2013 International Conference on Computing, Networking and Communications (ICNC) Workshop on Computing, Networking and Communications

Improving Service Differentiation of Immediate and Advance Reservation in Resource-Partitioned Optical WDM Networks

Derek R. Rousseau*, Joan Triay†*, and Vinod M. Vokkarane*
*Department of Computer and Information Science, University of Massachusetts, Dartmouth, MA, USA
†Department of Telematics Engineering, Universitat Politècnica de Catalunya (UPC), Castelldefels, Spain
u_drousseau@umassd.edu, {joan.triay,vvokkarane}@ieee.org

***Abstract*—As a result of the rise of Cloud and Grid computing, network operators are requested to improve their resource provisioning systems. In such scenarios, advance reservation (AR) and immediate reservation (IR) are gaining importance. While the former is capable of delivering low connection blocking for delay-tolerant applications, the latter is used by delay-sensitive ones. Even when both IR and AR requests share the same network substrate, service differentiation is needed. Additionally, under network resources partitioning, not only can the overall network capacity be under-utilized, but service differentiation can also be degraded. In this paper, we assess hybrid IR/AR sharing and partitioning mechanisms in optical wavelength-division multiplexing (WDM) networks, and propose efficient schemes for providing relative quality of service (QoS) differentiation among traffic classes. In order to overcome the poor resource utilization of strict partitioning, we develop and evaluate a preemption-based flexible partitioning framework aiming to provide relative QoS in hybrid IR/AR environments still supporting resource partitioning. Two preemption policies are proposed: switch first then preempt (SFTP) and preempt first then switch (PFTS). Through extensive simulations, we show that the proposed flexible partitioning framework can improve resource utilization, lower the overall blocking, and achieve well-differentiated relative QoS compared to strict partitioning.**

***Index Terms*—immediate reservation; advance reservation; optical WDM network; resource partitioning.**

I. INTRODUCTION

In large-scale experimentation scenarios, applications such as Grid computing and off-site backups require guaranteed resources, but the exact start time does not need to be assured provided a certain deadline is met, i.e., they are delay-tolerant to provisioning. For these applications, an advance reservation, wherein the request is sent much before the time the actual resources need to be allocated, can be used (see Fig. 1(b)). AR gives the network flexibility on reserving the resources, and as a result, increases the probability that the request can be provisioned. On the contrary, certain delay-sensitive applications request the resources to be reserved and allocated for immediate use. These applications are handled by immediate reservation mechanisms (see Fig. 1(a)). Both AR and IR are needed to ensure such diversity of applications are efficiently provisioned and service differentiation is provided.

Mixing IR and AR traffic can become an issue in optical WDM networks for Grids. For instance, if resources are reserved in advance by an AR request, fewer resources are available later for IR requests thus increasing connection blocking of the latter. To cope with such a problem, some methods include IR/AR admission control models [1] using

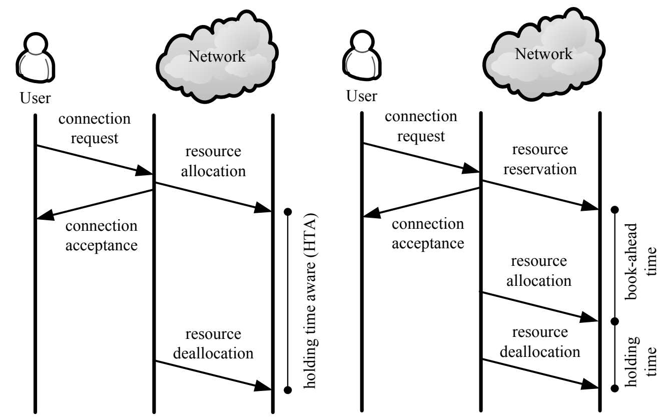

(a) Immediate reservation. (b) Advance reservation.

Fig. 1. Reservation types.

application-aware look-ahead mechanisms to predict IR demands. Also, Ahmad et al. [2] proposed to reduce the blocking of IR requests by using a dynamic look-ahead time controller. Another method of handling resource sharing is by rerouting requests that have been preempted [3]. All of these works consider only a single IR and a single AR class sharing resources. We argue that this configuration cannot provide full service differentiation with diverse IR/AR requests. This motivated our proposal of an IR/AR relative QoS framework [4].

It can also happen that network operators are forced to manage partitioned network domains [5]. This can occur while upgrading the capacity of the network [6], supporting different transport technologies (i.e., packet or circuit-based) [7] and because of network resiliency/backup provisioning [8]. Partitioning of network resources can even arise when considering the management of multiple virtual networks in parallel [9]. Despite such scenarios, network operators would still be required to provide some level of service differentiation among IR/AR classes. Furthermore, depending on the input traffic, partitioning does limit and inefficiently utilize the available network resources. This is due to the fact that a request blocked in its own partition may have been satisfied if it had the available resources from the other partition. Previously, some works have dealt with the allocation of certain requests to specific network resources for resource optimization [10] and resiliency [11], but not for QoS and service differentiation. To the best of our knowledge, the only work that has considered partitioning issues under hybrid IR/AR is [12]. However,

978-1-4673-0009-4/13/$26.00 ©2013 IEEE

TABLE I
IR/AR SERVICE MAPPING AND APPLICATIONS.

Reservation class	Request type	Algorithm time complexity	Service delay	Service deadline	Blocking probability	Typical applications
IR0	IR HTA	$O(\|W\|\|V\|\tau)$	Delay-sensitive	Strict	Unspecified	Best-effort traffic, file-sharing, FTP
IR1	IR LPS	$O(\|W\|k\|V\|\tau)$	Delay-sensitive	Strict	Medium	Grid applications, IPTV
AR0	AR STSD-fixed	$O(\|W\|\|V\|\tau)$	Delay-tolerant	Strict	Medium	Delay-tolerant Grid co-scheduling
AR1	AR STSD-flexible	$O(\|W\|\|V\|(\omega-\alpha))$	Delay-tolerant	Flexible	Low	Data-storage synchronization

the revenue-based scheme proposed by the authors aimed at improving the IR acceptance rate without considering service differentiation capabilities.

Motivated by the aforementioned scenario, in this paper we evaluate full-sharing, strict-partitioning, and flexible-partitioning frameworks to provide relative QoS differentiation for hybrid IR/AR in optical WDM networks. Firstly, we show the benefits of fully sharing the network resources to provide service differentiation and decrease the average total blocking probability. Then, we show how partitioning (strict with and without shared overlap) changes the IR/AR class differentiation and increases the total blocking. To overcome the higher blocking probability due to partitioning resources strictly, we evaluate two preemption-based policies that allow requests from partitions other than those originally assigned to reserve resources from other partitions present in the network. Finally, we propose a partitioning assignment to meet IR/AR service differentiation while using the flexible-partitioning framework. Results show that we can improve resource utilization, lower the overall blocking, and achieve well-differentiated relative QoS by using the proposed flexible-partitioning framework.

The paper is structured as follows: Section II introduces the IR/AR relative QoS framework. Section III evaluates the blocking performance and service differentiation of strict partitioning with and without shared overlap. Section IV presents the flexible partitioning framework while Section V evaluates its simulation results. Finally, Section VI concludes this paper.

II. SERVICE-AWARE HYBRID IR/AR QOS FRAMEWORK

In [4] we introduced a hybrid IR/AR QoS framework making use of full-sharing of network resources. The goal of the framework was to guarantee a relative service differentiation among diverse delay-sensitive and delay-tolerant applications.

We considered the case of a wavelength-routed optical network represented by $G = (V, E, W)$, where V, E, and W denote the set of nodes, links, and wavelengths, respectively. Moreover, time is divided up into time-slots. Applications request to set up a circuit between the source and the destination nodes for a specified duration or number of time-slots τ. There are no wavelength converters on the network; therefore, the wavelength-continuity constraint applies.

The framework is part of a centralized network resource provisioning system. This model is extensively used, especially in hybrid immediate and advance reservation capable optical Grid networks. Production networks like ESnet [13], and others devised in recent projects make use of this approach [14]. This architecture has also less implications in terms of management and control functions. However, it requires to set up extra capabilities for resiliency and/or protection.

In the framework, requests are handled by the network service layer which formats them into the proper setup. The setup is then forwarded to the service-aware adaptation layer. This layer maps the request onto an existing reservation class based on its requirements. Four reservation classes are specified: 2 IR classes (IR0 and IR1) and 2 AR classes (AR0 and AR1). Each one of these classes utilizes a specific scheduling algorithm that best matches the requirements of the application request and provides blocking differentiation. Table I summarizes the classes, scheduling complexity, and typical applications.

The low priority IR (IR0) is for regular applications that need to reserve resources immediately and do not specify any strict QoS requirements. On the contrary, IR1 is for immediate higher priority traffic. IR1 uses the *lightpath switching* (LPS) [15] scheduler to help achieve lower blocking probability by allocating slots on multiple wavelengths (channels) and routes. The number of candidate routes chosen is specified by k, from k-shortest paths. With LPS, the biggest difference is that virtual circuits can switch to another lightpath if there are no more resources available on the existing one. This additional flexibility allows higher priority IR requests to achieve a better blocking probability, especially when used in conjunction with AR requests.

The low priority AR (AR0) can be used for applications that can schedule in advance but with strict deadline QoS requirements. AR0 uses specified start and specified deadline (STSD)-fixed scheduling [16]. STSD-fixed is similar to basic IR; the only difference is that the former introduces a book-ahead time. This book-ahead time, denoted by α, increases the chances of allocating resources since they are requested far in advance of the connection request arrival. STSD-fixed has a strict delay-tolerant provisioning requirement. Finally, the high priority AR (AR1) is used with delay-tolerant applications that have a low blocking QoS requirement. AR1 uses STSD-flexible [17] which introduces the sliding scheduling window, ω, giving greater flexibility to this type of request, i.e., the request can be allocated anywhere within the specified allocation window.

III. HYBRID IR/AR SERVICES IN FULL-SHARING AND PARTITION-BASED FRAMEWORKS

As we introduced, there are a number of situations wherein network operators have their available network resources partitioned. Without loss of generality, we will consider the case of an optical WDM network where the partitioned resources

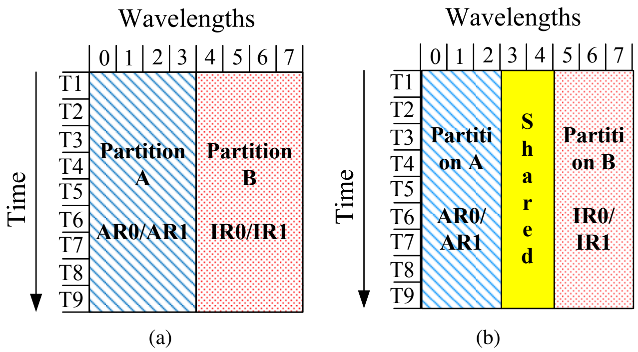

Fig. 2. (a) Example of strict partitioning (AR vs. IR), and (b) example of resources with an overlapping strict partitioning (AR vs. IR).

are wavelengths. That is, wavelengths are split into different partitions; for instance, wavelengths zero to three belong to partition A, and wavelengths four through seven to partition B. In strict partitioning, wavelength partitions are disjoint. Formally, let A and B be the resource sets, then, $A \cap B = \emptyset$.

We can also have the case where some of the available resources on the network can be used by any kind of service request. That is, some resources are only available to requests assigned to group A, others to group B, and some resources can be used by requests both of groups A and B. Again, formally, let A and B be the two sets of resources, with shared overlap, $A \cap B \neq \emptyset$, that is, there exist resources x so that $x \in A$ and $x \in B$. Additionally, there exist resources, y and z, which are not shared, hence $B \setminus A = \{y \in B | y \notin A\}$ and $A \setminus B = \{z \in A | z \notin B\}$ ($\setminus$ is the relative complement).

Upon partitioning, classes can be assigned. An example of strict partitioning is shown in Fig. 2(a). In this case, AR requests (AR0 and AR1) are assigned to partition A, whereas immediate reservation requests can only use resources of partition B, i.e., wavelengths λ_4 through λ_7. Fig. 2(b) shows an example of partitioning with shared overlap. In comparison with strict partitioning, here λ_3 and λ_4 can be used by any IR/AR request, i.e., these wavelengths overlap both partitions.

A. Full-Sharing and Strict-Partitioning Evaluation

We simulated the two strict partitioning strategies on the 14-node NSF network topology with 8 wavelengths per link. In all cases, results were averaged over 30 simulation runs for 10^6 connection requests each. Four IR/AR request types (IR0, IR1, AR0, and AR1) are run on the same network with both partitioned and non-partitioned frameworks. For both AR classes, the book-ahead is set to 1440 slots and a flexibility of $f_{AR1} = \lfloor \frac{\omega - \alpha}{\tau} - 1 \rfloor = 1$ is used for the higher priority AR (AR1). $f = 1$ means the scheduling window is double the holding time of the connection, denoted by τ and that follows an exponential distribution with a mean value of 180 slots. Parameters ω and α denote the latest deadline to allocate the connection and its book-ahead time, respectively. Requests arrive following a Poisson process. For IR1, k, the number of alternate paths of the LPS algorithm, is set to 3. In all scenarios we assume a traffic distribution of 25% assigned to each of the four request types. We have chosen this traffic configuration for being a representative case where all request classes contribute equally to the network load. We tested configurations with other traffic class percentages and the results led to similar conclusions.

We first analyze the results for the full-sharing case. Fig. 3(a) shows the blocking probability as a function of the offered load. As presented in [4], a clear IR/AR service differentiation remains along the whole load range between both delay-sensitive and delay-tolerant applications' requests. Specifically, the higher priority AR (AR1) has the lowest blocking probability; on the contrary, the lower priority IR (IR0) has the highest. At low loads, the blocking differentiation is not significant due to the greater availability of resources. One may argue that the blocking probability for IR requests is always higher. However, the aim in this scenario was to prioritize requests that were delay-tolerant still allowing higher priority IR to differentiate themselves from lower priority IR0. It is worth noting that since the model uses a centralized resource broker, the delay on signaling the optical cross-connects has no or little effect on the blocking probability as long as the time-slot is sufficiently large.

With partitioning we consider the case where wavelengths are split into two equal groups. We assign the classes defined previously in Table I to the newly defined partitions as shown in Fig. 2(a). Utilizing this grouping, sets are completely disjoint; hence, the IR and AR requests only compete for resources within their own type.

Fig. 3(b) shows the results for this strict partitioning scenario. The higher priority classes from IR and AR achieve nearly identical blocking probabilities, especially at high loads, and decreased blocking probability than their lower priority counterparts due to their scheduling improvements (LPS and STSD-flexible, respectively). At low loads, IR1 improves over AR1 due to the capabilities of the LPS scheduling algorithm used by the former. Both lower priority classes from AR and IR also achieve nearly identical blocking. In both cases, IR and AR only compete within their own partitions and only against their own types (i.e., IR0 vs. IR1 and AR0 vs. AR1). As seen, the precise service differentiation shown in the full-sharing scenario is no longer present in the strict partitioning case. Moreover, as shown in Fig. 4, the average total blocking probabilities (computed from all the blocking that occurred in the network irrespective of its original class) are much higher than in the full-sharing case. This is one of the main motivations to develop a flexible-partitioning framework.

We also examine a scenario where the partitions overlap. In this case, the AR classes have wavelengths zero through four and the IR classes use wavelengths three through seven (refer to example in Fig. 2(b)). Therefore, wavelengths three and four can be scheduled by any of the classes. This configuration balances the number of fixed partitioned wavelengths "lost" by each group with one extra "shared" wavelength. As we observe in Fig. 3(c), the lower priority AR (AR0) and the higher priority IR (IR1) have nearly identical blocking probabilities. The higher priority AR is able to decrease its blocking probability using the additional shared wavelengths. This would allow

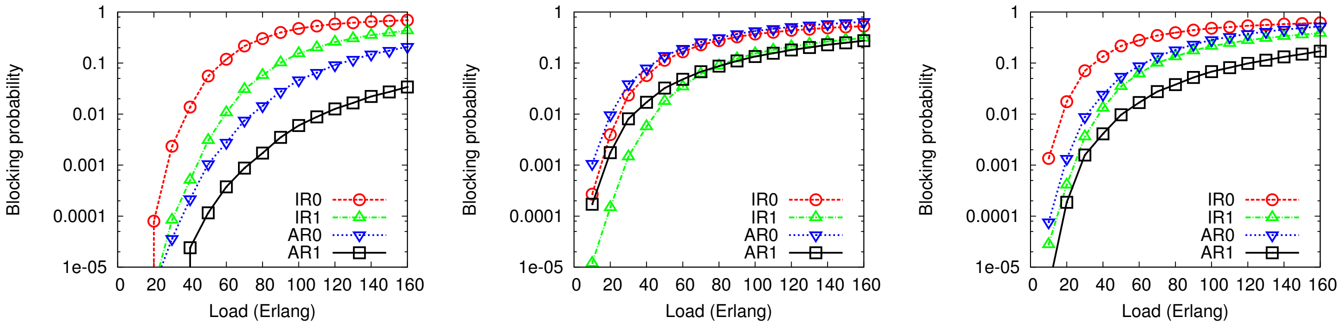

(a) (b) (c)

Fig. 3. Blocking probability versus offered load for different sharing/partitioning frameworks with traffic distribution IR0=25%:IR1=25%:AR0=25%:AR1=25%: (a) full sharing, (b) strict partitioning, and (c) strict partitioning with shared overlap.

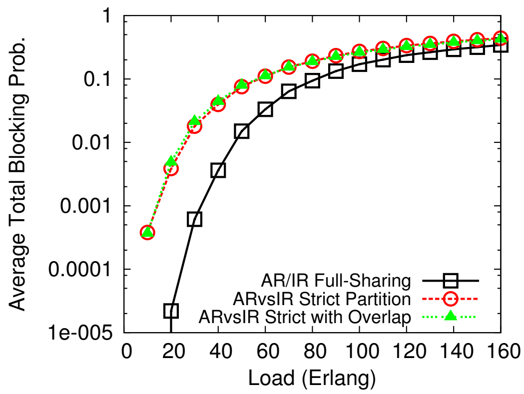

Fig. 4. Total average blocking probability with traffic distribution IR0=25%:IR1=25%:AR0=25%:AR1=25% for full sharing, strict and strict with overlap partitioning.

network operators to provide a QoS where the relative priority in terms of blocking performance is IR0<AR0=IR1<AR1. Nonetheless, the total average blocking is still higher than in the full-sharing case as shown in Fig. 4.

As illustrated, strict partitioning (with and without overlap) see an increased average blocking probability. Therefore, without additional mechanisms to improve the performance, these two schemes do not only change the class differentiation behavior, but also increase the total blocking on the network. To overcome this situation, we propose a flexible partitioning scheme in the following section.

IV. Flexible-Partitioning for Improving Service Differentiation of Hybrid IR/AR

Since strict partitioning is unable to utilize all the network resources, we propose a preemption-based flexible partitioning framework that allows the network operator to handle different network resource partitions. We define two possible policies: *preempt first then switch (PFTS)* and *switch first then preempt (SFTP)*. The first policy promotes the preemption of displaced requests (i.e., those using resources other than the primary partition), whereas SFTP stimulates the sharing of resources.

Network resources will be partitioned by wavelengths in the same manner introduced in Section III. Based on the policies defined below, the request may switch to the other partition in an attempt to allocate resources if it was unable to find in its primary partition. If it is able to allocate resources in the other partition, it will be set a preemption *flag* to this request, allowing it to be preempted by requests attempting to schedule within their primary partition. As an example, if both AR classes are mapped to partition A and both IR classes are mapped to partition B, when referencing an AR class, partition A will be referred to as the primary and partition B will be referred to as secondary. Note that, since IR1 (LPS) is split into segments possibly over different lightpaths, both policies require that all segments be reserved using resources from a single partition.

1) Preempt First Then Switch: In this policy (see Algorithm 1), we first try to schedule within the primary partition. If resources are not found to be allocated, the algorithm checks the primary partition again this time including resources taken by decreased priority requests (i.e., with preemption flag set). If resources are found, the request is scheduled by preempting the requests with decreased priorities that it conflicts with. If no resources are found either, it will then switch to the secondary partition and attempt to schedule there. If resources are found in the secondary partition then the request is flagged and the resources are reserved. If no resources are found at this point, the request is blocked.

Algorithm 1 Preempt First Then Switch (PFTS)

```
Check primary partition for available lightpath.
if Lightpath found then
    Reserve wavelength on lightpath.
else
    Check primary partition for available lightpath including re-
    sources taken by reservations with the preemption flag set.
    if Lightpath found then
        Reserve wavelength on lightpath.
    else
        Check secondary partition for available wavelength and
        lightpath.
        if Lightpath found then
            Set preemption flag.
            Reserve wavelength on lightpath.
        else
            Block request.
        end if
    end if
end if
```

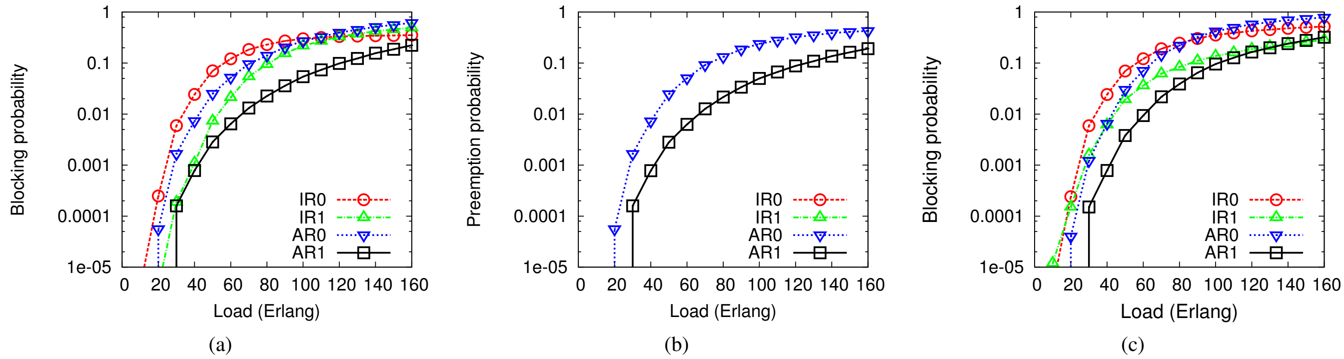

Fig. 5. Results in network partitioned with flexibility with AR vs. IR partitioning. The traffic distribution is IR0=25%:IR1=25%:AR0=25%:AR1=25%: (a) blocking probability versus offered load in PFTS, (b) preemption rate in PFTS, and (c) blocking probability versus offered load in SFTP.

Formally, let A be the primary partition for an incoming request r. We define A^0 as the free resources in partition A. Similarly, A^+ and A^- are the resources taken (already reserved) by requests assigned to this partition and requests that switched into it with preemptable flag, respectively. As such, $A^0 = A \setminus (A^- \cup A^+)$. The same relation does exist for partition B. The order in which the resources are inspected in PFTS is: i) A^0, ii) $A^0 \cup A^-$, and iii) B^0.

2) Switch First Then Preempt: This policy (see Algorithm 2) tries to promote the sharing of resources as opposed to PFTS. In the algorithm, the request to be scheduled first checks the primary partition. If no resources are available, then the policy switches to the secondary partition and checks for free resources. If resources in the secondary partition are found, the request is set a preemption flag (i.e., it can get preempted) and the resources are reserved. If no resources are found in this secondary partition, the scheduler will then check the primary partition again, but this time, any reservations that have the preemption flag set will be included and be preempted if the request can be successfully allocated. If no resources are found at this point, the request is blocked.

Algorithm 2 Switch First Then Preempt (SFTP)

Check primary partition for available lightpath.
if Lightpath found **then**
  Reserve wavelength on lightpath.
else
  Check secondary partition for available lightpath.
  if Lightpath found **then**
    Set preemption flag.
    Reserve wavelength on lightpath.
  else
    Check primary partition for available lightpath including resources taken by reservations with the preemption flag.
    if Lightpath found **then**
      Reserve wavelength on lightpath.
    else
      Block request.
    end if
  end if
end if

Following the same notation introduced for PFTS, in SFTP the ordered list of resources inspected is now: i) A^0, ii) B^0, and iii) $A^0 \cup A^-$.

The flexible-partitioning requires keeping track of the reserved resources and whether these are taken or not by preemptable requests. This would slightly increase the complexity of the scheduling system in addition to the extra signaling to deallocate and reconfigure the switching devices, but the latter is implicit to any preemption-based system.

V. Simulation Results of Flexible and Preemption-Based Partitioning

As we have seen previously, strict partitioning caused the average total blocking probability to increase compared to the results achieved in Fig. 3(a). In this section, we introduce the results of the two flexible partitioning and preemption policies (PFTS and SFTP) when resources are split between AR and IR requests. We then show that the proposed resource partitioning policies not only lower the total blocking probability, but also achieve better differentiation among traffic classes.

Because of the preemptive mechanisms, the blocking probability results also include preempted requests. Graphs showing the preemption probability are also included. As showcased before, we assume a scenario wherein all IR/AR classes contribute with the same percentage to the total traffic load (25% each). Such traffic configuration helps simplify the analysis and focus more on the performance of the policies themselves instead on the traffic contributions. The network under testing (NSFNET) and number of wavelength per link (8) is the same, so are the simulation runs and number of requests in each run.

We start off by examining the PFTS policy. As shown in Fig. 5(a), when AR is in one partition and IR is in the other, advance reservation requests rarely preempt IR requests due to the large book-ahead time of AR. However, IRs preempt AR requests frequently, especially the AR0 requests which, due to its less flexible scheduling capabilities, tend to switch to the secondary partition, hence being prone to having its preemption flag set more often. As a result, the AR requests are the only ones that experience preemption as shown in Fig. 5(b) (IR preemption rate was almost non-existent, so it does not show in the graph). Also, we observe that mostly all AR blocking is due to preemption if we set side by side Figs. 5(a) and 5(b).

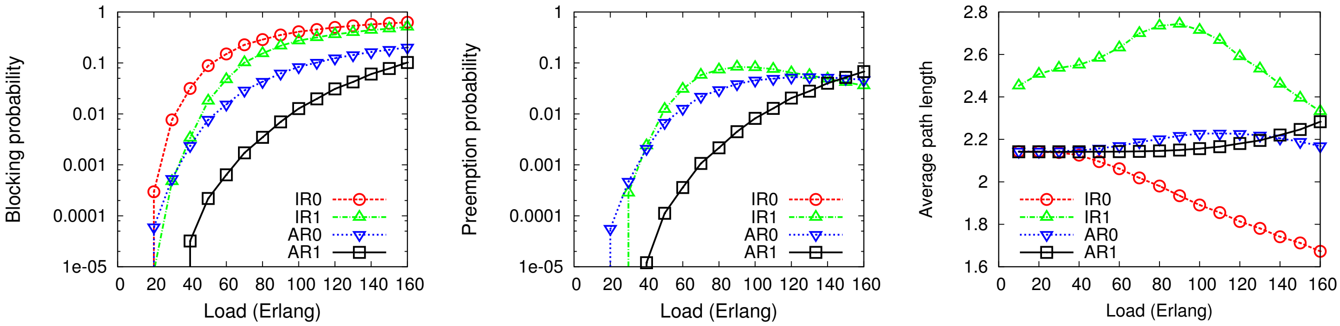

(a) (b) (c)

Fig. 6. Results on network partitioned with flexibility, with two partitions (AR1+IR0 vs. AR0+IR1) for PFTS policy. The traffic distribution is IR0=25%:IR1=25%:AR0=25%:AR1=25%: (a) blocking probability versus offered load, (b) preemption rate, and (c) average path length.

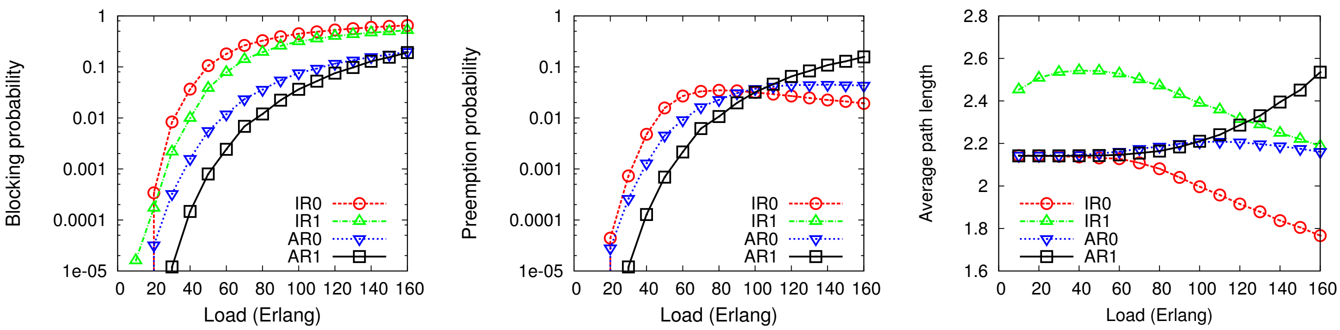

(a) (b) (c)

Fig. 7. Results on network partitioned with flexibility, with two partitions (AR1+IR0 vs. AR0+IR1) for SFTP policy. The traffic distribution is IR0=25%:IR1=25%:AR0=25%:AR1=25%: (a) blocking probability versus offered load, (b) preemption rate, and (c) average path length.

Fig. 5(c) shows the blocking probability with the same AR vs. IR partitioning but now for the SFTP policy. In this case, IR and AR blocking performance is differentiated within their own respective reservation type, i.e., IR0 vs. IR1 and AR0 vs. AR1. However, as in the previous case with PFTS, there is no clear differentiation among all the classes. Finally, in both cases, PFTS and SFTP, results showed that the average total blocking probability was lower than in the strict partitioning frameworks (the graph is not attached due to space restrictions, but it is similar to Fig. 8).

A. Partition Assignment for Improving Service Differentiation

In spite of decreasing the total average blocking probability with the new flexible partitioning framework, the results in Fig. 5 have shown that service differentiation among the four IR/AR classes is not clear enough. We have assessed different partitioning assignments of IR and AR classes and we found that the AR1+IR0 vs. AR0+IR1, i.e., AR1 and IR0 assigned to partition A, and AR0 and IR1 assigned to partition B, is able to deliver more service differentiation than the IR vs. AR assignment evaluated previously. Partition A is the primary for AR1 and IR0, and B is the secondary, and viceversa for AR0 and IR1 requests.

The plots in Fig. 6 illustrate the results with the new partitioning assignment for the PFTS policy. As seen in Fig. 6(a), the level of differentiation among the four IR/AR classes improves over the previous PFTS assignment of Fig. 5(a), especially from medium to high loads. An even better blocking differentiation is obtained when making use of the SFTP policy. Fig. 7(a) shows the blocking probability as a function of the offered load to the network for this case. Only at very low loads, the IR0 blocking tends to decrease abruptly due to the higher availability of wavelength resources. Nonetheless, the differentiation is clear for the remaining load range.

The differentiation performance between SFTP and PFTS for this particular partitioning assignment (AR1+IR0 vs. AR0+IR) can be explained by the preemption results of Figs. 6(b) and 7(b). We can see how in PFTS, the AR0 preemption probability is increased in comparison to SFTP. That is, with PFTS, AR0 requests are preempted more often by requests from the primary partition, including AR1 and in a greater extent IR0 requests. The inferior performance of the IR0 scheduling algorithm diminishes the capabilities of this type of immediate reservation to find free resources in its own partition, thus preempting more often requests that moved into it (AR0 and IR1). The behavior is even more stressed by the fact that IR0 requests compete for resources with the most flexible reservation type, AR1. Alternatively, with the SFTP policy, AR0 requests are not preempted so often (refer Fig. 7(b)). Instead, now IR0 requests are heavily preempted by other IR1 requests from the primary (B) partition when the former have moved into this partition. Additionally, AR0 requests tend to be preempted by both IR0 and AR1 requests when AR0 have switched to the secondary partition (A).

To sum up this part, we can conclude that SFTP promotes the “sharing” of resources among partitions, whereas PFTS

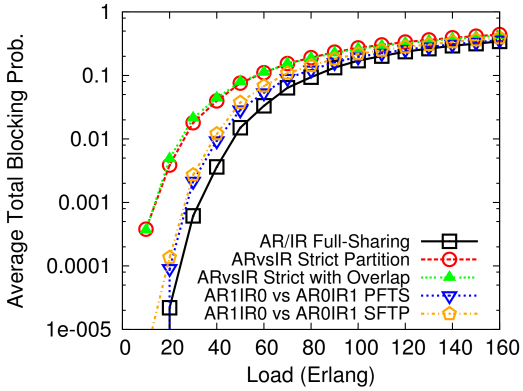

Fig. 8. Average total blocking probability with traffic distribution IR0=25%:IR1=25%:AR0=25%:AR1=25%.

tends to penalize the sharing of resources. This is why the results of SFTP resemble more those obtained with the full-sharing framework of Fig. 3(a).

The results about the average path length as a function of the offered load for both PFTS and SFTP policies are shown in Figs. 6(c) and 7(c), respectively. In both cases, the path length of IR1 requests is generally longer than the rest of classes. This is due to using the k-shortest path with lightpath switching scheduling algorithm ($k = 3$). The inferior scheduling algorithm for IR0 shortens the path length for successfully reserved requests when load on the network increases and free resources are running short.

With the proposed framework, we have not only improved the service differentiation while using partitioning, but also reduced the average total blocking probability. Fig. 8 shows the gain in blocking performance obtained by the flexible partitioning with both preemption policies. It does not match the performance of full sharing, but it improves more than an order over the strict frameworks at low loads.

VI. Conclusion

Service differentiation is a requirement in today's networks, even when network operators have to deal with different resource partitions. In turn, immediate and advance reservation are capable to satisfy delay-sensitive and low connection blocking needs, respectively. In this paper, we have shown that in hybrid IR/AR scenarios with full sharing of resources we were able to provide a well-defined service differentiation using different scheduling algorithms. However, switching to a strict partitioning or a strict partitioning with shared overlap has shown that not only is the service differentiation diminished, but the average total blocking probability on the network increased. To overcome such behavior, we have proposed two preemption-based flexible partitioning frameworks, preempt first then switch (PFTS) and switch first then preempt (SFTP). Results have demonstrated that we can decrease the total blocking probability on the network, thus improving the network resource utilization. Furthermore, by applying a specific IR/AR partition assignment, we have proved that the flexible policies do not only lower the blocking probability on the network, but also achieve better service differentiation of IR and AR requests and similar to the full-sharing scenario.

As future work, we plan to introduce probabilistic preemption to decrease AR preemption and provide absolute QoS.

Acknowledgment

Work partially supported by the Department of Energy (DOE) COMMON project under grant DE-SC0004909 and the NSF CARGONET project under grant CNS-1218973. Joan Triay was a visiting researcher at University of Massachusetts, Dartmouth, sponsored by a Fulbright fellowship.

References

[1] Y.-D. Lin, C.-H. Chang, and Y.-C. Hsu, "Bandwidth brokers of instantaneous and book-ahead requests for differentiated services networks," *IEICE Trans. on Communications*, vol. 85, no. 1, pp. 278–283, 2002.

[2] I. Ahmad, J. Kamruzzaman, and S. Aswathanarayaniah, "A dynamic approach to reduce preemption in book-ahead reservation in QoS-enabled networks," *Computer Communications*, vol. 29, no. 9, pp. 1443–1457, May 2006.

[3] I. Ahmad, J. Kamruzzaman, and D. Habibi, "Rerouting in advance for preempeted IR call in QoS-enabled networks," *Computer Communications*, vol. 31, no. 17, pp. 3922–3928, Nov. 2008.

[4] J. Triay, D. Rousseau, C. Cervelló-Pastor, and V. M. Vokkarane, "Dynamic service-aware reservation framework for multi-layer high-speed networks," in *Proc. of 20th IEEE ICCCN 2011*, Maui, HI, USA, Aug. 2011, pp. 1–6.

[5] "Energy Sciences Network," Department of Energy (DoE). [Available online: http://www.es.net].

[6] R. Roy, A. Nag, and B. Mukherjee, "Telecom mesh network upgrade to manage traffic growth," *IEEE/OSA J. Opt. Commun. Netw.*, vol. 2, no. 5, pp. 256–265, May 2010.

[7] C. Vadrevu, M. Liu, M. Tornatore, C. Guok, E. Chaniotakis, I. Monga, and B. Mukherjee, "Hybrid circuit/packet networks with dynamic capacity partitioning," in *2010 ITU-T Kaleidoscope: Beyond the Internet? - Innovations for Future Networks and Services*, Dec. 2010, pp. 1–7.

[8] C. S. Vadrevu and M. Tornatore, "Survivable IP topology design with re-use of backup wavelength capacity in optical backbone networks," *Optical Switching and Networking*, vol. 7, no. 4, pp. 196–205, 2010.

[9] S. Figuerola and M. Lemay, "Infrastructure services for optical networks [invited]," *IEEE/OSA J. Commun. Netw.*, vol. 1, no. 2, pp. A247–A257, 2009.

[10] M. M. Hasan and J. P. Jue, "Online clustering for hierarchical WDM networks," in *Proc. of 2008 OFC/NFOEC*, San Diego, CA, USA, Feb. 2008, pp. 1 –3.

[11] ——, "2-SRLG-connected partitioning in optical networks," in *2010 Int. Conf. on High Performance Switching and Routing (HPSR)*, Dallas, TX, USA, Jun. 2010, pp. 183 –188.

[12] D.-H. Yi and J. Kim, "A revenue-based bandwidth sharing for advance and immediate reservation calls," *IEICE Transactions on Communications*, vol. E91-B, no. 11, pp. 3617–3626, Nov. 2008.

[13] I. Monga, C. Guok, W. Johnston, and B. Tierney, "Hybrid networks: lessons learned and future challenges based on ESnet4 experience," *IEEE Commun. Mag.*, vol. 49, no. 5, pp. 114–121, may 2011.

[14] N. Charbonneau, V. M. Vokkarane, C. Guok, and I. Monga, "Advance reservation frameworks in hybrid IP-WDM networks," *IEEE Commun. Mag.*, vol. 49, no. 5, pp. 132–139, may 2011.

[15] N. Charbonneau and V. M. Vokkarane, "Dynamic circuits with lightpath switching over wavelength routed networks," in *4th IEEE Int. Conf. ANTS*, Mumbai, India, Dec. 2010.

[16] J. Zheng and H. Mouftah, "Routing and wavelength assignment for advance reservation in wavelength-routed WDM optical networks," in *IEEE Int. Conf. on Communications (ICC) 2002*, vol. 5, New York, NY, USA, Aug. 2002, pp. 2722–2726.

[17] E. S. Jung, Y. Li, S. Ranka, and S. Sahni, "Performance evaluation of routing and wavelength assignment algorithms for optical networks," in *IEEE Int. Conf. on Communications (ICC) 2008*, Beijing, China, Jul. 2008, pp. 62–67.