Preemption-Enabled Setup of Optical Connections Coupled with Event-Driven Rerouting

Wissam Fawaz, Zahi Nakad, and Maroun Aad

Abstract—**In an era of uncertain business tendencies, realizing profitable optical networks is emerging as the most urgent task for optical operators. This task is aggravated by the need to support quality of service (QoS) differentiation, especially with the continuous development of new applications, each having different QoS requirements. This letter proposes to tackle these challenges through the definition of a novel QoS-aware connection setup strategy. The proposed setup approach enables network operators to ensure both efficient utilization of network resources and connection blocking probability differentiation. Resource efficiency is achieved via event-driven rerouting, whereas differentiated blocking probability is realized by preemption. Specifically, high priority connections are equipped with the capability of preempting lower priority connections, which through event-driven rerouting experience smaller blocking probability. Performance of this novel connection setup strategy is analyzed in this letter by simulating its operation in the context of the National Science Foundation Network (NSFNET). The obtained simulation results show that the proposed strategy can minimize the overall blocking probability while achieving QoS differentiation.**

Index Terms—**Optical networks, connection setup management, quality of service.**

I. Introduction

OPTICAL networks are in the process of evolving from single service networks to multi-service ones, in which various kinds of services with different quality of service (QoS) requirements are to be supported. Under such circumstances, it becomes of extreme importance for optical operators to improve the conditions related to the provisioning of the ever-emerging different types of applications. According to [1], one of the parameters that greatly affects the QoS perceived by optical clients during the connection establishment phase is *the blocking probability* parameter. Blocking probability indicates the fraction of connection requests that are denied access to the optical network due to lack of resources. This parameter is likely to become a potential service differentiator and it can be used to achieve QoS differentiation during the course of connection setup. Therefore, numerous research efforts ([2], [3]) have envisaged achieving blocking probability differentiation by reducing the blocking probability experienced by the high priority connections. However, privileging high priority connections without regard to low priority connections does not constitute an optimal solution as it degrades the QoS perceived by the low priority connections. This is especially true since lower priority connections end up suffering an unfair blocking probability increase. As such, the design of a new connection setup strategy that smoothes the impact of high priority connections on low priority ones is desirable. Inspired by this observation, this letter proposes a novel connection setup management strategy whose aim is twofold: first, to serve the largest number of high priority connections; and second, to prevent low priority connections from being unnecessarily disrupted by the high priority ones.

The proposed setup strategy can be viewed as a combination between a preemption policy and an event-driven rerouting strategy. The preemption policy seeks to increase the number of high priority connections that are granted access to the optical network by allowing incoming high priority connections to preempt the already established lower priority ones. In addition, the event-driven rerouting strategy strives to protect the low priority connections against the greediness of the high priority ones by reducing the number of preemptions taking place in the network. The preemption policy is what primarily distinguishes the proposed setup management strategy from existing solutions, since preemptions, as explained later, are restricted to a subset of the low priority connections that exist in the network.

II. The Proposed Connection Setup Strategy

As mentioned earlier, optical operators must deal not only with the challenge of maximizing the network throughput but also with the need to minimize the overall access blocking probability. The setup strategy under study addresses these issues by making sure that the following requirements are met:

- The maximum number of high priority clients is accepted into the network;
- The minimum number of low priority clients is influenced by the accepted high priority connections.

The connection setup strategy discussed in this letter proposes to cater for each of the requirements above as follows. A novel preemption strategy is designed to meet the first requirement and an event-driven rerouting policy is defined so as to match the second requirement.

A. Improved Hard Preemption Strategy

In what follows, a preemption strategy is proposed to serve the purpose of maximizing the number of high priority connections accommodated by the network. This is achieved by enabling the blocked high priority connections to preempt existing low priority ones according to a well-defined strategy. It is important to point out in this respect that multiple preemption strategies were studied in [4], yet the one investigated in this letter presents a better performance as it incurs less

Manuscript received October 27, 2010. The associate editor coordinating the review of this letter and approving it for publication was C. Assi.

The authors are with the Department of Electrical and Computer Engineering, Lebanese American University, Byblos, Lebanon (e-mail: {Wissam.Fawaz, Zahi.Nakad, Maroun.Aad}@lau.edu.lb)

Digital Object Identifier 10.1109/LCOMM.2011.122010.102058

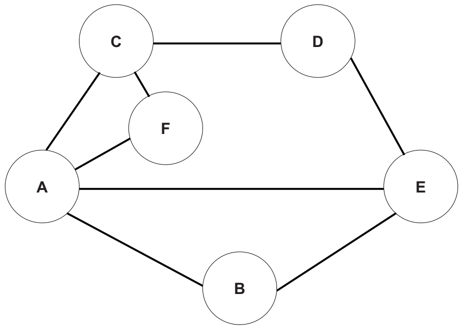

Fig. 1. Sample scenario: rerouting example.

preemptions. Since the proposed preemption strategy is an improved version of the hard preemption scheme found in [4], it will be referred to as the Improved Hard Preemption (IHP) strategy. IHP is described hereafter.

Let us illustrate the main idea behind IHP in the context of the sample network topology depicted in Fig. 1. Each connection request is assumed to be requiring a wavelength of bandwidth. It is also supposed that the capacity of each fiber link is limited to only one wavelength. Furthermore, 2 low priority connections L_1 and L_2 are assumed to be established from A to C and from A to D respectively. L_1 is routed along the single fiber path $A - C$ while L_2 is routed along the 3-fiber long path $A - F - C - D$. In addition to the low priority connections, 2 high priority connections H_1 and H_2 are considered to be in place between A and E along the paths $A - E$ and $A - B - E$ respectively.

IHP comes into play when the establishment of a high priority connection turns out to be impossible. For instance, let us assume that a setup request arrives at A and that this request corresponds to a high priority connection t_H addressed to D. It is clear that A cannot establish a lightpath to carry t_H due to resource shortage. Consequently, IHP is activated requiring A to scan through a list L that contains all of the low priority connections emanating from A sorted in a descending order of their hop count. Note that in this case L consists of the 2 connections L_1 and L_2. One by one the low priority connections stored in L are preempted and preemption continues until either A succeeds in serving t_H or A reaches the end of L in which case t_H is blocked. In the context of the considered example, this means that both L_1 and L_2 are preempted to make room for the insertion of t_H. An attempt is made then to re-provision the preempted low priority connections. That is, once t_H has been successfully routed along the path $A - C - D$, A tries to re-establish L_1 and L_2. Obviously, only L_1 can be re-provisioned along the path $A - F - C$ in the case of the example under study.

By trying to re-establish the preempted low priority connections, IHP enables more low priority connections to be admitted to the optical network without penalizing the high priority ones. Moreover, IHP has the merit of being applied solely to the low priority connections originating from the same optical node as the incoming high priority connection. This makes IHP a preemption strategy with a reduced complexity level as compared to the preemption strategies traditionally studied in the open literature.

B. Event-driven Rerouting

In order to avoid the starvation of low priority clients, and more specifically to minimize the number of preempted low priority connections, there is a need to supplement the preemption strategies with a rerouting strategy. The main purpose of the rerouting strategy is to smoothen the effect that preemption has on lower priority clients. In fact, a lower priority connection t_L is blocked when either of the following events occurs:

- There are not enough resources to establish t_L;
- A high priority connection took over the resources reserved for t_L and as a result t_L got preempted.

Hence, it is clear that in order to reduce the blocking probability of a low priority connection, resource consumption ought to be periodically minimized. In this manner, more resources are freed up and thus may be allocated to the incoming low and high priority connections. This higher availability of resources translates into a smaller number of preemptions and therefore improves the blocking probability of low priority clients. This letter optimizes resource consumption through a well-tailored event-driven rerouting policy. The authors in [4] considered a rerouting strategy that is driven by connection departures. As a distinguishing feature, the strategy proposed in this letter is driven by both connection departures and connection arrivals. The following example captures the essence of the rerouting strategy. Let us revisit the topology given in Fig. 1. Suppose that 3 connections t_1, t_2, and t_3 are established from A to E. t_1 is routed along the path $A - E$, t_2 along $A - B - E$, and t_3 along $A - C - D - E$. According to the considered strategy, rerouting is performed upon the occurrence of an arrival/departure event. Let us consider the case of a departure and assume that the connection t_1 traversing the fiber link $A - E$ terminates and frees up the fiber link between A and E. The proposed rerouting strategy establishes that on the occurrence of this event, the connections routed along paths with lengths greater than 2 fibers (t_2 and t_3 in this case) are rerouted according to an ascending order of their path lengths. So, t_2 will be rerouted along the freed up path $A - E$, then t_3 is rerouted to occupy the liberated path $A - B - E$. The deployment of the proposed rerouting strategy has thus the advantage of liberating 3 fiber links, namely $A - C$, $C - D$, and $D - E$ as opposed to the 1 fiber link that would have been freed up, had the rerouting strategy not been used. The freed up fiber links can be used by any subsequent connections, and the throughput would thus be improved.

In summary, the rerouting strategy discussed in this letter is triggered on the occurrence of two types of events, namely the departure and the arrival of connections. When a connection originating at a source node A departs from or arrives at the network, the rerouting strategy proceeds as follows. All of the connections emanating from A and whose lengths are greater than 2 are rerouted in an order consistent with their path lengths. That is, 2-fiber long connections are rerouted first to occupy the released resources, followed by the (3)-fiber long connections. This process continues until no further reroutings are possible and as such a better utilization of network resources is realized.

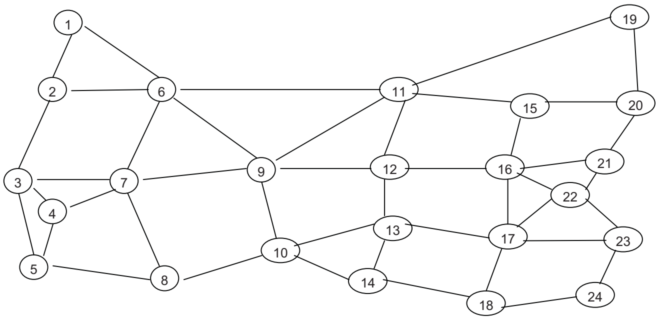

Fig. 2. Network topology used in simulation.

III. Simulation Study and Numerical Results

A Java-based discrete event simulator was built to study the performance of the proposed setup strategy in the context of the National Science Foundation Network (NSFNET) shown in Fig. 2, which has 24 nodes and 43 bidirectional fiber links. It is assumed that each node has a full wavelength conversion capability, and that each fiber link can support up to 8 wavelengths in each direction. Connection requests are arranged into 2 priority levels referred to as gold and silver and they arrive at and depart from the network in a dynamic way. The overall arrival process is Poisson and the connection holding time is exponentially distributed. Each connection request requires a wavelength. Dijkstra's shortest path algorithm is used to route the arriving connections, while wavelengths are assigned to the routed connections according to a first-fit strategy. 100000 connection requests are simulated per run of the simulator. The performance measure that will be used to evaluate the benefit of the proposed setup strategy is the blocking probability. Each obtained value is the result of multiple runs to make sure that a 95% confidence interval is achieved.

Different simulation scenarios were tested. Two scenarios investigated the impact of both IHP combined with rerouting and IHP without rerouting on the blocking probability of gold and silver connections. The simulation results pertaining to these scenarios are shown in the topmost figure of the composite Fig. 3. The reported results confirm that the coupling of the IHP preemption policy with the proposed rerouting strategy can greatly decrease the number of blocked gold and silver connections. One argument for such an observation is that the rerouting strategy helps to free up more resources for use by the incoming gold and silver connections and to reduce accordingly the number of preemptions taking place in the network. This underlines the fact that the rerouting strategy discussed in this letter can greatly benefit silver connections while boosting the quality of service perceived by the gold connections. To further evaluate the gain resulting from the proposed IHP with rerouting setup strategy, three additional scenarios featuring the proposed connection setup strategy, a preemption free rerouting free setup strategy, and a preemption free rerouting-based setup strategy were considered. The results corresponding to the aforementioned scenarios are reported in the bottom figure of the composite Fig. 3. It is clear from these results that the blocking probability of gold connections in the case of the IHP with rerouting setup strategy is much lower than those achieved by the 2 other traditional schemes. In addition, the results demonstrate that while the traditional schemes don't support blocking probability differentiation between gold and silver connections, the IHP with rerouting strategy has the luxury of enforcing such a differentiation. It follows from these observations that the proposed setup strategy allows optical operators to accept a large number of gold connections due to preemption while maintaining an acceptable blocking probability level for the silver ones.

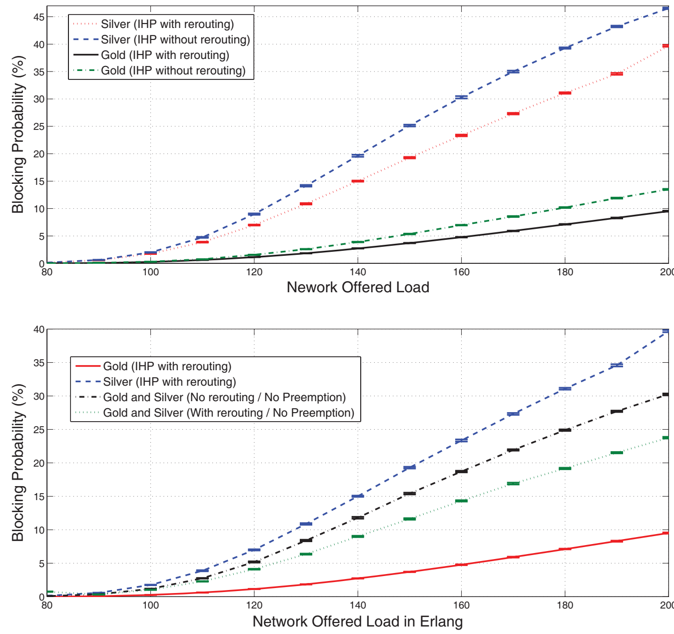

Fig. 3. IHP with/without rerouting setup strategies (uppermost); IHP with rerouting strategy vs. traditional setup schemes (bottom).

IV. Conclusion

This letter proposes to combine preemption-based connection setup with event-driven rerouting giving birth to an effective connection setup management strategy. The performance of the proposed setup strategy was analyzed in the context of a simulation model with a view to obtaining an estimate of its impact on the blocking probability of the different classes of connections. The obtained simulation results showed that the proposed strategy is capable of enforcing blocking probability differentiation while optimizing resource allocation.

References

[1] A. Szymanski, A. Lason, J. Rzasa, and A. Jajszczyk, "Grade-of-service-based routing in optical networks," *IEEE Commun. Mag.*, vol. 45, pp. 82–87, Feb. 2007.

[2] Y. Chen, M. Hamdi, and D. H. K. Tsang, "Proportional QoS over WDM networks: blocking probability," in *Proc. 6th IEEE Symp. Comp. and Commun.*, pp. 210–15, July 2001.

[3] T. Tachibana and S. Kasahara, "QoS-guaranteed wavelength allocation for WDM networks with limited-range wavelength conversion," *IEEE Trans. Commun.*, vol. 6, pp. 1439–50, June 2004.

[4] W. Fawaz, K. Chen, Z. Nakad, and C. Abou-Rjeily, "A simple QoS connection setup management approach for optical networks," in *Proc. ICC 2008*, May 2008.