like a chord on a circle. The key distinguisher of *p*-cycles as opposed to any kind of ring or graph cycle covers is the protection of straddling spans which themselves can each bear two units of working capacity and *zero* spare capacity. (Cycle double covers under the recent name of protection cycles still match every working unit of capacity with 100 percent redundant spare capacity and protect only on-cycle failures.)

Figure 5a shows an example of a *p*-cycle. In Fig. 5b an on-cycle failure occurs and the surviving arc of the cycle is used for restoration. Like rings (specifically the BLSR), a *p*-cycle protects against on-cycle failures through loop-back to protection. The failed signals effectively turn away from the break and go the other way around the cycle. Figure 5c, however, shows how the same *p*-cycle is accessed to support restoration of a straddling span failure as well. In fact, the efficiency of covering the case shown in Fig. 5c is double that of an on-cycle failure because two restoration paths are available from each unit of *p*-cycle protection capacity. In contrast, any ring yet proposed (UPSR, BLSR, FDDI) all provide at most one restoration path per unit of ring protection capacity and protect only on-cycle failures.

Initially this seems to be a minor technical difference, so how big can the network effects be? In fact they are rather dramatic. By allowing the same protection capacity as a ring to also protect straddling span failures, we have found that sets of *p*-cycles can cover all span failures with three to six times less capacity than required with rings. The initial report of these mesh-like capacity efficiencies [4, 13] has been confirmed by several groups

Notice there is hardly any difference between the working capacity requirements of any of the schemes. Path restoration is the most efficient scheme of all. It is remarkable to note that at its lowest, the path spare capacity corresponds to about 40 percent redundancy for 100 percent restorability.

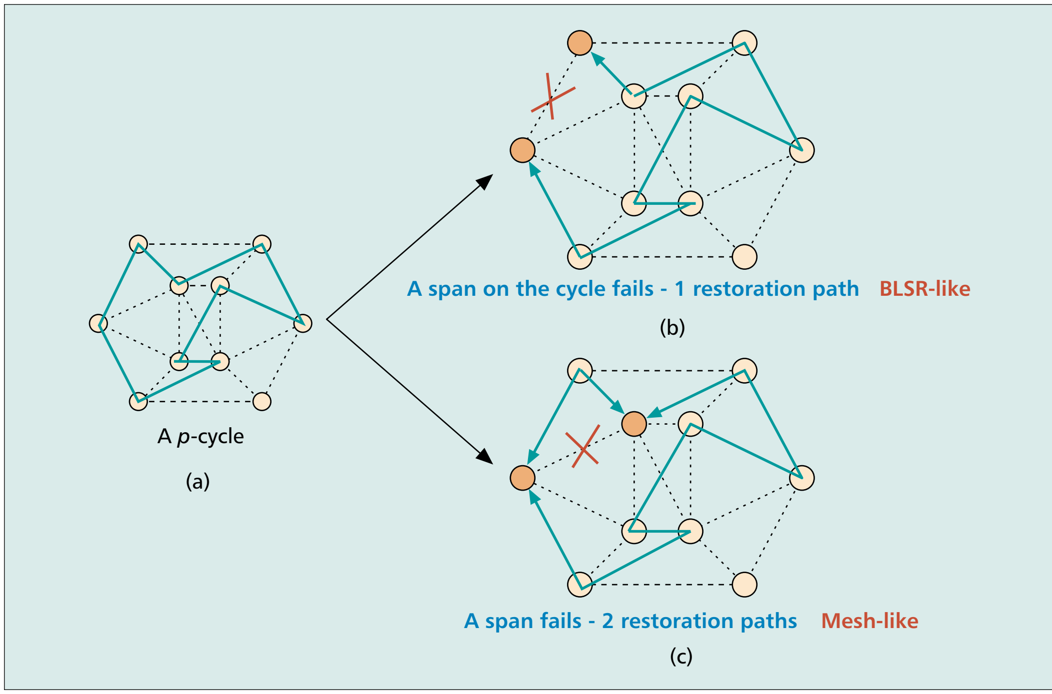

■ **Figure 5.** *A* p-*cycle protecting both on-cycle and straddling failures.*

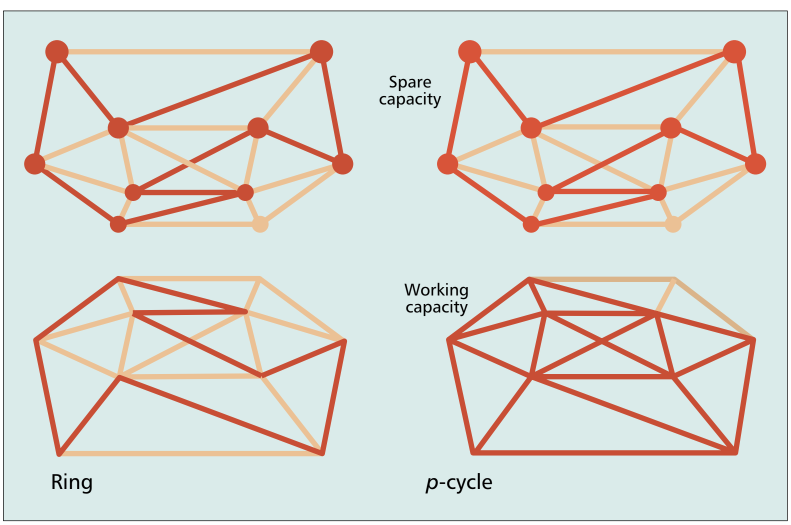

■ **Figure 6.** *Clamshell diagram illustrating ring and* p-*cycle span coverages for same spare capacity.*

in recent years, notably in [14] where WDM network design with *p*-cycles was considered with and without wavelength conversion. Reference [14] found that redundancy levels as low as 30 percent, and most often below 70 percent, could be achieved in the COST 239 European study network, depending on the demand pattern and the allowed circumference of the *p*-cycles.

In our own efforts to double-check results and to intuitively understand why *p*-cycles are so efficient, we produced the "clamshell" diagram in Fig. 6. The upper plane shows an identical investment of nine units of spare capacity, connected in a cycle. The corresponding spans that obtain working capacity protection are shown in the lower plane. On the left the spare capacity is interpreted as a nine-hop ring, so there is an identical top and bottom plane. On the right, the same investment in spare capacity is considered as a *p*-cycle where we see that it protects 19 spans (nine on-cycle relationships and ten straddling relationships). Moreover, the *p*-cycle provides two restoration paths for each of its straddling spans. The *p*-cycle thus protects 2*10 + 9 = 29 units of working capacity, 3.2 times that of the corresponding ring. In the limit of a full set of straddling span relationships, an *N*-hop *p*-cycle can protect up to $N(N-2)$ units of working capacity, making it up to $N-2$ times more efficient than a corresponding ring. If *N* could be up to 16 we can see why *p*-cycles can conservatively reach at least 3 to 6 times the efficiency of rings.

p-cycles have other advantages as well. One is that for straddling failures, protection paths are on average half the *p*-cycle circumference, whereas in rings protection paths are essentially the full circumference of the ring. Unlike rings, *p*-cycles can also be formed from individual units of spare capacity on optical cross-connect systems and logically rearranged to adapt to growth patterns as needed. In contrast rings commit a whole module of working and spare capacity and embody a structural association between the protection capacity and the working demands that they protect. Demands must be routed within the ring, whereas with *p*-cycles the working paths can be freely routed on shortest paths over the physical graph (or any other route) giving a significant reduction in *working* capacity as well as the considered reductions in spare capacity. Also, because *p*-cycles are formed in the spare capacity only, they can be adapted to suit the working path layer at any time, without any impact on working demands. In contrast, rings assert the routing that demands must take, rather than adapting to the routes they want to take. Jointly optimizing the working path routing with *p*-cycle placement (part of our current research) should yield even further capacity savings.

We envisage a process that continually adapts a set of "unseen" *p*-cycles to the actual accumulations of working demand flow. This is the opposite of the provisioning framework in rings, where new service paths must be explicitly threaded through the available rings. It would also take the need to make explicit protection arrangements away from the service provisioning process, unlike SBPP, where an explicit protection path must be arranged for every path provisioned. A service status indicator can simply tell the hidden *p*-cycle adaptation layer to include a demand flow in its protection scope or not. The planning of *p*-cycles for WDM can also be done in a way that only the p-cycle access points need wavelength conversion [14]. The working lightpaths below can be provisioned with a single wavelength assignment.

To delve further [15] explains the adaptation of *p*-cycles to the MPLS layer and introduces the additional concept of *node encircling p*-cycles. In [4] an ADM-like nodal device for *p*-cycle networking is proposed, and [13] describes a self-organizing protocol embedded in the physical layer that adaptively forms *p*-cycles for restorability of the current working demand flows.

Closing Comments

Research on survivable transport networking is conservatively at least 15 years old, originating in the mid-80s when fiber network outages reached some crisis proportions. A variety of ring and mesh-based approaches have been well developed since then, but it is interesting how innovations and wholly new paradigms are still being uncovered. Ring mining, meta-mesh, and *p*-cycles are examples. We believe the focus on *capacity*-efficient survivable design will continue to translate into *cost* efficiencies as user demand keeps growing. A specific future vision we propose is one in which lightpaths are established on demand between routers, but *without* the shareability and topology database exchanges required for explicit setup of protection paths at provisioning time. Instead, guaranteed and best-effort protection services are provided to each working path by an otherwise unseen but self-organizing adaptive layer of *p*-cycles.

We envisage a process that continually adapts a set of "unseen" p-cycles to the actual accumulations of working demand flow. This is the opposite of the provisioning framework in rings, where new service paths must be explicitly threaded through the available rings.

References

[1] M. Clouqueur *et al.*, "Mining the Rings: Strategies for Ring-to-Mesh Evolution," DRCN 2001, Budapest, Hungary, Oct. 2001, pp. 113–20.

[2] W.D. Grover and J. Doucette, "Increasing the Efficiency of Span-restorable Mesh Networks on Low-connectivity Graphs," *DRCN 2001*, Budapest, Hungary, Oct. 2001, pp. 99–106.

[3] J. Doucette and W.D. Grover, "Comparison of Mesh Protection and Restoration Schemes and the Dependency on Graph Connectivity," *Proc. 3rd Int'l. Wksp. Design Reliable Commun. Net. (DRCN 2001)*, Budapest, Hungary, Oct. 7–10, 2001, pp. 121–28.

[4] W.D. Grover and D Stamatelakis, "Bridging the Ring-mesh Dichotomy with P-cycles" (invited paper), DRCN 2000, Munich, Germany, Apr. 2000, pp. 92–104.

[5] W.D. Grover, "Self-organizing Broadband Transport Networks," *Proc. IEEE*, vol. 85, no. 10, Oct. 1997, pp. 1582–611.

[6] R. Kawamura, K. Sato, and I. Tokizawa, "Self-healing ATM Networks Based on Virtual Path Concept," *IEEE JSAC*, vol. 12, no. 1, Jan. 1994, pp. 120–27.

[7] W.D. Grover and Y. Zheng, "VP-based ATM Network Design with Controlled Over-subscription of Restoration Capacity," *DRCN 1998*, Brugge, Belgium, May 1998, paper O33.

[8] R.R. Iraschko and W.D. Grover, "A Highly Efficient Path-restoration Protocol for Management of Optical Network Transport Integrity," *IEEE JSAC*, vol. 18, no. 5, May 2000, pp. 779–93.

[9] B. Van Caenegem *et al.*, "Dimensioning of Survivable WDM Networks," *IEEE JSAC*, vol. 16, no. 7, Sept. 1998, pp. 1146–57.

[10] 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. Net.*, vol. 6, no. 3, June 1998, pp. 325–36.

[11] M. Herzberg and S. Bye, "An Optimal Spare-capacity Assignment Model for Survivable Networks with Hop Limits," *Proc. IEEE Global Telecommun. Conf. (GLOBECOM)*, 1994, pp. 1601–607.

[12] Y. Xiong and L.G. Mason, "Restoration Strategies and Spare Capacity Requirements in Self-healing ATM Networks" *IEEE/ACM Trans. Net.*, vol. 7, no. 1, Feb. 1999, pp. 98–110.

[13] W.D. Grover and D. Stamatelakis, "Cycle-Oriented Distributed Preconfiguration: Ring-like Speed with Mesh-like Capacity for Self-planning Network Restoration," *Proc. IEEE ICC '98*, Atlanta, June 1998, pp. 537–43.

[14] D.A. Schupke, C.G. Gruber, and A. Autenrieth, "Optimal Configuration of p-Cycles in WDM Networks," submitted to *IEEE ICC '02*.

[15] D. Stamatelakis and W.D. Grover, "IP-Layer Restoration and Network Planning Based on Virtual Protection Cycles," *IEEE JSAC*, vol. 18, no. 10, Oct. 2000, pp. 1938–49.

Biographies

Wayne D. Grover [F] (grover@trlabs.ca) is chief scientist/network systems at TR*Labs*, and a professor in electrical and computer engineering at the University of Alberta. He received his B.Eng. from Carleton University, Ottawa, Canada, his M.Sc. degree from the University of Essex, U.K., and his Ph.D. from the University of Alberta, Edmonton, all in electrical engineering. He was on the scientific staff and management at BNR (now Nortel) for ten years before joining the senior management of TR*Labs*, where he served as the founding technical VP. He is a recipient of the IEEE Baker Prize Paper Award for his work on Self-organizing Broadband Networks and the IEEE Canada Outstanding Engineer of the Year Award (1999), and he is the holder of an NSERC E.W.R. Steacie Memorial Fellowship for 2001-2002. He is a P.Eng. in the province of Alberta.

John Doucette (john.doucette@trlabs.ca) received a B.Sc. from Dalhousie University and a B.Eng. from the Technical University of Nova Scotia. He is a research engineer with TR*Labs* and is pursuing a Ph.D. at the University of Alberta, where he has been an instructor since 1998. His research interests include optical networks, restoration and protection, network topology, and network optimization.

Matthieu Clouqueur (matthieu.clouqueur@trlabs.ca) is a graduate from Ecole Nationale Supérieure de l'Electronique et de ses Applications (National School of Electronics and Electrical Engineering), France. Since joining TR*Labs* in 1998 he has worked on blind equalization for non-linear channels and on transport network survivability. He is now a research engineer at TR*Labs* and is pursuing a Ph.D. degree at the University of Alberta. His research is on the availability of service in mesh-restorable networks.

Dion Leung (dion.leung@trlabs.ca) received his B.Sc. degree in electrical engineering from the University of Alberta in 2000. Funded by the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Alberta Informatics Circle of Research Excellence (iCORE), he joined TR*Labs* as a M.Sc. candidate and transferred into a Ph.D. program in 2001. His current research is to investigate issues on mesh restorable transport networks with uncertainty in demand forecasts.

Demetrios Stamatelakis (Demetrios.Stamatelakis@NetworkPhotonics.com) received his B.Sc. and M.Sc. degrees in electrical engineering from the University of Alberta in 1994 and 1997, respectively. From 1997 to January 2001 he was a research engineer in the TR*Labs* Network Systems group, working on mesh-network preconfiguration concepts. He currently works for Network Photonics in Calgary, Alberta, an optical switching company, in the Optical Network Engineering group. He is a P. Eng. in the Province of Alberta.