A UNIFIED FRAMEWORK FOR SHARED PROTECTION SCHEMES IN OPTICAL MESH NETWORKS

Caroline Rocha*
DIRO – Université de Montréal
Montréal, QC, H3C 3J7, Canada
thennecy@crt.umontreal.ca

Brigitte Jaumard
CIISE – Concordia University
Montréal, QC, H3G 1M8, Canada
bjaumard@ciise.concordia.ca

* *Corresponding author* / autor para quem as correspondências devem ser encaminhadas

Abstract

While the advantages of *p*-cycles and FIPP *p*-cycles are well established, there has been no systematic analysis of how much bandwidth they consume in comparison with the classical shared link and path protection schemes. It was recently observed that, even enumerating a huge number of cycles, is not necessarily a guarantee for obtaining good quality solutions with the ILP models if tools for large scale programming are not used. We propose to investigate the bandwidth protection costs of *p*-cycles and FIPP *p*-cycles in comparison with those of shared link and path protection by applying the column generation technique to solve relaxed LP models for the four protection schemes, and then solving the resulting ILP models. Provably near-optimal solutions allow us to perform accurate quantitative comparisons on real-world networks.

Keywords: optical networks; shared protection; column generation.

Resumo

Embora as vantagens dos *p*-ciclos e FIPP *p*-ciclos estejam bem estabelecidas, não existe uma análise sistemática de quanta largura de banda eles consomem em comparação aos mecanismos clássicos de proteção compartilhada por caminhos e por *links*. Foi observado recentemente que a enumeração de um enorme número de ciclos não é necessariamente uma garantia de se obter soluções de boa qualidade se ferramentas para programação de larga escala não forem usadas. Nós propomos investigar os custos de proteção de largura de banda dos *p*-ciclos e FIPP *p*-ciclos em comparação com aqueles de proteção compartilhada por *links* e por caminhos, através da aplicação da técnica de geração de colunas para resolver modelos relaxados para os quatro esquemas de proteção. Em seguida, resolvemos os modelos de programação inteira resultantes. Soluções próximas da otimalidade nos permitem desempenhar comparações quantitativas exatas em redes reais.

Palavras-chave: redes óticas; proteção compartilhada; geração de colunas.

1. Introduction

A key driver for optical networking technology has been the sustainment of the Internet growth. Researchers have contributed with many advances in optical wavelength division multiplexing (WDM) equipment and networking architectures to meet these Internet traffic needs, leading to an optical technology that currently offers immense bandwidth scalability. As it has gained notable market traction over the last several years, enterprises rely more and more on their communication services to reach their business milestones. Given the immense scale of WDM networks and how much downtime can cost a business, service survivability issues are of paramount importance. In spite of an average of 3 to 13 cuts for every 1000 miles of fiber in long haul networks (Grover, 2004), failures have a strong impact on direct revenues.

The design of a mesh WDM network usually proceeds in two steps, firstly the establishment of the *working* (or routing) *paths* with the objective of minimizing the working capacity or the equipment cost, among others. Secondly, *protection paths* are set in order to offer resilience against failures. It is well known that fiber cuts are the dominant failure pattern, and that protection against single link failure is a reasonable assumption. Although some researchers have started to investigate the joint design of working and protection paths, our focus here is on the sequential approach, in which working paths are defined a priori, usually via shortest routes.

Several survivability strategies can be found in the literature, all based on a set of features that have an impact on the network operation and/or design (Grover, 2004). Firstly, a survivable network can either use a protection or a restoration scheme. In a *protection scheme*, the redundant resources are pre-computed and reserved in advance. On the opposite, *restoration schemes* take action in real time, including resource allocation and path cross-connections, based on the failure and the state of the network at the time of failure. While restoration schemes are usually more bandwidth efficient because they do not allocate spare capacity in advance, protection schemes have faster restoration time and can always guarantee recovery from failure. Secondly, one can choose either link or path protection. *Link protection/restoration* consists in protecting each link as one entity, regardless of the connection demands that go through it, while *path protection/restoration* protects each demand individually by providing a surviving protection path between its end nodes. Although path protection (restoration) schemes lead to an efficient utilization of backup resources, they also lead to a longer failure detection and recovery than link protection (restoration). Moreover, survivability mechanisms can either use *dedicated capacity*, where spare capacity for each link or path is exclusively allocated, or *shared capacity*, where spare capacity can be shared among several protection paths under the single failure assumption.

In this work, we propose to investigate the bandwidth protection costs of the classical shared link and path protection schemes against two well-known particular cases: *p*-cycles (Grover & Stamatelakis, 1998) and failure-independent path-protecting (FIPP) *p*-cycles (Kodian & Grover, 2005). The key advantage of pre-configured protection cycles, or *p*-cycles for short, lies in their switching speed and simplicity, similar to ring networks, as the protection paths around the surviving portions of the cycle are pre-connected at the outset and the only required switching actions take place at the end nodes of the failure. In spite of many existing studies on protection/restoration schemes (Kennington *et al*., 2007; Wang *et al*., 2002; Zhang & Mukherjee, 2004) and although it is known that *p*-cycles are less capacity efficient, there has been no systematic analysis of how much bandwidth *p*-cycles schemes require in

comparison with the basic shared link and path protection. Due to the highly combinatorial nature of *p*-cycle designs, nearly all studies are based on an explicit enumeration of cycles, resulting in difficulties for assessing the quality of the solutions provided by the resulting huge ILP models. Therefore, our goal is to accomplish an accurate comparison based on results provably close to optimality by using a specialized tool for solving large scale mathematical programs: the column generation technique. A unified mathematical programming framework is presented to minimize the overall protection costs for those protection schemes.

Column generation technique (see, e.g., Chvatal (1983), Lübbecke & Desrosiers (2005), Maculan *et al.* (2003), Loiseau *et al.* (2005) for general references, or Maculan *et al.* (2002) for a reference on network design) is a powerful tool for solving linear programs with a very large number of variables where constraints can be expressed implicitly. The method relies on a decomposition of the initial linear program into a *master problem* and a *pricing problem*. The master problem corresponds to a linear program subject to a first set of explicit constraints and a second set of implicit constraints expressed throughout properties of the coefficients of the constraint matrix. The pricing problem consists in the optimization of the so-called reduced cost subject to the set of implicit constraints: it either identifies favorable columns to be added to the master problem or indicates that no such column exists.

The paper is organized as follows. The next section presents an ILP model for each protection and the definition of their respective pricing problems Experimental results are depicted in Section 4, followed by the conclusion of the study.

2. Mathematical Models for Network Protection

Consider a WDM mesh network represented by an undirected graph $G = (V, E)$ where V is the set of nodes and E is the set of links, representing bidirectional fibers, indexed by e in a non-failure state and by f in a failure state. Each link e is associated with a cost c_e. Let K be a set of connection requests, where each request k requires b_k units of working traffic (channels) between nodes o_k and d_k. We assume that the working route of each connection request k is given and denoted by WP_k. The overall working traffic on a link e is denoted by $w_e = \sum_{k \in K : e \in WP_k} b_k$. In this section, we present the mathematical models, suitable to be solved by a column generation method, for establishing protection paths with enough spare capacity to protect all working traffic while minimizing the overall protection cost. Although the working traffic is non-bifurcated, it is assumed that it can be split in integer fractional parts (integer numbers of channels) during protection. For instance, the working traffic of a given request can be restored on different protection routes, each carrying an integer fraction (i.e., a multiple of bandwidth optical units) of the failed traffic, under a path protection scheme.

2.1 Link Protection

Under a link protection scheme, the interrupted traffic is rerouted only around the failed link. Thus, the total amount of working traffic on each link is considered for protection, regardless of the connections going through it. An example of the classical shared protection scheme is illustrated in Figure 1(b).
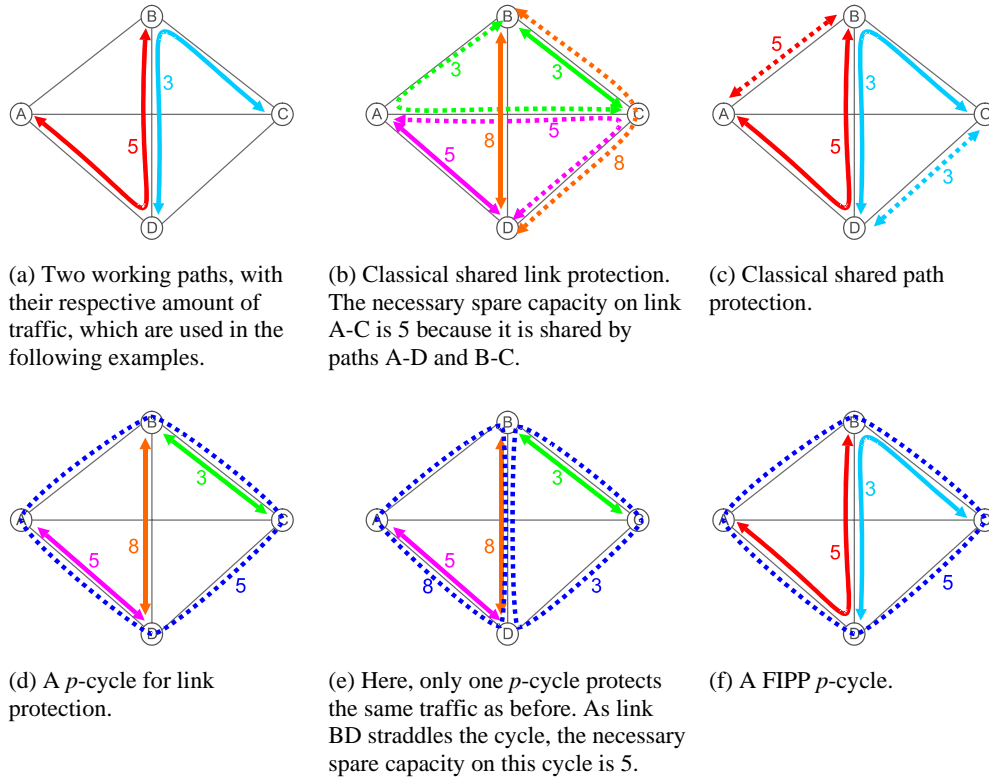
(a) Two working paths, with their respective amount of traffic, which are used in the following examples.

(b) Classical shared link protection. The necessary spare capacity on link A-C is 5 because it is shared by paths A-D and B-C.

(c) Classical shared path protection.

(d) A p-cycle for link protection.

(e) Here, only one p-cycle protects the same traffic as before. As link BD straddles the cycle, the necessary spare capacity on this cycle is 5.

(f) A FIPP p-cycle.

Figure 1 – Examples of the four protection schemes. Working and protection paths are respectively represented by solid and dashed lines, and numbers represent the amount of traffic.

2.1.1 Shared Link Protection Model

Let P_f^e be the overall set of potential protection routes for link f that use link e, and let us define set $P_f = \bigcup_{e \in E} P_f^e$, both indexed by p. The variables of model are the amount of spare capacity allocated (number of spare channels) on each link e, denoted by s_e, and the amount of traffic restored through each route p for protecting link f, denoted by n_f^p. The Shared Link Protection Model (SLP-M) is defined as follows:

$$\min \sum_{e \in E} c_e \cdot s_e \tag{1}$$

subject to:

$$\sum_{p \in P_f} n_f^p \geq w_f \qquad f \in E \tag{2}$$

$$\sum_{p \in P_f^e} n_f^p \leq s_e \qquad e \in E, f \in E \setminus \{e\} \tag{3}$$

$$n_f^p \in Z_+ \qquad\qquad f \in E, \, p \in P_f \qquad\qquad (4)$$

$$s_e \in Z_+ \qquad\qquad e \in E \qquad\qquad (5)$$

The objective function minimizes the overall protection cost. The link costs c_e can represent information such as equipment cost at end nodes, link length, etc. For instance, if $c_e = 1$, for every $e \in E$, the total spare capacity is minimized. Constraints (1) ensure that the overall working traffic on each link is protected. Constraints (2) determine the total number of spare channels that must be allocated on each link.

This model has a huge number of variables n_f^p as there is an exponential number of potential protection paths for each link. To avoid an explicit representation of all paths, a column generation approach is applied so that a solution of the LP-relaxed SLP-M model (SLP-RM) can be obtained by generating only a very small fraction of the potential paths without preventing from reaching the optimal solution. The column generation algorithm initially solves a restricted SLP-RM, which contains only a small subset of paths, and generates additional paths when needed. At each iteration, the pricing problem is set and solved for each link f: it consists in finding a path with minimum reduced cost. If a path with negative reduced cost is found, the solution of the restricted SLP-RM can be further improved and the path is added to the restricted constraint matrix. Otherwise, the optimal solution of the restricted SLP-RM is also the optimal solution of SLP-RM and, consequently, we can claim that the LP relaxation of SLP-M has been solved to optimality. For each link f, the pricing problem consists in minimizing the reduced cost of variable n_f^p, denoted as follows:

$$\bar{c}_f^p = -\mu_f - \sum_{e \in p} \pi_f^e \, ,$$

where $\mu_f \geq 0$ and $\pi_f^e \leq 0$ are the dual prices of constraints (1) and (2), respectively. For a given link f, the term $-\mu_f$ is a constant, and therefore the path with the minimum reduced cost can be found as the shortest path in the modified graph $G' = (V, E')$, where $E' = E \setminus \{ f \}$ and the link costs are defined by $c'_e = -\pi_f^e$.

2.1.2 *p*-Cycle Model

Introduced in Grover & Stamatelakis (1998), *p*-cycles are fully pre-connected cyclic protection structures with pre-planned spare capacity. When a link failure occurs, only the two end nodes of the failed link perform protection switching. Unlike rings, *p*-cycles protect against straddling link failures, enabling two protection paths, one on each half of the cycle, with only one unit of spare capacity. *p*-Cycles also provide protection against failures on links over the ring itself, as illustrated by Figures 1(d) and 1(e).

The following notation is introduced. Let *P* be the overall set of candidate *p*-cycles, indexed by *p*. The decision variables are denoted by n_p and represent the amount of traffic circulating through *p*-cycle *p*, i.e., the number of unit copies of *p*-cycle *p*. The coefficients $a_f^p \in \{0, 1, 2\}$ define the protection provided by *p*-cycle *p* for link *f*: 1 if link *f* lies on cycle *p*, 2 if link *f* straddles cycle *p* (*f* is a chord of the cycle), and 0 otherwise. The *p*-cycle model (PC-M) for link protection can be written as follows:

$$\min \sum_{p \in P} COST_p \cdot n_p \qquad\qquad (6)$$

subject to:

$$\sum_{p \in P} a_f^p \cdot n_p \geq w_f \qquad\qquad f \in E \qquad\qquad (7)$$

$$n_p \in Z_+ \qquad\qquad p \in P. \qquad\qquad (8)$$

The objective function calculates the overall protection cost, where $COST_p$ is the unit cost of p-cycle p, given by $COST_p = \sum_{e \in p} c_e$. Constraints (5) ensure that the overall working traffic on each link is protected.

Because we also aim to solve the linear relaxation of the above model with on-line generation of the cycles, by using a column generation technique, we need to define the formulation of the pricing problem. The pricing problem to be solved looks for the lowest cost cycle protecting the most profitable links. It corresponds to an NP-hard optimization problem, the Quadratic Selective Traveling Salesman Problem, described in details and solved using a branch-and-cut algorithm in Thomadsen & Stidsen (2003). Let us introduce three sets of variables. The first set contains variables x_e such that $x_e = 1$ if e defines one of the links supporting the sought cycle, and 0 otherwise. The second set is made of variables y_v such that $y_v = 1$ if v belongs to the cycle, and 0 otherwise. Last, the third set is made of variables z_e such that $z_e = 1$ if link e is protected by the cycle, and 0 otherwise. Constants μ_e are the dual prices from constraints (5) of master problem. The pricing problem can be formulated as follows:

$$\min \sum_{e \in E} (c_e + \mu_e) \cdot x_e - 2 \sum_{e \in E} \mu_e \cdot z_e \qquad\qquad (9)$$

subject to:

$$\sum_{e \in \omega(v)} x_e \leq 2 \cdot y_v \qquad\qquad v \in V \qquad\qquad (10)$$

$$z_e \leq y_v \qquad\qquad v \in V, e \in \omega(v) \qquad\qquad (11)$$

$$z_e \geq y_v + y_{v'} - 1 \qquad\qquad v, v' \in V, e = \{v, v'\} \in E \qquad\qquad (12)$$

$$\sum_{e \in \omega(V')} x_e \geq 2 \cdot (y_v + y_{v'} - 1) \qquad V' \subset V, 3 \leq |V'| \leq |V| - 3, v \in V', v' \in V \backslash V' \qquad (13)$$

$$y_v \in \{0, 1\} \qquad\qquad v \in V \qquad\qquad (14)$$

$$x_e, z_e \in \{0, 1\} \qquad\qquad e \in E. \qquad\qquad (15)$$

The objective function (9) is composed of two terms: one corresponding to the compound cost of the cycle under construction, and another one associated with the reward obtained from the protection provided to the links. Constraints (10) establish design conditions to build cycles. Constraints (11) and (12) link variables y and z. Constraints (13) are the Subtour Elimination Constraints, which ensure to build only one cycle at a time, as building more than one cycle would entail difficulty to identify the straddling links. And finally, we have binary domain constraints. Note that the domain of variables z can be relaxed to [0,1], without loss of generality, as they follow the binary status of variables y.

The PC-M model can be seen as an adaptation of that proposed in Stidsen & Thomadsen (2004). In there, the authors present a column generation model for the joint optimization of working routing and *p*-cycle protection, where two pricing problems are solved at each iteration: one for generating candidate working paths and another one for generating protecting cycles. Here, only protection is considered as we only investigate the sequential approach in this paper. In such a case, the pricing problem for generating *p*-cycles is the same as in Stidsen & Thomadsen (2004).

2.2 Path Protection

Path protection consists in protecting each demand individually by providing a protection path, diversely routed from its working path. In case of a link failure, a notification signal is sent to the end nodes of each connection traversing the failed link in order for them to switch the traffic over from the working path to the protection path. We discuss two failure-independent schemes: the Shared Backup Path Protection (SBPP) and the Failure-Independent Path-Protecting (FIPP) *p*-cycles, illustrated by Figures 1(c) and 1(f) respectively.

2.2.1 Shared Path Protection Model

Known in the literature as Shared Backup Path Protection (SBPP) (Grover, 2004) and Global Backup Path Protection (Calle *et al.*, 2004), this is a failure independent path protection scheme where the connections on the affected working paths are switched to predefined and diversely routed protection paths. Cross-connection operations to set up the protection paths are performed at the time of the failure. Unlike 1+1 protection, SBPP allows the spare capacity allocated to protection routes to be shared over failure-disjoint working paths.

Let us introduce the model for the optimal design of SBPP. Here, the set P_k^e is defined for each connection request k as the set of candidate protection routes for k using link e and the set P_k is the overall set of candidate protection routes for k, given by $P_k = \bigcup_{e \in E} P_k^e$. The decision variables are denoted by s_e and represent the number of spare channels allocated on link e, and by n_k^p, representing the amount of restored traffic through the protection route p for protecting demand k.

The optimal design of networks using SBPP as protection scheme is defined by the model SPPM below:

$$\min \sum_{e \in E} c_e \cdot s_e \tag{16}$$

subject to:

$$\sum_{p \in P_k} n_k^p \geq b_k \qquad\qquad k \in K \tag{17}$$

$$\sum_{k \in K: f \in WP_k} \sum_{p \in P_k^e} n_k^p \leq s_e \qquad e \in E,\, f \in E \backslash \{e\} \tag{18}$$

$$n_k^p \in Z_+ \qquad\qquad k \in K,\, p \in P_f \tag{19}$$

$$s_e \in Z_+ \qquad\qquad e \in E \tag{20}$$

The objective function minimizes the overall protection cost. Constraints (17) state that all working traffic must be protected. Constraints (18) determine the total number of spare channels to be allocated on each link ensuring that there will be enough spare capacity to protect non-diversely routed demands.

A column generation algorithm is also applied to the LP-relaxation of the SPP-M model (SPPRM) in order to implicitly generate protection paths. For each demand k, the pricing problem that has to be solved consists in finding a path with minimum reduced cost. The reduced cost of protection path p for demand k can be written as:

$$\overline{c}_k^p = -\mu_k - \sum_{e \in p} \sum_{f \in WP_k} \pi_f^e \ ,$$

where $\mu_k \geq 0$ and $\pi_f^e \leq 0$ are the dual prices of constraints (19) and (20), respectively. For a given demand k, the term $-\mu_k$ is a constant, and therefore the path with the lowest reduced cost for k can be defined as the shortest path in the modified graph $G' = (V, E')$, where $E' = E \setminus \{e : e \in WP_k\}$ and the link costs are defined by $c'_e = -\sum_{f \in WP_k} \pi_f^e$.

In Stidsen *et al.* (2007), the joint optimization of working routing and SBPP in directed networks with asymmetric traffic is considered. The authors present an LP model solved by column generation. The subproblem consists in computing a pair of paths (working and protection paths) with minimum total cost. Although NP-hard, the subproblem is solved with a dynamic programming based label algorithm in pseudo-polynomial time. Due to the focus of our study on a sequential optimization process, the SPP-RM model can be very easily solved by column generation since the associated pricing problem is a well-known polynomial problem instead of a NP-hard one.

2.2.2 FIPP *p*-Cycle Model

Link-protecting *p*-cycles were extended with the goal of providing end-to-end path protection, originating the Failure Independent Path-Protecting (FIPP) *p*-Cycles (Kodian & Grover, 2005). Under FIPP *p*-cycles, the cyclic protection structures can be shared by a set of working paths for protection as long as the working paths in this set are mutually disjoint or, if they are not, their protection paths are mutually disjoint.

Similarly to the model presented in Section 2.1.1, the model for the optimal placement of FIPP *p*-cycles is based on variables n_p, which represents the amount of spare capacity to be allocated on each link of FIPP *p*-cycle p. However, there is an important difference between these models: here, we let variables n_p to be possibly associated with the same topological cycle but with different coefficient vectors a^p. This means that there may be several identical cycles providing different levels of protection. For a given request k, coefficient a_k^p is equal to 2 if request k is protected for two units by cycle p, 1 if one unit of protection is provided, and 0 if no protection is provided. Note that a request can have two protected units only if it straddles cycle p. However, at most one protection unit can be provided for concurrent straddling requests at the same time.

The FIPP *p*-cycle model (FIPP-M) is formulated as follows:

$$\min \sum_{p \in P} COST_p \cdot n_p \tag{21}$$

subject to:

$$\sum_{p \in P} a_k^p \cdot n_P \geq b_k \qquad k \in K \tag{22}$$

$$n_p \in Z_+ \qquad p \in P \tag{23}$$

Constraints (22) are the demand constraints, where $a_k^p \cdot z_p$ gives the amount of traffic of demand k protected by cycle p.

The pricing problem is the minimization of the reduced cost subject to the constraints for defining a cycle and identifying the requests protected by that cycle. The model uses two sets of binary variables: the first one contains variables x_e, such that $x_e = 1$ if the cycle crosses link e, and the second one contains variables x_e^k, such that $x_e^k = 1$ if link e is used to protect request k. Furthermore, for any $V' \subset V$, let us define sets $E(V') = \{e = \{v, v'\} : v, v' \in V, e \in E\}$ and $\omega(V') = \{e = \{v, v'\} : v \in V', v' \in V \setminus V', e \in E\}$. The pricing problem is then formulated as follows:

$$\min \sum_{e \in E} c_e \cdot x_e - \sum_{k \in K} \left(\mu_k \cdot \sum_{e \in \omega(o_k)} x_e^k \right) \tag{24}$$

subject to:

$$\sum_{e \in \omega(v)} x_e \leq 2 \qquad v \in V \tag{25}$$

$$\sum_{e' \in \omega(v): e' \neq e} x_{e'} \geq x_e \qquad v \in V, e \in \omega(v) \tag{26}$$

$$\sum_{e \in \omega(V')} x_e \geq 2(x_{e'} + x_{e''} - 1) \quad V' \subset V, 3 \leq |V'| \leq |V| - 3, e' \in E(V'), e'' \in E(V \setminus V') \tag{27}$$

$$\sum_{e \in \omega(o_k)} x_e^k = \sum_{e \in \omega(d_k)} x_e^k \qquad k \in K \tag{28}$$

$$\sum_{e' \in \omega(v): e' \neq e} x_{e'}^k \geq x_e^k \qquad k \in K, v \in V \setminus \{o_k, d_k\}, e \in \omega(v) \tag{29}$$

$$\sum_{k \in K: e' \in WP_k} x_e^k \leq x_e \qquad e \in E, e' \in E \setminus \{e\} \tag{30}$$

$$x_e^k = 0 \qquad k \in K, e \in WP_k \tag{31}$$

$$x_e^k \in \{0, 1\} \qquad k \in K, e \in E \tag{32}$$

$$x_e \in \{0, 1\} \qquad e \in E \tag{33}$$

Constraints (25)-(27) address the cycle construction. Constraints (28)-(29) determine the protection routes(s) for the requests over the cycle. Next, we have constraints (30) that prevent requests from using more protection capacity than what is provided by the cycle under construction, and constraints (31) that prevent requests from using a given span in both

working and protection routes. The objective function (24) also minimizes the reduced cost and is basically composed of two terms, as for PC-M model: one corresponding to the compound costs of the links used by the cycle, and another one associated with the reward resulting from the requests chosen to be protected.

2.3 Implementation Details

In order to solve the optimization models described in Section 2, some implementation issues must be addressed. First, let us explain how the working paths have been computed. Since working and protection routes must be link disjoint, we must ensure that there exists at least one simple cycle providing one link-disjoint protection route for each working route. Therefore, for each connection request, the working route is set to the first shortest path between end nodes, computed by an algorithm for the k shortest paths problem (Epstein, 1998), for which there is a node-disjoint alternate path. The same input set of working route is provided to all models.

The LP programs of the four models were solved using version 10.1.1 of CPLEX solver. The solution approach for the pricing problems varies according to the protection scheme. The pricing problems of the SLP-RM and SPP-RM models were both solved using Dijkstra's algorithm on their respective modified networks in order to find the most profitable protection routes. The pricing MIP programs of PC-RM and FIPP-RM were solved using the CPLEX MIP solver. However, they are not solved to optimality at each iteration: solver execution is stopped as soon as a solution with negative reduced cost is obtained. Note that this does not hamper the solution optimality; instead, it often speeds up the solution process of the master problem in practice. In addition, the solution of the two latter pricing problems requires the use of costly subtour elimination constraints, which are only introduced to the set of constraints when they are violated in the incumbent solution.

Initially, all column generation algorithms start with a set of artificial (dummy) columns. An artificial column is a path (in SLP-RM and SPP-RM models) or cycle (in PC-RM and FIPP-RM models) providing some protection but which is so costly that it is very unlikely to be part of the optimal solution. The column generation algorithms obtain optimal solutions for the LP relaxation of the protection models, which are not guaranteed to be integer. In order to obtain integer solutions, the MIP models are solved using CPLEX solver with the columns introduced during the column generation process. Although, it is not certain that doing so necessarily leads to optimal integer solutions, the gap against optimality can be accurately evaluated using the optimal lower bound from the column generation algorithm.

3. Computational Results

The four column generation approaches were tested on four network instances, all described in Table 1. For each network, Table 1 shows the number of nodes, the number of links, the average node degree, the number of connection demands and the working cost. The network topologies as well as the traffic matrices for the USA and Germany instances were obtained from Hulsermann *et al.* (2004), for Pan-European COST239 and Atlanta networks are available in Batchelor *et al.* (1999) and Orlowski *et al.* (2007), respectively. The link costs in all networks are considered to be proportional to the geographical distances between cities. Each element of the traffic matrices represents the number of communication channels required between each pair of nodes and is obtained by dividing the original traffic load by 2.5 Gbit/s. Tests were performed on a AMD 64-bit machine with 16GB of RAM under Linux operating system.

The solutions obtained for all models, using the algorithms discussed in Section 2.2, are compared in Tables 2 and 3. The comparison is based on the redundancy of the resulting networks. Redundancy is a measure of architectural efficiency for survivable networks and is measured by the ratio of spare to working cost (Grover, 2004), $R = \sum_{e \in E} c_e \cdot s_e / \sum_{e \in E} c_e \cdot w_e$, where w_e and s_e are the number of working and spare channels on link e, respectively. In other words, the more redundant is the design, the more protection cost it requires. Moreover, we also show the gap between the integer and the optimal linear-relaxed solutions obtained by the column generation algorithms, the number of generated columns and the total computation time (linear relaxation time + integer solution time) for each model. In all cases, we observed very small gaps, which proves that our solutions are very close to the optimality.

As expected, path protection schemes are more cost efficient than link protection schemes. For instance, FIPP p-cycles yield between 44.63% and 107.60% extra cost while link-protecting p-cycles require between 55.27% and 113.46%. Also from Tables 2 and 3, it is clear that the denser is the network, the less extra cost is required for protection.

Table 1 – Characteristics of real-world networks.

Networks	Nodes	Links	Avg. Node Degree	Demands	Working Cost
COST239	11	26	4.7	55	137,170
USA	14	21	3.0	91	5,926,306
Atlanta	15	22	2.9	105	151,019
Germany	17	26	3.1	136	407,130

Table 2 – Results for link protection.

Networks	SLP				PC			
	R (%)	Gap (%)	Columns #	Time (s)	R (%)	Gap (%)	Columns #	Time (s)
COST239	54.25	0.26	260	0.13	55.27	0.78	24	0.67
USA	108.28	0.00	49	0.01	113.46	0.00	10	0.19
Atlanta	86.80	0.00	43	0.01	90.22	0.01	15	0.43
Germany	96.53	0.00	71	0.02	111.95	0.00	15	0.49

Table 3 – Results for path protection.

Networks	SBPP				FIPP			
	R (%)	Gap (%)	Columns #	Time (s)	R (%)	Gap (%)	Columns #	Time (s)
COST239	41.87	0.30	357	3.74	44.63	2.82	260	5376.25
USA	84.22	0.00	153	0.29	95.24	0.10	145	2038.11
Atlanta	82.67	0.00	153	0.13	90.02	0.00	257	1799.99
Germany	79.94	0.01	268	1.05	107.60	0.01	334	2025.14

According to the results, in average, basic *p*-cycles yield 8.2% more spare cost than classic link protection while the extra spare cost required by FIPP *p*-cycles in comparison with classic path protection can raise to 20.1%. Although *p*-cycles and FIPP *p*-cycles schemes require some extra cost with respect to the basic link and path protection costs respectively, the difference is rather small on dense networks, such as Cost239, which offers a greater flexibility for choosing cyclic structures. In spite of a slightly higher cost, *p*-cycle designs remain attractive and offer operational advantages: they present the properties of ring "50 ms" switching as a consequence of pre-cross-connection of paths (Grover, 2004), while restoration times are significantly higher with shared link and protection schemes (Fumagalli & Valcarenghi, 2000).

Computational times are satisfactorily small for all models, except for FIPP-M. Indeed, this model is not scalable and consumes significantly more time as the network topology and traffic matrix grow. To overcome this issue, we need to devise an efficient solution approach for the pricing problem, which is responsible for 99.99% of time consumption of the linear relaxation.

We performed additional computational tests on a 15-node network family available in Doucette (2004). Figure 2 shows the redundancy of the designs obtained by the four survivability schemes versus the average node degree for the 15 networks in the test set. Once again, we can see that denser networks take more advantage of spare capacity sharing. We can also observe that, for most cases, the survivable designs require less protection cost than the working networks: this is the great benefit of spare capacity sharing. Also, in these experiments, *p*-cycle designs are more redundant than the respective classical schemes for all networks.
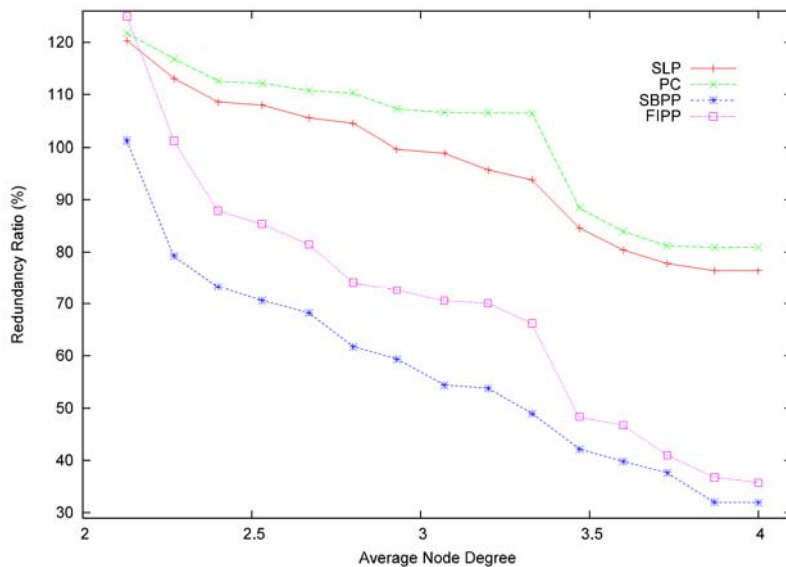

Figure 2 – Results for 15-node network family.

4. Conclusions

We proposed here to investigate the bandwidth protection costs of *p*-cycles and FIPP *p*-cycles in comparison with those of classical shared link and path protection, using column generation techniques for solving all four design models without explicit enumeration of cycles and paths. Our goal was to perform a fair and accurate comparison based on optimal or near optimal design solutions.

Results showed that classical protection schemes are more cost efficient than *p*-cycle schemes, which can be larger than the working costs for some networks. SBPP designs appear as the most cost efficient among all others. More redundant designs are less capacity efficient and, consequently, less costly. Although *p*-cycle designs are more redundant than classical protection schemes, it is an attractive choice from the operational point of view.

Acknowledgments

References

(1) Batchelor, P.; Daino, B.; Heinzmann, P.; Weinert, C.; Späth, J.*; Van Caenegem, B.; Hjelme, D.R.; Inkret, R.; Jäger, H.A.; Joindot, M.; Kuchar, A.; Le Coquil, E.; Leuthold, P.; de Marchis, G.; Matera, F.; Mikac, B.; Nolting, H.-P.; Tillerot, F. & Wauters, N. (1999). Ultra high-capacity optical transmission networks: Final report of action COST 239. Technical report, Faculty of Electrical Engineering and Computing, University of Zagreb.

(2) Calle, E.; Marzo, J. & Urra, A. (2004). Protection performance components in MPLS networks. *Computer Communications*, **27**, 1220-1228.

(3) Chvatal, V. (1983). *Linear Programming*. Freeman.

(4) Desaulniers, G.; Desrosiers, J. & Solomon, M.M. (2005). *Column Generation*. GERAD 25th Anniversary Series, Spring.

(5) Doucette, J. (2004). Advances on design and analysis of mesh-restorable networks. PhD thesis, University of Alberta, Edmonton, Canada.

(6) Epstein, D. (1998). Finding the *k* shortest paths. *SIAM J. Computing*, **28**(2), 652-673.

(7) Fumagalli, A. & Valcarenghi, L. (2000). IP restoration vs. WDM protection: is there an optimal choice? *IEEE Network*, **14**(6), 34-41.

(8) Grover, W.D. (2004). *Mesh-Based Survivable Networks*. Prentice Hall.

(9) Grover, W.D. & Stamatelakis, D. (1998). Cycle-oriented distributed preconfiguration: Ring-like speed with mesh-like capacity for self-planning network restoration. **In**: *IEEE International Conference on Communications (ICC 1998)*, 537-543.

(10) Hulsermann, R.; Bodamer, S.; Barry, M.; Betker, A.; Gauger, C.; Jager, M.; Kohn, M. & Spath, J. (2004). A set of typical transport network scenarios for network modelling. Technical report, ITG-Fachtagung Photonische Netze, Leipzig.

(11) Jaumard, B.; Rocha, C.; Baloukov, D. & Grover, W.D. (2007). A column generation approach for the design of networks using path-protecting *p*-cycles. **In**: *Proceedings of 6^{th} Workshop on Design of Reliable Communication Networks (DRCN 2007)*.

(12) Kennington, J.; Olinick, E. & Spiride, G. (2007). Basic mathematical programming models for capacity allocation in mesh-based survivable networks. *Omega*, **35**, 629-644.

(13) Kodian, A. & Grover, W. (2005). Failure-independent path-protecting *p*-cycles efficient and simple fully preconnected optical-path protection. *Journal of Lightwave Technology*, **23**(10), 3241-3259.

(14) Loiseau, I.; Ceselli, A.; Maculan, N. & Salani, M. (2005). Génération de colonnes en programmation linéaire en nombres entiers. **In**: *Optimisation combinatoire 1: concepts fondamentaux* [edited by V. Paschos], Hermès, 237-262.

(15) Lübbecke, M.E. & Desrosiers, J. (2005). Selected topics in column generation. *Operations Research*, **53**(6), 1007-2003.

(16) Maculan, N.; Passini, M.M.; Brito, J.A.M. & Lisser, A. (2002). Column generation method for network design. **In**: *Transportation and Network Analysis: Current Trends* [edited by M. Gendreau and P. Marcotte], Kluwer Academic Publishers, 167-179.

(17) Maculan, N.; Passini, M.M.; Brito, J.A.M. & Loiseau, I. (2003). Column generation in integer linear programming. *RAIRO-Operations Research*, **37**, 67-83.

(18) Orlowski, S.; Pióro, M.; Tomaszewski, A. & Wessaly, R. (2007). SNDlib 1.0 – Survivable Network Design Library. **In**: *Proceedings of the Third International Network Optimization Conference (INOC 2007)*, Spa, Belgium. <http://sndlib.zib.de>.

(19) Stidsen, T.; Petersen, B.; Rasmussen, K.; Spoorendonk, S.; Zachariasen, M.; Rambach, F. & Kiese, M. (2007). Optimal routing with single backup path protection. **In**: *Proceedings of International Network Optimization Conference (INOC 2007)*.

(20) Stidsen, T. & Thomadsen, T. (2004). Joint optimization of working and p-cycle protection capacity. Technical report, Technical University of Denmark.

(21) Thomadsen, T. & Stidsen, T. (2003). The quadratic selective travelling salesman problem. Technical report, Technical University of Denmark.

(22) Wang, J.; Sahasrabuddhe, L. & Mukherjee, B. (2002). Path vs. subpath vs. link restoration for fault management in IP-over-WDM networks: performance comparisons using GMPLS control signaling. *IEEE Communications Magazine*, **40**, 80-82.

(23) Zhang, J. & Mukherjee, B. (2004). A review of fault management in WDM mesh networks: basic concepts and research challenges. *IEEE Network*, **18**, 41-48.