Optimal Configuration of *p*-Cycles in WDM Networks

D.A. Schupke, C.G. Gruber, A. Autenrieth
Munich University of Technology
Institute of Communication Networks
Munich, Germany

***Abstract*—We investigate the deployment of *p*-cycles (preconfigured protection cycles) in WDM mesh networks with and without wavelength conversion. We develop optimization models for the configuration of the cycles and apply these on a case study for a pan-European network. The results in particular for wavelength converting networks show that *p*-cycles achieve high efficiency.**

I. Introduction

THE *p*-cycle concept was introduced by Grover and Stamatelakis in [1-3]. *p*-Cycles can be characterized as preconfigured protection cycles in a mesh network. With the hybrid cycle/mesh approach, the *p*-cycle concept is able to benefit from the advantages of both worlds.

Ring protection mechanisms offer very fast recovery times (about 50-60 ms), but the required spare to working resources ratio is at least 100%, in real networks sometimes more than 200% [2].

For mesh restoration (or protection) mechanisms, however, the required spare to working resources ratio can typically be in the range of only 50-70% for well-connected physical network graphs [2].

Since restoration involves a complex distributed signaling, mesh restoration is generally slower than ring recovery. Moreover, network operators rather use simple recovery mechanisms with a predictable behavior than a fully distributed restoration.

The concept of *p*-cycles utilizes the benefits of both alternatives: the efficiency of mesh restoration and the recovery speed of ring networks [1]. The main goal of this paper is to investigate the efficiency of *p*-cycles applied to WDM networks.

The next section summarizes the general concept of *p*-cycles. Section III discusses the deployment of *p*-cycles in WDM networks with and without wavelength conversion. In Section IV mathematical formulations for the optimal combination of *p*-cycles in WDM networks are presented. In Section V a case study for a pan-European network is performed, and the results are discussed in Section VI. Section VII concludes this paper.

II. The *p*-Cycle Concept

In this section we summarize the general concept of *p*-cycles applied to circuit-switched networks [2]. Two basic types of *p*-cycles exist. Link *p*-cycles protect the individual channels within a link. Node-encircling *p*-cycles are routed through all adjacent neighbor nodes of a node to be protected, but exclude the protected node itself, thus protecting all the connections traversing the node. In this paper we put an emphasis on link *p*-cycles, assuming WDM nodes to be very reliable, e.g., using internal redundancy.

D.A. Schupke is supported by the German Federal Ministry of Education and Research, project TransiNet. Email: Schupke@ei.tum.de.

Fig. 1 (a) depicts a network with one link *p*-cycle. The *p*-cycle is able to protect on-cycle links as shown in Fig. 1 (b). Furthermore, a *p*-cycle is also able to protect straddling links. A straddling link is an off-cycle link having *p*-cycle nodes as endpoints. In the case of a straddling link failure, each *p*-cycle can actually protect two working paths on the link by providing the two alternative paths around the *p*-cycle as shown in Fig. 1 (c)-(d).

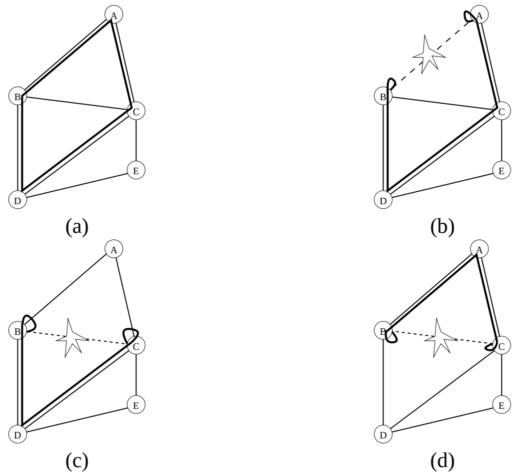

Fig. 1. A network with one link *p*-cycle (a) which can protect on-cycle links (b) and straddling links (c)-(d).

In our consideration we assume the *p*-cycles are configured by a network management system. Another approach is to find the *p*-cycles by distributed self-organization [1]. Although the autonomy of such a scheme can simplify network operations and reduce the dependency on an accurate network database, it can yield sub-optimal capacity utilization and needs higher signaling complexity.

The network configuration process of *p*-cycles in an existing (two-connected and capacitated) network can be sketched as follows. First, a given demand for connections is routed through the network, so that the links reserve (working) capacity for the demands. The spare capacity of the links is the remaining available capacity (which can be zero for individual links).

The *p*-cycles are formed in the spare capacity of the network. The set of link *p*-cycles is chosen such that for every link the working connections are protected by *p*-cycles of corresponding capacity. The routing of the demands has to be adapted, if a protecting set of *p*-cycles cannot be found. Like SDH/Sonet line-switched rings the *p*-cycles protection capacity can also be used for low-priority pre-emptible traffic.

p-Cycles have the outstanding property that the protection switching can be made very fast, since only the nodes neighboring the failure need to perform any real-time actions. Switching times in the order of some 10 ms comparable to SDH/Sonet line-switched rings can be achieved.

Therefore *p*-cycles can be configured efficiently by a (centralized, slow) network management system and the (distributed, node-internally processed) reaction to failures is very quick.

Even subsequent failures following a single failure can be survived by finding *p*-cycles in the remaining topology (without the failed element). For this the network configuration process has to be restarted based on the new topology and the new paths (which have been restored). The process can be started automatically and has to be finished before the next failure occurs. The active connections are not affected during the reconfiguration process, since the switching is done for the spare resources only. The same procedure can be used when new connection requests or releases occur.

So far, the applicability and efficiency of the *p*-cycle concept has been published with a focus on SDH/Sonet networks [1,2]. In the following we show how *p*-cycles can be incorporated in WDM networks with and without wavelength conversion and investigate their efficiency in these networks.

III. Deployment of *p*-Cycles in WDM Networks

We consider virtual wavelength path (VWP) and wavelength path (WP) WDM networks. A (unidirectional) link interconnecting the WDM nodes is a fiber dedicated to the transmission in one direction. A duct (or span) comprises all the fibers between a given pair of nodes.

The nodes in VWP WDM networks perform full wavelength conversion, i.e. lightpaths can be switched to a fiber output if there is some free wavelength channel. The nodes in WP WDM networks do not have any wavelength converters at all. Lightpaths sourced with a wavelength retain this wavelength entirely. Therefore the nodes can switch a lightpath to a fiber output if the lightpath's wavelength is free.

From the network dimensioning point of view VWP networks can be regarded as a special case of SDH/Sonet networks with a connection granularity of a single capacity unit (one lightpath).

WP networks can be treated as VWP networks with additional wavelength constraints in the fibers and wavelength continuity constraints in the nodes. Consequently *p*-cycles have a color in WP networks, since conversion is not possible. Furthermore there can be situations where routed demands and *p*-cycles should have different wavelengths in both directions. See for instance the situation in Fig. 2 where a pair of fibers (one for each direction) connects two nodes. The two *p*-cycles can protect the bidirectional *on-cycle* working connection only, if the wavelengths of the working connection are different in both directions and the *p*-cycles have the complementary wavelengths. This is the reason why we consider in the following working paths and *p*-cycles as unidirectional. By this asymmetric topologies and demands can also be treated properly.

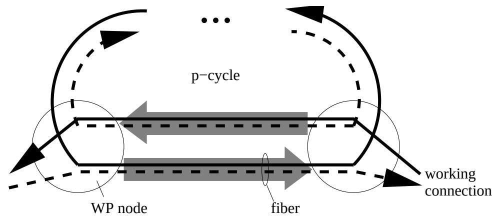

Fig. 2. A demand with two different wavelengths for each direction and a *p*-cycle with the complementary wavelengths.

A further restriction important for WP networks is the *p*-cycle length, because a lightpath can be attenuated too much if it becomes too long after protection switching. Length restrictions have also to be accounted for if the delay of a connection is limited.

IV. Optimal Combination of *p*-Cycles

Based on the network technologies of the previous section we give a mathematical formulation for the optimal combination of *p*-cycles in such networks. We assume the working connections to be routed in the network (in Section VI we will extend the influence of demand routing), consequently we obtain the spare capacity of the network.

Each duct in the network is represented by a pair of counterdirectional edges. The network is modeled as a directed graph $G = (V, E)$ where V represents the set of WDM nodes and E is the set of edges. Each edge j contains l_j fibers which in turn contain a set of wavelength channels K. $cost_j$ represents the cost for a capacity unit on edge j.

We find a set of cycles P in the spare capacity of the network by a breadth first search algorithm initiated from each duct. The cycles are subject to two characteristics: the cycles are simple (i.e. the nodes of the cycle path are pairwise different) and length-restricted (both physical length and number of hops). However, we also look at non-simple cycles (Section VI).

For the constraints of the mathematical models two incidence matrices are involved (computed after finding the cycles). $p_{i,j} \in \{0, 1\}$ indicates if edge j is element of *p*-cycle i. $x_{i,j} \in \{0, 1\}$ indicates if a working connection on edge j is protectable by *p*-cycle i. In contrast to [2] we use unidirectional *p*-cycles (Section III), therefore straddling fibers can be protected in one *p*-cycle direction only (and not in two, where $x_{i,j} \in \{0, 1, 2\}$, see Fig. 1 (c)-(d)). However, the advantage of having the choice between two protection paths is still maintained, since a straddling fiber can now be protected by two counterdirectional cycles.

A. VWP Networks

An edge j has a capacity of $c_j^{VWP} = l_j \times |K|$. On an edge j, w_j and s_j are the (given) number of working channels and the number of spare channels used by a *p*-cycle, respectively.

For the *p*-cycle configuration we are interested in the number of units n_i of a cycle i that is needed (i.e. the *p*-cycle capacity).

The problem for VWP networks can be formulated as follows:

$$\min \sum_{j=1}^{|E|} cost_j s_j \tag{1}$$

$$s_j = \sum_{i=1}^{|P|} p_{i,j}\, n_i, \quad \forall j \in E \tag{2}$$

$$w_j \leq \sum_{i=1}^{|P|} x_{i,j}\, n_i, \quad \forall j \in E \tag{3}$$

$$w_j + s_j \leq c_j^{VWP}, \quad \forall j \in E \tag{4}$$

$$n_i \in \{0, 1, 2, \ldots\}, \quad \forall i \in P \tag{5}$$

The objective (1) is to minimize the used spare resources, constraint (2) determines the protection capacity allocation, constraint (3) ensures the working capacity to be protected and by (5) we require integer *p*-cycle units. Note that (1)-(3) and (5) are according to [2]. The constraint (4) introduces the capacity restriction on an edge. By this the problem can become infeasible in which case we recommend an adapted demand routing or augmented network capacity.

B. WP Networks

Compared to VWP networks for the WP case we additionally need to take care of the individual wavelengths $k \in K$ of the system. This makes the problem more complex.

For an edge j containing $c_j^{WP} = l_j$ fibers, $w_{j,k}$ and $s_{j,k}$ are the (given) number of working channels with wavelength k and the number of spare channels used by a *p*-cycle with wavelength k, respectively.

For the *p*-cycle configuration we are interested in how many units $n_{i,k}$ of a cycle i with wavelength k are needed. There can be multiple *p*-cycle units using the same wavelength, if all *p*-cycle units traverse each edge on different fibers.

The problem for WP networks can be formulated as follows:

$$\min \sum_{j=1}^{|E|} cost_j \sum_{k=1}^{|K|} s_{j,k} \tag{6}$$

$$s_{j,k} = \sum_{i=1}^{|P|} p_{i,j}\, n_{i,k}, \quad \forall j \in E, k \in K \tag{7}$$

$$w_{j,k} \leq \sum_{i=1}^{|P|} x_{i,j}\, n_{i,k}, \quad \forall j \in E, k \in K \tag{8}$$

$$w_{j,k} + s_{j,k} \leq c_j^{WP}, \quad \forall j \in E, k \in K \tag{9}$$

$$n_{i,k} \in \{0, 1, 2, \ldots\}, \quad \forall i \in P, k \in K \tag{10}$$

The interpretation of the constraints is similar to the VWP case except for the further dimension introduced by the wavelengths. It can be observed that the WP problem can be solved for each wavelength k independently, which can reduce the calculation time and memory consumption very much.

V. CASE STUDY

The models of Section IV are formulated in AMPL and solved by CPLEX 6.6.0. In the study we optimize the pan-European COST 239 network [4] with 11 nodes and 26 ducts. We consider two demand patterns: `demand_1` is based on [4] where the traffic matrix entries are divided by 2.5 Gbit/s and interpreted then as lightpath demands (yielding 348 lightpath demands). `demand_2` is a centralized pattern where each node has a demand from and to node 7 in the network of [4] (yielding 800 lightpath demands).

The demands are routed on the shortest path, in the case of `routing_1` with equal link metrics and in the case of `routing_2` with metrics reciprocal to the free capacity of the link, which is calculated each time a demand has been routed. The former case tries to minimize the link resources and the latter case tries to achieve a balanced load on the links. In any case, a link with no remaining available capacity is removed for the routing.

Throughout the case study we use equal values for the number of fibers per edge $l_j = 2, \forall j \in E$ and fixed costs $cost_j = 1, \forall j \in E$, in order to minimize the link resources. The number of wavelengths of each fiber is $|K| = 128$.

VI. NUMERICAL RESULTS

We obtained the following results for VWP and WP in the case study of Section V. For simplicity we denote the maximum allowed physical *p*-cycle length as L_{max}, i.e. all cycles in P and thus all chosen cycles are not longer than L_{max}.

A. VWP Case

Fig. 3 shows the efficiency ratio $\sum_j s_j / \sum_j w_j$ over length L_{max} for the (scaled) pattern `demand_1` which is routed according to the two routing alternatives.

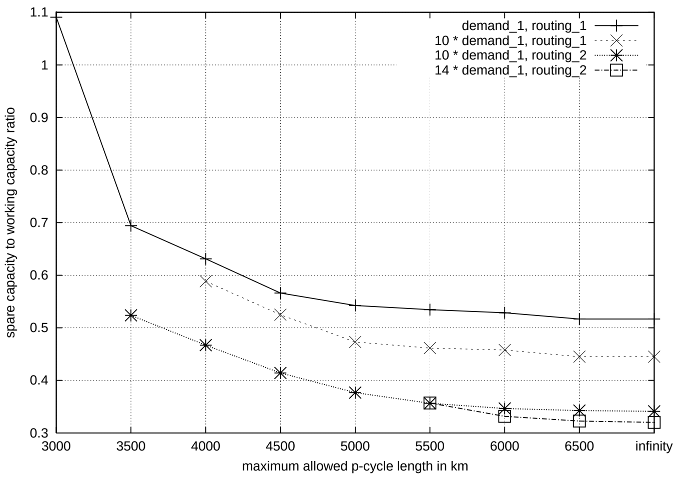

Fig. 3. VWP case: the efficiency ratio over the allowed maximum physical *p*-cycle length L_{max} for the scaled pattern `demand_1`.

If the pattern is not scaled, we obtain feasible solutions for L_{max} values greater than 3000 km. By restricting L_{max} to values greater than 4500 km, very good efficiency ratios lower than 60% are achieved. Nearly the same results are obtained for a scaling by 5 of pattern `demand_1`.

The efficiency ratio even improves in the situation where more traffic is offered to the network through a scaling of the demand pattern by 10. This is plausible, since existing cycles can cover a part of the working capacity of the additional demands. Although feasible solutions are only available for L_{max} values greater than 4000 km, protection capacity only needs 45% to 59% as much as working capacity. A substantial gain can be achieved by using `routing_2` instead of `routing_1`.

The network capacity limit is reached by a scaling of 14 of the pattern `demand_1`. For feasible solutions `routing_2` is needed and L_{max} greater than 5500 km. The working demands utilize 61% of the network capacity and the protection capacity takes not more than 36% of the working capacity. For protection not more than 22% of the total network capacity is needed!

Using this case study we are able to verify the efficiency of *p*-cycles as reported, e.g., in [2]. To achieve high total network utilization one has to take care of the demand routing. In general, `routing_2` yields better performance in efficiency ratio and network utilization than `routing_1`.

With `demand_2` we do not obtain these efficiency levels. For L_{max} greater or equal than 4500 km the efficiency ratio maintains a value of 84%. Similar to line-switched rings, *p*-cycles are less suitable to support centralized traffic patterns.

When allowing cycles to have more physical length, more cycles are available for the combination process as shown in Fig. 4. Therefore the efficiency ratio becomes better with growing length, since the solution space for the optimization procedure (the set of *p*-cycles in Section IV) enlarges and the objective must be steady or improve. The better performance of `routing_2` for 10 * `demand_1` also becomes obvious in Fig. 4, because many more cycles can be found than with `routing_1`. Some ducts are filled to 100% with working capacity by `routing_1` which makes these ducts useless for *p*-cycles.

We can proceed with the idea of increasing the set of *p*-cycles by using non-simple cycles, i.e. closed paths where a node can be traversed more than once. For functional *p*-cycles we restricted the cycles to visit an edge at maximum once. We note here that such non-simple cycles cannot protect every on-cycle duct, which has to be considered in the setting of the $x_{i,j}$'s in Section IV.

The number of such cycles grows much faster than the number of simple cycles, therefore we restricted L_{max} to 4000 km. In this case for 10 * `demand_1` and `routing_1` (`routing_2`) 13046 (14185) non-simple cycles are found in comparison to 864 (1098) simple cycles. At lower demand levels non-simple cycles can reasonably reduce the efficiency ratio.

With `routing_1` non-simple cycles improve the efficiency ratio from 59% (simple cycles) only to 58%. Virtually no improvement can be found with `routing_2`, where only one spare capacity unit can be saved.

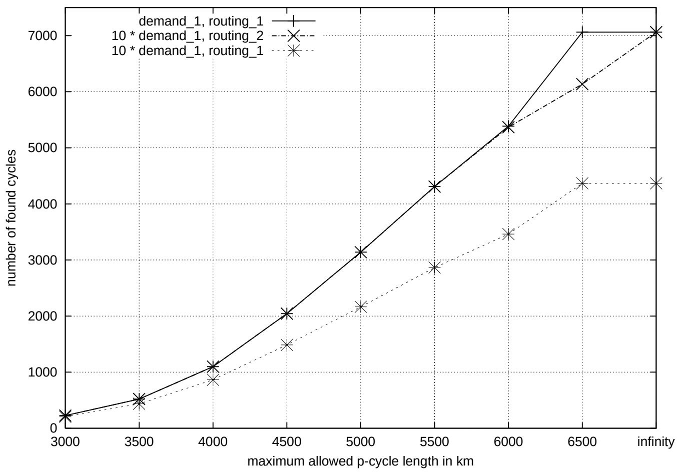

Fig. 4. VWP case: the number of found cycles over the allowed maximum physical *p*-cycle length L_{max}.

Fig. 5 depicts the minimum, mean and maximum number of *p*-cycle nodes over L_{max} (the drawing is similar if done over the maximum allowed *p*-cycle nodes). As the mean curve is situated nearer to the maximum curve, the optimization procedure (Section IV) tends to choose cycles with higher number of nodes. This is in accordance to the theoretical investigations of [5], stating that "there should be a significant number [...] of large *p*-cycles."

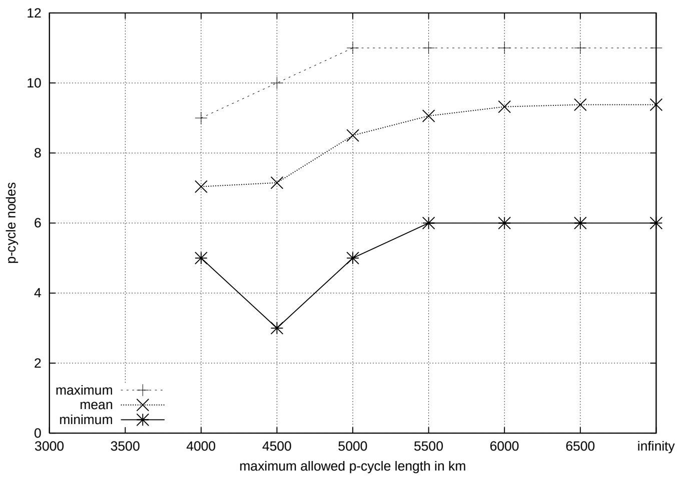

Fig. 5. VWP case with 10 * `demand_1` and `routing_1`: the minimum, mean and maximum number of *p*-cycle nodes over L_{max}.

B. WP Case

Fig. 6 shows the efficiency ratio $\sum_{j,k} s_{j,k} / \sum_{j,k} w_{j,k}$ over length L_{max} for the (unscaled) pattern `demand_1`. The working paths are routed with `routing_1` and their wavelengths are assigned first by order of the fiber and then by the first fit principle.

The behavior of the efficiency ratio over L_{max} is similar to the VWP case (Fig. 3), but the WP case performs slightly less efficient. By restricting L_{max} to values greater than 5000 km, efficiency ratios lower than 73% are achieved.

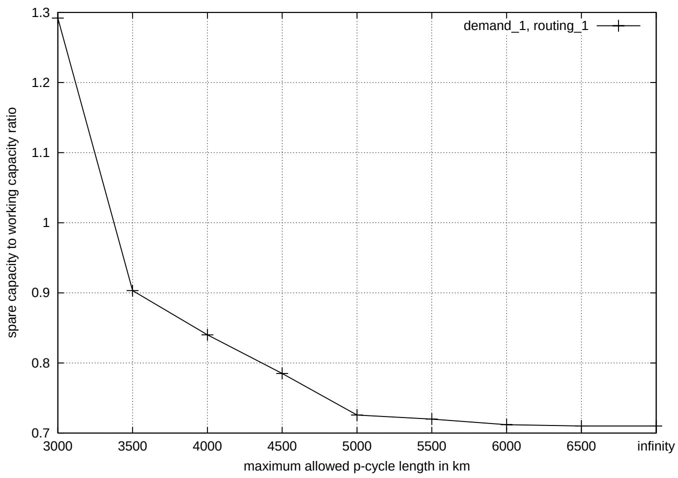

Fig. 6. WP case: the efficiency ratio over the allowed maximum physical *p*-cycle length L_{max} for the pattern `demand_1` routed by `routing_1`.

C. Computation Times

For the network planning process the computation time has to be rather short, in order to obtain design alternatives quickly. Moreover, during the operation of the network a fast reaction time upon failures has to be provided to recalculate a selection of *p*-cycles.

The total computation times are mostly much lower than three hours. The cycle search takes the biggest portion of the time as it can be seen in the typical case of Fig. 7. This is acceptable, since the search needs only to be performed once, in advance of any optimization procedures. Enhanced cycle search algorithms are also possible. The search of non-simple cycles takes even longer time (up to 20 hours) with marginal improvements obtained as described in the previous section.

The times for the other computations are negligible. When using `routing_2` instead of `routing_1` the routing computation time becomes slightly longer. Although CPLEX solves an integer linear program, the computation time is surprisingly short.

VII. Conclusions

We showed how *p*-cycles can be deployed in VWP and WP WDM networks and described the particularities of the latter. We further developed an integer linear program for the optimal configuration of the cycles in both cases. For a case study of a pan-European network we obtained the following results:

- For VWP networks the spare resources used for *p*-cycles of practical lengths (max. 4000-6000 km) consume about half of the capacity compared to the working resources. In our studies WP networks reach a spare to working capacity ratio of 71%. High capacity efficiency can be achieved.
- The *p*-cycle optimization performed better with a demand routing which tries to achieve balanced load on the links than a minimum hop routing.
- The longer *p*-cycles are allowed to be, the better the efficiency.
- The applicability of *p*-cycles is very dependent on the size of the network. High length values can introduce too much additional delay for a connection in protection state.
- The *p*-cycles selection tends to choose cycles with higher number of nodes.
- The deployment of non-simple link *p*-cycles is due to long computation times and low efficiency gains not recommendable.
- The configuration calculation time is very acceptable for a quasi-online computation, even if a reconfiguration after a failure is necessary.

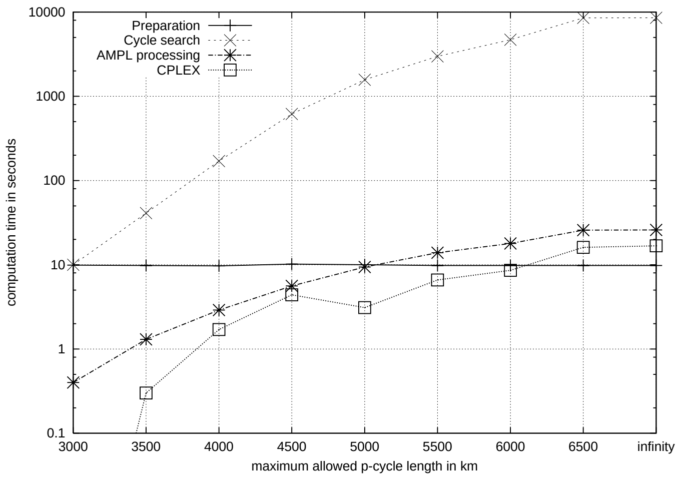

Fig. 7. Computation times for preparation (graph construction, routing, ...), for the cycle search, for the AMPL processing (data generation and processing), and for the execution of the solver CPLEX in the VWP case with 10 * `demand_1` and `routing_1`.

Another approach is to deploy partial wavelength conversion per node, providing a WP and a VWP part. Working paths can be routed in the WP part and *p*-cycles can be deployed in the VWP part, thus avoiding colored *p*-cycles.

The authors would like to acknowledge the suggestions of W.D. Grover.

References

[1] W.D. Grover and D. Stamatelakis, "Cycle-oriented distributed preconfiguration: ring-like speed with mesh-like capacity for self-planning network restoration," in *Proc. of ICC*, 1998.
[2] W.D. Grover and D. Stamatelakis, "Bridging the ring-mesh dichotomy with p-cycles," in *Proc. of DRCN Workshop*, 2000.
[3] D. Stamatelakis and W.D. Grover, "IP layer restoration and network planning based on virtual protection cycles," *IEEE JSAC*, vol. 18, Oct 2000.
[4] P. Batchelor et al., "Ultra high capacity optical transmission networks," Final report of action COST 239, 1999.
[5] D. Stamatelakis and W.D. Grover, "Theoretical underpinnings for the efficiency of restorable networks using preconfigured cycles ("p-cycles")," *IEEE Trans. on Comm.*, vol. 48, Aug 2000.