A Novel Event-Driven QoS-Aware Connection Setup Management Scheme For Optical Networks

Wissam Fawaz, Abdelghani Sinno, Raghed Bardan, and Maurice Khabbaz

Lebanese American University,
Electrical and Computer Engineering Department,
Byblos, Lebanon
{wissam.fawaz,abdelghani.sinno,raghed.bardan,
maurice.khabbaz}@lau.edu.lb
http://soe.lau.edu.lb

Abstract. This paper proposes a QoS-Aware Optical Connection Setup Management scheme that uses the Earliest Deadline First (EDF) queueing discipline to schedule the setup of the optical connections. The benefits of this EDF-based scheme are twofold: *a*) it reduces the blocking probability since blocked connection requests due to resource unavailability are queued for possible future setup opportunities and *b*) it realizes QoS differentiation by ranking the blocked requests in the EDF queue according to their connection setup time requirements, which are viewed as deadlines during connection provisioning. As such, pending lesser delay-tolerant requests are guaranteed to experience better QoS than the ones having longer setup time requirements. Extensive simulations are performed to gauge the merits of the proposed EDF-based scheme and study its performance in the context of two network scenarios, namely, the National Science Foundation Network (NSFNET) and the European Optical Network (EON).

Key words: Optical networks, connection setup management, Earliest Deadline First (EDF) scheduling, performance evaluation.

1 Introduction

Wavelength Division Multiplexing-based (WDM) optical networks are foreseen to support the numerous ever-emerging applications having distinct Quality of Service (QoS) requirements. The evolution of such networks along with all related technological developments are driven by the urgent need to cater for these requirements. QoS-Aware WDM-based optical networks come forth as a promising solution, the realization of which is, however, a major challenge. Over the last couple of years, a large body of researchers contended in addressing this challenge through various proposals (*e.g.* [1–3]). These proposals mainly aimed at enabling WDM-based optical networks to provide the so-called Predictable Quality of Transport Services (PQoTS) by monitoring a set of parameters that affect an established connection's data flow.

The *Connection Setup Time* (CST) defined as the maximum amount of time that elapses between the instant an optical connection is first requested and the instant the requested connection is setup, is an important parameter that may significantly improve the optical connections setup management. However, this parameter has been given little attention. In fact, CST can be interpreted as a *deadline* prior to which a received connection request must be established and thus provides an opportunity for network operators to carry out QoS differentiation during connection provisioning. According to the authors in [4, 5], CST is expected to become an integral part of an optical connection's service profile and is thus foreseen as a potential service differentiator in the Service Level Agreements (SLAs) established between optical operators and their clients.

Inspired by the above observation, we presents a novel QoS-Aware Optical Connection Setup Management strategy that uses CST as both an indicator of the priority of a connection request and a measure of the delay tolerance associated with that request. Whenever connection blocking occurs due to the lack of optical resources, the blocked connection setup requests are queued and scheduled in an order that is consistent with their deadlines. This objective is achieved by utilizing the well known *Earliest Deadline First* (EDF) queueing discipline whereby the connection request with the smallest setup time (i.e., earliest deadline) is served first. As such, what distinguishes this newly proposed strategy from others found in the open literature is the introduction of QoS differentiation among the incoming connection requests during the course of connection provisioning.

The rest of this paper is structured as follows: in Section II, the EDF-based connection setup scheme is described in details, then its relation to the literature is discussed. Section III introduces the simulation framework used to highlight the merits of the proposed setup strategy. The simulation results are given in Section IV. Finally, Section V concludes the paper.

2 Related Work

In [6–8], the authors investigated the problem of dynamic bandwidth allocation for Deadline-Driven Requests (DDRs). For this purpose, they designed several algorithms that enabled transmission rate flexibility throughout the DDR provisioning procedure in WDM-based optical networks. To this end, their followed approach differs from ours in that they considered the deadline to be the maximum connection holding time. Nonetheless, in order to increase the fraction of successfully provisioned DDRs the algorithms proposed in [6–8] can be easily supplemented with the management technique developed in this paper.

The authors of [9] studied the effectiveness of using an EDF-based optical for managing the setup of point-to-point connections in optical networks with single-wavelength fiber links. The work in [10] appears as an extension to [9] where the

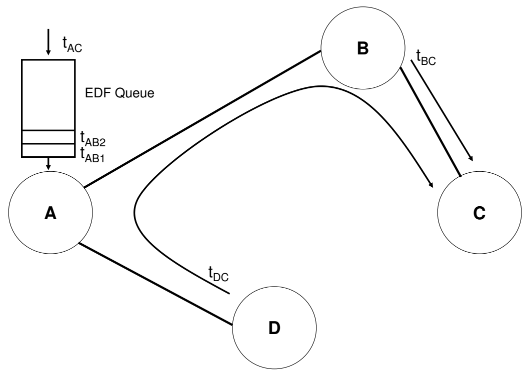

Fig. 1. A sample 4-node optical network.

authors considered the multi-wavelengths fiber links' case. However, both of the aforementioned studies lacked generality as they did not assess the performance of their corresponding EDF-based connection setup management schemes when applied to a wavelength routed optical network. Indeed, it was proven in [11–13] that the EDF-based strategies proposed in [9, 10] may not be as efficient as expected when utilized in such networks. Instead, the authors presented improved strategies which they showed to be more suitable for wavelength routed optical networks. Nevertheless, those improved strategies also suffered from two major limitations: *a*) they were driven only by connection departures and *b*) they solely served the head-of-line pending request when a departure event occurred.

In this paper, we alleviate the above observed deficiencies by introducing an improved event-driven EDF-based connection setup policy that, first, accounts for both a connection's arrival and departure, and second, ensures the setup of a wider spectrum of the pending requests upon the occurrence of a departure or an arrival event. This is achieved by having the proposed scheme target not only the head-of-line pending request but also a large portion of the other pending requests that may potentially be provisioned into the network.

3 Description of the Proposed Scheme

For illustration purposes, the sample network topology given in Figure 1 is used to explain the main idea behind the EDF-based setup scheme. The figure shows four optical cross-connects (OXCs), namely: A, B, C, and D, that are connected together by means of 3 fiber links. Each OXC serves an incoming connection request by attempting to establish an end-to-end lightpath connecting the source node of the connection request to its destination. When such an attempt fails, the connection request is said to be blocked. Historically, blocked connection requests used to be immediately dropped. Instead, we propose to insert such connections into an EDF queue and to arrange them in an ascending order of their Connection Setup Time (CST) requirements. This solution is motivated by the studies made in [4, 5], which stipulate that CST is a parameter that determines a connection request's class of service where the smaller a connection request's CST value is, the higher its associated class of service becomes. It is important to emphasize on the fact that, in the context of the proposed EDF-based connection setup policy, the priority of a connection request waiting in the EDF queue increases as that pending request approaches the deadline prior to which it must be established. If a CST expires before its corresponding connection request gets provisioned, then this request is called a *dead request*. The way an EDF queue treats a dead request depends on whether a *work-conserving* or a *non work-conserving* EDF policy is implemented. In particular, the non work-conserving variant immediately drops a dead request, whereas the work-conserving one further holds dead requests in the queue until they get served. In this paper, the more realistic non-work conserving EDF variant is considered.

Let us illustrate the main operation of the EDF-based setup scheme by considering the following scenario. Suppose that the capacity of each fiber link in Figure 1 is limited to 2 wavelengths and that 2 connections t_{BC} and t_{DC} are already established from B to C and from D to C respectively. For the sake of simplicity, each connection setup request is assumed to be requiring a full wavelength of bandwidth. Furthermore, let us assume that the EDF queue associated with A is finite and contains 2 previously blocked requests t_{AB1} and t_{AB2} destined for B, with t_{AB1} being the head-of-line pending request. t_{AB1} and t_{AB2} have deadlines of 1 and 2 units of time associated with them, respectively. Eventually, under such circumstances, a setup request t_{AC} addressed to C that arrives at A with a deadline of 3 would be blocked and consequently the event-driven EDF-based setup strategy would come into play. The event-driven aspect of the EDF-based setup scheme is highlighted by the fact that it is activated on the occurrence of an arrival event. Recall that the proposed scheme is driven by both arrivals and departures. Once the setup management scheme is activated, an attempt is made to serve the pending request occupying the front of the EDF queue. In the context of the scenario under study, t_{AB1} will hence be provisioned into the network. Then, the setup scheme turns to the next pending request attempting to serve it. This process continues until the setup strategy reaches a pending request that cannot be routed into the network. In this case, since the establishment of t_{AB2} turns out to be impossible, the setup scheme stops and inserts the blocked request t_{AC} into the EDF queue at the appropriate position relative to t_{AB2}. Given that t_{AB2}'s deadline is less than that of t_{AC}, t_{AC} ends up being queued behind t_{AB2}. As time evolves, the degree of urgency of t_{AB2} and t_{AC} increases. Ultimately, if one of the pending requests reaches its deadline, that request is immediately dropped out of the queue, in which case a *deadline mismatch* is said to have taken place. Subsequent connection requests whose deadlines are less than the deadline associated with t_{AC} are placed in front of t_{AC} in the EDF queue and as such are served prior to t_{AC}. Note that if the number of such connection requests is large enough, t_{AC} may end up being pushed out of the EDF queue.

Now consider what happens upon the occurrence of a connection departure event. Say the previously provisioned t_{AB1} connection departs from the network before t_{AB2}'s deadline is violated. On the occurrence of the departure event, the EDF-based setup scheme is activated requiring A to examine its associated EDF queue one request at a time and to establish each pending request in turn into the network. This process terminates when A encounters a pending request that cannot be provisioned. In the case of the considered scenario, A would succeed in provisioning t_{AB2}, but would fail in serving t_{AC}. As such, t_{AC} becomes the sole pending request in the EDF queue and would thus be obliged to wait until the next arrival or departure event occurs before retrying to gain access into the network.

In summary, the EDF-based connection setup strategy proposed in this paper is activated upon the occurrence of two types of events, namely the departure and the arrival of connections. When a connection emanating from an arbitrary source node A departs from or arrives to the network, the setup strategy proceeds as follows. It scans through the EDF queue associated with A aiming at provisioning as many pending requests as possible. This process continues until either all pending requests are provisioned or the setup strategy comes across a pending request whose setup is impossible, in which case the setup strategy stops its probing for possible connection setups. This suggests that the event-driven EDF-based setup strategy enjoys a wide setup probing scope.

4 Simulation Study

A Java-based discrete event simulator was developed to analyze the performance of the proposed event-driven EDF-based connection setup strategy in the context of two real-life optical networks, namely: *a*) the National Science Foundation Network (NSFNET) and *b*) the European Optical Network (EON). NSFNET's topology is shown in the bottom part of the composite Figure 2, while the one corresponding to EON is depicted in the topmost part of the Figure. NSFNET consists of 24 nodes and 43 bidirectional fiber links while EON is composed of 19 nodes together with 39 bidirectional links. The data relating to the physical topologies of both networks was taken from [14, 15]. Our simulation study is based on the following assumptions:

1. In both optical networks, it is assumed that each node has a full wavelength conversion capability.
2. Incoming connection requests are uniformly arranged into 3 service classes referred to as gold, silver, and bronze.
3. Each incoming request requires a full wavelength of bandwidth.
4. The overall arrival process is Poisson and the connection holding time is exponentially distributed with a mean normalized to unity.
5. Following the guideline presented in [4, 5], the parameters associated with the three service classes are as follows:

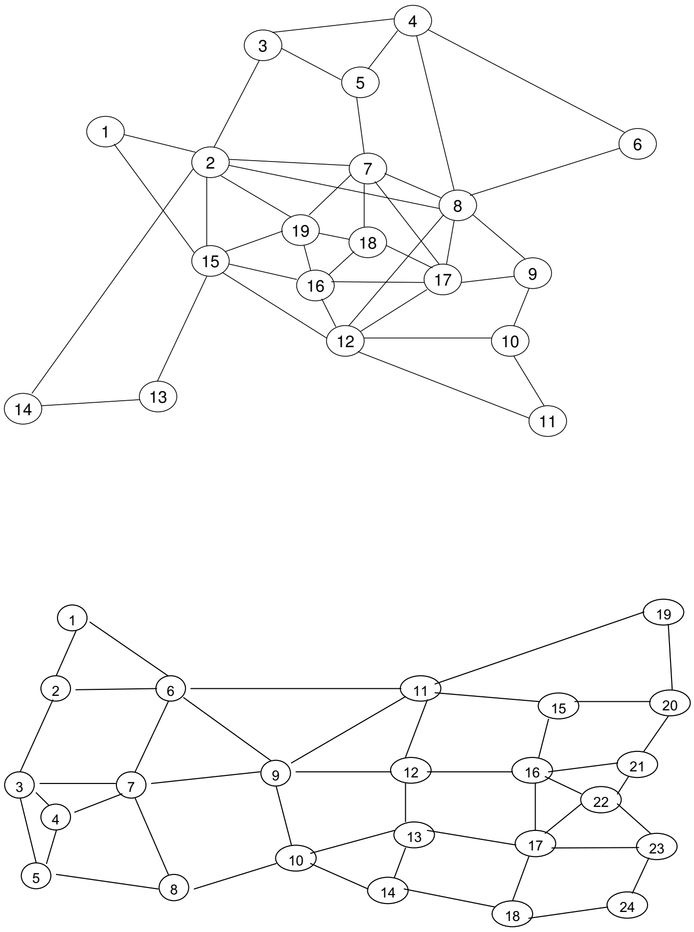

Fig. 2. Network topologies used in simulation.

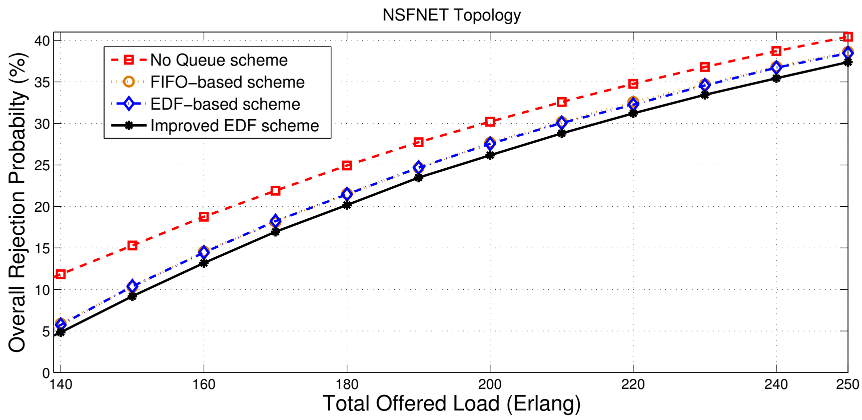

Fig. 3. Overall rejection probability for different setup strategies (NSFNET).

   – Gold connection requests arrive with an initial deadline of 6 units of time.
   – Silver requests have deadlines of 10 units of time associated with them.
   – The initial deadline of the bronze requests is set to 14 time units.
6. One EDF queue is deployed per optical node with a capacity to hold up to 20 pending connection requests.
7. Dijkstra's algorithm is used to find the shortest path for the arriving connections, while wavelengths are assigned to the provisioned connections according to a first-fit strategy.
8. The capacity of each fiber link is set to 8 wavelengths.

It is important to stress that 10^6 connection requests are simulated per run of the simulator and that each obtained value of the results is the average of the outcomes of multiple simulation runs to ensure that a 95% confidence interval is realized. The 10^6 simulated connection requests are uniformly distributed among the nodes of the considered optical networks.

4.1 Performance Metrics and Benchmarks

The performance metrics used to gauge the benefits of the proposed event-driven EDF-based connection setup strategy are: (i) the overall rejection probability and (ii) the rejection probabilities for gold, silver, and bronze connection requests. Note that the rejection probability is nothing else but the fraction of connection requests whose access to the network is blocked. Blocking could occur either due to buffer overflow or due to deadline mismatch which, as mentioned earlier, happens when a request's CST expires prior to its provisioning.

Three connection setup management approaches will serve as benchmarks:

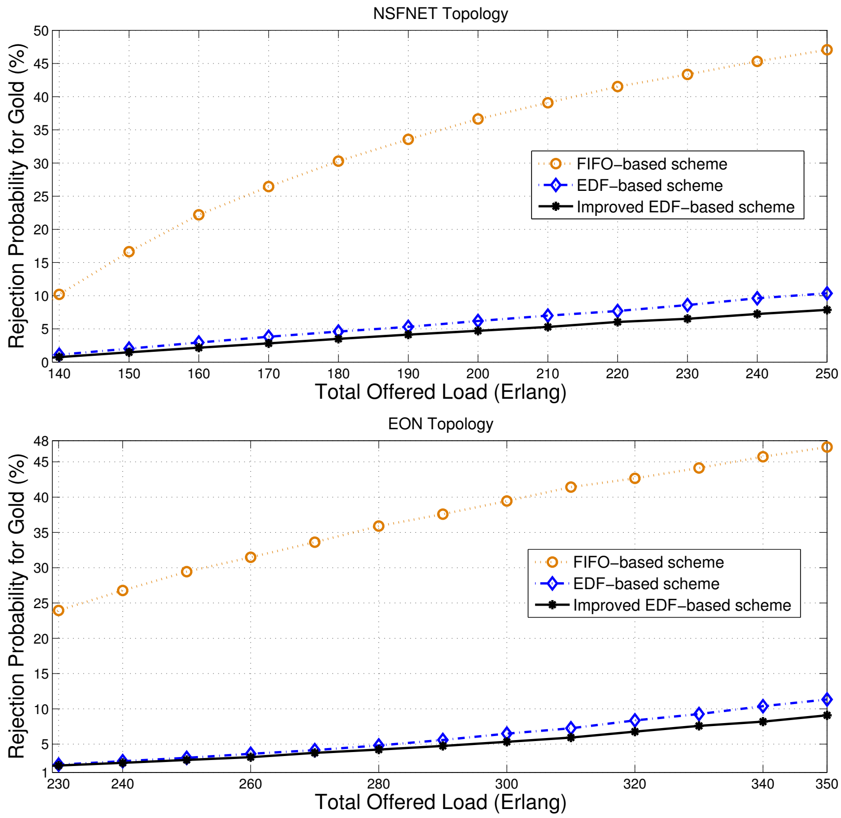

Fig. 4. NSFNET: Gold rejection probability for FIFO, EDF, and IEDF based setup schemes (uppermost); EON: Gold rejection probability for FIFO, EDF, and IEDF based setup schemes (bottom).

- A queue-free connection setup strategy, where no queues are used to store the blocked connection requests due to optical resource unavailability. This strategy will be referred to henceforth as the No-Queue strategy.
- A First In First Out (FIFO) queue-based connection setup scheme, where blocked connections requests are queued and then served according to the FIFO principle.
- The EDF-based connection setup mechanism studied in [11–13].

In order to distinguish our newly proposed event-driven EDF-based strategy from the one of [11–13], we will refer, in what follows, to our strategy as the Improved EDF-based (IEDF) connection setup strategy.

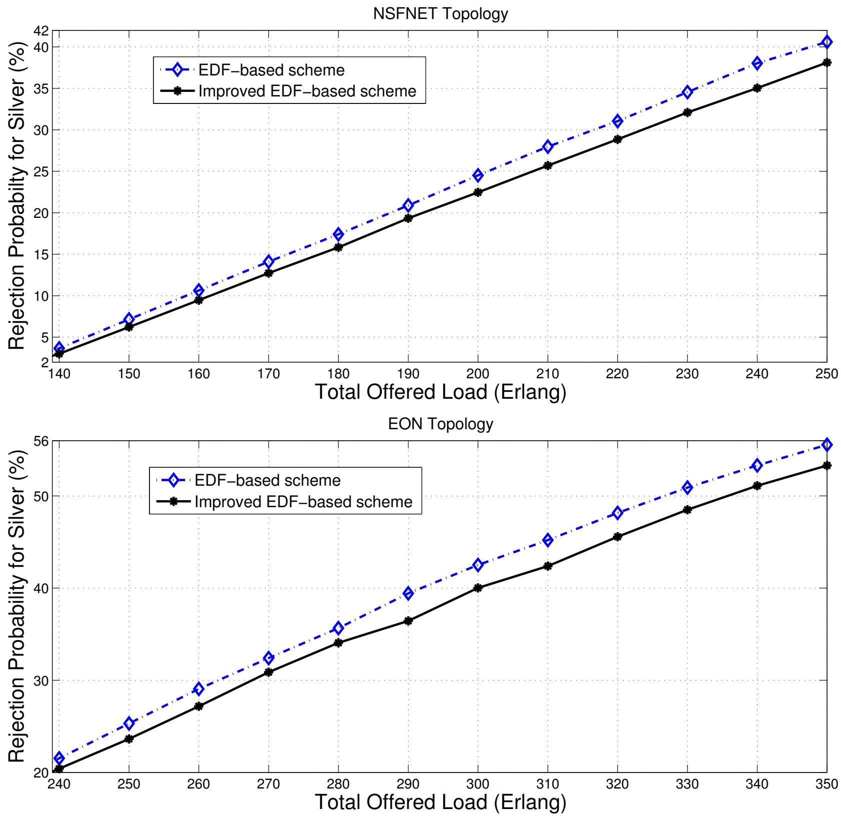

Fig. 5. NSFNET: Silver rejection probability for EDF and IEDF based setup schemes (uppermost); EON: Gold rejection probability for EDF and IEDF based setup schemes (bottom).

4.2 Numerical Results

Consider the NSFNET topology shown in bottom part of Figure 2. In this context, Figure 3 compares the overall rejection probabilities achieved by all of the IEDF and the three benchmark schemes as a function of the load offered to the network. IEDF clearly outperforms the other three strategies as it presents the lowest blocking probabilities. In contrast, a No-Queue scheme yields the worse performance in terms of the overall blocking probability since, simply, blocked connection requests are immediately dropped. Building on this observation, this scheme will not be considered in the subsequent set of results. It is worth mentioning that by limiting their setup probing scope to only the head-of-line pending request, the FIFO-based and the traditional EDF-based schemes achieved the same overall blocking probabilities and hence their blocking probability curves overlapped with each other.

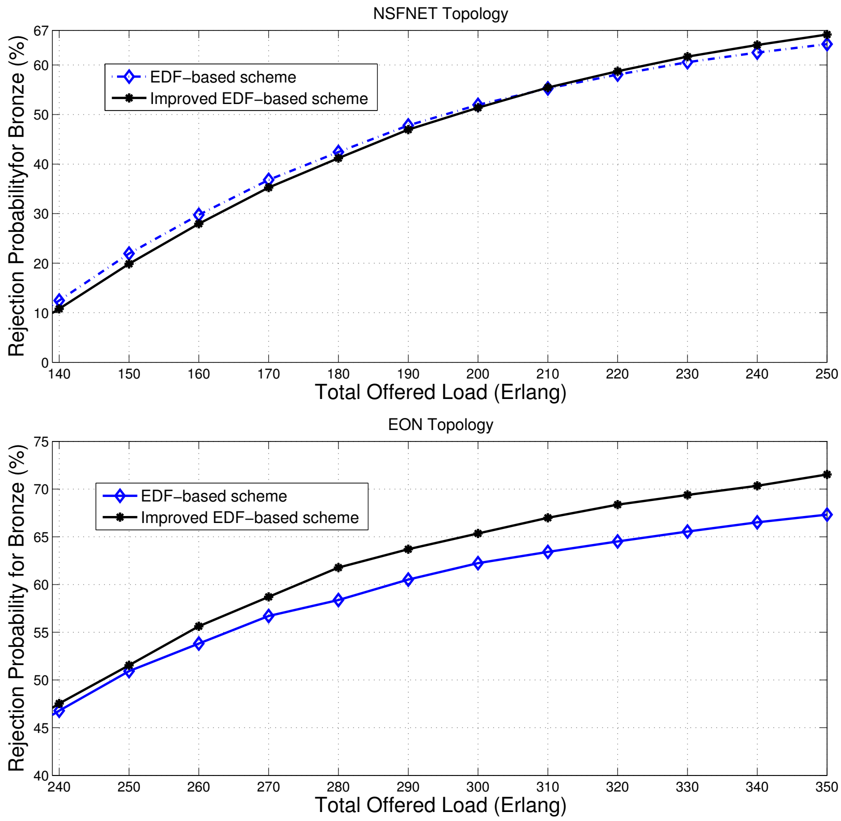

Fig. 6. NSFNET: Bronze rejection probability for EDF and IEDF based setup schemes (uppermost); EON: Bronze rejection probability for EDF and IEDF based setup schemes (bottom).

The rejection probabilities for gold connection setup requests resulting from the deployment of the FIFO-based, EDF-based, and IEDF connection setup schemes are graphed in Figure 4 in the context of both, the NSFNET and EON networks. Based on these reported results, smaller gold rejection probabilities are observed for the EDF-based and IEDF schemes relatively to the FIFO-based strategy. This finding can be justified by the fact that, in terms of access to the network, the EDF-based and IEDF schemes privilege the connections with the smallest deadline requirements (*i.e.* gold connections). Furthermore, in contrast to the traditional EDF-based scheme, targeting more than one of the front pending gold connection setup requests upon the occurrence of a departure or arrival event, enabled IEDF to provision a larger number of those requests. This is due mainly to the fact that gold connection setup requests are most likely to be found towards the front of the EDF queue because of their small deadline requirements and thus have a higher chance of being provisioned on time under the proposed IEDF scheme. Also in the context of both NSFNET and EON, Figure 5 shows the rejection probability associated with silver connections for different values of the network's offered load. The results demonstrate that IEDF is also hard to beat when it comes to the provisioning of silver connection setup requests in comparison to the traditional EDF-based scheme. This is again due to the fact that silver connection setup requests occupy the middle of the EDF queue and thus can benefit from the wider setup probing scope characterizing IEDF. This feature causes more silver setup requests to be provisioned on time and accordingly reduces the silver rejection probability.

Finally, Figure 6 compares the performance of IEDF to the that of the EDF-based strategy in terms of the rejection probability corresponding to bronze requests as a function of the network offered load in the context of the NSFNET and EON networks. The fact that IEDF privileges gold and silver connection setup requests in terms of network access comes at the expense of bronze requests. This explains the slightly degraded performance that the bronze requests experience under IEDF compared to the EDF-based strategy.

5 Conclusion

This paper proposes IEDF, an event-driven EDF-based connection setup scheme that has the luxury of triggering the setup of pending connection requests upon the occurrence of either a connection departure or the arrival of a new connection. This reduces the likelihood of incidence of deadline mismatches. An additional characterizing feature of IEDF is its ability to provision multiple pending connection setup requests per single event occurrence as opposed to existing EDF-based policies that consider provisioning only the head-of-line pending request. Extensive simulations were conducted to evaluate the performance of IEDF and to measure its impact on the quality of service perceived by the end clients. Throughout the simulation study, the performance of IEDF was contrasted with that of multiple other benchmark schemes, including the traditional EDF-based setup scheme. The reported simulation results proved that IEDF has the upper hand when it comes to rejection probability improvement. Moreover, these results inferred that IEDF is able to simultaneously achieve QoS differentiation and lower blocking probability of the incoming connection requests without compromising the privilege of higher priority clients with respect to network access.

References

1. Szymanski, A., Lanson, A., Rzasa, J., Jajszczyk, A.: Performance evaluation of the Grade-of-Service routing strategies for optical networks. In: IEEE ICC 2008, pp. 5252-5257. IEEE Press, New York (2008)
2. Martinez, R., Pinart, C., Cugini, F., Andriolli, N., Valcarenghi, L., Castoldi, P., Wosinska, L., Comellas, J., Junyent, G.: Challenges and requirements for introducing impairement awareness into the management and control planes of ASON/GMPLS WDM networks. IEEE Communications Magazine. 44, 76–85 (2006)

3. Jukan, A., Franzl, G.: Path selection methods with multiple constraints in service-guaranteed WDM networks. IEEE/ACM Transactions on Networking. 12, 59–72 (2004)
4. Fawaz, W., Daheb, B., Audouin, O., Berde, B., Vigoureux, M., Du-Pond, M., Pujolle, G.: Service Level Agreement and provisioning in optical networks. IEEE Communications Magazine. 36–43 (2004)
5. Sambo, N., Pinart, C., Le Rouzic, E., Cugini, F., Valcarenghi, L., Castoldi, P.: Signaling and Multi-layer Probe-based Schemes for guaranteeing QoT in GMPLS Transparent Networks. In: IEEE OFC 2009, pp. 22-26. IEEE Press, New York (2009)
6. Andrei, D., Batayneh, M., Martel, C.U., Mukherjee, B.: Provisioning of Deadline-Driven Requests with flexible transmission rates in different WDM network architectures. In: IEEE OFC 2008, pp. 1-3. IEEE Press, New York (2008)
7. Andrei, D., Batayneh, M., Martel, C.U., Mukherjee, B.: Deadline-Driven bandwidth allocation with flexible transmission rates in WDM networks. In: IEEE ICC 2008, pp. 5354-5358. IEEE Press, New York (2008)
8. Andrei, D., Tornatore, M., Batayneh, M., Martel, C.U., Mukherjee, B.: Provisioning of Deadlin-driven requests with flexible transmission rates in WDM mesh networks. IEEE/ACM Transactions on Networking. 353–366 (2010)
9. Fawaz, W., Chen, K.: A Shortest Setup Time First optical connection setup management approach with quantifiable success rate. In: IEEE Globecom 2006, pp. 1-5. IEEE Press, New York (2006)
10. Fawaz, W., Chen, K., Abou-Rjeily, C.: A novel connection setup management approach for optical WDM networks. IEEE Communications Letters. 998–1000 (2007)
11. Fawaz, W., Ouaiss, I., Chen, K., Perros, H.: Deadline-based Connection Setup in Wavelength-routed WDM Networks. Elsevier Journal of Computer Networks. 1972–1804 (2010)
12. Cavdar, C., Tornatore, M., Buzluca, F., Mukherjee, B.: Dynamic Scheduling of Survivable Connections with Delay Tolerance in WDM Networks. In: IEEE Infocom Workshops, pp. 1-6. IEEE Press, New York (2009)
13. Cavdar, C., Tornatore, M., Buzluca, F., Mukherjee, B.: Shared-Path Protection with Delay Tolerance in Optical WDM Mesh Networks. IEEE Journal of Lightwave Technology. 2068–2076 (2010)
14. Miyao, Y., Saito, H.: Optimal Design and Evaluation of Survivable WDM Transport Networks. IEEE Journal on Selected Areas in Communications. 1190–1198 (2004)
15. Fumagalli, A., Cerutti, I., Tacca, M., Masetti, F., Jagannathan, R., Alagar, S.: Survivable Networks Based on Optimal Routing and WDM Self-healing Rings. In: IEEE Infocom, pp. 726-733. IEEE Press, New York (1999)