

LSP Request Bundling in a PCE-Based WDM Network

J. Ahmed, P. Monti, L. Wosinska
 The Royal Institute of Technology (KTH), Stockholm, Sweden
 {jawwad, pmonti, wosinska}@kth.se

Abstract: LSP requests may be bundled to improve the network optimization process at the expense of an increased connection setup delay. A detailed study is conducted to evaluate the pros and cons of the bundling approach.

© 2008 Optical Society of America

OCIS codes: (060.0060) Fiber optics and optical communications; (060.4250) Networks

1. Introduction

In optical networks, automatic setup and teardown of connections, also referred to as Label Switched Paths (LSPs), can be supported by a GMPLS-based control plane. LSP setup includes both path computation and resource reservation. Path computation is computationally expensive, and may be subject to multiple constraints such as wavelength continuity, physical impairments, and QoS requirements (e.g., delay, bandwidth, and load balancing). Typically, in a GMPLS-based network, path computation is performed at ingress nodes in a distributed manner.

A distributed LSP path computation approach enables, on one side, network scalability, but it assumes, on the other, that all nodes in the network have sufficient computational resources to compute complex multi-constrained paths. This assumption, however, may not always be true. Computational power may be limited at some nodes, and/or some legacy equipment may not support the required control plane path computation functionality.

To cope with such scenarios, a central entity referred to as Path Computation Element (PCE) was proposed [1]. The PCE serves as an entity specialized in solving complex multi-constrained LSP path computation requests. 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 one (or several) PCE is specified by the Path Computation Element communication Protocol (PCEP) [4]. The PCEP protocol defines communication between the Path Computation Client (PCC) and the PCE, where such interactions include path computation requests and path computation replies. Among other options, the PCEP protocol allows the PCC to send to the PCE more than one LSP (path computation) request at a time. In other words, multiple LSP requests can be *bundled* before being sent to the PCE. 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), (ii) be highly beneficial (i.e., reduce race conditions) whenever a large set of traffic requests are competing, almost simultaneously, for the same resources (e.g., in the event of catastrophic failures and/or with highly dynamic traffic patterns), and (iii) reduce the overall control plane overhead which in turns allow a more efficient management of control plane resources, especially when the control plane is overloaded. All these advantages come at a cost in terms of connection setup delay. A tradeoff, that to the best of the authors' knowledge has never been studied, needs to be found to fully exploit the benefits provided by the PCE architecture.

This paper presents pros and cons, measured in terms of setup-time versus communication overhead, of enabling the bundling of LSP requests by the PCEP. The purpose of the study is three fold: (i) estimate the beneficial effects of bundling, in terms of overhead reduction, (ii) evaluate the tradeoff between connection setup delay and communication overhead via a time-threshold-based bundling mechanism and (iii) identify possible effects bundling may have on the network blocking probability.

Presented results show that, by choosing an appropriate time threshold for the bundling procedure, a significant reduction in communication overhead can be achieved without a noticeable increase of the LSP setup-time and/or the overall network blocking probability.

2. A time-threshold-based bundling approach

In the framework of PCEP, a *PCReq* message is used to send a LSP request from a PCC to a PCE, and a *PCRep* message is used to send the computed LSP request 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 in a single *PCRep* message, to be sent back to the PCC. Note that in the presented study bundled LSP requests are assumed to be *non-synchronized* and *independent*. With this assumption the SVEC (Synchronization Vector) object, as specified for PCEP format, is not required.

A time-threshold-based approach is used to bundle LSP requests at each node, where PCC generates a bundled *PCReq* message only after a timer (equal to the threshold value) expires. An alternative approach could be to employ a connection threshold value, where each ingress node waits for a specific number of LSP requests to arrive before bundling all of them in one *PCReq* message. However, a time-threshold-based approach can be beneficial since it ensures an upper bound on the waiting time for each LSP request at the ingress node.

3. Simulation setup and environment

Results are obtained extending the POSE discrete event-driven simulator [5] to implement a detailed simulation model for a PCE-based optical network. Simulation results are collected using the European Optical Network (EON) topology, which comprises 19 node and 39 links. Each link in the network is bidirectional with one fiber and 20 wavelengths for each direction. Wavelength conversion is not allowed, i.e., the wavelength continuity is enforced. Routes at the PCE are computed using the WLCR algorithm [6], while First-Fit is used for wavelength assignment. LSP requests' arrival rate is assumed to follow a Poisson distribution and the service time is assumed to be exponentially distributed. A Destination Initiated Reservation (DIR) signaling scheme similar to RSVP-TE is employed for resource reservation after the path computation is finalized.

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.

For each performance result a *Baseline* scenario is also identified, in which LSP request bundling is not used. 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 total number of LSPs required to be established is 10,000. The mean LSP service time is equal to 60 s, while results are collected for 3 different arrival rate (λ) values: (low) λ , equal to 1/120 arrivals per second, (medium) λ , equal to 1/90 arrivals per second, and (high) λ , equal to 1/60 arrivals per second.

4. Simulation results

Fig. 1 shows the effect of bundling on the blocking probability as the time-threshold is varied from 1 to 100 seconds. Results are also reported for the Baseline case (no bundling of LSP requests), for each of the three arrival rate cases. Results show that blocking probability is almost non-existent until the time threshold value reaches 20 seconds, under both low and medium arrival rate cases. Blocking increases rapidly after 60 seconds because of the congestion in the network, i.e., too many LSP requests need to be setup at approximately the same time. This effect is even more pronounced in the higher arrival rate case, when the blocking probability curve is even steeper.

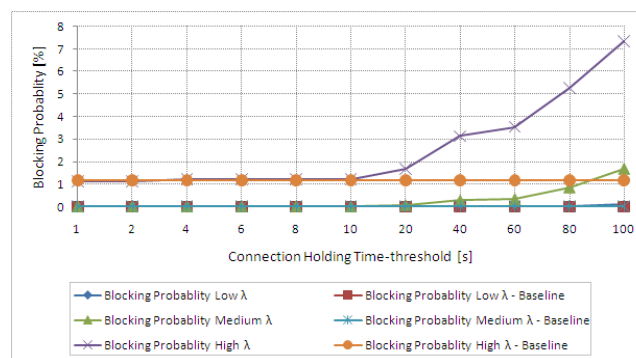


Fig. 1. Blocking Probability [%] as a function of connection holding Time-threshold [s]

Fig. 2 presents the reduction of PCEP overhead. About 42% reduction in control overhead can be achieved when the time threshold is equal to 20 seconds with medium arrival rate. For a threshold of 60 seconds this value is about

49% without much increase in blocking. Also the reduction in PCEP overhead, when the arrival rate goes from medium to high, is not very significant.

From Fig. 3 it can be observed that the LSP setup-time increases more rapidly after a time threshold of 20 seconds, but this increase is linear. With high arrival rate, the average LSP setup-time is significantly higher even for the Baseline scenario because the single PCE in the network becomes the bottleneck and the set-up delay due to bundling is only a fraction of the total setup-time. By evaluating these performance figures, it can be concluded that, under the given network scenario, employing a time threshold of 20 seconds can result in significant reduction in control overhead without a major increase in the average LSP setup-time and no increase in blocking probability.

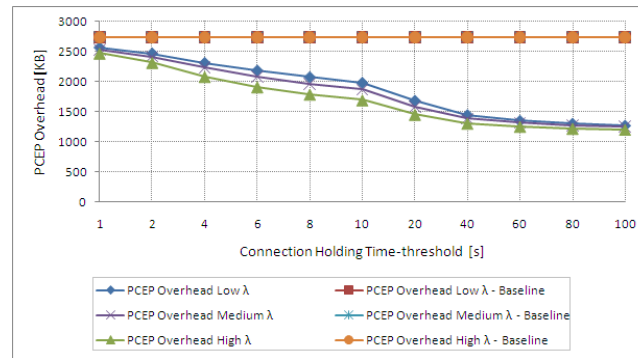


Fig. 2. PCEP Overhead [KB] as a function of connection holding Time-threshold [s]

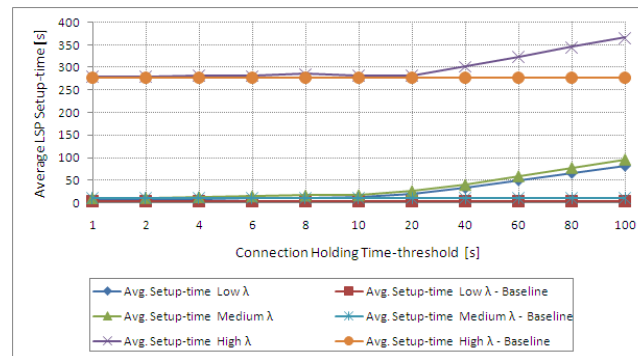


Fig. 3. Average LSP setup-time [s] as a function of connection holding Time-threshold [s]

5. Conclusion

In the paper a study was presented where the performance of LSP request bundling was evaluated using a time-threshold-based mechanism. Results show that, by carefully choosing the time threshold a significant reduction in communication overhead can be achieved without a noticeable increase of the LSP setup delay and/or the overall network blocking probability.

References

1. A. Farrel, J.-P. Vasseur, and J. Ash, "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 and S. Nawaz, "POSE: A New Discrete Event Optical Simulator for the Java Platform," NCET 2004.
6. X.-W. Chu, et al., "A Dynamic RWA Algorithm in a Wavelength-Routed All-Optical Network with Wavelength Converters," IEEE INFOCOM 2003.