On Hamiltonian Cycles as Optimal *p*-Cycles

Dominic A. Schupke, *Member, IEEE*

***Abstract*— Using Hamiltonian *p*-cycles, it can be shown that *p*-cycle design is able to reach the logical redundancy bound of $1/(\bar{d}-1)$ where $\bar{d}$ is the average node degree. We formulate two conditions on which the design is able to reach this bound if and only if Hamiltonian *p*-cycles are used.**

I. Introduction

p-Cycles represent a both fast and capacity-efficient recovery mechanism [1]. The special case of Hamiltonian *p*-cycles has recently attracted much interest [1], [2], [3], [4], since it provides an insight into lower bounds on capacity. In this paper, we prove a basic lower bound along with finding conditions on which Hamiltonian *p*-cycles are necessary and sufficient to reach this bound. Besides enhancing prior analytical results on the topic, the results of this paper are practically useful for the design of *p*-cycle networks, since it can facilitate cycle selection, topology analysis, and capacity estimating.

We can summarize the concept of *p*-cycle protection as follows (for more detail see [1], Chap. 10). Figure 1 illustrates the protection principle of *p*-cycles for link protection. The *p*-cycle in Figure 1(a) is preconfigured as a closed connection on the cycle B-C-D-F-E-B, requiring one protection capacity unit (e.g., a WDM channel) on its links. Preconfiguration means that the configuration is done before a failure occurs.

The *p*-cycle is able to protect working capacity on its own links, called *on-cycle* links, as shown in Figure 1(b). Upon failure of on-cycle link B-C, the *p*-cycle offers protection by the route on the remainder of the cycle (C-D-F-E-B). The protectable capacity on on-cycle links is thus one capacity unit. *p*-Cycles also protect links outside the *p*-cycle: Each link which has both its end points on the *p*-cycle can also be protected. Figure 1(c) shows the protection of such a link (E-D) which is called *straddling* link. We can provide two protection routes for straddling links, in the example, routes E-B-C-D and E-F-D. In effect, we can protect two working capacity units of straddling links.

Since the nodes adjacent to the failure perform protection switching, fast recovery times are possible. Multiple *p*-cycles can be used to cover the network.

II. Non-Joint Optimization Model for *p*-Cycle Selection

We review the basic mathematical optimization model for the capacity-optimal selection of link-protecting *p*-cycles in circuit-switched networks with given working capacity [1]. Non-joint optimization refers to the methodology that protection capacity optimization (here, for *p*-cycles) is separated from working capacity optimization. A standard way for the latter is, e.g., that demands are routed on the shortest paths (with a desired link metric); we then obtain the working capacity on a link as the sum of all demands routed through it.

To assess the efficiency of a network design solution, the concept of redundancy comes in place, which measures the additional resources needed to protect the working resources. We define the *logical redundancy*, or simply *redundancy*, as the ratio of the protection capacity to the working capacity. The logical redundancy is often called "efficiency ratio." The *cost-weighted redundancy* is the ratio of the cost-weighted protection capacity to the cost-weighted working capacity, where the weights are given cost-values per link. Conclusions about bounds on general cost values are hard to draw, since we can often construct arbitrary special cases (see [1]). Therefore, we study the more general logical redundancy measure. The virtue is that this redundancy measure holds for any optimization model with arbitrary cost optimization.

The network to be protected is modeled by an undirected and two-connected graph $G=(V,E)$ without self-looping or parallel edges (links). Associated with the graph is an edge cost vector $\mathbf{z}$, a working capacity vector $\mathbf{w}$, and a protection capacity vector $\mathbf{p}$; all vectors are of size $|E|$. Based on a set of (precomputed) cycles C, the two matrices $\mathbf{\Pi}$ and $\mathbf{\Phi}$, both of size $|E|$ by $|C|$, are given. A matrix entry $\pi_{e,k} \in \{0,1\}$ of $\mathbf{\Pi}$ indicates whether edge $e \in E$ is element of cycle $k \in C$ or not. The matrix entry $\phi_{e,k} \in \{0,1,2\}$ of $\mathbf{\Phi}$ indicates how many working capacity units on edge $e \in E$ are protectable by a single *p*-cycle on cycle $k \in C$. The values of zero, one, and two represent non-protectable links, on-cycle links, and straddling links, respectively. At this point we also define the straddling link matrix $\mathbf{\Sigma} = \frac{1}{2}(\mathbf{\Phi} - \mathbf{\Pi})$. A matrix entry of $\mathbf{\Sigma}$ indicates whether edge $e \in E$ is a straddling link of cycle $k \in C$ or not. For the *p*-cycles configuration, we are interested in the number of *p*-cycles n_k for each cycle $k \in C$ that is needed. The corresponding vector of size $|C|$ is denoted by $\mathbf{n}$.

The basic problem for non-joint optimization can be formulated in vector/matrix notation [1]:

$$\min \quad \mathbf{z}^T\mathbf{p} \tag{1}$$
$$\mathbf{p} = \mathbf{\Pi}\mathbf{n} \tag{2}$$
$$\mathbf{w} \leq \mathbf{\Phi}\mathbf{n} \tag{3}$$
$$\mathbf{p} \in [0,\infty)^{|E|} \tag{4}$$
$$\mathbf{n} \in \{0,1,2,\ldots\}^{|C|} \tag{5}$$

The operation "$\leq$" applies to each entry of the vector. The Objective (1) minimizes the cost-weighted protection capacity

The author is with Siemens AG, Corporate Technology, Information and Communications, Otto-Hahn-Ring 6, D-81730 Munich, Germany (dominic.schupke@siemens.com). This work is part of the work while the author was with the Institute of Communication Networks at Technische Universität München, Munich, Germany.

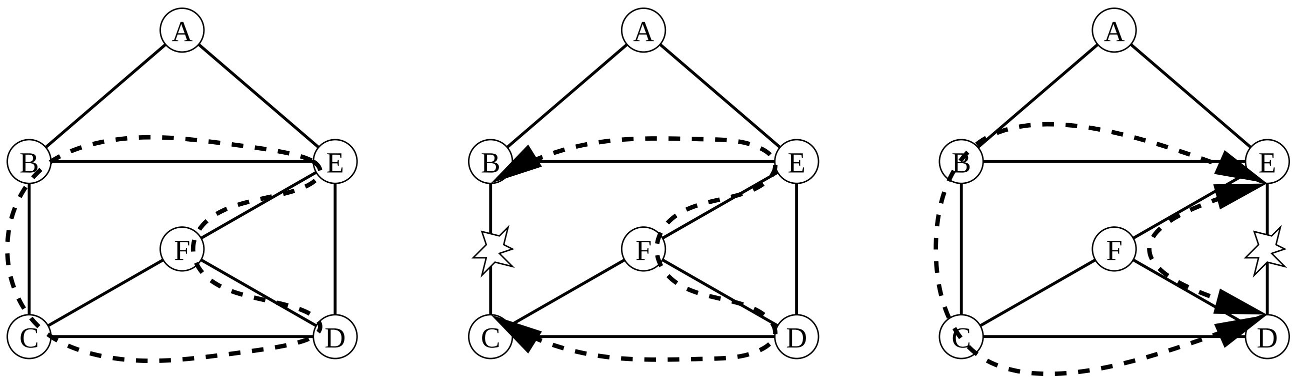

(a) A *p*-cycle (dashed line) in a network. (b) Protection of an on-cycle link. (c) Protection of a straddling link.

Fig. 1. Protection principle of *p*-cycles for link protection.

(which is equivalent to minimizing the cost-weighted redundancy), Constraint (2) determines the protection capacity allocation, and Constraint (3) ensures the working capacity to be protected. The protection capacity variables are defined in (4) and the integer number of *p*-cycles are defined in (5).

III. Relationship to Hamiltonian Cycles

For link recovery mechanisms, which include *p*-cycles as a special case, we can state the well-known lower bound $1/(\overline{d}-1)$ as bound on the logical redundancy, i.e., the protection capacity to working capacity ratio [1]. The bound is expressed as function of the average node degree $\overline{d}$.

A Hamiltonian cycle is a cycle which traverses all nodes of the graph exactly once. The network in Figure 1, e.g., has exactly two Hamiltonian cycles (A-B-C-F-D-E-A and A-B-C-D-F-E-A).

For *p*-cycles, reference [1] shows that a *p*-cycle reaches the lower bound by construction of a Hamiltonian cycle. This implies that *p*-cycles are not hindered structurally from reaching the bound (note that this is, e.g., unlike link-protection rings for which the redundancy cannot be lower than one). We can also learn from such optimal constructions how to guide the *p*-cycle design for reaching or coming close to the bound. The general ability to reach the bound is, however, no guarantee that *any* design instance can be realized with low redundancy.

Against this background, we aim to find an answer to the question: On which condition do *p*-cycles reach the $1/(\overline{d}-1)$ lower bound only by *p*-cycles which are Hamiltonian?

We partition the cycle set into Hamiltonian (C_h) and non-Hamiltonian ($C_{\overline{h}}$) and rewrite (2)-(3)

$$\mathbf{p} = \mathbf{\Pi}_h \mathbf{n}_h + \mathbf{\Pi}_{\overline{h}} \mathbf{n}_{\overline{h}} \quad (6)$$

$$\mathbf{w} + \Delta\mathbf{w} = \mathbf{\Phi}_h \mathbf{n}_h + \mathbf{\Phi}_{\overline{h}} \mathbf{n}_{\overline{h}} \quad (7)$$

where $\Delta\mathbf{w} \in [0,\infty)^{|E|}$ denotes a slack vector variable. For a feasible set of *p*-cycles, $\Delta\mathbf{w}$ is the capacity that could be protected in excess of the working capacity.

We ask when the redundancy is equal to the lower bound, i.e.,

$$r = \frac{\mathbf{1}^T \mathbf{p}}{\mathbf{1}^T \mathbf{w}} \overset{!}{=} \frac{1}{\overline{d}-1} \quad (8)$$

where $\mathbf{1}$ is a vector, of suitable size, which has only entries of one. With (6)-(7) and $\overline{d} = \frac{2|E|}{|V|}$ we obtain

$$(8) \Leftrightarrow \frac{\mathbf{1}^T \mathbf{\Pi}_h \mathbf{n}_h + \mathbf{1}^T \mathbf{\Pi}_{\overline{h}} \mathbf{n}_{\overline{h}}}{\mathbf{1}^T \mathbf{\Phi}_h \mathbf{n}_h + \mathbf{1}^T \mathbf{\Phi}_{\overline{h}} \mathbf{n}_{\overline{h}} - \mathbf{1}^T \Delta\mathbf{w}} = \frac{|V|}{2|E|-|V|}. \quad (9)$$

It is easy to see, that a Hamiltonian *p*-cycle has $|V|$ on-cycle edges while the other $|E|-|V|$ edges can become straddling edges. Thus, such a *p*-cycle is able to protect $|V| + 2(|E|-|V|) = 2|E|-|V|$ working capacity units. Therefore, we can simplify as follows

$$\mathbf{1}^T \mathbf{\Pi}_h = |V| \mathbf{1}^T \text{ and } \mathbf{1}^T \mathbf{\Phi}_h = (2|E|-|V|)\mathbf{1}^T. \quad (10)$$

This results into expression

$$(9) \Leftrightarrow \frac{\mathbf{1}^T \mathbf{n}_h + \frac{1}{|V|} \mathbf{1}^T \mathbf{\Pi}_{\overline{h}} \mathbf{n}_{\overline{h}}}{\mathbf{1}^T \mathbf{n}_h + \frac{1}{(2|E|-|V|)} \mathbf{1}^T \mathbf{\Phi}_{\overline{h}} \mathbf{n}_{\overline{h}} - \frac{1}{(2|E|-|V|)} \mathbf{1}^T \Delta\mathbf{w}} = 1. \quad (11)$$

At this point we can already conclude, that—if Hamiltonian cycles exist—just by using the Hamiltonian *p*-cycles (i.e., $\mathbf{n}_{\overline{h}} = \mathbf{0}$) we can reach the bound if there is no slack (i.e., $\Delta\mathbf{w} = \mathbf{0}$). In [1], [2] this has already been found for *p*-cycles based on a *single* Hamiltonian cycle leading to the proposed concept of "semi-homogoneous networks." As such, many network examples with Hamiltonian *p*-cycles and no slack can be constructed; e.g., start from a network with a cycle having working capacity 1 and arbitrarily add edges with working capacity 2 to non-connected node-pairs of the cycle. Also superpositions of such constructed networks belong to this group, given their number of nodes is equal.

We solve (11) for the slack variables

$$\mathbf{1}^T \Delta\mathbf{w} = \mathbf{1}^T (\mathbf{\Phi}_{\overline{h}} - (\overline{d}-1)\mathbf{\Pi}_{\overline{h}}) \mathbf{n}_{\overline{h}}$$

$$= \mathbf{1}^T (2\mathbf{\Sigma}_{\overline{h}} - \overline{d}\mathbf{\Pi}_{\overline{h}}) \mathbf{n}_{\overline{h}} = \mathbf{1}^T Q_{\overline{h}} \mathbf{n}_{\overline{h}} \quad (12)$$

where $\mathbf{\Sigma}_{\overline{h}}$ denotes the straddling-link matrix (Section II) of all non-Hamiltonian cycles. As $\triangle\mathbf{w} \geq \mathbf{0}$ and $\mathbf{n}_{\overline{h}} \geq \mathbf{0}$, we now ask: On which condition is each column-sum of $Q_{\overline{h}}$ less than zero, such that (12) can only hold for $\mathbf{n}_{\overline{h}} = \mathbf{0}$? We easily see that the sum of the entries of a column of $Q_{\overline{h}}$ can be expressed by the number of on-cycle links ν_k and the number of straddling links μ_k that the respective non-Hamiltonian cycle k has. Therefore we can write the condition

$$2\mu_k - \overline{d}\nu_k < 0,\ \forall k \in C : \nu_k < |V|. \tag{13}$$

If this condition holds, only Hamiltonian cycles can achieve the lower bound. Let d_{max} be the maximum degree in the network. Then every node of a cycle has at most $(d_{max} - 2)$ incident straddling links. If we sum the number of straddling links per node of a cycle (the number of straddling links are then counted twice), we can derive the following upper bound on a cycle's number of straddling links

$$\mu_k \leq \frac{1}{2}(d_{max} - 2)\nu_k,\ \forall k \in C. \tag{14}$$

Combining (13) and (14) leads to a degree-based condition

$$d_{max} - 2 < \overline{d}. \tag{15}$$

Hence, if the average degree is greater than the maximum degree diminished by two, only Hamiltonian *p*-cycles can achieve the lower bound. On the contrary, on condition (13), or even (15), we cannot achieve the best possible solution, if the graph is not Hamiltonian or if Hamiltonian cycles are not feasible for a network (e.g., because of length restrictions to avoid overly signal degration in optical networks).

It is interesting that condition (15) is fulfilled for many practical optical networks, since these are often (two-connected) networks with $d_{max} = 3$ or $d_{max} = 4$ (implicitly $\overline{d} > 2$). We also inspect that for an optimal combination of Hamiltonian *p*-cycles exactly covering working capacity (optimal solution $\hat{\mathbf{n}}$), homogeneous link capacity applies, since $\mathbf{w}+\mathbf{p} = \mathbf{\Pi}\hat{\mathbf{n}} + \mathbf{\Phi}\hat{\mathbf{n}} = 2\mathbf{1}\mathbf{1}^T\hat{\mathbf{n}} = a\mathbf{1}$, with some constant scalar a.

Reference [3] suggests that for homogenous working capacity networks ($\mathbf{w} = w\mathbf{1}$) *p*-cycles on a *single* Hamiltonian cycle will always attain the optimal solution. On condition (13), however, we can only conclude that *p*-cycles on one Hamiltonian cycle or on several (different) Hamiltonian cycles are optimal. A counterexample is a fully meshed network with five nodes having $w = 3$ working capacity units per link. A single Hamiltonian cycle, requiring three capacity units, has redundancy $r = \frac{3\times5}{3\times10} = \frac{1}{2}$. Less redundancy is obtained with two *p*-cycles, where one is routed link-disjoint from the other. Since the network is four-edge-connected, this routing is possible, and we reach $r = \frac{1\times5+1\times5}{3\times10} = \frac{1}{3} = \frac{1}{\overline{d}-1}$.

IV. Conclusion

A *combination* of Hamiltonian *p*-cycles is able to reach the redundancy limit given by the network-associated bound $1/(\overline{d} - 1)$. On one of two further conditions, Hamiltonian *p*-cycles are the only *p*-cycles to reach the bound. The first condition is that every non-Hamiltonian *p*-cycle has less straddling links than half of the average node degree times the number of its on-cycle links. The second (more restrictive) condition is that the average degree is greater than the maximum degree diminished by two.

Acknowledgment

The author would like to thank the reviewers for their invaluable comments.

References

[1] W. D. Grover. *Mesh-based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking*. Prentice Hall, Upper Saddle River, 2003.

[2] A. Sack and W. D. Grover. Hamiltonian p-Cycles for Fiber-Level Protection in Homogeneous and Semi-Homogeneous Optical Networks. *IEEE Network*, 18(2):49–56, March/April 2004.

[3] H. Huang and J. Copeland. Hamiltonian Cycle Protection: A Novel Approach to Mesh WDM Optical Network Protection. In *Proceedings of the IEEE Workshop on High Performance Switching and Routing (HPSR)*, (Dallas, TX, USA), May 2001.

[4] H. Huang and J. A. Copeland. A Series of Hamiltonian Cycle-Based Solutions to Provide Simple and Scalable Mesh Optical Network Resilience. *IEEE Communications Magazine*, 40(11):46–51, November 2002.