Joint Routing and Protection Using p-cycles

Thomas Stidsen* Tommy Thomadsen†

Informatics and Mathematical Modelling
Technical University of Denmark
DK-2800 Kongens Lyngby, Denmark

May 13, 2005

Abstract

Today people rely heavily on electronic communication systems like the Internet, telephone systems, etc. Hence, it is important to ensure reliable electronic communication. The bulk of the electronic communication today is carried by circuit switched networks, thus protection against failures in these networks is paramount. Protection is possible by rerouting the electronic communication, bypassing the failed network component. In order to be able to reroute, extra capacity is, nevertheless, needed.

This article considers the recently suggested fast protection method, p-cycles. We develop a method for minimizing the capacity needed for protection using p-cycles. The routing of traffic influence the amount of extra capacity needed, thus we consider joint optimization of routing and protection.

An integer linear programming model is presented and a column generation algorithm is developed. The algorithm is faster and obtains better bounds and solutions than existing methods. The algorithm enables an experimental study of the capacity efficiency of p-cycles. The results show that p-cycles are comparable to any other protection method, with respect to the capacity usage. The results also show that substantial capacity savings can be achieved when routing and protection is performed jointly.

Based on the integer linear programming model, we discuss how protection costs can be taken into account in routing methods. We also discuss an alternative efficiency measure of the p-cycles, which takes into account the interaction with existing p-cycles.

Keywords: networks, p-cycles, routing, protection, column-generation.

*Email: tks@imm.dtu.dk
†Email: tt@imm.dtu.dk

1 Introduction

Reliable communication networks are important in society today because of the increasing dependency on electronic communication. However, most communication networks are vulnerable to equipment failures, cable cuts, electric outages, etc. Furthermore, is difficult, if possible at all, to avoid such failures. Alternatively, the traffic may be reestablished by rerouting the traffic around the failed network components.

Most of the high capacity communication networks today are circuit switched, i.e. a connection is established prior to sending actual data. In this article we consider bidirectional connections which enables two way communication via bidirectional links. In Figure 1(a), a bidirectional connection is established between node A and node D. Given a network of nodes and bidirectional links and a communication demand defined by a set of bidirectional connections, routing is the optimization task of deciding the paths which should be used by the bidirectional connections. The path of the connection in Figure 1(a) use the links AF, FE, and ED. Furthermore, the connection occupies a certain bandwidth on the links AF, FE, and ED of its path. The required capacity of a link is the sum of the bandwidths of all connections which pass the link. The *working capacity* of the network is the summed capacity of all the links.

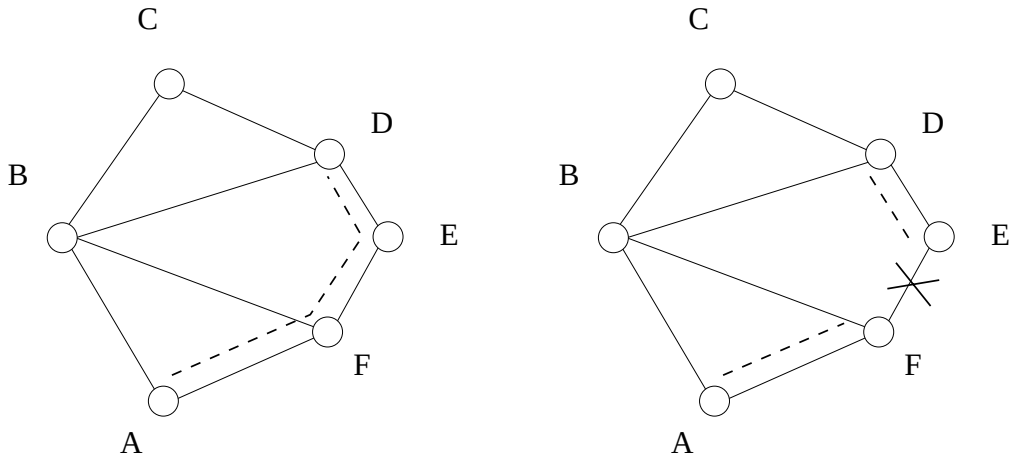

(a) A bidirectional connection

(b) A broken bidirectional connection

Figure 1: A circuit switched network

In Figure 1(b), the effect of a link failure of link FE is shown: The link between E and F fails and the connection between A and D is lost. When a link failure occurs, the connections may be reestablished by rerouting the connections passing the failed link. Rerouting is applied in protection methods to recover failures. A number of different protection methods has been suggested: Span protection, path

protection, ring protection, global rerouting [16, 19], and p-cycle protection. For a comprehensible review of the different protection techniques we refer to [6]. Generally these protection methods protect against any single link failure. We consider cycle protection [13] as an abstract model of ring protection.

In order to be able to reroute in link failure situations, capacity is needed in addition to the working capacity. Determining what additional capacity should be installed in order to be able to handle any single link failure is an optimization task. The additional capacity necessary is denoted the *protection capacity*.

Recently the so-called p-cycle protection method (Pre-configured Protection Cycle), has been suggested [8]. The authors claim that the p-cycle protection method is both capacity efficient and offer fast protection, leading to the claim that p-cycles provide "ring-like speed with mesh-like capacity".

Routing is usually done prior to protecting the network. For p-cycles, it is, however, beneficial to optimize routing and protection jointly. This article considers the joint routing and protection problem, where protection is performed using p-cycles. The main contribution is the development of a column generation algorithm which determines close to optimal solutions for the joint routing and protection problem.

The remainder of the article is organized as follows. In Section 2, the general problem of routing and p-cycle protection is described. Previous work on optimization of p-cycle protection is briefly reviewed in Section 3. In Section 4, the column generation algorithm for joint routing and p-cycle allocation is described. This algorithm is tested on six networks and in Section 5, the results are presented and discussed. Finally concluding remarks are given in Section 6.

2 The p-cycle Protection Method

The p-cycle protection method uses additional capacity allocated in cycles to protect the links. The allocated cycles are denoted p-cycles. The same amount of capacity is required on all links of the p-cycle. The capacity is pre-configured such that in case of a link failure, the only nodes that need to do rerouting are the end nodes of the failed link. Thus no signaling is required. The p-cycle protects two types of links, *on-cycle* links, see Figure 2(a) and *straddling* (chord) links, see Figure 2(b). In the figures, the thick solid lines indicate the pre-configured capacity of the p-cycle. The failed links, in Figure 2(a) link EF and in Figure 2(b) link BF, are marked with a cross.

On-cycle protection uses the fact that there is always one other way around the cycle, in case of a link failure. In case a link on the cycle fails, the connections on that link are rerouted over the remaining links of the cycle. The rerouted connections are illustrated with the dashed line from node E through $D\ C\ B\ A$ to node F. Naturally

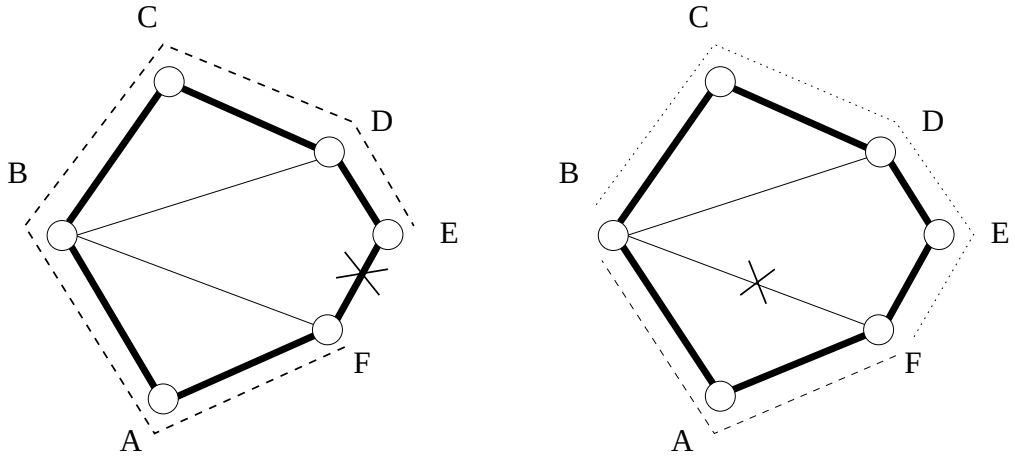

(a) On-cycle link protection

(b) Straddling link protection

Figure 2: p-cycle protection

the maximal number of connections which can be protected for the on-cycle links corresponds to the pre-configured capacity of the p-cycle.

Figure 2(b) shows how a straddling link can be protected. Because the end-nodes are on the cycle but the *link* is not on the cycle, the cycle has two routes between the end-nodes (B and F) of the failed link, illustrated with the dashed line and the dotted line in Figure 2(b). The p-cycle can hence protect twice the pre-configured capacity of the p-cycle. For a more comprehensive description of p-cycles we refer to chapter 10 in [6].

A link may be protected by several p-cycles, i.e. if a link fails, the connections using that link may be protected by rerouting them along several different p-cycles protecting that link.

This article studies the problem of jointly planning routing and p-cycle protection. Given a network and a connection demand, the sum of the working capacity and the protection capacity of the network is minimized.

3 Previous Work on p-cycle Planning

p-cycles were first suggested in [8] and it was claimed that p-cycles provide "ring-like speed with mesh-like capacity". Since then, a number of articles have been published regarding different aspects of p-cycles. In this section we briefly review those which are most relevant in connection with the joint routing and p-cycle protection problem considered here.

In [15] theoretical arguments are given for the efficiency of p-cycles. Bounding-type

arguments are given for the claim that the p-cycle method is the most capacity efficient pre-configured protection method. They do, however, base the argumentation on fully connected networks. This seems far from the rather sparsely connected telecommunication networks.

A Mixed Integer Program (MIP) model for planning p-cycle protection is presented in [8]. A prerouted network is assumed, i.e. the protection capacity requirement is minimized given the working capacity. One problem with the MIP model is that it requires enumeration of all possible p-cycles. Furthermore, it may be capacity inefficient since routing and protection are performed separately. Because the number of p-cycles grows exponentially, only networks of moderate size may be solved to optimality. By pre-selecting "promising" p-cycles, the size of the networks which can be handled may be increased albeit sacrificing the optimality guarantee. This is considered in [4, 9]. In [4] two measures for evaluating p-cycles are suggested. In Section 4.4.2 we study these measures in more detail. In [9], the effect of preselecting p-cycles of different size is investigated.

The problem of *joint* routing and p-cycle protection is studied in [7, 12]. In [7] a number of paths and p-cycles are pre-selected, making optimization of networks of medium size possible, again sacrificing the optimality guarantee. In [12], column generation is applied to implicitly represent all paths and p-cycles. The column generation subproblem is, however, not solved to optimality and thus no bounds can be derived, see Section 4.4.1. Still, it is in [12] demonstrated that low capacity requirements can be achieved using p-cycle protection. In [13], the related problem of joint routing and protection using cycles is studied.

A different approach is taken in [14]. Here a complex MIP model, which does *not* require enumeration of all possible p-cycles, is formulated. The number of binary variables of the formulation is $O(|N| \cdot |L| \cdot |C|)$, where $|N|$ is the number of nodes, $|L|$ is the number of nodes and $|C|$ is the number of different p-cycles which are actually used. While this is certainly an improvement compared to an exponential number of variables in the MIP formulation from [8], the size of the formulation still grows significantly making optimal solution methods intractable for networks of medium size. Instead, an elaborate method for stepwise optimization of gradually refined models is suggested. The approach is verified by application to full meshed networks with up to 25 nodes.

4 Solution Methodology

As discussed in Section 3, the MIP model suggested in [8] requires enumeration of all possible p-cycles to achieve an optimal solution. In this section, we describe how the use of a column generation algorithm allows us to solve a relaxation of the MIP model through *implicit* enumeration of the p-cycles. This enables solution of the LP-relaxed MIP model to optimality generating only a fraction of the possible

p-cycles.

Section 4.1 describes the MIP model for joint routing and p-cycle protection. In Section 4.2 the column generation algorithm which is needed to solve the relaxed MIP model is described. The column generation algorithm requires the solution of two sub-problems: The path generation problem described in Section 4.3, and the p-cycle generation problem described in Section 4.4. Finally, Section 4.5 describes how to use the generated paths and p-cycles to find near optimal solutions to the original MIP model.

4.1 The Joint Routing and Protection Planning Problem

Consider a network consisting of a set of nodes N and a set of links L. Furthermore, a set of connection demands D, indexed by unordered node pairs, $k, l \in V$ are defined. The constant $d_{kl} \in N_0$ is the number of connections demanded between nodes k and l. A set of paths P_{kl}, exist for each demand kl and a set of p-cycles R are given. Let c_{ij} be the capacity cost for allocating one unit of capacity on link $ij \in L$. The capacity cost of a path $p \in P_{kl}$ is $c_p^{kl} = \sum_{ij \in p} c_{ij}$. The capacity cost of a p-cycle $r \in R$ is $c_r = \sum_{ij \in r} c_{ij}$, i.e. the sum of the capacity cost of the *on-cycle* links of the p-cycle. We assume that the capacity unit of the required connections d_{kl} is equal to the capacity units of the links.

The constants $PATH_{p,ij}^{kl}$ have value 1 if path $p \in P_{kl}$ use link $ij \in L$ and 0 otherwise. The constants $PCYC_{r,ij}$ have value 1 if link ij of p-cycle r is on-cycle, 2 if link ij is straddling, and 0 otherwise. The $PCYC_{r,ij}$ constants define the protection offered by the p-cycle.

The variables $v_p^{kl} \in Z^+$ are the number of connections of demand $kl \in D$ that use path $p \in P_{kl}$ and the variables $u_r \in Z^+$, are the pre-configured capacity of p-cycle $r \in R$. Then a MIP model for the Joint routing and p-cycle protection problem, henceforth called the JP model can be formulated:

minimize:

$$\overbrace{\sum_{r \in R} c_r \cdot u_r}^{protection\ cost} + \overbrace{\sum_{kl \in D} \sum_{p \in P_{kl}} c_p^{kl} \cdot v_p^{kl}}^{routing\ cost} \tag{1}$$

subject to:

$$(\xi_{kl}) \qquad \sum_{p \in P_{kl}} v_p^{kl} = \quad d_{kl} \quad \forall\ kl \in D \quad (2)$$

$$(\pi_{ij}) \quad \sum_{r \in R} PCYC_{r,ij} \cdot u_r - \sum_{kl \in D} \sum_{p \in P_{kl}} PATH_{p,ij}^{kl} \cdot v_p^{kl} \geq \quad 0 \quad \forall\ ij \in L \quad (3)$$

$$v_p^{kl} \in Z^+ \qquad \forall\ p \in P_{kl} \quad (4)$$

$$u_r \in Z^+ \qquad \forall\ r \in R \quad (5)$$

The objective function (1) calculates the combined routing and protection cost. The constraints (2) ensure that all demands are satisfied by routing exactly the required connections along one or more of the available paths. The constraints (3) ensure that each link is protected against failure by allocation of enough protection capacity along p-cycles which offers protection to the link. Notice the difference between on-cycle link protection and straddling link protection is included in the $PCYC_{r,ij}$ constant. The dual variables of constraints (2) are ξ_{kl} and the dual variables of constraints (3) are π_{ij}.

The JP model is a generalization of the MIP model suggested in [8], which arises when P_{kl} contain exactly one path, the shortest, for each demand kl. The same model as the above is used in [7, 12]. The main problem with the JP model is that the number of paths and p-cycles grows exponentially with the number of nodes (and links) in the network. In order to ensure optimality, all paths and p-cycles must be considered explicitly or implicitly. To avoid explicit representation of paths and p-cycles, column generation is applied. The Relaxed JP model, R-JP is created by relaxing the integer domain constraints (4) and (5) of the variables v_p^{kl} and u_r, i.e. $v_p^{kl}, u_r \in R^+$. The R-JP model is an LP model, which can then be solved using column generation where the paths and p-cycles are taken into account implicitly. The column generation algorithm solves a R-JP model of reduced size where only a small subset of paths P and p-cycles R are included. We denote this the R-JP(P,R) model. Paths and p-cycles are then generated when needed.

4.2 Column Generation Algorithm

The idea of a column generation algorithm is to only generate the variables when needed, i.e. when the *reduced cost* of a variable is negative. For each iteration of the column generation algorithm the paths (one for each demand) with the minimal reduced cost is found and the p-cycle with the minimal reduced cost is found. If the reduced cost of a path or a p-cycle is negative, they are called *improving*. If no improving paths or p-cycles are found, the algorithm terminates and the R-JP model has been solved to optimality using only a subset of possible paths and p-cycles. The column generation algorithm is given in pseudo-code in Figure 3.

P = Shortest path for each demand node pair kl
R = one dummy p-cycle for each link ij
do
  Solve the R-JP(P,R) problem
  Solve routing subproblems searching for improving paths
  if improving paths found **then**
    Add improving paths to P
  Solve p-cycle subproblem searching for an improving p-cycle
  if improving p-cycle found **then**
    Add improving p-cycle to R
while improving path or improving p-cycle is found

Figure 3: *The joint routing and p-cycle protection column generation algorithm*

Initially the column generation algorithm is started with a set of shortest paths, one for each demand, and a set of *dummy* p-cycles one for each link ij. A dummy p-cycle is a (non-existent) p-cycle which has the ability of protecting just one link and which is so expensive that it will never show up in the optimal solution. Then the R-JP(P,R) model is solved based on the current set of paths P and the current set of p-cycles R. Based on the dual variables from equation (2) (ξ_{kl}) and equation (3) (π_{ij}), improving paths and p-cycles are found. This process continues until no improving paths or p-cycles are found.

4.3 Subproblem I: Path Generation

The path generation problem is the problem of generating paths with negative reduced cost. The reduced cost of a variable $\hat{c}_p^{kl}$ can be calculated based on the dual variables ξ_{kl} and π_{ij} from equation (2) and equation (3) in the R-JP(P,R) model and the link cost c_{ij} as follows.

$$\hat{c}_p^{kl} = \sum_{ij \in p} c_{ij} - \xi_{kl} + \sum_{ij \in p} \pi_{ij} \tag{6}$$

Each reduced cost contains three terms, the sum of the link costs c_{ij}, a reward term ξ_{kl} for providing an additional path to route the demand kl and a sum of the link protection costs π_{ij}. The term ξ_{kl} appear in the reduced cost for all kl-paths. Therefore the path with the lowest reduced cost for a demand kl can be found as the shortest path in a network with link costs defined as follows.

$$\overline{c}_{ij} = c_{ij} + \pi_{ij} \tag{7}$$

By duality $\pi_{ij} \geq 0$ and by assumption $c_{ij} \geq 0$, this means that $\overline{c}_{ij} \geq 0$. Thus we can apply the Floyd-Warshall algorithm [3] and the shortest paths for all demands kl can be calculated in $O(|N|^3)$. The running time may be improved using iterated Dijkstra, but since running time for generating paths is insignificant, this has not been implemented.

For all node pairs kl, if a path exists with $\hat{c}_p^{kl} < 0$ it is an improving path and it is included into the set of paths P in the R-JP(P,R) model.

Column generation is a standard approach used for the multi commodity flow problem, hence the pricing problem of paths has been extensively studied. The problem is studied in connection with network restoration in [11] and in connection with p-cycles in [12].

When the column generation algorithm terminates, the price $\overline{c}_{ij}$ is the price for using that link, including both routing costs *and* protection costs. Often joint routing and p-cycle protection is unrealistic because, as is argued in [14], introduction of new p-cycles in a network is a strategic decision, whereas routing is an operational decision. We suggest that after the strategic choice of p-cycles, based on a *forecast* of the demand, routing is performed as shortest path routing based on $\overline{c}_{ij}$ prices. This is not optimal because new p-cycles may be needed, but an estimation of the protection costs is utilized in the routing.

4.4 Subproblem II: p-cycle Generation

The second subproblem is the p-cycle generation problem, i.e. the problem of generating p-cycles with negative reduced costs. The reduced costs of the p-cycles depend only on the π_{ij} dual values. The reduced costs may be calculated as given below.

$$\widehat{c}_r = \sum_{ij \in r} c_{ij} - \sum_{ij \ straddling \ r} 2 \cdot \pi_{ij} - \sum_{ij \in r} \pi_{ij} = \sum_{ij \in r} c_{ij} - \sum_{ij \ straddling \ r \ or \ ij \in r} 2 \cdot \pi_{ij} + \sum_{ij \in r} \pi_{ij} \tag{8}$$

The last equality sign follows immediately by including both straddling and on-cycle links into the second sum and afterwards correcting by adding the third sum.

The p-cycle generation problem is an NP-hard optimization problem [12, 17] which we have previously termed the Quadratic Selective Travelling Salesman problem. This problem is described in detail in [17] and we refer to this article for an in-depth treatment. In [10] a polyhedral study of the problem is carried out.

The following MIP model of the P-Cycle Generation problem, henceforth called the PCG model, uses three types of variables: The variables $y_i \in \{0, 1\}$ represent the nodes which are part of the p-cycle, 1 for being included 0 otherwise. The

variables $x_{ij} \in \{0,1\}$ for $ij \in L$ represents the links where the protection capacity is pre-configured, 1 for being included 0 otherwise. Finally the variables $z_{ij} \in R^+$ represents all node pairs $ij \in L$ which are included in the p-cycle, 1 for the node pair being included 0 otherwise. The PCG is then expressed as follows.

minimize:

$$\sum_{ij \in L} (c_{ij} + \pi_{ij}) x_{ij} - \sum_{ij \in L} 2\pi_{ij} z_{ij} \tag{9}$$

subject to:

$$\sum_{j \in V} x_{ij} = 2y_i \qquad \forall\, i \in N \tag{10}$$

$$z_{ij} \leq y_i \qquad \forall\, ij \in L \tag{11}$$

$$z_{ij} \leq y_j \qquad \forall\, ij \in L \tag{12}$$

$$z_{ij} \geq y_i + y_j - 1 \qquad \forall\, ij \in L \tag{13}$$

$$\sum_{i \in S, j \notin S, ij \in L} x_{ij} \geq 2(y_k + y_l - 1)$$

$$\forall\, S \subset N, 3 \leq |S| \leq |N| - 3\ , k \in S, l \notin S \tag{14}$$

$$x_{ij}, y_i \in \{0,1\} \quad z_{ij} \in R^+ \tag{15}$$

The objective equation (9) calculates the reduced cost as described above in equation (8). All nodes which are in the p-cycle are required to have to two incident links, which is ensured by equation (10). For each $ij \in L$ where i and j are included in a p-cycle, i.e. $y_i = 1$ and $y_j = 1$, the variable $z_{ij} = 1$. This is ensured by equation (11), (12) and (13). Since $\pi_{ij} \geq 0$, constraints (13) are implied and are thus not necessary in the formulation. Sub-tour elimination constraints are added in order to ensure that connected cycles are constructed (14). Finally the domain constraints (15) ensure integer values for the x and y variables which in turn force integrality of the z variables.

The PCG model is solved using the branch-and-cut algorithm described in [17]. Solving the PCG model is *the* bottleneck of the column generation algorithm. This is validated by computational tests, see Table 2 in Section 5.1. However, the total computation time is acceptable, thus we have deemed improvements unnecessary. If a speed up of the column generation algorithm is needed, heuristic generation of improving p-cycles could be applied. Inspiration for this could be sought in algorithms for the TSP problem and pricing problems for cycles [12]. However, to ensure optimality of the column generation algorithm, guarantee of non-negative reduced costs are required, thus, ultimately optimal solution of the PCG model is required.

4.4.1 Reduced cost of cycles

For comparison we modify the algorithm to deal with the Joint routing and Cycle protection (JC) model. We use the same column generation algorithm as for the JP model. The only difference is the removal of straddling protection from the $PCYC_{r,ij}$ constant, i.e. $PCYC_{r,ij} = 1$ if link ij of p-cycle $r \in R$ is on-cycle and 0 otherwise. The path generation problem, see Section 4.3, remains the same, but the MIP model for the cycle generation problem is slightly different from the PCG model. The objective is to find the cycle with the most negative reduced cost, hence equation (9) is changed to equation (16) below, where the reward for the straddling links have been removed.

$$\text{minimize} \quad \sum_{ij \in L} (c_{ij} - \pi_{ij}) x_{ij} \tag{16}$$

Only the objective function is changed to find improving cycle instead of improving p-cycles. However, at the same time the z_{ij} variables and the constraints in equation (11), (12) and (13) become obsolete. This indicates that cycle generation is easier than p-cycle generation, and in fact cycles can be generated in polynomial time. Applying the Bellman-Ford algorithm [3], cycles with negative reduced costs, negative cycles, can be found in $O(|L| \cdot |N|^2)$ time.

4.4.2 p-cycle Efficiency

As mentioned in Section 3, one way to reduce the problem of the large number of possible p-cycles is to pre-select a fraction of promising p-cycles. In [4] two different measures of the p-cycles efficiency for p-cycle pre-selection is suggested: "A Priori p-cycle Efficiency" $AE(r)$, see equation (17); and "Demand-weighted p-cycle Efficiency" $EW(r)$, see equation (18).

$$AE(r) = \frac{\sum_{ij} PCYC_{r,ij}}{\sum_{ij \in r} c_{ij}} \tag{17}$$

$$EW(r) = \frac{\sum_{ij} CAP_{ij} \cdot PCYC_{r,ij}}{\sum_{ij \in r} c_{ij}} \tag{18}$$

The efficiency measure $AE(r)$ counts the number of protected links, divided by the cost of the p-cycle. In $EW(r)$ the offered protection capacity is weighted with the working capacity which needs to be protected for each link, CAP_{ij}. Hence this measure assumes that the demands are already routed.

To compare $AE(r)$ and $EW(r)$ measures with the reduced cost ($\widehat{c}_r$) from equation (8), the reduced costs of the p-cycles is divided by the cost $\sum_{ij \in r} c_{ij}$ (assuming $\sum_{ij \in r} c_{ij} \neq 0$) of the p-cycle and equation (19) is obtained.

$$\widehat{c}'_r = 1 - \frac{\sum_{ij \in r} \pi_{ij} \cdot PCYC_{r,ij}}{\sum_{ij \in r} c_{ij}} \tag{19}$$

Given a p-cycle r, the sign of $\widehat{c}'_r$ is the same as $\widehat{c}_r$, because we assume $c_{ij} \geq 0$, i.e. $\widehat{c}_r < 0 \Rightarrow \widehat{c}'_r < 0$. However, the division may have changed the order of the p-cycles with negative reduced costs, hence the best p-cycle according to equation (8) is not necessarily the best p-cycle according to equation (19). The division effectively makes the shorter p-cycles more attractive. If we ignore the constant term and change the sign of the fraction, we obtain a new measure which should be maximized.

$$\widehat{c}''_r = \frac{\sum_{ij} \pi_{ij} \cdot PCYC_{r,ij}}{\sum_{ij \in r} c_{ij}} \tag{20}$$

It is interesting to compare the optimal p-cycles according to the three different measures: $AE(r)$ (equation (17)), $EW(r)$ (equation (18)) and $\widehat{c}''_r$ (equation (20)). The optimal p-cycle according to the $AE(r)$ measure, is the p-cycle with the lowest average cost for link protection. The main problem is that it does not take into account the actual need for protection, i.e. the working capacity which needs to be protected. The $EW(r)$ measure weighs the importance the protection of the links according to the working capacity CAP_{ij} of each link. The main problem with the $EW(r)$ measure is that it does not take the interplay of the different p-cycles into account, i.e. a link may not be very interesting to protect, even though CAP_{ij} is high, because the link might already be cheaply covered by other efficient p-cycles. Given a network, a set of existing p-cycles and a demand, we conjecture that the measures defined in equation (8) and equation (20) are the best measures of future p-cycles to include into the network. Furthermore, these measures seems most appropriate when choosing p-cycles to add in response to increased demand.

4.5 Getting Integer Solutions

The column generation algorithm obtains an optimal solution to the R-JP model, but it is *not* guaranteed to return an integer solution, i.e. the optimal solution to the JP model. In this article we have chosen the simple solution of solving the JP model using a standard MIP solver with the paths P and p-cycles R collected during the column generation algorithm. Hence it is really the JP(P,R) model which is solved and it is important to acknowledge that the MIP solver only returns the optimal solution given the available paths and p-cycles and *not* the optimal integer solution to the full JP model. But because we have an optimal lower bound from the column

generation algorithm, the solution to the R-JP model, we can quantify the worst case optimality gap. As the results clearly illustrates in Section 5.3 this approach is fully sufficient to achieve close to optimal performance for all the networks tested. In order to get the real optimal solution a branch-and-price algorithm is needed [1, 18].

5 Results and Discussion

Our column generation algorithm and the integer heuristic is tested on six networks, see Table 1. The objective of the tests and discussions in this section is twofold: To examine the efficiency of the column generation algorithm and to compare the capabilities of p-cycles with cycle protection and the global rerouting lower bound [16, 19]. The global rerouting lower bound is achieved by allowing rerouting of all connections in case of any single link failure. Note that global rerouting is a lower bound for any protection method. In [19] a heuristic for global rerouting is suggested but here we report results obtained by a column generation algorithm which guarantees the lower bound [16].

	Nodes	Links	Avg. Node Degree	Working Capacity	Global Rerouting Abs.	Global Rerouting Rel.
Cost239 [2]	11	26	4.73	86	11.6	13 %
Europe	13	21	3.23	158	90.0	57 %
USA [4]	28	45	3.21	1273	641.2	50 %
Italy [5]	33	68	4.12	1718	581.4	34 %
France [4]	43	71	3.3	3473	1604.0	46 %
France 2 [4]	43	71	3.3	4043	3156.3	78 %

Table 1: The tested networks

The columns in Table 1 contain (in order): The number of nodes, the number of links, the average node degree, the working capacity i.e. capacity after shortest path routing of all demands, and the global rerouting lower bound both in absolute extra capacity and the percentage extra compared to working capacity.

For all networks, one connection is requested for all node pairs, except for the network France 2 where the same (sparse) demand pattern as in [4] was used. Otherwise, the networks France and France 2 are identical. In the tests we assume unit costs for the links, i.e. $c_{ij} = 1$, as have been done previously in [4].

Based on the test networks, Section 5.1 presents results regarding computational efficiency of the algorithm. Section 5.2 compares the protection capacity of p-cycles with the protection capacity of cycles, using shortest path routing and joint routing, and with the global rerouting lower bound. In Section 5.3, the integer solutions are compared with the bound. Finally, in Section 5.4, the importance of straddling link protection offered by the p-cycle method is studied.

5.1 Computational Efficiency

Table 2 presents data regarding the running time of the column generation algorithm. For each network, the total running time, the percentage of the time spent on solving the R-JP problem (including time spent on initialization and path generation), the percentage of the time spent on generating p-cycles, the percentage of the time spent on obtaining integer solutions, the number of p-cycles generated (Gen.), the number of p-cycles used in the integer solutions (Used), and the time spent on generating one p-cycle on average. The CPLEX 9.0 solver is used both to solve the R-JP model in the column generation algorithm, to solve the linear programs in the branch-and-cut algorithm for the PCG model, and to obtain the integer solutions of the JP(P,R) model. The MIP solver generates the integer solutions as described in Section 4.5, but if a provably optimal solution is not found after 30 seconds, the MIP solver is terminated and the best feasible solution found is returned. Preliminary tests show that 30 seconds was sufficient to obtain good heuristic solutions. The computer used was a 1200 MHz SUN Fire 3800 machine.

	Total	JP	PCG	Integer	#p-cycles		Avg. PCG
	Time (sec.)	Time (%)	Time (%)	Time (%)	Gen.	Used	Time (sec.)
Cost239	0.4	25.0 %	75.0 %	0.0 %	8	3	0.04
Europe	1.0	10.0 %	90.0 %	0.0 %	10	5	0.09
USA	44.3	2.0 %	30.2 %	67.7 %	17	11	0.79
Italy	160.6	3.1 %	78.1 %	18.7 %	44	14	2.85
France	360.1	2.1 %	96.3 %	1.6 %	43	22	8.07
France 2	239.2	0.7 %	99.1 %	0.2 %	41	19	5.78

Table 2: Computational efficiency

The column generation algorithm terminates in less than 361 seconds for all the test networks. The main part of the running time is spent on generating p-cycles, and to some extend finding an integer solution. The number of generated p-cycles is low, always less than 50, even though the number of possible p-cycles for example in the France network is at least 500000 [4]. Furthermore only about half of these are used in the integer solutions. The running time may be improved by pre-generating a number of p-cycles e.g. by using the pre-selection methods suggested in [4, 9].

5.2 Protection Capacity Efficiency

In this section, we compare the (integer) solutions for the JP model, the Shortest path routing p-cycle protection (SP) model, the JC model and the Shortest path routing Cycle protection (SC) model. The SP protection method and the SC protection method only allows the demands to be satisfied using one path: The shortest. Hence, the JP model is reduced to contain only the shortest path in the set of paths P_{kl} for each demand. As mentioned earlier in Section 3, other articles have

considered p-cycle protection assuming a prerouted demand. Thus the comparison between JP and SP is interesting. The cycle protection models are solved using the same column generation algorithm described in Section 4.2, with the modifications described in Section 4.4.1.

For each network in Table 3, the working capacity is given in the first column. The protection capacity in absolute number and relative to the working capacity is given for the integer solutions for the four different models.

		p-cycle Protection				Cycle Protection			
	Working	JP		SP		JC		SC	
	Abs.	Abs.	Rel.	Abs.	Rel.	Abs.	Rel.	Abs.	Rel.
Cost239	86	30	35%	37	43%	90	105%	94	109%
Europe	158	112	71%	147	93%	162	103%	182	115%
USA	1273	861	68%	1064	84%	1328	104%	1472	116%
Italy	1718	868	51%	1206	70%	1774	103%	1892	110%
France	3473	2255	65%	2904	84%	3732	107%	3954	114%
France 2	4043	3345	83%	3470	86%	4774	118%	4848	120%

Table 3: p-cycle and cycle protection efficiency

From Table 3 it is clear that the joint routing and p-cycle protection method is the most efficient. This method requires between 35% and 83% protection capacity to protect the network. The corresponding cycle protection method requires between 103% and 118% protection capacity. p-cycles are most capacity efficient for the networks with the highest node degree: Cost239 and Italy. The higher density of the networks enables better use of straddling links, which is shown in Section 5.4.

In [14] a number of good arguments against joint routing and protection are given. While we acknowledge these, we find it interesting that savings of $3\% - 22\%$ of the required protection capacity is possible for p-cycles. A possible explanation of the improved efficiency of joint routing and protection is offered in Section 5.4. For cycle protection, the total capacity savings are, however, only $2\% - 12\%$.

In [4], it is suggested to solve the SP model using p-cycles generated prior to optimization. Results are presented for USA and France 2. For USA, all p-cycles can be enumerated, thus the optimal solution of 1064 is obtained, which coincide with the solution we have obtained. For France 2, all p-cycles cannot be generated, but by generating 15000, a heuristic solution of 3675 is obtained. For comparison, we obtain a heuristic solution of 3470 using the SP model and a solution of 3345 if joint optimization is applied.

It could be argued that the comparison in Table 3 is not fair, since the percentages are given compared to no protection at all. In Table 4, the global rerouting lower bound [16, 19] is compared to p-cycle protection with and without joint routing. The first column contains the working capacity and the second column the addi-

tional capacity needed for global rerouting protection. Then follows the protection capacity, the extra capacity compared to the global rerouting lower bound, and the extra capacity in percent of the working capacity for JP and SP.

	Working	Global	JP			SP		
		Rerouting	Abs.	Extra	Extra %	Abs.	Extra	Extra %
Cost239	86	11.6	30	18.4	21 %	37	25.4	30%
Europe	158	90.0	112	22.0	14 %	147	57.0	36%
USA	1273	641.2	861	219.8	17 %	1064	422.8	33%
Italy	1718	581.4	868	286.6	17 %	1206	624.6	36%
France	3473	1604	2255	651	19 %	2904	1300	37%
France 2	4043	3156.3	3345	188.7	5 %	3470	313.7	7.8%

Table 4: Global rerouting vs. p-cycle protection

It is interesting to observe, that at most 21% extra capacity is needed to ensure protection using p-cycles as compared to the global rerouting lower bound. Furthermore, France 2 is only 5% from the global rerouting lower bound. This is due to the sparse demand pattern.

5.3 Integer Solution Quality

Table 5 shows the gap between the lower bound found by the column generation algorithm and the integer solution found by the MIP solver.

	p-cycle Protection		Cycle Protection	
	JP	SP	JC	SC
Cost239	2.33 %	4.65 %	5.45 %	2.50 %
Europe	0.00 %	0.59 %	0.75 %	0.99 %
USA	0.35 %	0.15 %	0.09 %	0.00 %
Italy	0.81 %	0.11 %	0.08 %	0.14 %
France	0.15 %	0.01 %	0.00 %	0.00 %
France 2	0.08 %	0.08 %	0.04 %	0.03 %

Table 5: Integer gap (%) to Column generation lower bound

As can be seen from Table 5, the solutions obtained are within 1% from optimum for all variants of the algorithms for all networks, except Cost239 where the integer solutions are up to 5.45% from the lower bound. Thus in general the algorithm produces close to optimal solutions.

The Cost239 network is small, thus all p-cycles can be generated. Furthermore, the optimal integer solution can be obtained, however the gap is still substantial.

Thus, the heuristic solution obtained is good, but the lower bound for Cost239 is significantly worse than for the other networks. This may be caused by the high density of the Cost239 network compared to the other networks.

5.4 Straddling Link Protection and Surplus Capacity

The difference between p-cycles and cycles is in essence the possibility of protecting straddling links. In this section we investigate how much of the protection capacity is straddling protection compared to on-cycle protection. The p-cycles may be able to protect more working capacity in the link than is actually present. This is denoted *surplus capacity*. The surplus capacity for a link ij corresponds to the value of the left hand side of inequality (3) and the total surplus capacity is the sum of surplus capacity for all links.

Table 6 compare the on-cycle protection capacity, the straddling protection capacity and the surplus capacity for JP and SP. All capacities are given in percent of the total protection capacity, i.e. on-cycle plus straddling protection capacity. It can

	JP			SP		
	On-cycle	Straddling	Surplus	On-cycle	Straddling	Surplus
Cost239	29 %	71 %	17 %	31 %	69 %	29 %
Europe	52 %	48 %	25 %	52 %	48 %	44 %
USA	52 %	48 %	20 %	56 %	44 %	33 %
Italy	39 %	61 %	18 %	43 %	57 %	39 %
France	50 %	50 %	17 %	58 %	42 %	30 %
France 2	56 %	44 %	30 %	55 %	45 %	36 %

Table 6: Pre-configured capacity: On-cycle, straddling and surplus

be seen in Table 6 that for the networks Cost239 and Italy, more than 50% of the protection capacity is straddling protection. For the rest of the networks 40% to 50% is straddling protection. The higher amount of straddling capacity in the networks Cost239 and Italy is due to the higher density of these networks. This also explains the higher capacity efficiency of the p-cycle protection method for these networks in Table 3.

The effect of joint routing and protection only slightly increases the amount of straddling capacity. On the other hand, joint routing and protection significantly decreases the amount of surplus capacity and this seems to be the main reason for the improved capacity efficiency of joint routing and protection.

6 Conclusion

In this article we have described an integer linear programming model for the problem of jointly routing and protecting a network using p-cycles. A column generation algorithm is implemented to obtain lower bounds. Based on the columns generated, heuristic solutions are found. The gap between the lower bound and the heuristic solution is insignificant.

An experimental study shows that the algorithm obtains better bounds and solutions faster than previously used algorithms. Lower bounds and solutions are found for networks with up to 43 nodes and 71 links in at most six minutes.

The experiments further show that straddling link protection is a valuable addition to cycle protection. Also, joint routing and protection reduce the total capacity usage compared to when routing is predetermined. The gap between the joint routing and protection for p-cycles and the global rerouting lower bound is only $5\% - 21\%$, which is quite remarkable since the p-cycle protection method is fast.

Based on the integer linear programming model, it is discussed how the protection cost can be taken into account in routing methods. Finally, a new measure of p-cycle efficiency is discussed, which takes the interplay of existing p-cycles into account.

References

[1] C. Barnhart, E.L. Johnson, G.L. Nemhauser, M.W.P. Savelsbergh and P.H. Vance. Branch-and-price: column generation for solving huge integer programs. *Operations Research*, 46(3):316–29, 1998.

[2] P. Batchelor, B. Daino, P. Heinzmann, D.E. Hjelme, P. Leuthold, R. Inkret, G.D. Marchis, H.A. Jager, F. Matera, M. Joindot, B. Mikac, A. Kuchar, H.-P. Nolting, E.L. Coquil, J. Spath, F. Tillerot, B.V. Caenegem, N. Wauters, and C. Weinert. Study on the implementation of optical transparent transport networks in the european environment-results of the Research Project COST 239. *Photonic Network Communication* 2(1):15–32, 2000.

[3] T.H. Cormen, C.E. Leiserson and R.L. Rivest. *Introduction to Algorithms*. McGraw-Hill Book Company, 2th edition, 2001.

[4] J. Doucette, D. He, W. Grover, and O. Yang. Algorithmic approaches for efficient enumeration of candidate p-cycles and capacitated p-cycle network design. *Fourth International Workshop on the Design of Reliable Communication Networks*, 212–220, 2003.

[5] W. Grover, J. Doucette, M. Clouqueur, D. Leung, and D. Stamatelakis. New options and insights for survivable transport networks. *IEEE Communications Magazine*, 40(1):34–41, 2002.

[6] W.D. Grover. *Mesh-Based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking*, Prentice Hall, 2003.

[7] W.D. Grover and J. Doucette. Advances in optical network design with p-cycles: joint optimization and pre-selection of candidate p-cycles. *IEEE/LEOS Summer Topical Meeting*, WA2–49, 2002.

[8] W.D. Grover and D. Stamatelakis. Cycle-oriented distributed preconfiguration: ring-like speed with mesh-like capacity for self-planning network restoration. *IEEE International Conference on Communications*, 1:537–543, 1998.

[9] J. Kang and M. Reed. Bandwidth protection in MPLS networks using p-cycle structure. Proceedings of the Fourth International Workshop on the Design of Reliable Communication Networks, 356–362, 2003.

[10] V. Mak and T. Thomadsen. Facets for the Cardinality Constrained Quadratic Knapsack Problem and the Quadratic Selective Travelling Salesman Problem. *IMM-Technical Report-2004-19*, 2004.

[11] K. Murakami, and H.S. Kim. Joint optimization of capacity and flow assignment for self-healing ATM networks. *IEEE International Conference on Communications*, 1:216–220, 1995.

[12] D. Rajan, and A. Atamtürk. Survivable network design: routing of flows and slacks. In Anandalingam, G and Raghavan, S., editors, *Telecommunications Network Design and Management*, Kluwer Academic Publishers, 65–81, 2003.

[13] D. Rajan, and A. Atamtürk. A directed cycle-based column-and-cut generation method for capacitated survivable network design. *Networks*, 43(4):201–211, 2004.

[14] D.A. Schupke. An ILP for optimal p-cycle selection without cycle enumeration. *Eighth Working Conference on Optical Network Design and Modelling (ONDM)*, 2004.

[15] D. Stamatelakis and W.D. Grover. Theoretical underpinnings for the efficiency of restorable networks using preconfigured cycles (p-cycles). *Communications, IEEE Transactions on* 48(8):1262–1265, 2000.

[16] T. Stidsen. P. Kjærulff. Protection lower bounding through global rerouting. *Work In Progress.*

[17] T. Thomadsen and T. Stidsen. (2003) A Branch-and-Cut Algorithm for the Quadratic Selective Travelling Salesman Problem. *Submitted for Telecommunication Systems.*

[18] F. Vanderbeck and L.A. Wolsey. An exact algorithm for IP column generation. *Operations Research Letters* 19(4):151–159, 1996.

[19] J. Yamada. A spare capacity design method for restorable networks. *IEEE Global Telecommunications Conference* 2:931–935, 1995.