Directed vs. Undirected p-Cycles and FIPP p-Cycles

Caroline Rocha* Brigitte Jaumard⋆ Pierre-Etienne Bougué⋄

*DIRO, Université de Montréal, Canada

⋆CIISE, Concordia University, Canada

⋄ENSIIE, France

Abstract

While it is acknowledged that the carried Internet traffic is very asymmetric, most work on p-cycle and FIPP p-cycle protection schemes have been conducted for symmetric traffic. The focus of this paper is to provide estimations of how much we loose in terms of usage of network resources when using symmetric links in spite of asymmetric traffic. To reach our goal, we propose a unified presentation of directed and undirected models for the p-cycle and FIPP p-cycle protection, using column generation formulations. Results show that the use of undirected models can be very cost ineffective under asymmetric traffic scenarios.

Keywords: *traffic asymmetry, directed p-cycles, undirected p-cycles*

1 Introduction

Current deployed networks, whether metropolitan or backbone, are still inherently symmetric, although the traffic that has to be carried out is asymmetric [16]. This means that large portions of network resources may be wasted in one transmission direction under highly asymmetric traffic. Although it entails more complex processes in network configuration and management in order to handle asymmetry, the resulting gain is worth the effort, as investigated in [14]. Not many studies were carried out on the influence of the asymmetric traffic on the underlying optical layer cost, and on the comparison in terms of cost/bandwidth between the use of bidirectional and unidirectional line systems [11]. While most provisioning models for operation and protection networks have been conducted for both directed and undirected models, the p-cycle protection scheme and its variants have been only studied with undirected models, except for very few studies [1]. Hence, the objective of this paper is to further investigate the concern about directed vs. undirected models for the p-cycle-based protection schemes.

p-Cycles are fully pre-connected cyclic protection structures that protect against straddling link failures, in addition to common ring-like failures [5]. What is specially interesting with p-cycles is their pre-cross connected feature. Indeed, while p-cycles can be viewed as a particular case of the shared link protection (SLP) scheme, they offer much faster recovery delays, comparable to those of ring protection. 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 [9]. Under FIPP p-cycles, the cyclic protection structures can be shared by a set of requests for protection as long as their working paths are mutually disjoint or, if they are not, their protection paths over the cycle are mutually disjoint.

While several different solution methods have been proposed for p-cycles, most models rely on a similar ILP model with the only difference that the set of possible cycles is either generated off-line, as in [5], or generated on-line, most of the time using a column generation algorithm, as in [15] and [1].

In [1], the authors propose an exact approach and a polyhedral study of the directed p-cycle placement problem. However, they assume that the spare capacity installed on the p-cycles is fractional, while in this paper we assume this quantity to be integer. In [13], the authors present a collection of mathematical models, following a uniform fashion, for comparing p-cycles and FIPP p-cycles against classical shared link and path protection. The undirected models discussed here were previously presented in [13], and an improved column generation algorithm as well as efficient algorithms for getting integer solutions for undirected FIPP p-cycles are provided in [8].

In order to investigate the behavior of directed systems under asymmetric scenarios, we describe mathematical models for the design of directed and undirected networks using both p-cycle and FIPP p-cycle protection schemes. As exact methods relying on column generation tend to reduce the scalability issue, we invest in those techniques. Please refer to [3] and [10] for a nice background on linear programming and column generation respectively. Moreover, we define two expressions for measuring traffic asymmetry, which allow us to assess the impact on the protection cost caused by the use of directed versus undirected models under different levels of asymmetry.

The paper is organized as follows. First, we present the various models, directed and undirected, for p-cycle and FIPP p-cycle protection, with column generation formulations. All models are described with capacity constraints. Thereafter, we present and discuss the experimental results obtained with those models on different network and traffic instances. Results clearly show off more attention should be given to directed models for p-cycle-based networks.

2 Notations

In order to simplify the exposition of the different formulations, we will use the same notations for directed and undirected models, except that their definitions may vary depending on whether they are used in a directed or in an undirected model. For instance, let $G = (V, L)$ be a graph associated with a WDM (wavelength-division multiplexing) network, where V is the set of nodes and L the set of links. Let ℓ refer to an undirected link in the undirected models, and to a directed link in the directed models. Note that, when G is a directed graph, each pair of connected nodes is associated with a pair of directed links, one in each direction, denoted by ℓ and $-\ell$. We denote by $\omega(v)$ the set of links incident to v, and, if G is directed, by $\omega^+(v)$ (resp. $\omega^-(v)$) the set of outgoing (incoming) links of v. Let us also define $\omega(W), W \subset V$, as the cut induced by W, i.e, the set of spans incident to a node in W and to another one in $V \backslash W$.

The traffic is defined by a set K of requests, indexed by k. Bandwidth requirement for each request k is denoted by b_k. The overall working traffic on a link ℓ is denoted by $w_\ell = \sum_{k \in K : \text{WP}_k \ni \ell} b_k$, where WP_k is the working route of request k, known beforehand. The total capacity reserved for protection on a link is subject to spare capacity constraints. Let c_ℓ denote the capacity installed on link ℓ. Then, the spare capacity on link ℓ is given by $c_\ell - w_\ell$.

Futhermore, whether we address a directed or undirected model, we define a set $\mathcal{T}$ of (FIPP) p-cycles, indexed by t, and corresponding variables n_t, which represent the amount of traffic circulating through p-cycle t, e.g., the number of unit copies of p-cycle t or the amount of pre-configured spare capacity of p-cycle t. The cost of a p-cycle t can be decomposed as $\text{COST}_t = \sum_{\ell \in t} \text{LCOST}_\ell$, where LCOST_ℓ is the cost of link ℓ (e.g., the cost of its endpoint interconnection equipment).

3 p-Cycles

The most relevant difference between directed and undirected p-cycles is that, in the first case, at most one protection unit can be provided by a cycle for each link, while with undirected p-cycles, straddling links have two units under protection.

Before setting the ILP models for p-cycles, we need to introduce some more notations, common to both directed and undirected cases. Let coefficients α_ℓ^t define the protection provided by p-cycle t for link ℓ. As for undirected p-cycles, α_ℓ^t is equal to 2 if link ℓ straddles cycle t; 1 if if link ℓ is on cycle t; and 0 otherwise. In the directed case, although the expression of the master problem is identical, the definition of coefficients α_ℓ^t is different: $\alpha_\ell^t = 1$ if link ℓ has both end nodes on cycle t but is not crossed by t; 0 otherwise. Coefficients $\beta_\ell^t \in \{0,1\}$ define the cycle topology: β_ℓ^t is equal to 1 if link ℓ lies on cycle t, and 0 otherwise, in both directed and undirected cases.

The p-cycle model in an undirected (pC_SYM) or directed (pC_ASYM) network can be written as follows:

$$\min \quad \sum_{t \in \mathcal{T}} \text{COST}_t \, n_t$$

$$\begin{aligned} \text{subject to} \quad & \sum_{t \in \mathcal{T}} \alpha_\ell^t \, n_t \geq w_\ell && \ell \in L && (u_\ell^1) \\ & \sum_{t \in \mathcal{T}} \beta_\ell^t \, n_t \leq c_\ell - w_\ell && \ell \in L && (u_\ell^2) \\ & n_t \in \mathbb{Z}_+ && t \in \mathcal{T} \end{aligned}$$

where the first set of constraints ensures that the overall working traffic is protected on each link, and the second set of constraints guarantees that the requested bandwidth for establishing working paths and protection cycles does not exceed the link transport capacity.

Because we aim at solving the model above with on-line generation of the cycles, by using a column generation technique, then we need to establish the formulation of the pricing problem, i.e., of the problem that allows the on-line generation of the most promising cycles. By promising cycle, we mean one with a negative reduced cost (in terms of the linear programming theory). Let us introduce three sets of variables. First set contains variables x_ℓ such that $x_\ell = 1$ if ℓ defines one of the links supporting the sought cycle, and 0 otherwise. Second set is made of variables y_v such that $y_v = 1$ if node v belongs to the cycle, and 0 otherwise. Last, third set is made of variables z_ℓ such that $z_\ell = 1$ if link ℓ is protected by the cycle, and 0 otherwise. Constants $u_\ell^1 \geq 0$ and $u_\ell^2 \leq 0$ are the dual prices from master problem constraints. Note that, in order to provide the pricing problem with dual prices, linear relaxation of the master problem is solved. In Section 5, we explain how integer solutions are obtained.

The pricing problem for pC_SYM model can be formulated as described below. The objective function minimizes the reduced cost, which is composed of the following terms: the compound costs of the links used by the cycle ($\sum_{\ell \in L} \text{LCOST}_\ell \, x_\ell$), the reward resulting from the requests chosen to be protected ($\sum_{\ell \in L} u_\ell^1 (2 z_\ell - x_\ell)$), and the price for using spare capacity ($\sum_{\ell \in L} u_\ell^2 x_\ell$). The first three sets of constraints establish conditions to build cycles. The fourth set of constraints ensures to build only one cycle at a time, as building more than one cycle would entail difficulty to identify the straddling links. They are the Generalized Subtour Elimination Constraints in outer form (or cut form) [4] and are taken into account in the solution process along with the so-called lazy constraints feature (cf. [7]). These constraints state that, if a node is selected in both sides of a cut, then at least two links crossing the cut must be part of the cycle. Relations between the variables of the pricing problem and the coefficients of the master problem can be defined as $\alpha_\ell = 2\, z_\ell - x_\ell$ and $\beta_\ell = x_\ell$.

As for the pricing problem for directed p-cycles, more differences appear, as it can be seen in the formulation described in Table 1. The first four sets of constraints also establish conditions to build cycles and determine which links are under protection. Note that a link ℓ crossed by a cycle cannot be protected, but the reverse link $-\ell$ can. Next, we have the subtour elimination constraints for the directed case. Relations between the variables of the pricing problem and the coefficients of the master problem can be defined as $\alpha_\ell = z_\ell$ and $\beta_\ell = x_\ell$.

4 FIPP p-Cycles

There is an important difference between p-cycle and FIPP p-cycle models concerning the variable definition in the master problem. As for FIPP p-cycles, we let variables n_t, $t \in \mathcal{T}$, to be possibly associated with the same topological cycle but with different coefficient vectors α^t. This can be justified by the

Pricing problem for undirected p-cycles	*Pricing problem for directed p-cycles*
$\min \quad \sum_{\ell \in L} (\text{LCOST}_\ell + u^1_\ell - u^2_\ell)\, x_\ell - 2 \sum_{\ell \in L} u^1_\ell z_\ell$	$\min \quad \sum_{\ell \in L} (\text{LCOST}_\ell - u^2_\ell)\, x_\ell - \sum_{\ell \in L} u^1_\ell \, z_\ell$
$\sum_{\ell \in \omega(v)} x_\ell = 2\, y_v \qquad v \in V$	$\sum_{\ell \in \omega^+(v)} x_\ell = y_v \qquad v \in V$
$z_\ell \leq y_v \qquad v \in V,\ \ell \in \omega(v)$	$\sum_{\ell \in \omega^-(v)} x_\ell = y_v \qquad v \in V$
$z_\ell \geq y_v + y_{v'} - 1 \qquad v, v' \in V,\ \ell = \{v, v'\} \in L$	$z_\ell \leq y_v - x_\ell \qquad v \in V,\ \ell \in \omega(v)$
$\sum_{\ell \in \omega(V')} x_\ell \geq 2\,(y_v + y_{v'} - 1)$	$z_\ell \geq y_v + y_{v'} - x_\ell - 1 \qquad v, v' \in V, \ell = (v, v') \in L$
$V' \subset V,\ 3 \leq \lvert V' \rvert \leq \lvert V \rvert - 3,\ v \in V',\ v' \in V \setminus V'$	$\sum_{\ell \in \omega^+(V')} x_\ell \geq y_v + y_{v'} - 1$
	$V' \subset V,\ 3 \leq \lvert V' \rvert \leq \lvert V \rvert - 3,\ v \in V',\ v' \in V \setminus V'$
$y_v \in \{0,1\} \qquad v \in V$	$y_v \in \{0,1\} \qquad v \in V$
$z_\ell, x_\ell \in \{0,1\} \qquad \ell \in L$	$z_\ell, x_\ell \in \{0,1\} \qquad \ell \in L$

Table 1: Pricing Problems for p-cycles

fact that a FIPP p-cycle may not be able to protect every request with endpoints over it, since requests may be concurrent, i.e., they may not be link disjoint. Therefore, there may be several identical cycles providing different levels of protection.

The master problems of both directed and undirected models (respectively, FIPP_ASYM and FIPP_SYM models) are again the same except for the definition of the coefficient matrix. In the undirected case, the first set of coefficients are defined as follows: $\alpha^t_k = 2$ if request k is protected for 2 units by t; 1 if request k is protected for 1 units; and 0 otherwise. Indeed, note that a request can be protected for 2 units only if it straddles cycle t. However, at most one protection unit can be provided for concurrent straddling requests at the same time, see [8] for details. As for directed FIPP p-cycles, because they can offer at most one unit of protection for a request, coefficients α^t_k are now defined as follows: $\alpha^t_k = 1$ if request k is protected by cycle t; 0 otherwise. Coefficients β^t_ℓ remain the same as for p-cycle models in both directed and undirected cases.

The master problem for directed and undirected FIPP p-cycles is then formulated as follows:

$$\begin{aligned} \min \quad & \sum_{t \in \mathcal{T}} \text{COST}_t n_t \\ \text{subject to} \quad & \sum_{t \in \mathcal{T}} \alpha^t_k \, n_t \geq b_k && k \in K && (u^1_k) \\ & \sum_{t \in \mathcal{T}} \beta^t_\ell \, n_t \leq c_\ell - w_\ell && \ell \in L && (u^2_\ell) \\ & n_t \in \mathbb{Z}_+ && t \in \mathcal{T} \end{aligned}$$

Regarding the pricing problems, there are more differences between the two models, including the advantage for the directed model to be able to avoid the exponential number of subtour elimination constraints, while overall containing about twice the number of variables. In both cases, 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. There are two sets of binary variables: the first one contains variables x_ℓ, such that $x_\ell = 1$ if the cycle crosses link ℓ, and the second one contains variables x^k_ℓ, such that $x^k_\ell = 1$ if link ℓ is used to protect request k. The pricing problem is then formulated as bellow.

The formulation of the pricing problem for FIPP_SYM model is described in Table 2. Briefly, the first three sets of constraints address the cycle construction. The following two sets of constraints determine the protection path(s) for the requests over the cycle. Next, we have constraints that prevent requests from using more protection capacity than what is provided by the cycle under construction, and constraints that prevent requests from using a given span in both working and protection paths. For more details on this formulation, we recommend to see [8]. The objective function, which also minimizes the reduced cost, is basically composed of two terms: 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. Relations between the variables of the pricing problem and the coefficients of the master problem can be expressed as $\alpha_k = \sum_{\ell \in \omega(s_k)} x_\ell^k$ and $\beta_\ell = x_\ell$.

Pricing problem for undirected FIPP p-cycles		*Pricing problem for directed FIPP p-cycles*	
$\min \sum_{\ell \in L} (\text{LCOST}_\ell - u_\ell^2)\, x_\ell - \sum_{k \in K} \left(u_k^1 \sum_{\ell \in \omega(s_k)} x_\ell^k \right)$		$\min \sum_{\ell \in L} (\text{LCOST}_\ell - u_\ell^2)\, x_\ell - \sum_{k \in K} \left(u_k^1 \sum_{\ell \in \omega^+(s_k)} x_\ell^k \right)$	
$\sum_{\ell \in \omega(v)} x_\ell \le 2$	$v \in V$	$\sum_{\ell \in \omega^+(v)} x_\ell - \sum_{\ell \in \omega^-(v)} x_\ell = 0$	$v \in V$
$\sum_{\ell' \in \omega(v): \ell' \neq \ell} x_{\ell'} \ge x_\ell$	$v \in V, \ell \in \omega(v)$	$\sum_{\ell \in \omega^+(v)} x_\ell \le 1$	$v \in V$
$\sum_{\ell \in \omega(V')} x_\ell \ge x_{\ell'} + x_{\ell''} - 1$	$V' \subseteq V, \ell' \in L(V'), \ell'' \notin L(V')$	$\sum_{\ell \in \omega^-(s_k)} x_\ell^k = 0$	$k \in K$
$\sum_{\ell \in \omega(s_k)} x_\ell^k = \sum_{\ell \in \omega(d_k)} x_\ell^k$	$k \in K$	$\sum_{\ell \in \omega^+(v)} x_\ell^k - \sum_{\ell \in \omega^-(v)} x_\ell^k = 0$	$k \in K, v \in V \setminus \{s_k, d_k\}$
$\sum_{\ell' \in \omega(v): \ell' \neq \ell} x_{\ell'}^k \ge x_\ell^k$	$k \in K, v \in V \setminus \{s_k, d_k\}, \ell \in \omega(v)$	$\sum_{k \in K: \{\ell', -\ell'\} \cap \text{WP}_k \neq \emptyset} x_\ell^k \le x_\ell$	$\ell \in L, \ell' \in L \setminus \{\ell, -\ell\}$
$\sum_{k \in K: \ell' \in \text{WP}_k} x_\ell^k \le x_\ell$	$\ell \in L, \ell' \in L \setminus \{\ell\}$	$x_\ell^k + x_{-\ell}^k = 0$	$k \in K, \ell \in \text{WP}_k$
$x_\ell^k = 0$	$k \in K, \ell \in \text{WP}_k$	$x_\ell^k \in \{0,1\}$	$k \in K, \ell \in L$
$x_\ell^k \in \{0,1\}$	$\ell \in L, k \in K$	$x_\ell \in \{0,1\}$	$\ell \in L$
$x_\ell \in \{0,1\}$	$\ell \in L$		

Table 2: Pricing Problems for FIPP p-cycles

Next, let us describe the pricing problem for FIPP_ASYM model. The first two sets of constraints address the construction of the directed cycle, followed by the constraints that prevent a protection path from terminating at the orign node of a request. Next, we have flow conservation constraints on the protection paths, and constraints preventing requests from using more protection capacity than provides the cycle under construction. Finally, we have again constraints preventing requests from using a given span in both working and protection paths. The objective funtion is expressed similarly to the undirected case. Relations between variables of the pricing problem and coefficients of the master problem can be expressed as $\alpha_k = \sum_{\ell \in \omega^+(s_k)} x_\ell^k$; and $\beta_\ell = x_\ell$.

Since the pricing problem is allowed to provide columns composed of more than one cycle, there is no subtour elimination constraint in this model. This comes from the use of double-indexed variables together with directed flows, which allow us to properly identify protection paths for straddling requests. Note that, while it is possible to get a formulation without the exponential number of subtour elimination constraints for the other models, we have favored the presented formulations for their simplicity, overcoming this last difficulty with the use of the lazy constraints when solving the pricing problem.

5 Computational Experiments

All four mathematical models (pC_SYM, pC_ASYM, FIPP_SYM and FIPP_ASYM) were implemented under the assumption of uncapacitated links, as most work in the literature, and solved as proposed using CPLEX 10.1.1 solver. Integer solutions were obtained by providing the solver with the columns generated during the column generation algorithms. Although we cannot claim that this is an exact approach, we obtained optimal or nearly optimal solutions, as shown later. We first describe the traffic and networks instances used in the experiments, as well as how we measure the traffic asymmetry. Thereafter, we present the cost discrepancies obtained for the various instances.

We consider six network instances, all available in [2], [6] and [12], and their characteristics are listed in Table 3. In particular, we provide the number of requests and the working cost for the undirected models, assuming each working path is routed via a shortest path from source to destination. Working routes are carefully assigned so that they are exactly the same in both directions as well as for both directed and undirected models.

Instance	n	m	Node Degree	$\lvert K_{\text{SYM}}\rvert$	Working Cost (SYM)	Asymmetry (%) ASYM$_{sd}$	ASYM$_\ell$
COST239 [2]	11	26	4.7	55	137,170	17	15
NSF [6]	14	21	3.0	91	5,926,306	20	7
ATLANTA [12]	15	22	2.9	105	151,019	16	13
GERMANY [6]	17	26	3.1	136	407,130	20	13
NEW-YORK [12]	16	49	6.1	240	1,437	18	12
EUROPE [6]	28	41	2.9	378	1,796,669	18	8

Table 3: Characteristics of the data sets

A measure of traffic asymmetry needs to be defined in order to assess the effect of asymmetry on the bandwidth cost. Two different ratios are then proposed. The first one attempts at measuring the traffic asymmetry, in terms of sources and destinations, weighted by the amount of bandwidth to be carried out, and is given by $\text{ASYM}_{sd} = 1 - \text{MIN}_{\mathcal{SD}}/\text{MAX}_{\mathcal{SD}}$, where $\text{MIN}_{\mathcal{SD}} = \sum_{sd \in \mathcal{SD}} \min\{b_{sd}, b_{ds}\}$ and $\text{MAX}_{\mathcal{SD}} = \sum_{sd \in \mathcal{SD}} \max\{b_{sd}, b_{ds}\}$. The second one aims at measuring the resulting traffic asymmetry in the carried bandwidth in the network and is given by $\text{ASYM}_\ell = 1 - \text{MIN}_\ell/\text{MAX}_\ell$, where $\text{MIN}_\ell = \sum_{\ell \in L} \min\{w_\ell, w_{-\ell}\}$ and $\text{MAX}_\ell = \sum_{\ell \in L} \max\{w_\ell, w_{-\ell}\}$. Once we have defined how to measure the asymmetry of a traffic instance, we developed a traffic generator which takes as input the desired sd-asymmetry percentage and a symmetric traffic instance for a given network. For each symmetric request, the traffic generator randomly calculates reverse request pairs, each one composed of a request with the symmetric demand value and the other one with a random demand value in the opposite direction. The random values are selected in such a manner that the final sum of generated values meet the desired asymmetry ratio. Asymmetric traffic instance for ATLANTA network is available in [12], all others were obtained by our generator with traffic asymmetry close to 20% in average, as shown in Table 3. It is observed that link asymmetry is often reduced with respect to traffic asymmetry. Most results have been obtained for an average of 10% link asymmetry. In [11], authors propose a different asymmetry ratio, more precisely, $\text{ASYM}_{sd} = (\text{MAX}_{\mathcal{SD}} - \text{MIN}_{\mathcal{SD}})/(\text{MAX}_{\mathcal{SD}} + \text{MIN}_{\mathcal{SD}})$.

In Table 4, we provide the results related to the number of cycles and the redundancy ratios (protection over working cost) of the (nearly) optimal protection solutions obtained for directed and undirected models with p-cycles and FIPP p-cycles. For each protection scheme we also provide, between parenthesis, the overall number of generated p-cycles. For each network instance, we mention the traffic and link asymmetry ratios, which are meaningful only for the directed models. It is shown that the redundancy ratios decrease by about 10% with directed p-cycles, meaning that the protection cost is reduced by about 10% when using a directed model instead of an undirected one. Note that, when comparing directed against undirected models, working costs for the former are assumed to be double those for undirected models. Results followed by * correspond to optimal solutions. As for the remaining results, gaps against optimality are smaller than 1%, except with undirected FIPP p-cycle result for COST239 network, whose gap increases to about 3%.

With the purpose of evaluating the reduction in protection costs while increasing traffic asymmetry, additional experiments were performed and the obtained results are illustrated in Figure 1. Five different traffic instances with fixed traffic load were generated and with asymmetry ratios ranging from 0% (pure symmetric) to 100% (pure asymmetric, i.e., unidirectional traffic between every pair of source and

Instance	pC_SYM		pC_ASYM		FIPP_SYM		FIPP_ASYM	
	# cycles	RR (%)	# cycles	RR (%)	# cycles	RR (%)	# cycles	RR (%)
COST239	9 (24)	55.3	16 (53)	51.2	11 (260)	44.6	25 (1176)	40.5
NSF	8 (10)	113.5	15 (21)	102.8⋆	63 (315)	95.2	135 (1019)	86.2
ATLANTA	8 (13)	90.2	18 (21)	84.4	82 (257)	90.0⋆	179 (557)	83.8
GERMANY	15 (18)	111.9	16 (29)	100.6⋆	84 (334)	107.6	170 (969)	96.6
NEW-YORK	18 (65)	40.0	26 (117)	37.2	-	-	-	-
EUROPE	15 (47)	109.0⋆	29 (109)	100.6⋆	-	-	-	-

Table 4: Number of (FIPP) p-cycles and Redundancy Ratios (RR)

destination) for COST239 network. In the graphic of Figure 1, plotted points correspond to the average ratio of asymmetric over symmetric results obtained using p-cycles or FIPP p-cycles for a given asymmetry percentage. For instance, with an asymmetry ratio of 10%, there is an average gain of almost 5% in the protection cost over the symmetric case when using directed p-cycles. Indeed, the use of asymmetric links can yield an average gain of nearly 45% under pure asymmetric traffic.

It can be also observed that there is an almost linear reduction in protection costs as asymmetry increases, with both p-cycles and FIPP p-cycles. This shows that the use of directed models for protecting asymmetric traffics with p-cycle-based schemes may be of great interest.

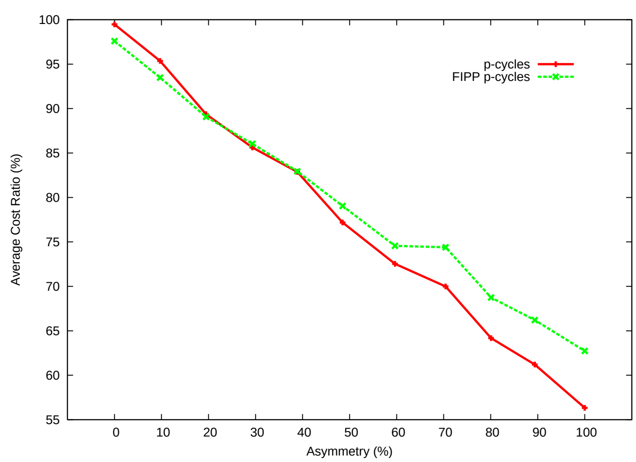

Figure 1: Accuracy of the symmetric model for COST239.

6 Conclusion

In this paper, we presented mathematical models for the design of survivable networks using directed p-cycles and FIPP p-cycles. The models were solved by means of column generation algorithms and the obtained solutions were compared against those obtained for undirected models. For that, we defined a measure of traffic asymmetry and developed a generator of asymmetric traffic instances.

Our goal was to evaluate the impact on the cost of p-cycle-based networks under asymmetric traffic scenarios. We observed a quite surprising efficiency gain when asymmetric links are used: reductions of up to 45% in the cost over undirected models. Indeed, the efficiency gain increased linearly as traffic asymmetry augments. In conclusion, the use of asymmetric links is very cost effective under asymmetric traffic scenarios and the difficulty implied by the realization of asymmetry in transport netowrks may be worthwhile.

Some possible improvements in the formulations are being investigated. For instance, the subproblem for directed p-cycles can be formulated without subtour elimination constraints, as for directed FIPP

p-cycles. But the benefits of this approach still need to be evaluated, since it implies more variables and constraints.

Acknowledgments

Financial support to the first author from the Brazilian National Council for Scientific and Technologic Development (CNPq) is gratefully acknowledged. The second author has been supported by a Concordia University Research Chair (Tier I) on the Optimization of Communication Networks and by an NSERC (Natural Sciences and Engineering Research Council of Canada) grant.

References

[1] A. Atamtürk and D. Rajan. Partition inequalities for capacitated survivable network design based on directed p-cycles. *Discrete Optimization*, 5:415–433, 2008.

[2] P. Batchelor *et al.* Ultra high-capacity optical transmission networks: Final report of action COST 239. Technical report, Faculty of Electrical Engineering and Computing, University of Zagreb, 1999.

[3] V. Chvatal. *Linear Programming.* Freeman, 1983.

[4] M. Fischetti, J.J. Salazar Gonzalez, and P. Toth. The symmetric generalized traveling salesman polytope. *Networks*, 26(2):113–123, 1995.

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

[6] R. Hülsermann. A set of typical transport network scenarios for network modelling. Technical report, Beitrge zur 5. ITG-Fachtagung Photonische Netze, Leipzig, 2004.

[7] ILOG S.A. and ILOG, Inc. *ILOG CPLEX 10.0 Documentation*, 2006.

[8] B. Jaumard and C. Rocha. Efficient computation of FIPP p-cycles. *Submitted to publication*, 2008.

[9] A. Kodian and W.D. Grover. Failure-independent path-protecting p-cycles efficient and simple fully preconnected optical-path protection. *Journal of Lightwave Technology*, 23(10):3241–3259, October 2005.

[10] M. E. Lübbecke and J. Desrosiers. Selected topics in column generation. *Operations Research*, 53(6):1007–1023, 2005.

[11] S. De Maesschalck, M. Pickavet, D. Colle, and P. Demeester. Influence of the IP traffic asymmetry on the cost of the optical network layer. *Photonic Network Communications*, 4(2):133–150, 2002.

[12] S. Orlowski, M. Pióro, A. Tomaszewski, and R. Wessäly. SNDlib 1.0–Survivable Network Design Library. In *Proceedings of the Third International Network Optimization Conference (INOC)*, Spa, Belgium, April 2007. http://sndlib.zib.de.

[13] C. Rocha and B. Jaumard. Revisiting p-cycles / FIPP p-cycles vs. shared link / path protection. In *17th International Conference on Computer Communications and Networks (ICCCN)*, St. Thomas, USA, 2008.

[14] D.A. Schupke. Link and path asymmetry issues in IP over WDM transport networks. In *Fifth Working Conference on Optical Network Design and Modelling (ONDM)*, Vienna, Austria, February 2001.

[15] T. Stidsen and T. Thomadsen. Joint optimization of working and p-cycle protection capacity. Technical report, Technical Univ. of Denmark, 2004.

[16] K. Thompson, G.J. Miller, and R. Wilder. Wide-area Internet traffic patterns and characteristics. *IEEE/ACM Transactions on Networking*, pages 10–23, November 1997.