Efficient burst reservation protocol: a hybrid signaling protocol for efficient burst-level reservations and quality-of-service differentiation in optical burst switching networks

Kostas Christodoulopoulos, Kyriakos Vlachos, and Emmanouel Varvarigos

Research Academic Computer Technology Institute and Department of Computer Engineering and Informatics, University of Patras, GR-26500, Rio, Greece

kvlachos@ceid.upatras.gr

Leontios Stampoulidis and Efstratios Kehayas

Photonics Communications Research Laboratory, Department of Electrical and Computer Engineering, National Technical University of Athens, GR-15773, Athens, Greece

lstamp@mail.ntua.gr

In this paper we present the *efficient burst reservation protocol* (EBRP) suitable for bufferless optical burst switching (OBS) networks. The EBRP protocol is a two-way reservation scheme that employs *timed* and *in-advance* reservation of resources. In the EBRP protocol *timed* reservations are relaxed, introducing a *reservation time duration* parameter that is negotiated during call setup phase. This feature allows bursts to reserve resources beyond their actual size to increase their successful forwarding probability and can be used to provide quality-of-service (QoS) differentiation. The EBRP protocol is suitable for OBS networks and can guarantee a low blocking probability for bursts that can tolerate the round-trip delay associated with the two-way reservation. We present the main features of the proposed protocol and describe in detail the timing considerations regarding the call setup phase and the actual reservation process. Furthermore, we show evaluation results and compare the EBRP performance against two other typical reservation schemes, a tell-and-wait and a tell-and-go (just-enough-time) like protocol. EBRP has been developed for the control plane of the IST-LASAGNE project. © 2006 Optical Society of America

OCIS codes: 060.0060, 060.4510.

1. Introduction

Optical burst switching (OBS) has been introduced as a true broadband solution offering a higher degree of data transparency and providing sufficient bandwidth for the establishment of the fast Internet and other advanced, high-speed services [1]. The recent advances in network techniques, such as the introduction of generalized multiprotocol label switching (GMPLS), have favored the development of new OBS concepts such as labeled OBS (LOBS), securing higher network performance standards by enabling a subwavelength, burst-level granularity [2]. The requirement for optimizing network resources and simultaneously minimizing buffering in the core has imposed severe constraints on designing simple and efficient OBS networks. Suitable protocols have been proposed to cope with

key issues such as connection establishment and bandwidth reservation. In this context two main classes of signaling protocols have been distinguished, usually referred as tell and wait (TAW) and tell and go (TAG) with a number of variants, which exhibit complementary performance characteristics [3].

In the case of TAW, a two-way reservation of resources is performed and an end-to-end connection is fully established before the transmission of data. Recent research efforts, such as wavelength-routed OBS (WR-OBS) [4], have shown that such reservation schemes enable the implementation of a bufferless core network with limited wavelength-conversion capability by moving processing and buffering functions at the edge. However, the establishment of end-to-end lightpaths is a time-consuming process that adds considerable delays, whereas the end-to-end round trip affects the bandwidth utilization, since resources are reserved immediately upon the arrival of the setup message. To reduce burst holding time at the edge nodes, burst size prediction-estimation techniques [5] can be applied.

In the TAG approach, the pre-establishment of a full virtual path before transmission is not required. The setup message, the optical burst, and (optionally) the release message are sent sequentially; thus while the connection establishment is in progress, the burst transverses the previous core nodes. A number of one-way reservation schemes have been proposed, including the ready-to-go virtual circuit protocol [6], just enough time (JET) [7], Horizon [8], and just in time (JIT) [9, 10]. Although very promising, these signaling schemes result in high burst loss ratio under heavy network load and rely on wavelength conversion combined with sophisticated void-filling channel scheduling algorithms to resolve contentions [11, 12]. Although contention resolution in the wavelength domain is considered as a straightforward process, it poses specific hardware requirements, raising issues such as node scalability, size, and cost. Alternative contention-resolution techniques have been also proposed employing fiber delay line structures [13–15] or deflection routing [16]. However, the former is impractical because of the lack of scalable optical buffers, whereas deflection routing affects the network load and cannot guarantee packet arrival in the correct order.

Having identified the major advantages and weaknesses of the two complementary classes of protocols, hybrid signaling schemes have been especially designed to combine features from both classes. In Ref. [18], a scheme that employs TAW-type reservation to an intermediate node, followed by unacknowledged one-way reservation until the egress, is proposed. The proposed hybrid protocols controls the trade-off between burst loss and delay (by selecting the intermediate node) and thus enables QoS differentiation. In this paper, we propose the *efficient burst reservation protocol* (EBRP). EBRP is suitable for bufferless optical burst switching networks and exploits the advantages of both classes of protocols, relaxing the strict time requirements to achieve efficient burst-level reservations and contention-free operation. EBRP resembles TAW in the sense that a burst can enter the core network only after having reserved its route, but avoids the wasteful reservation of resources. EBRP resembles the TAG feature of delayed-timed reservations, but relaxes the strict time requirements of the one-way reservation schemes. Its operation relies on the *timed* and *in-advance* reservation mechanisms [19], and further provides the capability of negotiating the starting time and the duration of the reservation to achieve a finer granularity over the reserved periods as well as a higher burst acceptance probability. Hence the algorithm schedules the requests efficiently in the time domain, detecting the earliest possible time instance that is available for reservation. The capability to negotiate the reservation horizon enables QoS differentiation with higher priority flows to request a larger reservation period during resource negotiation.

In what follows, we describe the protocol's main features including call establishment, reservation process, and timing considerations, and we provide performance evaluation results by comparing the performance of the EBRP scheme with two typical TAW and

TAG type protocols.

EBRP protocol has been developed as part of the control plane of the LASAGNE all-optical label-swapping network [20].

2. Protocol Features

The requirement for no or limited buffering in the core network challenges the network efficiency and throughput that can be achieved. In one-way reservation schemes the risk of not finding the appropriate resources is high, especially for large bursts, resulting in a high loss ratio in the core. On the other hand, two-way reservation mechanisms waste resources, as these are reserved for time intervals longer than the actual data transmission (capacity is reserved at an intermediate node upon the arrival of the setup packet, while it is actually needed at least one round-trip delay later). Thus the end-to-end round-trip time not only determines the pretransmission delay but also affects the resource consumption. The challenge is how to achieve an efficient usage of resources, that is, to consume the reserved resources only when they are actually needed with short pretransmission times, while simultaneously avoiding collisions in the core nodes.

These deficiencies can be overcome using a two-way reservation scheme employing timed and in-advance reservations to schedule the bursts. Specific features of the EBRP signaling protocol are:

- **Timed reservations**: Outgoing capacity is reserved only for the duration of the burst, and capacity is released after the data traversing through the node. EBRP negotiates the reservation duration during the downstream setup phase, which may exceed burst holding time. In that case, strict timed requirements are relaxed and restored during the acknowledgment phase.

- **In-advance reservations**: If the capacity is not available at the requested time, it is reserved in the future at the first time it becomes available. If this does not satisfy the maximum delay requirements, the bandwidth request is rejected.

Timed reservations are important for high-speed networks since they allow a greater number of requests to be served. EBRP also avoids the wasteful repetition of the call establishment process, because it enables a transmission to reserve the required capacity in its first attempt, probably at a time later than the requested time. If the time at which the required resources become available is unacceptable (delay requirements of the transmission), the call is blocked and is reattempted later, probably using a different path. In order to employ timed and in-advance reservations, the actual burst size or an accurate estimation has to be communicated during the connection establishment process. Moreover, it is crucial for each node in the core network to be aware of its own resource availability (reserved capacity of its outgoing links as a function of time) in order to grant or reject new requests.

The EBRP signaling protocol employs two messages for call setup, namely:

- a SETUP packet,

- an ACKNOWLEDGMENT/REJECT (ACK/REJ) packet.

The SETUP packet is transmitted from source to destination and is used for resource negotiation. If the reservation is successful, a confirmation message (ACK) is sent back to the source to confirm the timed reservation. The ACK packet is a replica of the SETUP packet and communicates to intermediate nodes the (agreed) time that the resources are allocated. If the reservation process is blocked at an intermediate node, a REJECT packet is generated and sent backward to release capacity and inform the source. The use of timing

information to schedule bursts relieves the network from additional control signal overhead associated with the tearing down of a reservation.

Figure 1 shows the fields of the SETUP packet. The SETUP packet is sent to the nodes across the path to communicate the necessary information for the in-advance and timed reservation. The path of the SETUP could be specified as a sequence of link identifiers $L_1, L_2, \ldots, L_h$, corresponding to the links that this packet must traverse. Each node reads the first link identifier to determine the outgoing link to which it should be routed, and cyclically rotates the link identifiers so that the one just read becomes last.

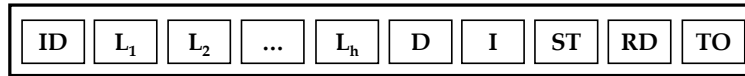| ID | L_1 | L_2 | ... | L_h | D | I | ST | RD | TO |

Fig. 1. Contents of SETUP packet.

Basic parameters of the SETUP packet that are communicated to all nodes are:

- The start time ST that specifies the time at which the reservation of capacity for the specific outgoing link should begin. ST is relative to the time that the SETUP packet is received by the node. The ST field is initially set equal to the round-trip delay time (T_{RTT}) and is updated at the intermediate nodes according to their resource availability, as presented in the following section.

- The time-offset field TO that contains the time, following the reception of the acknowledgment packet at the source, after which the source should start the transmission. The TO field is updated at every node in a way to be described later.

- The information field, I, that specifies the amount of information, i.e., the burst size that will be transmitted.

- The maximum delay field, D, that specifies the maximum allowable delay for this burst at the edge node (which is a QoS parameter). Clearly, we must have $D > T_{RTT}$; otherwise, the requested transmission cannot be served within the desired deadline over that path.

- The reservation duration time field, RD, that specifies the maximum time period following ST, during which the specific outgoing link should be reserved. Field RD provides a control over the allowed degree of the timed-reservation mechanism. For example, if initially RD is chosen equal to burst-transmission time, then resources are reserved exactly for the time needed, while when RD exceeds the transmission duration, a more relaxed timed reservation is made. If a SETUP packet reserves bandwidth at a node for a larger duration than the burst-transmission time, it is given more flexibility and a higher probability of reserving at least the minimum required duration at subsequent, downstream nodes. Therefore, the RD field can be used for QoS differentiation and also for treating more equitably bursts that traverse a large number of hops. How the RD field is updated is presented in the following section.

If the reservation is successful, a message is sent backwards (ACK) to the source to confirm the timed reservation. Since the ACK is used to notify the source of the exact status of the reservations, this message will be an updated version of the SETUP. The difference is that in the acknowledgment phase the link identifiers are resequenced and the ST, TO, and RD fields are not updated.

3. Call Setup and Reservation Process

For call establishment a SETUP packet is generated at the source node after the assembly of the burst, containing all the necessary routing and flow specification values. In particular, at the first hop, ST is set equal to T_{RTT}, TO equal to zero, and RD to a dynamic value per burst.

Since scheduling is required at each intermediate node, bursts traversing a longer path, and especially those requesting bandwidth for a longer duration, will have a higher risk of not finding the appropriate resources. In order to increase the burst acceptance probability, the key idea is to define an efficient reservation-time duration (RD) in the SETUP packet that is based on the class of service (CoS) of the traffic, the burst destination, and the burst size. The SETUP packet tries to reserve bandwidth for duration equal to RD at every hop, and if this is not possible, it tries to reserve bandwidth at least equal to the burst duration. If it cannot make a reservation even for that minimum duration, the SETUP packet is blocked. In this paper, we have considered the RD parameter proportional to number of hops and burst size as follows: $RD = khT_{data}$, where k is a constant, h is the number of hops on the path to be followed, and T_{data} stands for burst transmission duration. With this function, requests that traverse a large number of hops are granted a longer reservation time duration and thus have more flexibility in reserving resources and a higher forwarding probability. Moreover, larger bursts (which are more important than small ones in terms of throughput) are preferentially served. The increase of the acceptance probability for bursts that exhibit a high loss ratio introduces a certain amount of fairness. A detailed examination of the RD effect in network performance is provided in Subsection 4.B.

In general, RD can be used for burst QoS differentiation in the sense that high-priority bursts are allowed to reserve resources for time durations longer than their actual holding time, and thus to experience a lower probability of being blocked at subsequent links.

Figure 2 illustrates the timing considerations of the EBRP protocol where a setup process is instantiated between nodes S_0, S_1, S_2, and S_h. In particular, Fig. 2(a) shows the case when a request is blocked at an intermediate node, and Fig. 2(b), 2(c) illustrates the setup and acknowledgment phase of a successful process.

Assuming that a burst is routed through nodes $S_0, S_1, \ldots, S_h$, we denote with ST_{i-l}, TO_{i-l}, and RD_{i-l} the values of the fields ST, TO, and RD, when the SETUP packet traverses node $S_i (i = 1, \ldots, h-1)$. When node S_i receives the SETUP packet, it finds the first time t_{start}^i relative to the SETUP packet arrival that $t_{start}^i \geq ST_{i-1}$ and enough residual capacity is available to accommodate the burst. In order to do so, capacity should be available for a time period of $\lfloor t_{start}^i, t_{start}^i + T_{data} \rfloor$ and $t_{start}^i + T_{data} \leq ST_{i-1} + RD_{i-1}$.

- If $t_{start}^i + T_{data} > ST_{i-1} + RD_{i-1}$, the transmission is blocked and a REJECT packet is returned to free reserved resources at intermediate nodes and to notify the source.

- If t_{start}^i satisfies the above requirements, node S_i reserves the resources starting from t_{start}^i up to $t_{end}^i = \min \lfloor (ST_{i-1} + RD_{i-1}), t_{available}^i \rfloor$, where $t_{available}^i$ is the time that capacity is not available due to reservation made by another burst.

In the case of a successful reservation, node S_i updates the SETUP packet and forwards it to the next node. In particular, it updates the reservation starting time (ST) and time offset (TO):

- $ST_i = t_{start}^i$, $TO_i = TO_{i-1} + \delta_i$, where $\delta_i = t_{start}^i - ST_{i-1}$.

  Similarly, the reservation duration time (RD) is decremented:

- $RD_i = RD_{i-1} - (\sigma_i + \delta_i)$, where $\sigma_i = |(ST_{i-1} + RD_{i-1}) - t_{available}^i|$ if $(ST_{i-1} + RD_{i-1}) - t_{available}^i > 0$ or else $\sigma_i = 0$.
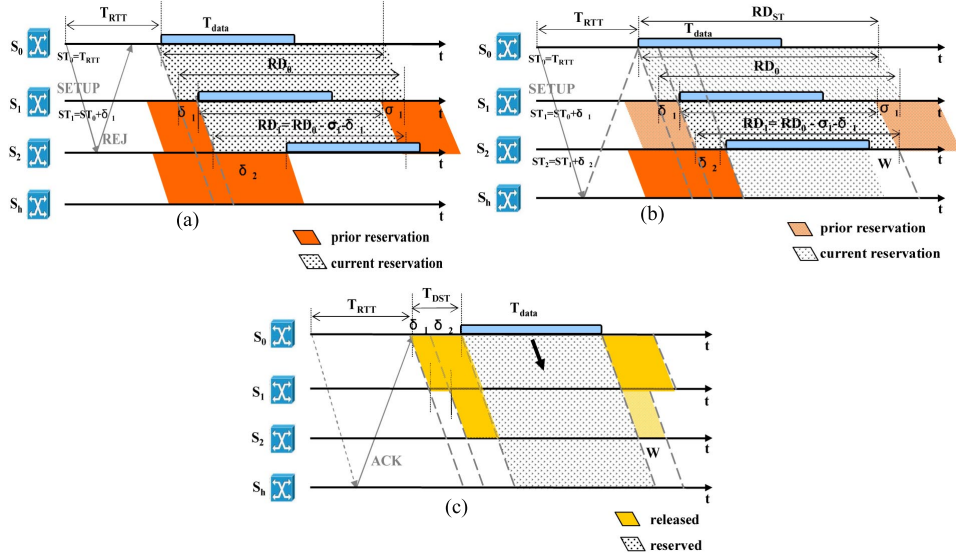
Fig. 2. Timing considerations in the EBRP signaling protocol: (a) blocked call setup phase at an intermediate node, (b) successful end-to-end call setup, and (c) acknowledgment phase with abundant resources release.

With respect to Fig. 2(a), node S_2 finds the earliest time that resources are available at an offset time δ_2. However, void filling cannot be performed since T_{data} size does not fit, and thus S_2 drops the call setup. In the case depicted in Fig. 2(b), S_2 finds adequate resources beyond the δ_2 offset and grants the request. During the acknowledgment phase (Fig. 2(c)) the abundant reserved resources in S_1 and S_0 nodes are released.

When the SETUP packet arrives at the last core node S_{h-1} (before the egress node S_h), there is no need to reserve resources beyond the burst duration (assuming that the egress edge router is commissioned only to buffer and disassemble the bursts). Thus, the node S_{h-1} (S_2 node in Fig. 2) reserves resources only for a duration equal to the burst transmission time, that is, $[t_{\text{start}}^{h-1}, t_{\text{start}}^{h-1} + T_{\text{data}}]$. Following the previous procedure, the SETUP message that reaches node S_{h-1} has accumulated all the time offsets δ_i issued by the intermediate nodes, and therefore ST_{h-1} determines the earliest transmission starting time that resources are available. Destination node S_h, after checking for the availability of adequate memory to store the specified burst size, sends an ACK packet back to the source with the following fields:

$$TO_{h-1} = \sum_{i=0}^{h-1} \delta_i \text{ and } ST_{h-1} = T_{\text{RTT}} + \sum_{i=0}^{h-1} \delta_i.$$

Upon receiving the ACK packet, the intermediate nodes retrieve the agreed transmission starting time ST_{h-1} and update their reservations to exactly match the burst duration time, T_{data}. In this way the abundant resources are released, and the strict time limits are restored. The source node upon receiving the ACK packet waits for time equal to TO (TO_{h-1}) and then begins transmission.

In the case that buffering at the destination egress node is also a limited resource, a timed buffer reservation may also have to be made. Thus, the buffer at the destination node can be treated in EBRP in the same way that bandwidth is treated, and can be viewed as the last leg of the reservation.

4. Performance Evaluation

In order to assess the performance of the resource reservation protocol in a network environment and compare it to that of other typical resource reservation protocols, we have developed a discrete-event simulation based on the ns-2 platform. For the sake of comparison, apart from the proposed signaling scheme, we also simulated (i) a typical tell-and-wait (TAW) protocol and (ii) a typical tell-and-go (TAG) protocol, namely, the just enough time (JET) with void filling. The simulation was conducted on the NSFnet backbone network topology, shown in Fig. 3, where all links were assumed to be bidirectional employing a single wavelength per direction. The SETUP and ACK packets processing delay was set to 0.01 ms. In the simulation experiments, we have used the Pareto ON/OFF traffic generator embodied in the ns-2 tool to generate packet streams, and further we modeled two edge router (ER) architectures, one employing virtual output queues, that is, a separate first-in-first-out (FIFO) queue per burst destination, and one with a single FIFO serving all requests.
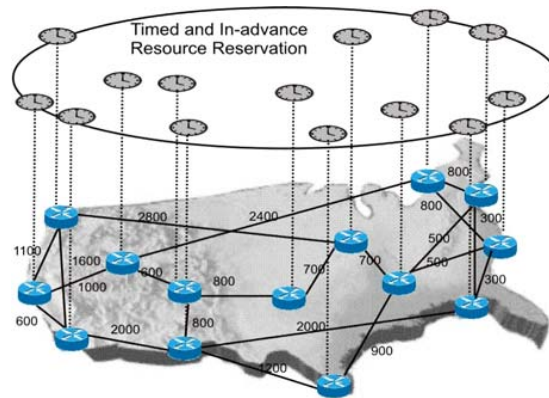

Fig. 3. 14-node NSFnet backbone network topology (the shown distances are in kilometers).

All Pareto sources contribute equally to the ER load, and packet destinations are evenly distributed among all network nodes. The core network is assumed to be free of blocking (bursts are transmitted after lightpath establishment), and thus throughout this study we have assumed that blocking refers to those bursts lost because of ER buffer overflow or the expiration of the maximum allowable delay at the edge [4].

The whole burst aggregation cycle can be described as follows. Each burst assembly queue (separate queue per source–destination pair) has an individual time-out signal denoted as $T_{\text{time-out}}$. In these simulations, burst aggregation cycles of the order of tens of milliseconds are considered, so that they are similar to the network propagation delays. Data packets from the N attached sources are collected, sorted, and buffered at the corresponding queues. When the burst size exceeds a size limit, denoted by B_{MAX} or when a time limit equal to $T_{\text{time-out}}$ has elapsed, the assembler requests to transmit a SETUP packet. An overall burst manager controller is responsible for setup transmissions, which in the case of virtual output queues (VOQ) selects requests from the assemblers in a simple, round robin way, while in the single FIFO, ER serves only the first request.

We used the burst blocking probability as the main metric for assessing protocol performance. Additional performance metrics that were used were the average number of retrials until a successful reservation and the average holding time of the burst at the network edge node. We define the holding time of a burst as the time that elapses between the assem-

bly of the burst and the time its first bit is transmitted in the network. This time includes the round-trip delay associated with the two-way reservation mechanism. Furthermore, the effect of the reservation-time duration, RD, on blocking performance is investigated for various loads. It must be noted here that if the connection-establishment process is blocked and the traffic delay requirements allow, the burst manager retries to make a reservation until either a successful reservation is made or the expiration of the allowed delay. The FIFO property of each virtual queue in the case of multiple FIFOs is maintained, and thus until the final successful reservation or rejection of a request, the manager does not proceed to the next burst residing in the queue.

4.A. Single and Multiple FIFO Experiments

In this subsection, simulation results for the EBRP signaling protocol are presented and compared with corresponding results obtained for the TAW and JET protocols. For these experiments, we have chosen the following parameters: $C_{core} = 10$ Gbits/s, $N = 10$ Pareto sources per ER, each with incoming rate $C_{access} = 1$ Gbit/s, and Pareto shape parameter $a = 1.2$. The minimum burst size is set equal to 400 bytes, which corresponds to a Mean_ON burst time of 19.2 µs, and thus the Mean_OFF (idle time) can be calculated from the desired load p [22]. The total offered load to the network from all sources is $14NpC_{access} = 140p$ Gbits/s, since the NSF network employs 14 edge routers. The actual network load can be calculated by the following equation [23]:

$$L = N_{IE} \overline{h} r / C_{core} (2L),$$

where N_{IE} is the number of ingress–egress pairs, $\overline{h}$ is the mean number of hops per path, r is the incoming-requesting flow rate (per ingress–egress pair), and L is the total number of the bidirectional links in the network. For the case of the NSF network we have $N_{IE} = 182$, $\overline{h} = 2.26$, $L = 21$, and for the particular ER and Pareto generators models, $r_{Pareto} = N/13pC_{access}$.

Figures 4(a), 4(b) show the blocking probability of the EBRP protocol as well as of the typical TAW and JET schemes for single and multiple FIFOs, respectively.
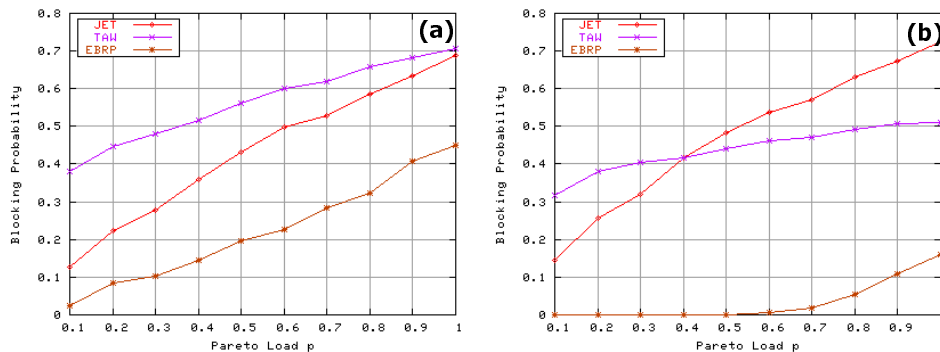

Fig. 4. Blocking probability of the EBRP, TAW, and JET protocol for the case of (a) single FIFO and (b) multiple FIFO ER architecture.

For these experiments the constant k of the RD parameter was set equal to $k = 4$, the burst aggregation time $T_{time-out}$ equal to 0.3 s, the burst time delay D equal to 0.3 s, and the ER buffer size equal to 256 Mbytes.

As shown in Fig. 4, EBRP protocol outperforms the TAW and JET schemes for both ER models. Especially in the case of multiple FIFOs, the blocking probability of EBRP is negligible for loads less than $p = 0.6$. As expected, the performance of the JET scheme remains

the same for both ER architectures. On the other hand, in the case of EBRP and TAW protocols, blocking performance improves when moving from single to multiple FIFOs. This improvement is significant for the EBRP and small for the TAW. The low performance of single FIFO was expected, since it exhibits the so called head-of-line blocking effect that results in consecutive burst deadlines expirations.

Figures 5(a), 5(b) show the holding times at the network edge points, when D is set to a very high value in order to better compare the EBRP scheme with the TAW. To this end, no burst blocking occurs, and therefore the worst-case holding time is measured. We can observe that lower holding times are achieved with the EBRP protocol, since timed and in-advance reservations better utilize the available resources. It is worth noting that holding time in the EBRP protocol increases faster than the TAW with respect to the offered load, primarily due to the in-advance mechanism that schedules burst transmission in the future, dispensing of the need for repeating call setup. JET holding times are not illustrated in these graphs, since a burst is released after an offset (depending on the number of traversed hops) that is negligible compared to the two-way scheme delays.
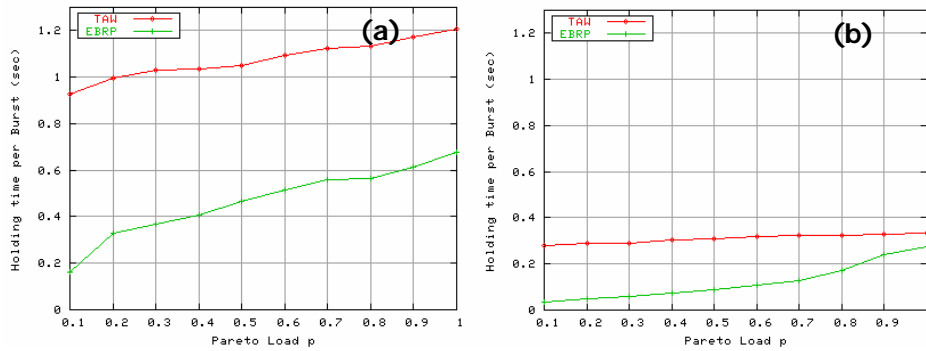

Fig. 5. Average holding time at the network ingress point for the EBRP and TAW protocols for the case of (a) single and (b) multiple FIFO ER architectures. Bursts of maximum allowable delays are set to infinity.

4.B. Effect of Reservation Time Duration

In this subsection, we investigate the performance of the EBRP protocol and the dependence of the reservation time duration period on the blocking probability. The parameter RD can be chosen based on the class of service (CoS), the burst destination, and the burst size. In the simulation experiments, we used a separate FIFO per burst destination and experimented with a linear RD function that is proportional only to number of hops and the burst size. For the implemented ER architectures, under light/medium load, the burst assembly functions in pure time-out mode. Thus average burst size ($B_{\text{aver-size}}$) is given by

$$B_{\text{aver-size}} = \frac{N}{13} p C_{\text{access}} T_{\text{time-out}}.$$

We conducted experiments with two different $T_{\text{time-out}}$ values to investigate indirectly the effect of the average burst size in protocol performance via the RD parameter.

Figures 6(a), 6(b) show the corresponding results for three different source loads ($p = 0.6, \ldots, 0.8$) that correspond to average burst size equal to $B_{\text{aver-size}} = 28.8p$ Mbytes and $B_{\text{aver-size}} = 14.4p$ Mbytes, respectively. For these experiments the burst time delay D was set again to 0.3 s.
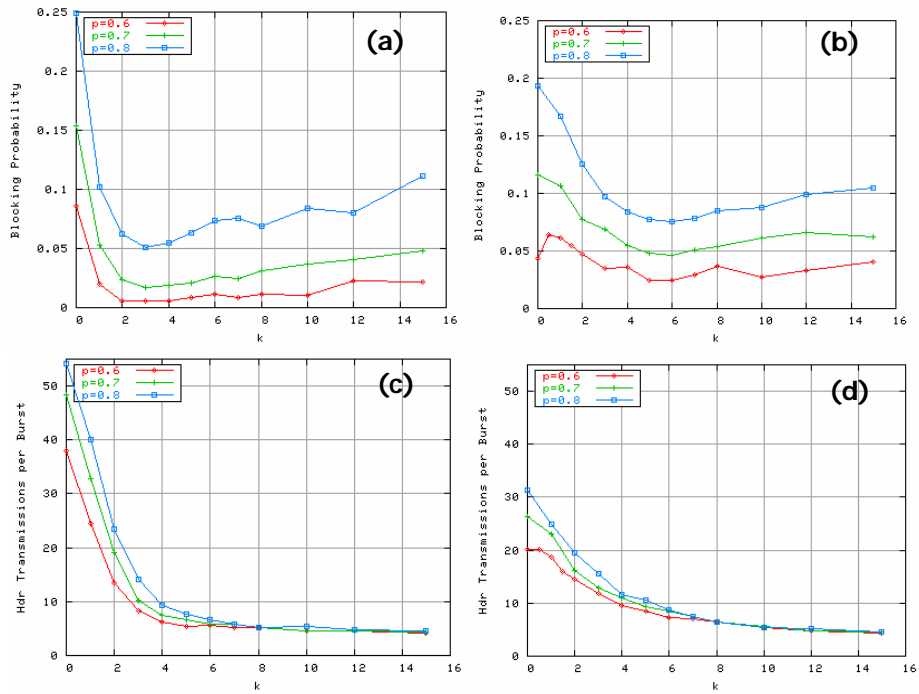
Fig. 6. Blocking probability for various reservation time duration values for average burst size equal to (a) $28.8p$ Mbytes and (b) $14.4p$ Mbytes. (c), (d) Corresponding number of SETUP retransmissions. The *x* axis corresponds to the choice of the constant *k* and reservation time duration is calculated on a per burst basis as $RD = khT_{\text{data}}$.

From these figures, it can be deduced that blocking probability initially decreases with the increase of the actual RD parameter. However, above a critical value, the abundant reservation of resources prevents other setup packets from performing successful reservations, and thus blocking starts to increase. This is evident for values of k higher than 10. In that case EBRP resembles TAW operation. Further, as k increases, nodes reserve abundant resources, and therefore more and more requests are shifted and scheduled to be served far in the future after other committed reservations. Since tight scheduling of bursts is not ensured, we end up with a low utilization of resources and increased holding times.

To further prove the usefulness of the RD parameter, we have also measured the SETUP packet retransmissions for all successful requests. Figures 6(c), 6(d) show the corresponding average number of retrials for the same average burst sizes. It is obvious that the number of SETUP retransmissions reduces rapidly, alleviating the network from processing numerous signaling packets.

It is worth noting here that in the above experiments, and especially in the case of Figs. 6(b), 6(d), burst transmission times are one order of magnitude smaller than the mean round-trip time, $E\{T_{RTT}\}$. Therefore a large number of small bursts are produced, a fact that makes efficient resource allocation a complicated task, and for k values smaller than 2, instabilities are observed. From Figs. 6(a), 6(d), it can be seen that for the simulated traffic scenario the optimum value of k is between 2 and 6.

In principle, other RD functions that depend logarithmically or as a polynomial on the number of hops or the transmission time can be applied as well. The simple linear function that was examined here proves the efficiency of the proposed scheme.

5. Conclusions

In this paper we have presented a hybrid burst reservation protocol suitable for bufferless, OBS networks. The protocol employs a two-way reservation mechanism that uses in-advance and timed reservations in order to detect the earliest available time instance to schedule the flows and block the capacity only for the duration that is needed. A key feature that the EBRP protocol introduces is the reservation time duration that adds the flexibility of negotiating the reservation horizon during the call setup phase to increase forwarding acceptance probability and can be used to provide QoS differentiation. In this paper, we have evaluated a reservation duration function that depends linearly on the number of hops and the burst size. We have shown that EBRP exhibits superior performance in terms of blocking probability and resource utilization for bursts that can tolerate the round-trip time delay in a full-scale network simulation. Other RD functions are currently investigated in order to optimize protocol performance.

Acknowledgments

References and Links

[1] C. Qiao and M. Yoo, "Optical burst switching (OBS)—a new paradigm for an optical internet," J. High Speed Netw. **8**, 69–84 (1999).

[2] C. Qiao, "Labeled optical burst switching for IP-over-WDM integration," IEEE Commun. Mag. **38**(9), 104–114 (2000).

[3] C. Qiao and M. Yoo, "Choices, features and issues in optical burst switching," SPIE Opt. Netw. Mag. **1**(2), 36–44 (2000).

[4] M. Dueser and P. Bayvel, "Analysis of a dynamically wavelength-routed optical burst switched network architecture," J. Lightwave Technol. **20**, 574–585 (2002).

[5] Jingxuan Liu, N. Ansari, and T. J. Ott, "FRR for latency reduction and QoS provisioning in OBS networks," IEEE J. Sel. Areas Commun. **21**, 1210–1219 (2003).

[6] E. A. Varvarigos and V. Sharma, "The ready-to-go virtual circuit protocol: a loss-free protocol for multigigabit networks using FIFO buffers," IEEE/ACM Trans.Netw. **5**(5), 705–718 (1997).

[7] M. Yoo and C. Qiao, "Just-enough-time (JET): a high speed protocol for bursty traffic in optical networks," in *IEEE/LEOS Technologies for a Global Information Infrastructure* (IEEE, 1997), pp. 26–27.

[8] J. S. Turner, "Terabit burst switching," J. High Speed Netw. **8**(1), 3–16 (1999).

[9] J. Y. Wei and R. I. MacFarland, Jr., "Just-in-time signaling for WDM optical burst switching networks," J. Lightwave Technol. **18**, 2019–2037 (2000).

[10] I. Baldine, G. N. Rouskas, H. G. Perros, and D. Stevenson, "JumpStart: a just-in-time signaling architecture for WDM burst-switched networks," IEEE Commun. Mag. **40**(2), 82–89 (2002).

[11] Y. Xiong, M. Vandenhoute, and H. C. Cankaya, "Control architecture in optical burst-switched WDM networks," IEEE J. Sel. Areas Commun. **18**, 1838–1851 (2000).

[12] J. Xu, C. Qiao, J. Li, and G. Xu, "Efficient channel scheduling algorithms in optical burst switched networks," in *Proceedings of 22nd Annual Joint Conference of the IEEE Computer and Communications Societies* (IEEE, 2003), Vol. 3, pp. 2268–2278.

[13] F. Masatti, "Fiber-delay lines optical buffer for ATM photonic switching applications," in *Proceedings of 12th Annual Joint Conference of the IEEE Computer and Communications Societies* (IEEE, 1993), pp. 935–942.

[14] I. Chlamtac, A. Fumagalli, L. G. Kazovsky, P. Melman, W. Nelson, P. Poggiolini, M. Cerisola, A. Choudhury, T. Fong, R. Hofmeister, C. Lu, A. Mekkittikul, D. Sabido IX, C. Suh, and E. Wong, "CORD: contention resolution by delay lines," IEEE J. Select. Areas Commun. **14**, 1014–1029 (1996).

[15] K. Vlachos, E. Varvarigos, and C. Bintjas, "Performance evaluation of an optically interconnected scheduling switch network for Pareto traffic," J. Opt. Netw. **3**, 760–768 (2004).

[16] C. F. Hsu, T. L. Liu, and N. F. Huang, "Performance analysis of deflection routing in optical burst-switching networks," in *Proceedings of 21st Annual Joint Conference of the IEEE Computer and Communications Societies* (IEEE, 2002), pp. 1803–1812.

[17] M. Yoo, C. Qiao, and S. Dixit, "QoS performance of optical burst switching in IP-over-WDM networks," IEEE J. Sel. Areas Commun. **18**, 2062–2071 (2000).

[18] R. Karanam, V. Vokkarane, and J. Jue, "Intermediate node initiated (INI) signalling: a hybrid reservation technique for optical burst-switched networks," in *Proceedings of Optical Fiber Communication Conference* (Optical Society of America, 2003), Vol. 1, pp. 213–215.

[19] E. A. Varvarigos and V. Sharma, "An efficient reservation connection control protocol for gigabit networks," Comput. Netw. ISDN Syst. **30**, 1135–1156 (1998).

[20] F. Ramos, E. Kehayas, J. Martinez, R. Clavero, J. Marti, L. Stampoulidis, D. Tsiokos, H. Avramopoulos, J. Zhang, P. Holm-Nielsen, N. Chi, P. Jeppesen, N. Yan, I. Tafur Monr Koonen, M. Hill, Y. Liu, H. Dorren, R. Van Caenegem, D. Colle, M. Pickavet, and B. Riposati, "IST-LASAGNE: towards all-optical label swapping employing optical logic gates and optical flip-flops," J. Lightwave Technol. **23**, 2993—3011 (2005).

[21] Q. Zhang, V. M. Vokkarane, J. P. Jue, and Biao Chen, "Absolute QoS differentiation in optical burst-switched networks," IEEE J. Sel. Areas Commun. 22, 1781–1795 (2004).

[22] K. Fall and K. Varadhan, "The ns Manual" Chap. 35, pp. 320–324; http://www.isi.edu/nsnam/ns/doc.

[23] J. Li, G. Mohan, and K. C. Chua, "Dynamic load balancing in IP-over-WDM optical burst switching networks," Comput. Netw. **47**, 393–408 (2005).