1st Int'l Conference on Optical Communications and Networks (ICOCN'02), Singapore, Nov.11-14, 2002, pp. 305-308.

Understanding *p*-Cycles, Enhanced Rings, and Oriented Cycle Covers

(invited paper)

Wayne D. Grover, *Fellow IEEE*

TR*Labs* and Department of Electrical and Computer Engineering, University of Alberta,
c/o TR*Labs* #800, 10611-98th Avenue, Edmonton, Alberta, Canada T5K 2P7 Edmonton, Alberta, Canada
Contact: grover@trlabs.ca

***Abstract*-This paper explains the important conceptual and technical differences between the method of *p*-cycles and two other recent advances involving a cyclic orientation to protection. These are enhanced rings and cycle double covers. The most fundamental difference that is unique to *p*-cycles is the aspect of straddling span failure protection. This enables mesh-like efficiency levels at well under 100% redundancy. In contrast enhanced rings and advanced cycle cover methods are both seeking to reduce span overlaps in what is otherwise a purely ring-like logical paradigm in which 100% redundancy remains the best that can possibly be achieved.**

I. Introduction and Objective

In recent years independent research efforts have been advancing a number of new approaches to survivable network architecture. One of these, the technique of *p*-cycles, has been of primary interest to the author. *p*-Cycles are remarkable and interesting because they combine the simplicity and speed of ring-like protection switching and yet, under fairly simple design optimization, they result in networks that as a whole exhibit mesh-like capacity efficiency. The key concepts and properties of *p*-cycles will shortly be reviewed. In parallel with the work on *p*-cycles, however, other advances have been made in aspects of ring-based protection itself. Notable amongst these is a recent patent by Fee on what we call "enhanced rings" and the theoretical research by Ellinas et al. on unidirectional or "oriented" cycle covers of a graph. All three of these advances in one way or another involve some form of cyclic layout of protection capacity on a network graph and are easily confused at a superficial level. Indeed, coincidentally similar names used by Grover and Ellinas has inadvertently even contributed to the confusion. So it is fair to ask: are these three advances really the same? If not, how do they relate? This paper hopes to provide a timely explanation of the important ways in which these schemes differ and, in particular, why neither of the other schemes are equivalent conceptually, functionally, or in terms of efficiency, to *p*-cycles. Such clarification is needed to ensure the industry and research community interested in *p*-cycles avoids confusion and false assumptions of equivalence arising from the superficial similarities or naming between these schemes.

II. *P*-Cycles

In protecting a network with *p*-cycles, one forms cyclic pre-connected closed paths of spare capacity while allowing working paths to take the shortest direct route over the facilities graph. *p*-Cycles are formed in advance of any failure and the switching actions required in real time are completely pre-planned and no more complex than that of a line-switched ring. Despite the similarity to rings in terms of switching function and the fact that both use a cycle on the network graph for their topology, *p*-cycles are unlike any survivable ring-based system heretofore developed (including 1+1, UPSR, BLSR and FDDI) in that they protect both *on-cycle* and *straddling span* failures. This initially seems to be a rather minor technical difference. When its implications are fully worked through it turns out that the aspect of straddling spans, and the direct routing of working paths, are the keys to obtaining mesh-like network efficiency while employing ring-like protection structures.

To illustrate the *p*-cycle concept, and convey an appreciation of why such high capacity efficiency is reached, Fig.1(a) shows a typical transport network example of 11 nodes and 24 spans. In Fig. 1(b) a cycle of ten hops is configured as an example. If this cycle is used as the layout for a BLSR ring then each unit of protection capacity on the ring protects the same amount of working capacity against failure of the ten spans underlying the ring itself, shown in (c). These are called the "on-cycle" failure spans. But if we use the same unit-capacity cycle in (b) as a *p*-cycle it means that we will also access it for protection against the failure of a span such as (6-10) which is said to "straddle" the *p*-cycle. A straddling span has both its end nodes on the *p*-cycle, but is not itself part of the cycle. In the example, the *p*-cycle has eight such straddling span relationships, shown in (d). Straddling span relationships have twice the leverage of an on-cycle span in terms of capacity efficiency because when they fail, the cycle itself remains intact and can thereby offer *two* protection paths for each of unit of protection capacity. Note, in the case of a span such as (9-8), that the straddling spans are not limited to being only inside the *p*-cycle perimeter on the graph.

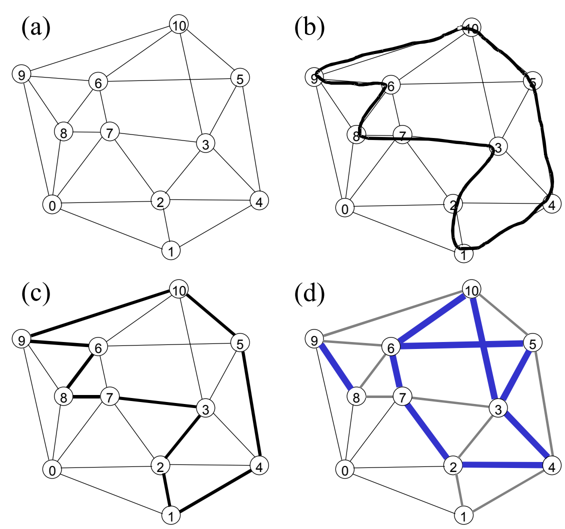

Fig. 1. - (a) a network, *G*, (b) a 10-hop cycle of unit-spare capacity, *x*, on *G*, (c) ten spans obtaining ring-like protection from *x*, (d) eight additional spans each receiving protection for two units of working capacity if is *x* is a *p*-cycle.

Fig. 2(a) shows the resultant impact on efficiency (or conversely the redundancy). The (w, s) indications note the working and spare composition of the spans, respectively. The example *p*-cycle can achieve 38% redundancy (if all its straddling spans have two units of working capacity equipped). The corresponding ring is 100% redundant.[1] In Fig. 2(b) and (c) the real time operation is illustrated. In (b) a span on the cycle breaks and the surviving arc of the cycle is used for protection. This is functionally like a unit-capacity BLSR.

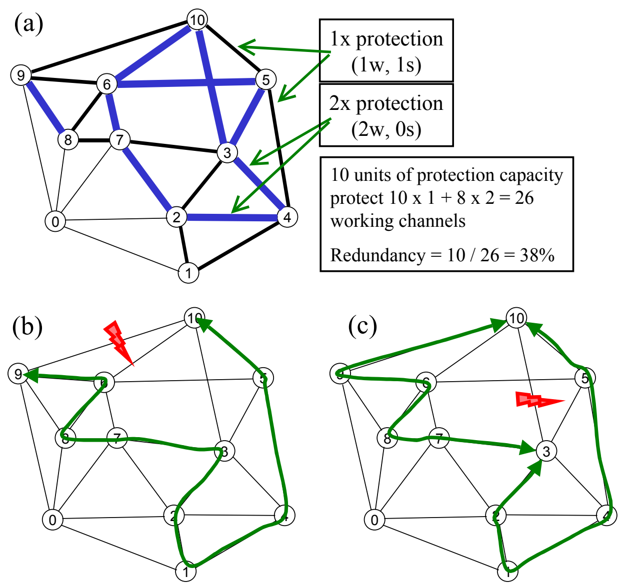

Fig. 2. - (a) the compete set of protection relationships offered by *p*-cycle, *x*, and corresponding protection redundancy, (b) example of an on-cycle failure, (c) example of a straddling-span failure.

In (c), however, the same *p*-cycle is accessed to protect a straddling span failure. Note the significant increase in protection coverage provided by using the same investment in spare capacity as a ring, but configuring it as a *p*-cycle. In Fig. 2 the single *p*-cycle of ten hops provides protection to 18 spans and since it protects two times its own capacity on each straddling span it actually covers 26 units of working capacity, compared to the ring which protects only 10. Note also, a "signature" feature of *p*-cycles: spans that have working capacity but *no* spare capacity. Spans (3-4), (2-4), and all other straddling spans, may bear either one or two units of working capacity, but in either case require *zero* units of protection capacity. In a complete design of many *p*-cycles such spans may in practice bear protection capacity associated with other *p*-cycles overlying them, but the very possibility of spans with no protection capacity is a unique property of *p*-cycle based networks. In fact this gives one of the simplest "acid tests" to determine, regardless of other complexities, whether a given scheme is or is not equivalent to *p*-cycles. Most notably, no form of enhanced ring, cycle cover, (or generalized loop-back design - to follow) exhibit "working only" spans which bear *zero* associated protection capacity.

There are a number of other important properties of *p*-cycles. *p*-Cycles can either be cross-connect based or based on an ADM-like "capacity slice" nodal device structure [1]. In the OXC-based approach *p*-cycles are formed from individual spare wavelength channels, offering the greatest flexibility to adapt and evolve the *p*-cycle configuration. On the other hand, the ADM-like *p*-cycle nodal element offers the "pay as you grow" advantage of conventional rings. In either case, *p*-cycles avoid the structural association between the routing of working demands and the configuration of protection capacity, unlike rings. *p*-Cycles are formed only within the spare capacity layer of the network, leaving the working paths to be routed freely on shortest paths, or any other route desired. The configuration of the *p*-cycles is adapted to the working flow, not the other way around. It is important to note that when *p*-cycle network redundancy values are stated in various studies, this is already stated with respect to the most efficient shortest-path routing of working demands. In other words, relative to rings there remains a further uncharacterized capacity saving just due to replacing ring-routing with mesh-like shortest path routing. Some in the industry have attributed another 30% working capacity savings overall due to this effect alone. Note also that the average length of protection paths in a *p*-cycle is half that of the corresponding ring for straddling spans, and the same as a BLSR ring for on-cycle spans.

The literature on *p*-cycles is growing. Since its introduction in 1998 [2] the basic theory has been developed [3,4], and there have been studies on self-organization of the *p*-cycle sets [2, 5], application of *p*-cycles to the MPLS/IP layer [6], application to DWDM networking [7] and studies on joint optimization of working paths and spare capacity [8]. Notably in [7] it was found that full survivability against any span cut could be achieved with as little as 39% total network redundancy in the COST 239 European community network model. Recent work is also extending the basic concept of *p*-cycles for span protection to flow-segment protection, which is in effect the path-protection version of *p*-cycles [9].

III. Enhanced Rings

Having stressed how *p*-cycles differ from conventional rings, we now need to make the distinction between *p*-cycles and what are referred to as "enhanced rings." The basic idea of enhanced rings, in the subject US Patent by Fee [10], is for two otherwise conventional BLSR-type rings to be able to share a common protection channel on spans where the rings run in parallel; that is to say, if the rings were tiling the plane, where two tiles are edge to edge.

To appreciate the motivation for this enhancement to normal ring-based networking, consider first the view of a capacity cross-section along the right-of-way indicated in the conceptual example of Fig. 3, which imagines a North American IXC with two rings that have facing edges on the Salt Lake to L.A. facility route. Each ring individually comprises a 100% matching of working and protection fiber capacity. If, as is not unusual, the working fiber "fill" is itself say only 40% (it may tend to be even lower than this on a North-South mid-west US route due to demography), then the overall redundancy of the capacity investment on this facility route may be 400%.[2] The enhanced

[1] Even this interpretation is generous to the ring if we recognize that usually to load the working spans of a ring we have already had to deviate working paths from their true shortest routes over the network graph.

[2] This is the ratio of total unused plus protection capacity to used working capacity, i.e., here: [2(0.60) + 2] / 2(0.40) = 4.

ring idea is to recognize that in a sense the situation is "doubly redundant" because there two completely separate but side-by-side ring systems, each with their own self-contained protection. More generally this situation is called a *span overlap* in ring network design. It arises commonly in even the most efficient ring-based network designs and is widely known as one of the most inefficient aspects of ring-based networking. Our example only illustrates one of its most "painful" real-world forms: overlapped spans, thousands of km long, with low load.

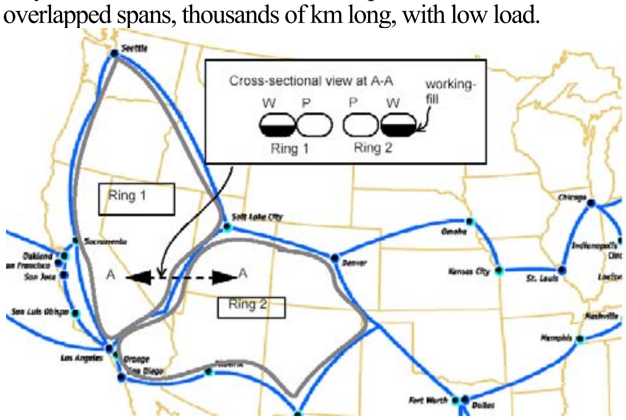

Fig. 3. - Example of two BLSR rings with a span overlap in a context aggravated by light loads and long distances.

Enhanced rings use arrangements to improve the situation by allowing the two facing rings to share a single allocation of protection bandwidth. This is done through provision of a 2:1 selector switch at each of the sites bordering the common span, and arranging an exchange of signaling between the two rings so the availability of the common protection span can be coordinated. See [10] for more details. In the example the net effect is that the overall redundancy is improved to 275% (by elimination of one complete protection channel) and the revised cross-section diagram becomes:

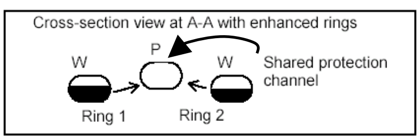

Fig. 4. - Cross-section view at A-A with enhanced rings.

Some have wondered (importantly, patent examiners) if this amounts to being a *p*-cycle. The answer is no because these remain coupled BLSR rings with an allocation of working and protection capacity on every span and still with the ability only to protect against on-cycle failures with an on-cycle loop-back. The key difference is again the aspect of straddling spans. If a true *p*-cycle was formed it would be on the outer perimeter of the two rings in Fig. 3, and the corresponding cross-sectional view in the same example would become:

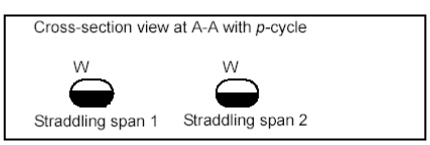

Fig. 5. - Cross-section view at A-A with *p*-cycle

In fact, the *p*-cycle approach allows an even further efficiency in this example that is categorically not available to the coupled ring scheme because each ring in the latter case must retain its complete cyclic integrity. Thus, even if all the Salt Lake - L.A. traffic was groomed over to one or the other ring only, the overall redundancy would not change because the other working fiber would still have to be there, completely intact, even if empty, to complete its ring.

In contrast, however, while the *p*-cycle can protect up to two line-rate straddling spans, there is no need to do so if demand does not require it. Thus, in the same example, the further option available to us on the Salt Lake - L.A. route when protected by a *p*-cycle is to have only one working line system, straddling the *p*-cycle between Salt Lake and L.A. So the final view of the cross-sectional capacity investment becomes:

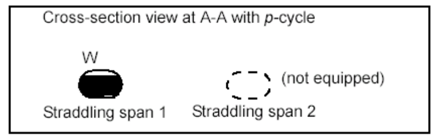

Fig. 6. - Potential cross-section view at A-A with *p*-cycle

The corresponding redundancy of this span, which was initially 400% with the overlapped rings, and reduced to 275% by enhanced rings, now reaches (1.0-0.8) / 0.8 = 25% under the *p*-cycle approach.[3]

A separate argument against *p*-cycles may be that in the example the rings should exhibit better availability because each individually has a smaller perimeter distance than the single *p*-cycle which is implied. However, with the sharing of access to a single protection span of the enhanced rings, the coupled rings in essence have the same total mileage exposure to a dual failure situation on the outer perimeter of themselves as a pair, as does the *p*-cycle. But in general, there may be circumstances where a single *p*-cycle could be made so efficient, through having a large number of straddlers, that we put separate limits in design on maximum *p*-cycle circumference or number of straddling spans, to manage the dual failure availability implications.

IV. Oriented Cycle Double Covers

Let us now explain the work of Ellinas [11] on oriented cycle double covers. Given a graph *G* a cycle cover is any decomposition of *G* into elementary cycles such that each edge is present in at least one cycle. To a large degree conventional ring network planning is an exercise in finding a suitably low cost cycle cover. The cycles have to take on a capacitated nature in the general view, and in some cases demands may be re-routed to off load an edge from the cycle-coverage requirement, a practice known as span elimination. But fundamentally every ring-based transport network comprises a form of cycle-cover.

In practice where the ring systems actually operate at the OC-n or DWDM λ-managed level, the problem is not just one of finding a set of logical cycles that cover all spans of the facilities graph. One needs also to provide for different capacity requirements on each edge, and wants to do so also generally with

[3] Note that these redundancy measures take unused working capacity into account. In other contexts one would say that the pure protection-to-working fiber-count redundancies were 100%, 50%, and zero in Fig.s 3,4,5, respectively

a minimum number of ring systems arranged in a near minimum-cost configuration.

But there is a technical context where the survivability design problem can be reduced to that of simply finding a logical cycle cover without further considerations of different specific capacity levels on each cycle. This is the paradigm of "four-fiber" networks (found in both [11,12]) where every span of the network is assumed to be comprised of exactly one bidirectional pair of working fibers and a completely matched pair of protection fibers. (DWDM allows the consideration of such "flat" fixed-fiber-count networks, at least theoretically, by assuming that any needed number of lightwave channels can be supported on each fiber pair). When one is working in this framework, it can be appreciated that the only thing that makes a cycle cover more than exactly 100 redundant (in the sense of footnote 2) is the unavoidability of span overlaps; exactly the problem illustrated in Fig. 3. Somehow, if these span overlaps could be avoided, a ring or cycle-cover network design could become more efficient. In the best case it would get to where exactly 1-for-1 matching of protection and working fiber (or channel counts) are needed.

What Ellinas et al. have shown is that if one works with individual *unidirectional* cycles (as opposed to trying to plan a cover of bidirectional pairs of cycles) you can always derive a layout of unidirectional or "oriented" cycles on a planar graph so that all span overlaps are avoided! Others have also exploited the same basic principle under the heading of "generalized loopback networks" [12]. Both approaches enable design solutions having exactly one working and one protection fibers on every span, for exactly 100% fiber-count redundancy by avoiding span overlaps through an oriented planning paradigm.

The problem which ring network planners know well, is that in general one cannot find a set of (bidirectional) covering cycles that cover every span only once. Some spans unavoidably wind up having two or more rings overlapping on them in order to support the set of cycles that cover other spans. In fact it can be fairly easily appreciated that if a graph is non-Eulerian, then more than one protection fiber will be needed on at least one span because the covering cycles are degree-2 structures and yet at least one node in a non-Eulerian network is of odd degree. Fig. 7 illustrates intuitively why odd degree nodes lead unavoidably to span overlaps in ordinary ring planning.

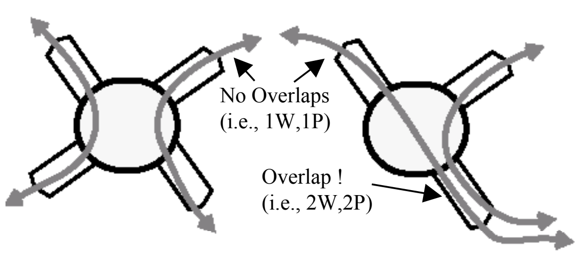

Fig. 7. - With bidirectional cycles a cover without overlaps is possible at even degree nodes but not odd degree nodes.

Ellinas et al. observes that the bi-directional nature of the actual transmission systems is almost always in practice realized by two fibers, one transmitting in each direction. Thus, at the single-fiber level one could take a unidirectional or oriented view to the formation of single-fiber cycles. They were then able to prove that as long as the graph is planar it is always possible to find a set of unidirectional cycles that covers every edge of the graph *exactly twice* (i.e., equivalent to once with a bidirectional cycle). Such a cycle set is called an *oriented cycle double cover* (O-CDC). Fig. 8 is a simple demonstration of what is more generally claimed mathematically in [11]. Working with oriented cycles, it is easy to find a set of such cycles to cover each span even at the odd degree node in Fig. 8, with exactly one working and one protection fiber per span, whereas this was not possible with bidirectional cycle covers.

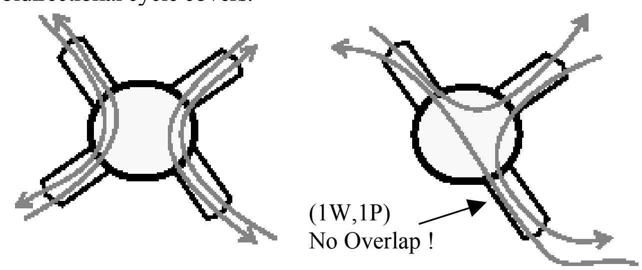

Fig. 8 - With oriented cycles the even-degree node cover is equivalent but at an odd degree node the overlap can be avoided!

Rather coincidentally this principle for achieving 100% fiber-level protection redundancy appeared in the same special issue of JSAC [11] which described *p*-cycle applications to the IP networking layer [6]. Unintentionally, both papers used the term "protection cycles"[4] in their titles, sowing the seeds for confusion. But O-CDCs are clearly not *p*-cycles. As elegant as the O-CDC approach is, it still only achieves the goal of avoiding worse than 100% overall redundancy. Protection remains logically ring-like in all cases and the result applies only to planar graphs. There are no straddling spans, which would bear zero protection capacity anywhere in an O-CDC design. In contrast *p*-cycles are far more efficient, involve the use of straddling spans, and can be efficiently planned on any form of facilities graph and are not limited to the "4-fiber network" paradigm. They can easily and fully taking differences in span fiber capacities into account.

REFERENCES

[1] W. D. Grover, D Stamatelakis, "Bridging the ring-mesh dichotomy with *p*-cycles", *Proc. DRCN 2000,* Munich, April 2000, pp. 92-104.
[2] W.D. Grover, D. Stamatelakis, "Cycle-oriented distributed preconfiguration: Ring-like speed with mesh-like capacity…," *Proc. ICC'98*, 1998, pp. 537-543.
[3] D. Stamatelakis, W.D. Grover, "Theoretical Underpinnings …," *IEEE Trans. Comm*, v.48, no.8, Aug. 2000, pp. 1262-65.
[4] D. Stamatelakis, …Preconfiguration of Spare Capacity in Mesh Restorable Networks, M.Sc. Thesis, Univ. of Alberta, 1997.
[5] D. Stamatelakis, W.D. Grover, "OPNET Simulation of Self-organizing Restorable SONET Mesh Transport Networks", *Proc. OPNETWORKS '98* (CD-ROM), Wash. D.C., April 1998, paper 04.
[6] D. Stamatelakis, W. D. Grover, "IP layer restoration … based on virtual protection cycles," *IEEE JSAC*, Oct. 2000, pp. 1938 - 1949.
[7] D. A. Schupke, C. G. Gruber, A. Autenrieth, "Optimal configuration of *p*-cycles in WDM networks," *Proc. ICC'02*, NYC, 2002.
[8] W. D. Grover, J. E. Doucette, "Advances in optical network design with *p*-cycles: Joint optimization and pre-selection …," in *Proc. IEEE-LEOS Topical Meetings*, Quebec, July 15-17, 2002.
[9] W.D. Grover, G. Shen, "Extending the *p*-Cycle Concept to Path-Protection," submitted to ICC 2003, Aug. 2002 (patent pending).
[10] J. A. Fee, *Method and System for Optical Restoration Tributary Switching in a Fiber Network*, U.S. Patent #5,884,017, Mar. 1999.
[11] G. Ellinas, A. G. Hailemariam, T.E. Stern, "Protection Cycles in Mesh WDM Networks," *IEEE JSAC,* Oct.2000, pp. 1924-37.
[12] M. Medard et.al., "Generalized loop-back recovery in optical mesh networks," *IEEE Trans. Networking,* Feb.2002, pp. 153-164.

---

[4] The "*p*" in *p*-cycles was initially chosen to suggest aspects of both protection and pre-connection concepts in general, but in early literature *p*-cycles were also called "protection cycles."