A Quantitative Evaluation of Reordering in OBS Networks and its Impact on TCP Performance

Jordi Perelló[†], Sebastian Gunreben[§] and Salvatore Spadaro[†]

[†]Optical Communications Group (GCO), Signal Theory and Communications Department (TSC)
Universitat Politècnica de Catalunya (UPC), Jordi Girona 1-3, 08034 Barcelona, Spain
Tel: (+34) 93 401 7179, Fax: (+34) 93 401 7200
Email: jperello@tsc.upc.edu, spadaro@tsc.upc.edu
[§]Institute of Communications Networks and Computer Engineering
University of Stuttgart, Pfaffenwaldring 47, 70569 Stuttgart, Germany
Tel: (+49) 711 685 67968, Fax: (+49) 711 685 57968
Email: gunreben@ikr.uni-stuttgart.de

Abstract—**This paper presents a layered framework to measure the reordering introduced by contention resolution strategies in OBS networks. In particular, characterization is based on the reordering metrics proposed by the IETF IPPM WG. The obtained results are twofold. First, they quantify the impact of burst reordering on TCP throughput performance. Second, they give insight into solving burst reordering by well dimensioned buffers.**

I. INTRODUCTION

Optical Burst Switching (OBS, [1]) has emerged as a promising transport technology for next-generation Internet. As a matter of fact, OBS networks become a combination of packet and circuit switched networks, where packets are firstly aggregated in edge routers and, then, are sent as bursts along bufferless optical networks. This provides the benefit from statistical multiplexing in the optical domain, which allows better adaptation than circuit-switched networks (OCS) to higher layer dynamics. Besides, it lessens technology requirements compared to all-optical packet switching (OPS) networks.

OBS relies either on one-way or two-way resource reservation. However, one-way resource reservation is usually implemented in wide area networks, where propagation delays may become much longer than burst transmission times. This means, that unlike in OCS networks, data transmission is not delayed until the reception of the reservation acknowledgement packet. Conversely, a burst is subsequently transmitted after an offset time since the reservation request control packet was sent. Hence, there is no sureness about the proper transmission of the bursts, which could be lost due to contention at intermediate nodes.

Several contention resolution strategies [2] have been proposed to minimize burst loss probability, namely deflection routing, Fiber Delay Lines (FDLs) as optical buffers, wavelength conversion, and even combinations of them. It has been extensively demonstrated that these strategies succeed in decreasing burst blocking probability. Nonetheless, an additional degree of reordering can be introduced to the network. In fact, in order burst delivery cannot be guaranteed, as long as the

extra time introduced in the FDLs or in the deflection routes exceed the burst inter-arrival time (IAT).

Current studies reveal that the majority of Internet traffic is carried by Transfer Control Protocol (TCP), which provides applications with reliable data transfer support. Likewise, although User Datagram Protocol (UDP) real-time applications are gaining ground (e.g., VoIP), it is expected that TCP will remain dominant in the foreseeable future. Therefore, insight into its performance on top of next-generation transport network technologies undertakes special significance.

It is widely known the effect of packet loss on TCP. In TCP Reno [3], the sender of a TCP session is notified of a packet loss by means of duplicate acknowledgements. In this context, the TCP *fast retransmit* algorithm is invoked, whether the duplicate acknowledgement threshold (DUP-ACK) is reached. As a result, the missing packet is retransmitted and the sender's congestion window is halved, which decreases TCP throughput significantly. A similar situation occurs whether a packet becomes reordered. Note, that in the event of reaching the DUP-ACK threshold, TCP may consider a reordered packet as lost, even though it is only delayed.

This paper focuses on the viability of OBS as a carrier technology for TCP. As mentioned earlier, optical bursts can be dropped or even reordered due to contention resolution. Whenever a burst is dropped, some or even the completeness of the contained packets are also lost. Additionally, the reordering introduced at the OBS layer is propagated to higher layers (e.g., TCP), affecting these ones as well.

Heretofore, the impact of OBS on end-to-end TCP throughput has been studied in detail. Mainly, previous works focus on the consequences on TCP of burst loss due to contention, assuming an integrated TCP/OBS scenario modeled as a monolithic block [4] [5] [6]. This integration, however, complicates the identification of key parameters at the OBS layer that become significant in the final TCP performance.

In this work, we follow a layered approach to model the impact of burst reordering on TCP. Firstly, we quantify the introduced reordering at the OBS layer under several contention resolution strategies. With such purposes in mind,

we apply the reordering metrics presented by the IETF in [7], which provide us extensive information. On the one hand, they quantify the buffer size, which should be placed at edge nodes, to solve reordering at the OBS layer. This would permit the sending of already ordered packets to the IP layer, so that burst reordering would remain transparent to TCP. On the other hand, whether reordering is left to the TCP layer, they provide information about the violation of the DUP-ACK threshold due to reordering, which allows TCP performance estimation.

The rest of this paper is structured as follows. Section II overviews the considered contention resolution strategies. Section III introduces the IETF reordering metrics. Section IV describes the scenario under study. Then, section V presents the obtained burst reordering results, whose repercussion on TCP is discussed on section VI. Finally, section VII concludes this work.

II. STRATEGIES FOR CONTENTION RESOLUTION IN OBS NETWORKS

As mentioned before, OBS provides the benefit of statistical multiplexing directly in the optical domain. Nonetheless, in order to achieve low burst loss probability, contention resolution emerges as a key issue in OBS networks. In particular, contentions occur whenever two or more bursts try to leave the OBS switch on the same output port at the same time, which leads to burst loss.

Up to now, several contention resolution strategies have been proposed in the literature, whose effectiveness has been extensively demonstrated. This makes widely accepted that future OBS transport networks will be enhanced with contention resolution capabilities. Therefore, for the sake of generality, we quantify reordering in OBS networks under several contention resolution strategies. With this aim, as in [8], we deal with basic strategies and combinations of them.

As basic strategies, we consider *Conv*, *Defl* and *FDL*. Specifically, *Conv* assumes that wavelength conversion without limitations is available at each node. In turn, *Defl* refers to deflection routing. In this case, whether contention exists, the burst is sent to an alternative output fiber of the node, following the shortest path towards the destination. Finally, *FDL* contemplates buffering by means of a single shared feedback FDL employing WDM.

Departing from the aforementioned strategies, combinations of them can be also applied. Because the order of application of each strategy is essential, combined strategies are named by a concatenation of the former's acronyms. In this way, we could have combined strategies like *ConvFDL*, *ConvDefl*, *ConvFDLDefl* or even *ConvDeflFDL*. For instance, whether *ConvDeflFDL* was used, conversion would be firstly attempted. If conversion was unavailable, deflection would be subsequently tried. If deflection also failed, buffering would be tried afterwards. If, unfortunately, none of them worked, the burst would be finally dropped. Note that conversion is always applied first. This assumption agrees with some previous works' results [9] [10].

III. REORDERING METRICS

A complete characterization of reordering becomes noteworthy, specially when assessing a protocol's viability over a given network. With this objective in mind, the IETF IPPM working group has recently standardized a set of metrics [7] to characterize reordering effects in generic packet networks (e.g., OBS networks).

In this section, we introduce some of these metrics, which have been used in the following sections to provide a broad view of reordering in the scenario under study.

A. Reordering Ratio

The detection of reordering is done at the destination, looking at the sequence number $s[i]$ of each packet, where i numbers the arriving packet order at destination. This sequence number is set at the source node, following a consecutive integer sequence. In turn, the destination node maintains a counter $s'[i]$, which identifies the sequence number of the following expected packet. Under normal conditions, $s'[i]$ is equivalent to the sequence number of the last received in order packet plus 1. When packet i arrives, the packet is considered as reordered whether $s[i] < s'[i]$. Conversely, whether $s[i] \geq s'[i]$, the packet is considered in order and $s'[i+1] = s[i] + 1$.

In this context, the reordering ratio quantifies, given a certain data stream, the ratio of reordered packets. This figure is easily obtained as the number of reordered packets divided by the number of received packets. It is noteworthy, that in case of duplicate packets, only the first copy is considered.

B. Reordering Extent

The reordering extent quantifies the extent to which packets are reordered. As a main application, it provides information about the minimal storage (i.e., buffer size) at the receiver, which would be needed to restore packet order at destination.

Let us suppose that $s[1]$, $s[2]$,..., $s[L]$ is the sequence of incoming packets' numbers at destination. Furthermore, imagine that a given packet i is considered as reordered. This would mean that a set of indexes j ($1 \leq j < i$) exist, which accomplish $s[j] > s[i]$. In this scenario, the reordering extent can be obtained as $i - j$, for the smallest value of j such that $s[j] > s[i]$.

C. n-Reordering Ratio

Finally, the n-reordering ratio is intended to provide information about the impact of reordering on TCP. Formally, a received packet with sequence number $s[i]$ is considered as n-reordered whether $i - j \leq n < i \wedge s[j] > s[i]$.

In other words, an n-reordered packet is a reordered packet which causes n DUP-ACK to TCP. Note that, this metric could be useful to quantify the percentage of reordered packets which would trigger TCP *fast retransmit* procedures. Indeed, by setting $n = 3$, we would obtain the percentage of reordered packets, which trigger *fast retransmit* in TCP Reno [3]. As a further application, this metric can be helpful for matching the duplicate ACK threshold for a given path, so that unnecessary retransmissions are avoided.

Fig. 1. Scenario under study

Fig. 2. Burst loss probability

Fig. 3. Burst reordering ratio

IV. SCENARIO UNDER STUDY

With evaluation purposes, we implement the 16-node COST 266 reference network (Fig. 1). For simplicity, all links have the same length of 200 km, which introduces a link propagation delay of 1 ms. Network resources are dimensioned according to a static traffic demand matrix, obtained from a 2006 European population model [11]. Particularly, a total demand of 9.9 Tbps is offered to the network, which corresponds to 990 Erlangs for a 10 Gbps line rate. Then, wavelength capacity is distributed in the network, so that shortest path routing leads to equal blocking probabilities on all links (i.e., dimensioning according to the Erlang model [12]). In this context, different network load situations can be achieved by overdimensioning wavelength capacity by a given factor (denoted as *overdimensioning factor* in the figures).

Regarding traffic characteristics, the burst departure process follows a Poisson process and burst length is exponentially distributed with mean 100 kbit [8]. In turn, in OBS nodes, the number of add/drop ports is unlimited and the switching matrix is non-blocking. Besides, the delay for burst control packet processing is compensated by a short extra FDL of appropriate length at the input of the node.

With contention resolution purposes, we assume one FDL per node with a certain number of wavelengths. The length of this FDL equals the mean burst transmission time, defined as the time needed to transmit an average sized burst (i.e., 10 μs for 100 kbit bursts over 10 Gbps data links). Note, that the wavelengths of this FDL are shared and the number of wavelength converters per node is unlimited. Nonetheless, if all wavelengths are occupied upon burst arrival in the FDL, the burst is discarded.

V. SIMULATION RESULTS

In this section, we evaluate the performance of the strategies *Conv*, *ConvFDL*, *ConvDefl*, *ConvFDLDefl* and *ConvDeflFDL* in an OBS scenario. With this aim, we focus not only on burst loss probability, but also on introduced reordering, which harms TCP performance as well. First of all, we compare, for each contention resolution strategy, burst loss probability improvements against introduced reordering. Then, we investigate the possibility to solve reordering at the OBS layer.
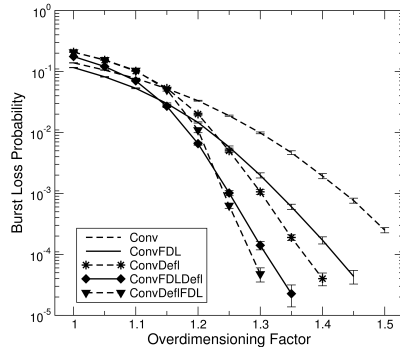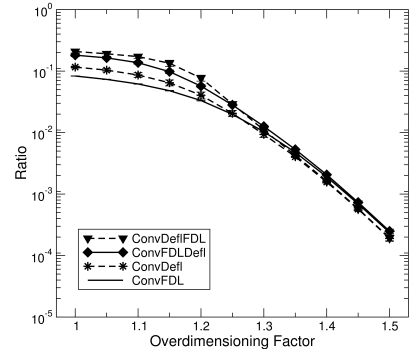
Finally, we conclude on the 3-reordering ratio at the OBS layer. With reordering quantification purposes, we measure burst reordering between each demand's source-destination pair and we provide global network statistics. Note, that if no wavelength conversion would have been feasible in the network, our conclusions on reordering would be still valid, as *Conv* is applied first in all schemes.

In our studies we assume, for cost reasons, 8 wavelengths in the FDLs. Moreover, to avoid unnecessary load and high propagation delays in the network, we limit the number of deflections to 1. Previous works demonstrate that the improvements due to further deflections are marginal [8], as long as a reasonable amount of flexibility is allowed in the network. The results have been obtained using the event-driven simulation library IKRSimlib [13].

A. Burst Loss Probability vs. Reordering Ratio

In Fig. 2, we depict the behavior of each considered strategy in terms of burst loss probability. Besides, Fig. 3 quantifies the reordering ratio introduced by each one. Note that *Conv* has not been included in Fig. 3. When applying this strategy, all bursts travel along the same path and no buffering is used. Therefore, no reordering is introduced.

Looking at Fig. 3, for high and medium loads, *ConvDeflFDL* introduces the highest reordering, followed by *ConvFDLDefl*,*ConvDefl* and *ConvFDL*. However, towards low loads, all strategies behave similarly in terms of reordering ratio. Observing now Fig. 2, for high loads the performance of all the strategies which use deflection routing (i.e., *ConvDefl*, *ConvFDLDefl* and *ConvDeflFDL*) is poor, as they overload an already highly loaded network. Nonetheless, towards lower loads, deflection (alone or combined) decrease burst loss probability rapidly, as enough network resources become available.

The majority of studies coincide that in a realistic OBS scenario, burst blocking probabilities should range from 10^{-3} to 10^{-6}. Fig. 3 shows that in this operating range, all strategies introduce the same reordering to the network. However, notice in Fig. 2 that *ConvDeflFDL* provides the best performance regarding burst loss probability. Therefore, this leads to the conclusion that this strategy may provide the best compromise between burst loss and introduced reordering.
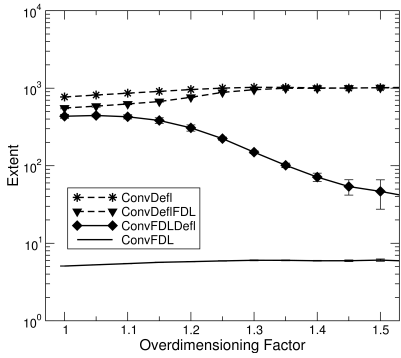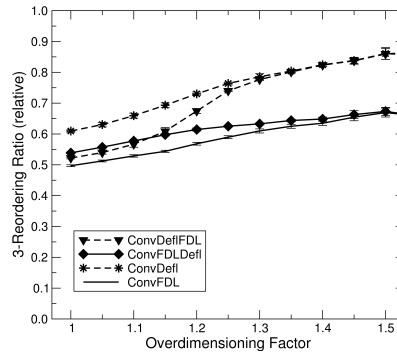
Fig. 4. Reordering Extent
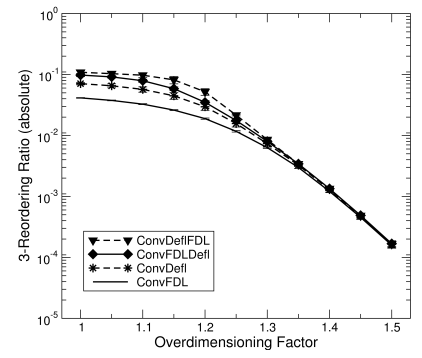

Fig. 5. 3-Reordering burst ratio (relative)

Fig. 6. 3-Reordering burst ratio (absolute)

B. Reordering Extent Quantification

In this subsection, we analyze the possibility to restore burst order directly at the OBS layer. Then, already ordered packets could be sent to the IP layer, so that burst reordering would remain transparent to TCP. With these purposes, a possible solution is the placement of buffers, on a per flow basis, at OBS edge nodes. Such buffers would store incoming out-of-order bursts, waiting for the expected one to be received. In this context, the reordering extent metric provides information about the mean extent to which bursts are reordered. Therefore, this gives an idea of these buffers' size.

Fig. 4 shows the mean reordering extent for each strategy under consideration. Deflection routing technique introduces large extents, in the order of one thousand. In fact, deflected bursts transverse at least one more hop than those going through the direct path. This accounts for an additional propagation delay of 1 ms, which is two orders of magnitude greater than the mean burst transfer time (10 μs in our scenario). Conversely, the use of buffering, like in *ConvFDL*, introduces relatively low extents (as can be seen, lower than 10). Therefore, by means of relatively low buffering capacities, these strategies would enable the restoration of the burst order directly at the OBS layer.

It is noteworthy that towards low loads, the introduced extent by combined strategies tend to the former's one (e.g., towards low loads, *ConvDeflFDL* tends to *ConvDefl*). This is due to the fact that in a low loaded network, contentions can be solved in the first attempt in most situations.

C. 3-Reordering Ratio Evaluation

Until now, we have quantified, for each contention resolution strategy under consideration, the reordering ratio and introduced reordering extent. While the former provides a general view of what happens in the network, the latter evaluates the possibility to restore order directly in the OBS layer. On the one hand, such information provides understanding about the origins of reordering, and evaluates specific solutions to restore it (as seen in previous subsection). However, on the other hand, these metrics do not illustrate the direct implication of reordering on TCP. In this subsection, we quantify the n-Reordering ratio metric. Specifically, as in TCP Reno, we consider n = 3.

Fig. 5 shows the relative 3-Reordering burst ratio, understood as the ratio of reordered bursts, which become 3-Reordered or more. As seen, 3-Reordering ratio increases along with the overdimensioning factor. This could be due to several reasons. For low loads, deflected bursts have more possibilities to succeed, which would increase 3-Reordering ratio. Moreover, for higher loads, since more reordering exists, this could decrease 3-reordering. For instance, let us assume a reordered burst. It may happen, that the following ones become also reordered, which could cause this one not to be 3-reordered. Further looking at Fig. 5, buffering technique introduces less 3-Reordering ratio than deflection. Particularly, it can be seen how *ConvFDL* outperforms all the remainder strategies. In fact, all these remainder strategies involve deflection routing technique.

In order to allow a better illustration of 3-Reordering in the network, we provide absolute 3-Reordering results. Specifically, Fig. 6 depicts the ratio of received packets, which become 3-Reordered or more (i.e., absolute 3-Reordering ratio). Absolute values present a behavior inline with the reordering packet ratio. For high loads, differences between the strategies can be appreciated, outperforming *ConvFDL* the remainder ones. However, towards lower loads, in a more realistic OBS scenario, all strategies behave equally.

VI. Impact of Burst Reordering on TCP Performance

In this section, we quantify the impact of burst reordering on final TCP throughput. Taking into account the already measured 3-Reordering ratio at the burst layer, we derive a worst case situation for 3-Reordering packet ratio. Then, considering both burst reordering and burst loss pernicious effects, we provide a new figure of merit, called P_{FR}, which quantifies the probability to invoke *fast retransmit* algorithm in TCP Reno. Finally, as the key point of this work, we estimate the theoretical TCP throughput over the scenario under study, which allows us to conclude on its viability.

A. Worst Case Scenario for 3-Reordering Packet Ratio

For the the n-Reordering packet ratio, and according to the definition presented in [7], only the first packet contained in an n-Reordered burst is considered as n-Reordered. Intuitively,

this leads to think that an upper bound for the n-Reordering packet ratio is given when exactly 1 packet of the TCP flow under study is contained in each burst. One should remind, that the n-Reordering packet ratio is measured on a per TCP flow basis. Therefore, only those packets belonging to the TCP flow under study are considered (bursts can be composed of more packets arriving from different TCP sessions). An analysis to verify our hypothesis is proposed below.

Let $P(N_r \geq n_r)$ denote the Complementary Cumulative Distribution Function (CCDF) of a burst to become at least n_r-Reordered. As shown in [14], if there are exactly n_p packets of the same TCP flow in each burst, the probability of a packet to be at least n_r-Reordered is

$$P(N_r^* \geq n_r) = \frac{1}{n_p} P\left(N_r \geq \left\lceil \frac{n_r}{n_p} \right\rceil\right), \quad n_p, n_r \epsilon N. \quad (1)$$

On the other hand, if there is only one packet of the same TCP flow in n_b bursts, the probability of a burst to be at least n_r-Reordered is

$$P(N_r^* \geq n_r) = P(N_r \geq n_b \, n_r), \quad n_b, n_r \epsilon N. \quad (2)$$

Looking at (2), note that due to the non strictly monotone behavior of the CCDF, the upper bound for $P(N_r^* \geq n_r)$ is given when $n_b = 1$ (i.e., exactly one packet per burst) . This agrees with our assumption. Let us focus now on (1). Here, particularizing for $n_p = 1$, we have that $P(N_r^* \geq n_r) = P(N_r \geq n_r)$. Therefore, to ensure the worst case assumption, the following inequation must hold:

$$P(N_r \geq n_r) \geq \frac{1}{n_p} P\left(N_r \geq \left\lceil \frac{n_r}{n_p} \right\rceil\right), \quad n_p, n_r \epsilon N. \quad (3)$$

With demonstration purposes, we have obtained the CCDF of the n-Reordering ratio for each strategy under study. Indeed, we observe how, for $n_r = 3$ and $n_p \epsilon N$, the gathered results accomplish inequation (3). This shows that we truly contemplate the worst case scenario for the 3-Reordering ratio.

B. TCP Throughput Estimation

The aforementioned analysis allows us to estimate a worst case for the final TCP throughput in our scenario. With such purposes, according to the conclusion above, we assume that 1 packet of the TCP flow under study is contained in each burst. Moreover, we consider that, in the case of contention, a burst is dropped. Thus, packet loss probability P_L equals to burst loss probability P_B. In the case that the receiver does not use selective acknowledgements, and the sender uses the basic congestion control presented in [3], reordering has the same effect as packet loss. In fact, reordered packets which exceed the DUP-ACK threshold also trigger the fast retransmit algorithm (i.e., as if they would have been lost). Therefore, the probability to invoke fast retransmit algorithm can be stated as $P_{FR} = P(N_r^* \geq n_r) + P_L$.

In Fig. 7, we depict the upper bound for P_{FR} for a DUP-ACK threshold set to 3. In particular, it is obtained as
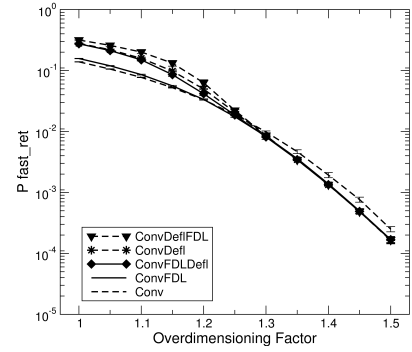

Fig. 7.   Probability to trigger TCP fast retransmit (worst case scenario)

$P_{FR} = P(N_r \geq 3) + P_B$, using the results presented in Fig. 2 and Fig. 6. As seen, for high loads, *Conv* and *ConvFDL* lead to better results, due to the lower reordering they introduce. However, for lower loads, all combined strategies provide similar performance. This is due to the fact, that along this range, 3 reordering ratio dominates in front of P_L. The fact that *Conv* alone provides substantially worse performance, demonstrates the need for additional contention resolution in OBS networks. Up to now, several analysis have been proposed in the literature to model the steady state throughput of a TCP connection. Among them, Mathis *et al.* [15] described the behavior of TCP by a simple model, which considered only the congestion avoidance phase of TCP. Alternatively, Padhye *et al.* [16] developed a more accurate model, considering both congestion avoidance and retransmissions caused by time out. In particular, the authors concluded that the theoretical throughput of a TCP session can be approximated as

$$B_{TCP} \simeq$$
$$\min\left(\frac{W_{\max}}{RTT}, \frac{MSS}{RTT\sqrt{\frac{2bp}{3}} + T_0 \min\left(1, 3\sqrt{\frac{3bp}{8}}\right) p(1+32p^2)} \right) \quad (4)$$

where W_{max} is the maximum receiver advertised window (typically 64 kB), RTT is the round trip time in seconds, MSS is the maximum segment size (typically 1460 bytes in Ethernet), T_0 is the mean duration of a TCP retransmission timeout (typically 1 s, as specified in [17]), b is the number of TCP segments an ACK refers to (typically $b = 2$ if delayed ACK [18]) and, finally, p is the total packet loss probability along the path (P_{FR} since, in this scenario, 3-Reordering has the same effect as packet loss). Fig. 8 illustrates, for different RTT values, the theoretical TCP throughput according to this model. Mainly, it depicts B_{TCP} and the limitation due to the receiver limitation window, both function of p. In this way, given a certain p, the theoretical TCP throughput will be the minimum of both curves.

As mentioned earlier, OBS networks are usually dimensioned to achieve burst loss probabilities ranging from 10^{-3} to 10^{-6}. Looking at Fig. 7, a network dimensioned to achieve these values (as seen in Fig. 2, overdimensioning the network
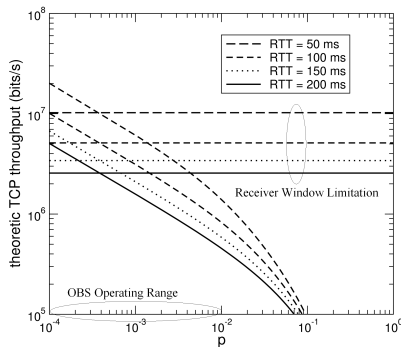
Fig. 8.   Theoretic TCP throughput (bits/s) according to the model proposed in Padhye *et al.* [16]

by 1.25 - 1.35) would experience P_{FR} values from 10^{-2} to 10^{-3}, depending on the strategy used. Observing now Fig. 8, we find that for these p values, the performance of TCP is highly affected by the reordering introduced at the OBS layer. In fact, to assure the proper performance of TCP, p should be lower than 10^{-3}, so that the limiting factor would be the receiver advertised window rather than the reordering introduced in the network. This demonstrates that reordering should be also considered when dimensioning an OBS network for TCP traffic. As seen, its impact on TCP is much more significant than P_L in the range of operation of typical OBS networks. Moreover, as far as TCP performance is concerned, almost all combined contention resolution strategies under study behave similarly. Although we mentioned earlier that *ConvDeflFDL* may outperform the remainder, such improvements are hidden by the fact that 3-Reordering dominates in front of P_L.

VII. Concluding Remarks

In this paper, we propose a layered framework to quantify the impact of burst reordering on TCP performance. First of all, we measure the reordering introduced by several contention resolution strategies. With such purposes in mind, we use the packet reorder metrics proposed by the IETF.

Two different approaches to tackle reordering in an OBS scenario have been highlighted and subsequently evaluated. On the one hand, reordering can be solved directly at the OBS layer, by means of well dimensioned buffers. On the other hand, reordering can be left to higher layers, expecting this one to be solved by them.

For the former strategy, we quantify the size of the buffers which should be placed at OBS edge nodes on a per flow basis. Following this line, we find that deflection routing prohibits this solution, since the introduced extents are extremely high. Conversely, we demonstrate that buffering introduces significantly lower extents, which would, a priori, enable this strategy. For the latter strategy, we focus on its impact on final TCP Reno performance. We propose a new figure of merit, named P_{FR}, which considers not only the pernicious effects from packet loss, but also the ones from caused by reordering. This allows us to conclude, based on the model proposed by Padhye *et al.*, that the usual OBS operating range

fits no more. On the contrary, network should be dimensioned taking into account not only burst loss probability, but also burst reordering introduced by contention resolution.

Further work would address the comparison of both strategies from the network planning point of view. This work would study which solution is more cost effective: the resource overdimensioning required for the proper performance of TCP or the placement of buffers in OBS edge nodes, making in this way reordering transparent to TCP.

References

[1] Y. Chen, C. Qiao, and X. Yu, *Optical Burst Switching: A New Area in Optical Networking Research*, IEEE Network May/June 2004 pp. 16-23.
[2] S. Yao et al., *A Unified Study of Contention-Resolution Schemes in Optical Packet-Switched Networks*, Journal of Lightwave Technology, March 2003.
[3] M. Allman, V. Paxson and W. Stevens, *TCP Congestion Control*, RFC 2581, April 1999.
[4] M. Schlosser, E. Patzak, and P. Gelpke, *Impact of deflection routing on TCP performance in optical burst switching networks*, In Proceedings of the 7th International Conference on Transparent Optical Networks (ICTON), pp. 220-223 vol.1, June 2005.
[5] S. Gowda *et al.*, *Performance evaluation of TCP over optical burst-switched (OBS) WDM networks*, In Proceedings of the IEEE International Conference on Communications (ICC), pp. 1433-1437 vol.2, 2003.
[6] A. Detti and M. Listanti, *Impact of segments aggregation on TCP Reno flows in optical burst switching networks*, In Proceedings of IEEE INFOCOM, 2002.
[7] A. Morton, *et al.*, *Packet Reordering Metrics*, RFC 4737, November 2006.
[8] C. M. Gauger, M. Köhn, and J. Scharf, *Comparison of contention resolution strategies in OBS network scenarios*, In Proceedings of the 6th International Conference on Transparent Optical Networks (ICTON), pp. 18-21 vol.1, June 2004.
[9] C. M. Gauger, *Performance of converter pools for contention resolution in optical burst switching*, In Proceedings of OptiComm, Boston, 2002.
[10] S. Yao, B. Mukherjee, S. Dixit, *Contention Resolution in Optical Packet Switching*, IP-over-WDM, Chapter 5, Wiley, 2003.
[11] S. Maesschalck *et al.*,*Pan-European Optical Transport Networks: An Availability-based Comparison*, Photonic Network Communications, vol. 5, pp. 203-225, 2003.
[12] M. Köhn and C. M. Gauger, *Dimensioning of SDH/WDM Multilayer Networks*, Proceedings of the 4. ITG Symposium on Photonic Networks, pp. 29-33, 2003.
[13] S. Bodamer *et al.*, *IND Simulation Library 2.3 User Guide Part I: Introduction*, September 2004.
[14] S. Gunreben and G. Hu, *A Multi-layer Analysis on Reordering in Optical Burst Switched Networks*, IEEE Communications Letters, vol 11.12, December 2007.
[15] M. Mathis, *et al.*, *The macroscopic behavior of the TCP congestion avoidance algorithm*, ACM Computer Communication Review, vol. 27, no. 3, pp. 67-82, July 1997.
[16] J. Padhye, *et al.*, *Modeling TCP reno performance: A simple model and its empirical validation*, IEEE/ACM Transactions on Networking, no. 2, pp. 133-145, April 2000.
[17] V. Paxson and R. Allman, *Computing TCP's Retransmission Timer*, RFC 2988, November 2000.
[18] R. Braden, *Requirements for Internet Hosts – Communication Layers*, RFC 1122, October 1989.