Benefits of Connection Request Bundling in a PCE-based WDM Network

(Invited paper)

J. Ahmed, P. Monti, L. Wosinska

Department of Microelectronics and Applied Physics, Royal Institute of Technology
KTH, Stockholm, Sweden
Tel: +46 87904083, Fax: +46 87527850, E-mail: {jawwad, pmonti, wosinska}@kth.se

The Path Computation Element (PCE) concept is considered to be beneficial in the network connection setup operations, especially in optical networks based on wavelength division multiplex (WDM) transport technology. In the PCE paradigm, communication between a node and the PCE is specified by the Path Computation Element Communication Protocol (PCEP). PCEP allows the PCC (Path Computation Client) to send to the PCE more than one LSP (path computation) request at a time, i.e., multiple LSP requests can be bundled together before being sent to the PCE. Enabling bundling, and consequently the concurrent optimization of a large set of LSP requests at the PCE, may result in significant improvements in terms of network optimization and reduced control plane overhead. However, these advantages come at a cost of increased connection setup-delay. This paper explores pros and cons of enabling bundling of LSP requests in terms of both control plane overhead reduction and benefits of sequential vs. concurrent path computation operations. A variety of scenarios are analyzed, including a WDM mesh network providing LSPs with both dedicated and shared path protection. Results demonstrate significant gains in terms of reduced control overhead using LSP bundling, and reduction in blocking probability using concurrent processing of bundled LSP requests at the PCE.

1. Introduction

Currently, a GMPLS based control-plane is one of the available options for the automated setup and tear down operations of LSPs (Label Switch Paths) in communication networks. Path computation operations require a lot of processing power, and may be subject to multiple constraints such as wavelength continuity, physical impairments, and QoS requirements (e.g., delay, bandwidth, and load balancing). On the other hand, the GMPLS based control plane is distributed in its nature and each node in the network is assumed to have enough power to perform the necessary path computation operations. However, this assumption may not always be true. In that case a Path Computational Element (PCE) [1] based network architecture can be deployed where the PCE serves as a centralized entity specialized in solving complex multi-constrained LSP path computation problems. The PCE concept was validated in an experimental network [2], where the performance of the path computation procedure was evaluated in both single and multi-area networks. Performance was then compared with a conventional, GMPLS-based, distributed path computation approach. Results demonstrated both efficient resource utilization and good network scalability of the PCE-based approach. A more detailed experimental evaluation of PCE-based versus distributed approach can be found in [3].

In the PCE paradigm, communication between a node and a PCE (as well as communication among multiple PCEs) is specified by the Path Computation Element communication Protocol (PCEP) [4]. The PCEP protocol defines the communication between the Path Computation Client (PCC) and the PCE, where such interactions include path computation requests and path computation replies.

An important feature of the PCEP protocol is the possibility to bundle multiple LSP requests in a single PCReq message to be sent to the PCE for path computation. Similarly, multiple LSP replies can be grouped together in a single PCRep message before they are sent back to the PCC. Enabling bundling, and consequently the concurrent optimization of a large set of LSP requests, may (*i*) produce substantial improvements in terms of overall network optimization (i.e., concurrent versus sequential optimization), and (*ii*) reduce the overall control plane overhead which in turn allows a more efficient management of control plane resources, especially when the control plane is overloaded. However, these advantages come at a cost of increased connection setup-delay. A first assessment of this trade-off was presented in [5], where the control plane overhead reduction was studied as a function of the setup time, in a scenario where path computation was performed sequentially by the PCE.

This paper extends the work presented in [5] by including a benefits analysis of enabling the PCE to optimize the path computation procedures by concurrently considering the entire LSP set in the bundle. The analysis is done in a WDM network scenario with unprotected, dedicated (path) protected and shared (path) protected LSPs. The objectives of the current study are: *(1)* estimate the beneficial effects of bundling, in terms of control overhead reduction, *(2)* study the extra overhead introduced in the PCEP protocol by the inclusion of path protection *(3)* evaluate the trade-off between connection setup-delay and communication overhead via a time-threshold based bundling mechanism *(4)* identify possible negative effects bundling may have on the network blocking probability, *(5)* study the benefits of using a concurrent RWA heuristic, *(6)* estimate the extent of increase in the LSP setup-time, when concurrent RWA is deployed on the PCE (in contrast to sequential RWA). With these objectives in mind a concurrent Routing and Wavelength Assignment (RWA) algorithm is proposed, and its performance is compared against a similar approach already available in the literature [6].

The paper is organized as follows. First the LSP bundling approach is described in Section 2. Then the proposed concurrent RWA heuristic algorithm is presented in Section 3. Finally, in Section 4, performance results for all three types of LSP requests (no protection, dedicated protection, shared protection) are presented to evaluate the benefits of concurrent processing of LSP requests. Concluding remarks are given in Section 5.

2. Bundling Approach

PCEP is used for the communication between a PCC and a PCE, or between two different PCEs. Two important messages defined for PCEP are: PCReq message that is used to send LSP requests from a PCC to a PCE and PCRep message that is used to send the computed LSP's back from the PCE to a PCC. PCEP allows sending multiple LSP requests in a single PCReq message, and similarly computed LSP requests from the PCE can also be bundled together in a single PCRep message, to be send back to the PCC. This feature is exploited in the presented study to reduce PCEP control overhead. Note that, since path protection capabilities are provided should one LSP require them, LSPs in a bundle are assumed to be both *synchronized* (from a concurrent optimization point of view) and *dependent* (because of the dependency of the protection path computation on the primary path). For the above mentioned reasons the SVEC (Synchronization Vector) objects need to be introduced in PCReq messages.

At each node bundling is enabled via a time-threshold based approach. An alternative would be to employ a connection- threshold approach where each ingress node (PCC) waits for a specific number of requests to arrive before bundling them together in PCReq message. However, time based threshold ensures an upper bound on the waiting time for the LSP requests on the ingress nodes before being sent to the PCE.

3. Proposed Heuristic for Concurrent RWA Processing

A greedy meta-heuristic to concurrently process a bundle of LSP requests that arrive at the PCE for path computation is proposed. It requires a pre-processing phase which works as follows:

1. Compute a set of K-shortest routes for each of the source destination pairs in the network to be stored at the PCE. These K-shortest routes will be used by the employed routing algorithm to select an appropriate primary route for each of the LSP requests.
2. Compute L link-disjoint routes for each of these K shortest routes to be used by the employed routing algorithm to select a suitable secondary route in case of shared/dedicated protection.

The proposed heuristic performs the following steps iteratively to process all the requests in a bundle:

1. Place all the LSP requests in the current bundle in an input queue.
2. For each of the LSP requests (LSP_Req) present in the queue compute a “temporary” RWA solution (using the selected sequential routing and wavelength assignment algorithms that are discussed later), and record the computed value of the objective function $F_{obj}(LSP\_Req)$ (described later).
3. Sort the LSP requests in the queue in ascending order to put LSP request with the lowest value of $F_{obj}(LSP\_Req)$ at the head of the queue.
4. Select the request at the head of the queue and remove it from the queue.
5. For the selected request choose the route for the primary path from the K-shortest routes, and select the link-disjoint protection route (if a protection mode is selected) from the L alternate routes (for that selected primary route) using the selected sequential routing algorithm.
6. Perform the wavelength assignment using the selected wavelength assignment algorithm for the primary and protection-route (if a protection mode is chosen).

7. Check if there are still any LSP requests remaining in the input queue. If yes go to Step 2, otherwise terminate and consider RWA computation for all the requests in the bundle to be completed.

It is easy to see that the dominant part in terms of computation time complexity is Step 2, where temporary routing and wavelength assignment is performed a number of times (equal to the number of LSP requests in the queue) per iteration. In total, it is required to perform temporary RWA solution $n(n+1)/2$ times until the heuristic terminates to find the concurrent solution, where n is the total number of requests in a bundle. It should be noted that Step 4 and 5 are required if temporary RWA solutions performed in Step2 are not stored in the PCE to reduce storage overhead.

Now the RWA algorithms to be used in Step 2, 4 and 5 of the proposed heuristic are described. First-Fit is used for the wavelength assignment. Note that in the shared-protection mode, First-Fit will always first try to find a feasible shared wavelength to be used for the protection-path before finally resorting to reserve a free wavelength. Enhanced Weighted Least Congested Routing (EWLCR) [7] is used for routing because of its very good performance in the dynamic-traffic scenario. EWLCR computes the weight function $W(R)$ for each route R of the input candidate K-routes, and a route with the highest weight value is selected. $F(R)$ is one of the input parameters to compute $W(R)$, denoting the total number of free wavelengths for the route R. In case of a network with no-wavelength conversion a wavelength is considered free only if it is available on "all" the links included in route R.

The EWLCR is used for the routing in dedicated protection and no-protection case as described in [7], but for the shared-protection case a minor but important modification to the algorithm is introduced. For the protection route computation in the shared-protection scheme the $F(R)$ parameter is replaced with another parameter $T(R)$, such that $T(R) = F(R) + S(R)$, where $S(R)$ denote the "number of common shared wavelengths" and $T(R)$ denote the "total number of available wavelengths" (including shared) on all the links of the route R. However, the shared wavelengths to be included in the $S(R)$ computation need to satisfy a certain shared-protection criteria, namely *"a shared wavelength in the protection path can be used only if its primary path is link-disjoint with the primary paths of rest of the protection paths using that shared wavelength".*

The objective function $\mathrm{F_{obj}(LSP\_Req)}$ that will return a value based on the primary and secondary paths computed by the routing (EWLCR) and wavelength assignment (First-Fit) algorithms for the temporary RWA solution of the current LSP request (LSP_Req) is defined as follows:

$$F_{obj}(LSP\_\mathrm{Re}q) = \begin{cases} F_{obj}(P_{pri}), & \text{If no-protection case} \\ F_{obj}(P_{pri}) + F_{obj}(P_{Sec}), & \text{otherwise} \end{cases} \tag{1}$$

$$F_{obj}(P_{pri}) = W_{New}(P_{pri}) \tag{2}$$

$$F_{obj}(P_{Sec}) = \begin{cases} W_{New}(P_{Sec}), & \text{If dedicated-protection case} \\ W_{New}(P_{Sec}) + W_{Resv}(P_{Sec}), & \text{otherwise} \end{cases} \tag{3}$$

where P_{pri} is the primary path for the current $\mathrm{LSP\_Req}$, P_{Sec} is the protection path for the current $\mathrm{LSP\_Req}$, $F_{obj}(LSP\_Req)$ is the objective function value for the current $\mathrm{LSP\_Req}$, $F_{obj}(P_{pri})$ is the objective function value for the primary path computed by the temporary RWA, $F_{obj}(P_{Sec})$ is the objective function value for the protection path computed by the temporary RWA, W_{New} is the number of wavelengths reserved by the current path that are NOT already used in the network and W_{Resv} is the "total" number of wavelengths reserved by the current path.

It can be observed that F_{obj} will prefer (compute a lower value) the temporary RWA solutions that reuse the already used wavelengths in the network (W_{New}). In addition, for the shared-protection case it will favor RWA solutions that reserve the minimum number of "new" wavelengths along the selected protection route, but rather use already shared wavelengths where possible(W_{Resv}). Inclusion of the objective function F_{obj} makes this heuristic flexible because it can be tailored to some specific requirements under the given network scenario.

4. Simulation Setup and Environment

In this work the POSE discrete event-driven simulator [8] is extended and a detailed simulation model for a PCE-based optical network is implemented. Simulation results are collected under EON (European Optical Network) topology. It is assumed that each link in the network is bidirectional (one fiber in each direction) with 20 wavelengths per fiber. Wavelength continuity constraint (WCC) is enforced in the network. For the sequential case, at the PCE routes are computed using the EWLCR [7] algorithm while First-Fit is used for wavelength-assignment. For the concurrent case, the heuristic described in Section 3 is used for both routing and wavelength assignment. LSP request arrivals follow a poison distribution and service times are exponentially distributed.

The following performance parameters are considered in the paper: blocking probability, PCEP overhead and average LSP setup-time. Performance is evaluated as a function of the time-threshold value. The starting time, for the threshold-based timers, is randomly offset at each node, to avoid synchronization. The estimation of the PCEP bandwidth overhead includes: TCP, IP and Ethernet overhead, assuming that the control plane is implemented over Ethernet. LSP setup-time includes not only path computation, communication and queuing time, but also the time necessary for reserving the computed path through the network (i.e., signaling time). The PCE is assumed to be *state-full* in this case, i.e. it maintains not only the free wavelength information for each link in the network, but also the state related to the currently established connections in the network (at-least for the shared-protection mode) where this information is required by the PCE to compute the protection paths for the LSP requests. The total number of LSPs required to be established is 10,000. The mean LSP service time (μ) is equal to 60 s, while the arrival rate (λ) is set to 1/80 arrivals per second per node-pair assuming a uniform traffic load per node-pair.

5. Results

This section presents simulation results for different performance metrics as the time-threshold is varied from 1 s to 80 s in the dedicated, shared and no path protection scenarios. For each scenario, four different curves are plotted, namely: a *Baseline* approach (where there is no bundling of LSP requests), a bundling approach employing *Sequential* RWA, a bundling approach employing the proposed *Concurrent* RWA, and finally a bundling approach employing a *Competing Heuristic* [6], at the PCE.

The dedicated path protection case is presented in Figure 1. Figure 1(a), shows an increase in blocking probability (20%) for the Sequential RWA when compared to the Baseline scenario for a time-threshold value of 20 s. The blocking probability increase can be reduced to just 10% when the Concurrent RWA approach is used. Blocking increases rapidly after this time-threshold because too many LSP requests need to be setup at approximately the same time. From Figure 1 (b), it can be seen that with the Concurrent approach a 25% reduction in control overhead can be achieved when the time-threshold is equal to 20 s. Control overhead can be reduced further, but at the expense of an increased blocking probability and LSP setup-time. Note that there is small overhead penalty when employing the Concurrent approach as compared to Sequential RWA because of the inclusion of SVEC objects in the PCReq message. However, this penalty brings a significant blocking probability reduction (Figure 1 (a)).

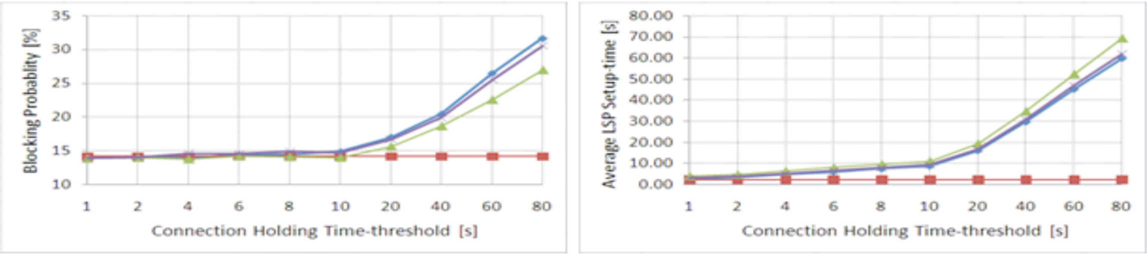

(a) Blocking Probability [%]

(c) Average LSP Setup-time [s]

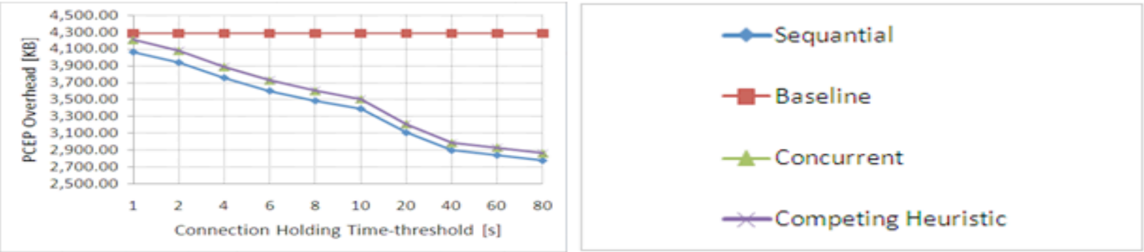

(b) PCEP Overhead [KB]

Figure 1: Performance results for the dedicated-protection case.

As for the LSP setup-time (Figure 1 (c)), it increases more rapidly after a time-threshold of 20 s. At this time-threshold an increase of 16% in the setup-time can be observed as compared to the Sequential RWA case. However this value is negligible when compared to the contribution made by the holding-time (20 s). From these performance figures, it can be concluded that the use of the Concurrent RWA approach in the given network scenario, employing a time-threshold of 20 s can result in significant reduction in control overhead without a major increase in the average LSP setup-time and minor increase in blocking probability. It is also interesting to note that at time-thresholds lower than 20 s, the Concurrent RWA approach can achieve a lower blocking when compared to the *Baseline* case (i.e., at a time-threshold of 10 s). Although this improvement is minimal in most cases, it is worthwhile considering that a noticeable reduction in control overhead is also achievable at the same time.

Figure 2 shows the shared path protection case. As expected the blocking probability (Figure 2 (a)) in general is much lower when compared to the dedicated-protection case. The trend is similar to the case with dedicated protection, but here the proposed approach is even more effective against the Sequential RWA approach. For a time-threshold of 20 s Concurrent RWA can reduce the incurred blocking probability penalty (compared to Baseline) from 58% to just 17%. This is partly due to the fact that the proposed heuristic facilitates wavelength sharing when computing protection paths. From Figure 2 (b), it can be seen that using the Concurrent RWA approach a 22% reduction in control overhead can be achieved at the time-threshold equal to 20 s. From (Figure 2 (c)) it can be noticed that the LSP setup-time increases more rapidly after a time threshold of 20 s. At this time-threshold a 20% increase in setup-time can be observed using the Concurrent approach as compared to the Sequential RWA. It is also interesting to note that the Competing heuristic also performs much better here than Sequential RWA but not as well as the proposed scheme, although computation overhead for the Competing RWA heuristic is lower, as shown in Figure 2(c).

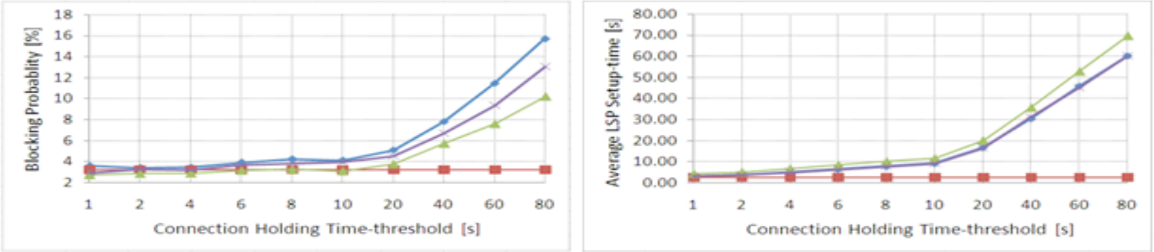

(a) Blocking Probability [%]

(c) Average LSP Setup-time [s]

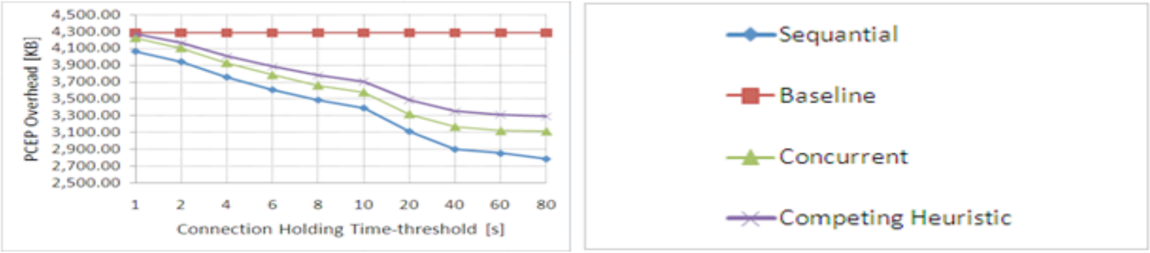

(b) PCEP Overhead [KB]

Figure 2: Performance results for the shared-protection case.

Figure 3 shows the no protection case. As expected, here the blocking probability (Figure 3 (a)) in general is much lower when compared to both the dedicated and shared-protection case. The trend is similar to the shared-protection case, but better reduction in blocking probability is achievable by using Concurrent RWA in contrast to the Sequential approach. For a time-threshold of 40 s the blocking probability is reduced by almost 52% using the Concurrent approach. From Figure 3 (b), it can be observed that by using the Concurrent approach the control overhead can be reduced by 50% when the time-threshold is equal to 40 s. Difference between the Sequential and the Concurrent approach is also negligible at this time-threshold. For the LSP setup-time (Figure 3(c)) one can see that at the time-threshold of 40 s a penalty of 20% increase in setup-time is associated with the Concurrent approach as compared to the sequential case.

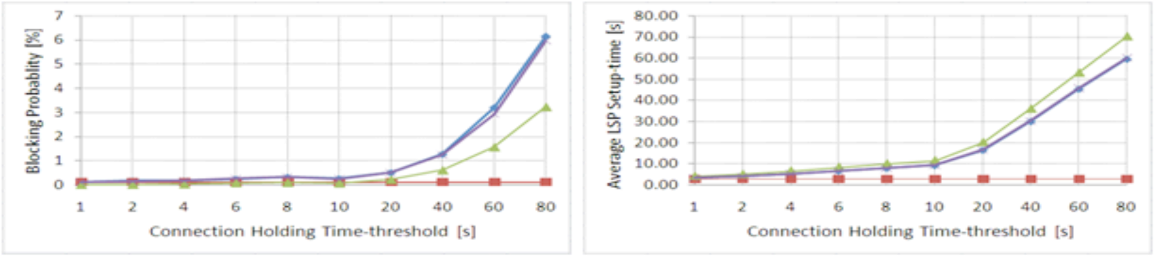

(a) Blocking Probability [%]

(c) Average LSP Setup-time [s]

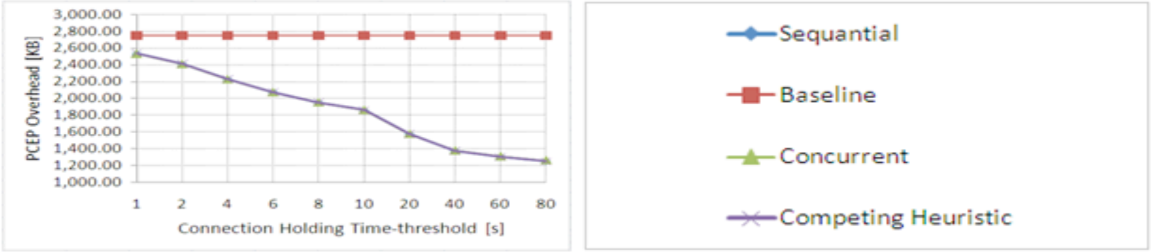

(b) PCEP Overhead [KB]

Figure 3: Performance results for the no-protection case.

6. Conclusion

This paper presents a performance study of a time-threshold based LSP request bundling approach. In particular, the attention is focused on the benefits analysis of enabling the PCE to optimize the path computation procedure by concurrently considering the entire LSP set in the bundle. A concurrent RWA approach was presented and analyzed in a WDM network scenario where LSPs require dedicated, shared or no protection. Results demonstrate that, by carefully choosing an appropriate time-threshold a significant reduction in communication overhead can be achieved without a noticeable increase of the LSP setup-time or overall network blocking probability under different traffic protection scenarios. Furthermore, a concurrent RWA heuristic tailored to these protection scenarios allows to significantly reduce blocking probability usually associated with sequential processing at the PCE.

References

[1] A. Farrel, et al., “A Path Computation Element (PCE) Based Architecture,” IETF RFC 4655, Aug. 2006.

[2] F. Cugini, et al., “Implementing a Path Computation Element (PCE) to encompass physical impairments in transparent networks,” *OFC/NFOEC 2007.*

[3] P. Castoldi, et al., “Centralized vs. Distributed Approaches for Encompassing Physical Impairments in Transparent Optical Networks,” ONDM 2007.

[4] J.-P. Vasseur and J. L. Le Roux, “Path Computation Element (PCE) Communication Protocol (PCEP)”, Internet draft, March 2008.

[5] J. Ahmed, et al., “LSP Request Bundling in a PCE-Based WDM Network,” *OFC/NFOEC 2009.*

[6] H. Zang, et al., "Path-protection Routing and Wavelength-Assignment in WDM Mesh Networks under Shared-Risk-Group Constraints". APOC 2001

[7] J. Ahmed, et al., “Dynamic routing in wavelength convertible WDM networks,” *INMIC 2004.*

[8] J. Ahmed and S. Nawaz, “POSE: A New Discrete Event Optical Simulator for the Java Platform,” *NCET 2004.*