

COMPARISON OF p -CYCLE CONFIGURATION METHODS FOR DYNAMIC NETWORKS

Dominic A. Schupke

*Dominic A. Schupke, Siemens AG, Corporate Technology, Information and Communications,
Otto-Hahn-Ring 6, 81730 Munich, Germany*
dominic.schupke@siemens.com

Abstract: Dynamic optical networks which are protected by p -cycles can be operated by different p -cycle configuration methods. We compare the blocking probability of two dynamic p -cycle configuration approaches and one static p -cycle configuration approach (protected working capacity envelope).

1. INTRODUCTION

In this paper we evaluate p -cycles [1,2,3] in optical networks with dynamic connections. A p -cycle is a cycle in a network and can protect working capacity of links which have both end points on the p -cycle. These links include the links covered by the p -cycle as well as those non-covered links which join nodes of the p -cycle.

Protected connections are an important transport network service not only for existing, but also for emerging networks [4]. Even connections with short holding times can require protection, since a failure of an established connection may slow down or disrupt the user's application. Thus, besides it is desirable to have low connection blocking, an established connection should also have high availability, which we ensure in our consideration by p -cycles.

This paper is organized as follows. In Section 2 we briefly recapitulate the p -cycles protection concept and discuss several connection management approaches for dynamic networks in Section 3. In Section 4 we describe the assumed dynamic demand model and in Section 5 we evaluate the connection management approaches using simulation. In Section 6 we draw several conclusions.

2. *p*-CYCLES

Figure 1 depicts the protection principle of *p*-cycles for link protection. The *p*-cycle in Figure 1(a) is preconfigured as a closed connection on the cycle B-C-D-F-E-B. Preconfiguration means that the configuration is done before a failure occurs. The *p*-cycle is able to protect working capacity on its own links, called on-cycle links, as shown in Figure 1(b). Upon failure of on-cycle link B-C, the *p*-cycle offers protection by the route on the remainder of the cycle (C-D-F-E-B). The protectable capacity on on-cycle links is thus one capacity unit. The protection of on-cycle links is logically equal to multiplex-section shared protection rings (MS-SPRings) in SDH and bidirectional line-switched rings (BLSRs) in SONET.




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

Unlike rings, however, *p*-cycles also protect links outside the *p*-cycle path: Each link which has both its end points on the *p*-cycle can also be protected. These links are called straddling links. Figure 1(c) shows the protection of such a link (E-D). We can provide two protection routes for straddling links, in the example, routes E-B-C-D and E-F-D. In effect, we can protect two working capacity units of straddling links.

3. CONNECTION MANAGEMENT APPROACHES

We investigate three approaches to management of dynamic connections. The first two approaches can change *p*-cycles per new connection (dynamic *p*-cycles). In the third approach we assume an unchangeable *p*-cycle configuration (static *p*-cycles). In any case, we do not reconfigure existing working connections to accommodate newly arriving connections. All operations on the *p*-cycles are hitless for the working paths, since they affect the protection configuration only.

We can describe the link resources of the network, at a decision point in time i , by their capacity vectors for working (w_i) and protection (p_i). In addition, we have

a vector for the installed capacity (c). For a new demand arrival at time instance i , we have to decide whether a connection (requiring working capacity Δw_i) can be admitted to the network or not.

We propose three methods in the next subsections.

3.1 Routing in Spare Plus Protection Capacity (RSPC)

In this method, we accept a new connection demand if it can be protected by p -cycles after it is virtually routed in the non-working capacity $c - w_{i-1}$, i.e., if a feasible p_i exists for $w_{i-1} + \Delta w_i$. If it is possible, we release the p -cycle configuration, set-up the new connection ($w_i = w_{i-1} + \Delta w_i$), and configure the new p cycles (p_i). Otherwise we block the request and leave the configuration ($w_i = w_{i-1}$, $p_i = p_{i-1}$).

3.2 Routing in Spare Capacity (RSC)

Now, we accept a new connection if it is routable in the spare capacity $c - w_{i-1} - p_{i-1}$ and can be protected by p -cycles, i.e., if a working path (Δw_i) exists in the spare capacity and we find a feasible p_i for $w_{i-1} + \Delta w_i$. If it is possible, we set-up the new connection and reconfigure the p -cycles. Otherwise we block the request and leave the configuration. This approach may achieve shorter times, during which the p -cycles are in the process of reconfiguration, than RSPC.

3.3 Routing in Protected Working Capacity Envelope (RPWCE)

In this principle (see [2,5]), we assume a static protected working capacity envelope (PWCE) in which connections are set up and torn down. While the PWCE can also be dynamic, with less frequent change actions than for connection requests, a static PWCE is clearly the easiest solution and can model a dynamic PWCE in quasi-static operation points.

In comparison with the above two approaches, the static PWCE is even simpler, since connection management deals only with configuration of paths and not with configuration of both paths and p -cycles.

For all time instances i , we have a fixed amount of protection capacity $p_i = p$ and a fixed protectable working capacity w' , i.e., the PWCE. We accept a new connection if it is routable in the PWCE, i.e., if a working path (Δw_i) exists in the PWCE which fulfils $w_{i-1} + \Delta w_i \leq w'$. Otherwise we block the request and leave the configuration. An accepted connection is automatically protected by the p -cycles.

4. DYNAMIC DEMAND MODEL

We derive a dynamic demand model from a static traffic matrix (which may be prescaled). For every demand unit, we implement one client dynamically requesting services, e.g., a general client interfacing by a user-to-network interface (UNI) or a requesting multilayer control instance setting up a connection for an IP router port pair. A single client follows an on-off-state model with a time distribution for the on-state and a time distribution for the off-state, regardless whether the network can provide a connection or not [4]. In other words, a client tries to set up a connection at off-to-on transition and releases an established connection for it at on-to-off transition.

Clients request connections independently from each other. If not enough available resources are found in the network, the request will be rejected. The estimation for the network-wide blocking probability using a sequence of connections requests is the ratio of rejected requests over generated requests.

5. EVALUATION

We evaluate the blocking performance of the hypothetical Germany network (17 nodes, 26 links) in [6]. We assume full wavelength conversion and length-based cost. We dimension it for the unscaled demand in [6] using firstly shortest path routing, yielding working capacity w_{dim} , and secondly p -cycle optimization [1,3], yielding protection capacity p_{dim} .

Dynamic and static p -cycles use this dimensioned capacity differently during network operation. For dynamic p -cycles (RSPC and RSC), we find connections and p -cycles in the aggregate capacity, i.e., $c = w_{\text{dim}} + p_{\text{dim}}$. Also for dynamic requests, we firstly route an admitted connection and secondly we configure the p -cycles. For static p -cycles, we find connections in the RPWCE, i.e., $w' = w_{\text{dim}}$. The overall capacity in the network is the same for all cases.

As in [4], the time distribution for the on-state is deterministic and the time for the off-state is exponentially distributed. The ratio of the on-state time to the mean time for the off-state is 1/12, i.e., if we take 12 dynamic traffic sources between a given pair of nodes, we will have the same total traffic demand in the dynamic network on average, as in the static traffic matrix (e.g., clients use their connections continuously for 1 hour out of 12 hours on average).




Figure 2: Blocking of dynamic connections over demand scaling factor.

Since an analytical treatment of dynamic configuration (paths and p -cycles) appears to be too complex, we evaluate the blocking performance by discrete event simulation. Although the network is dimensioned with the shortest paths, we use shortest available path routing with physical length as metric, since, except for high blocking levels, it yields lower blocking probabilities than shortest path routing. For RSPC and RSC, the computation of the p -cycles (p_i) has to finish within 10 s, which can be regarded as adequate time for an on-demand request [4]. We also allow an optimality gap of 1%.

Figure 2 depicts the simulation results for RSPC, RSC, and RPWCE. We show the blocking probability with confidence intervals (at 95% confidence level) over the demand scaling factor, i.e., the demand in [5] is prescaled by the value at the x-axis.

For an upper limit of 3% for the blocking probability, we can scale the demand matrix to integers of 6 for RPWCE, 9 for RSC, and 11 for RSPC. On the one hand, the RPWCE concept does not achieve as low blocking probabilities as dynamic p -cycles, since it is less flexible in finding capacity for a demand request. On the other hand, the higher complexity of RSPC and RSC (routing and p -cycle configuration) is opposed to the simplicity of RPWCE (routing only), which can justify to install more capacity for the RPWCE to reduce the blocking probability. For example, assume RPWCE follows the economy of scale rules of loss systems for which the Erlang-B formula apply; to shift the blocking performance curve of RPWCE in Figure 2 to the region of RSPC, i.e., to roughly half the blocking probability, by the economy of scale effect we need less than double the capacity

of the PWCE. An appraising simulation for the considered blocking range shows that RPWCE well outperforms RSPC if the PWCE of each link (and likewise its protection capacity) is scaled by 1.75.

RSPC is superior to RSC, since working paths are closer to the shortest paths, for which the network is dimensioned. This can also be explained by the average network load. While the load by RSPC is over 80% for the displayed values, the load by RSC drops to 63% for the demand scaling of seven. Hence, RSPC is able to utilize the network better than RSC. In RSC, working paths can deviate from the shortest paths, since protection capacity on links of the shortest paths prohibits using it. The relative outcome for RSPC and RSC is different from the study in [4]. In this study, however, the network is dimensioned for dedicated path protection, for which dispersion from the shortest paths can be beneficial (see also the results with balanced load routing in [3]).

6. CONCLUSIONS

We evaluated p -cycles in dynamic optical networks by simulation. Among the configuration approaches routing in spare plus protection capacity, routing in spare capacity, and routing in protected working capacity envelope, a case study has shown that the blocking probability increases in the order mentioned, given that the overall network capacity is the same for the three approaches. This is because the former two approaches use dynamically configured p -cycles which can adapt flexibly to the traffic situation. The latter approach, however, is significantly less complex in operation than the former two approaches.

Further research could investigate in detail how much capacity is needed for the protected working capacity envelope concept in excess to the one needed for an optimal assignment of dynamic p -cycles, given that the blocking probability is equal.

ACKNOWLEDGMENTS

This work is part of the work when the author was with the Institute of Communication Networks at Technische Universität München, Munich, Germany.

REFERENCES

- [1] W. D. Grover and D. Stamatelakis. Bridging the ring-mesh dichotomy with p-cycles. In Proceedings of the International Workshop on Design of Reliable Communication Networks (DRCN), (Munich, Germany), April 2000. Invited Talk.
- [2] W. D. Grover. Mesh-based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking. Prentice Hall, Upper Saddle River, 2003.
- [3] D. A. Schupke, C. G. Gruber, and A. Autenrieth. Optimal Configuration of p-Cycles in WDM Networks. In Proceedings of the IEEE International Conference on Communications (ICC), (New York City, NY, USA), April–May 2002.
- [4] D. A. Schupke, M. Jaeger, and R. Huelsermann. Comparison of Resilience Mechanisms for Dynamic Services in Intelligent Optical Networks. In Proceedings of the International Workshop on Design of Reliable Communication Networks (DRCN), (Banff, AB, Canada), October 2003.
- [5] W. D. Grover. The Protected Working Capacity Envelope Concept: An Alternate Paradigm for Automated Service Provisioning. IEEE Communications Magazine, 42(1):62–69, January 2004.
- [6] R. Hülsermann, S. Bodamer, M. Barry, A. Betker, C. Gauger, M. Jäger, M. Köhn, and J. Späth. A Set of Typical Transport Network Scenarios for Network Modelling. In Proceedings of ITG-Fachtagung Photonische Netze, (Leipzig, Germany), May 2004.