

A Simulation-Based Study of TCP Performance over an Optical Burst Switched Backbone with 802.11 Access

Isaias Martinez-Yelmo, Ignacio Soto, David Larrabeiti, and Carmen Guerrero

Universidad Carlos III de Madrid, 28911, Leganes (Madrid), Spain
{imyelmo, isoto, dlarra, guerrero}@it.uc3m.es

Abstract. The combined effect of optical and wireless subnetworks in an hypothetical future scenario where core networks have evolved to the still prototype Optical Burst Switching (OBS) technology is an open research issue.

This paper studies this hybrid scenario, in the particular case of 802.11 access, by reviewing the key aspects of OBS and 802.11 with an impact on the performance of TCP, and makes a simulation-based assessment of the relative influence of both technologies over the effective end-to-end behaviour of TCP.

1 Introduction

The exploitation of Dense Wavelength Division Multiplexing (DWDM) as a dynamic Optical Circuit Switching technology (OCS), where lightpaths are optically switched to interconnect the backbone edge nodes, has just started. This technology removes the opto-electronic conversion bottleneck in the core of the network and provides the full capacity of an optical carrier to transport data edge to edge. However, it yields no statistical time multiplexing gain of lightpath capacity and connectivity is constrained by the number of optical carriers provided by DWDM. This limitation will eventually be overcome when Optical Packet Switching (OPS) is technologically feasible. In the meantime, Optical Burst Switching [1] seems to be the key transition technology on the way to Optical Packet Switching (OPS).

On the other hand, wireless LAN (WLAN) access is becoming increasingly popular for personal broadband communications, and is envisaged as a real alternative to the widespread Third Generation (3G) Mobile System. One of the most widely deployed WLAN technologies is the IEEE Standard 802.11b. A lot of work has been devoted to the analysis of TCP/IP performance over 802.11 e.g. [2], [3]. However, up to the date, no study on the relative impact of OBS and 802.11 together on TCP has been realised in a scenario like the one shown in Fig.1. This scenario, interworking with electronically packet-switched subnetworks - assumed over-provisioned in our study to focus on the relative effects of WLAN and OBS -, might become a reality in the next future. A comparative study is hence relevant because the behaviour of both technologies is quite different with respect to delay and packet loss.

The structure of the paper is the following. Section 2 is a short review of factors driving the performance of TCP. Sections 3 and 4 are devoted to outline the effects of 802.11b and OBS over TCP respectively. Then, section 5 presents a simulation-based analysis of TCP performance on a Wireless LAN access network connected to an OBS core network. Finally, section 6 draws the conclusions.

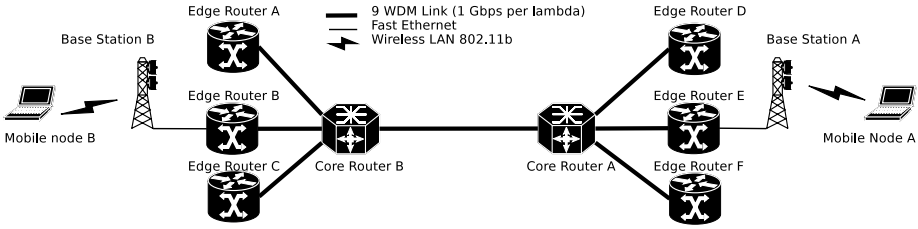


Fig. 1. OBS - WLAN simulation scenario

2 TCP Congestion Control

Basically, TCP has a byte-oriented sliding window mechanism to dynamically adapt the data rate to the network state. An End-to-End mechanism exists to avoid congestion in the network where only the transmitter and the receiver take part. The Congestion Control mechanism of TCP is ruled by several parameters [2]:

- **The lower window edge ($W_{min}(t)$):** All data numbered up to and including $W_{min}(t) - 1$ has been transmitted and ACKed. $W_{min}(t)$ is non-decreasing variable; the receipt of an ACK with sequence number $n > W_{min}(t)$ causes $W_{min}(t)$ to jump to n .
- **Congestion Window ($W(t)$):** The sender can send packets with a sequence number n where $W_{min}(t) \leq n < W_{min}(t) + W(t)$ where $W(t) \leq W_{max}$ where W_{max} is the maximum transmission window.
- **Slow Start threshold ($W_{th}(t)$):** This value serves to set an inflexion point in the growth of the congestion window from geometric to lineal (congestion avoidance region).

The behaviour of TCP congestion control defines two working regions:

- **Slow Start:** If $W < W_{th}(t)$, each first ACK implies that $W(t+1) = W(t) + 1$
- **Congestion Avoidance:** If $W \geq W_{th}(t)$, each first ACK implies that $W(t+1) = W(t) + 1/W(t)$.

TCP estimates with each first ACK received the maximum RTT of the ongoing connection (called Retransmission TimeOut - RTO) and with this value a timer is started. Depending on the version of TCP running on the transmitter, the congestion window W and the Slow Start threshold W_{th} are modified in a different way if the RTO time expires or duplicated ACKs are received. For a detailed explanation see [4], [5] and [6]. Depending on the version of TCP, different performance will be obtained in a particular scenario.

3 TCP on Wireless LAN 802.11

Most commercial Wireless LAN products are based on the IEEE 802.11 standards.

For the purpose of this paper, it is necessary to understand the basic behaviour of TCP over an 802.11b network. In-depth performance studies can be found in [7], [8]. The performance of TCP is severely hindered by packet losses due to interferences,

noise or collisions. This well-known fact is due to the assumptions made in the design of TCP, where an RTO timer expiration or the reception of duplicated ACKs are interpreted as network congestion. In wired scenarios the transmission error probability is very low, and packet loss is solely attributed to network congestion. However, in wireless technologies, losses are usually due to interferences and noise, and delay is not necessarily due to link congestion but to retransmissions following a frame error; thus, TCP reduces the transmission rate erroneously.

A good analysis of the performance of different versions of TCP over lossy links is provided in [3] and [2]. This latter work shows that the most recent implementations of TCP perform better on lossy links because the **congestion window** is not decreased drastically under packet loss. Other works focus on performance improvement [9], [10].

4 Optical Burst Switching

4.1 Overview

Next generation core networks are planned to be based on DWDM technology using Optical Circuit Switching featuring real lightwave switching. The problem comes in large core networks, where the number of available lightpaths becomes insufficient to build fully-meshed OCS networks. Statistical time multiplexing is hence required to achieve scalability and efficiency, either by means of optical traffic grooming or by Optical Burst Switching (OBS) [1].

The idea of OBS is achieving statistical time multiplexing at the optical layer but taking into account the current physical limits of the optical switching technology. Current optical switching laboratory prototypes take a few milliseconds to change the state of a light beam switch. This speed is insufficient to switch packets optically at Terabits per second. OBS is a practical trade-off between OCS and OPS whereby statistical multiplexing is obtained with a coarser granularity. The proposal in OBS is to use larger-than-packets switching units to keep up efficiency; this way most of the time is spent on transmitting data and less on changing the switch state. These large switching units are called "*bursts*": a group of packets that leave the OBS network at the same egress router; all the packets are encapsulated in the ingress router and the OBS burst is configured with a label to the destination address of the desired egress router. Finally, the egress router takes the received OBS burst, decapsulates the packets and forwards the packets electronically according to the routing table.

Since switching time is long in current optical switches, the transmission of a burst is preceded by a control packet named *Burst Header Packet* (BHP) whose purpose is reserving resources for the transmission of the burst through the OBS switches. This control packet is sent over a signalling channel to be electronically processed at each hop. If due to existing traffic, the burst cannot be scheduled to be transmitted in a free slot in one of the output wavelengths, the OBS burst is dropped. At each node a heuristic scheduling algorithm is used to allocate the necessary resources for the bursts to reach the next hop, [11].

Delay and jitter can be controlled as follows. Packets are buffered according to their egress router and they are transmitted all together in a burst. The transmission of a burst could be realised according two configuration configuration parameters:

- *Burst Threshold*: Maximum number of packets that can be transmitted in a burst. Thus, if this value is reached, the corresponding burst should be transmitted immediately. If this value is reached earlier, the waiting queue time will be smaller, and the delay will be also smaller.
- *Burst Timeout*: A burst will be transmitted if this timeout expires although the burst threshold has not been reached. Thus, this timeout guarantees a maximum value of delay introduced by the OBS network. A longer burst timeout implies less average delay but less network utilisation efficiency.

4.2 Key Issues on TCP over OBS

There are several works that analyse the performance of TCP in OBS networks, for instance [12]. This performance is closely related with the Retransmission timeout (RTO) of TCP and how the different versions of TCP treat this event [13]. Nevertheless, several qualitative general considerations can be made to understand the way OBS affects TCP.

The rate of a TCP connection can be approximated to the value of the TCP transmission Window Size divided by the Round Trip Time (RTT). The RTT of a connection crossing an OBS network is related to the value assigned to the Burst Size and Burst Timeout parameters. The RTT decreases with the value of these parameters, since it takes less time to send the OBS burst. The Burst size is the dominant parameter for a heavily loaded destination egress node because OBS bursts are sent as soon as the Burst size is reached. On the other hand, the Burst Timeout is the dominant parameter in a link with low load of traffic due to the fact that data is delivered only when the Burst Timeout timer expires.

Furthermore, it must be taken into consideration how the OBS network can affect the TCP Window Size. The size of the TCP Window is set according to different events depending on the TCP version used, upon two events: RTO and Duplicated ACKs, both events produced by packet losses. Losses of packets in an OBS network are due to OBS burst drops when these cannot be allocated in a free transmission time. This blocking probability depends on the available wavelengths for data and signalling, on the patterns of the incoming traffic and on the Burst Timeout and Burst Size parameters; but it is not easy to measure how all these parameters can affect the Window Size and throughput without simulation [12].

5 TCP Performance in an 802.11b Access Network Connected to an OBS Core Network

To the date, no previous study related with the relative impact of OBS and WLAN on TCP performance has been realised. Both technologies drop packets for different reasons: both for high load in the network and, in the case of WLAN, also for bad link conditions. A simulation study of the effects of both technologies is a first step to understand the interaction, eventual combined effects and overall performance achieved.

5.1 NS-2 Configuration

The simulation tool used to simulate the designed scenario is the NS-2 simulator [14] with an extension to simulate OBS networks used in [12] featuring the burst scheduler presented in [11]. The propagation model used in the wireless link is the *shadowing model* [15]. This model allows to approximate the Error Probability of a Wireless LAN link in different indoor and outdoor scenarios. Since the purpose of the simulation is comparing the relative effects of the WLAN and OBS losses together on TCP, the wired access network between the WLAN network and OBS core is assumed to be over-provisioned. Thus, the scenario under study is shown in Fig.1. Mobile nodes have 802.11b interfaces at 11 Mbps whereas the OBS network has nine wavelengths between each node working at 1 Gbps, two of them dedicated for signalling. WLAN losses happen on the link between base station A and mobile node A.

The methodology used is very similar to the one applied in [12]. Since the simulation of many background individual simultaneous connections has a high cost, its effect in the core has been modelled by means of self-similar background traffic, considering that Internet traffic has been widely characterised with this distribution probability [16]. The parameters used for the self-similar traffic have been 10000 batches per second, a Hurst exponent for arrival of 0.5 and a Hurst exponent size of 0.5. The values of the Hurst parameters has been set to be conservative. Different batch sizes have been used to increase the background traffic and OBS blocking probability (OBS Pb). These patterns have been configured between each edge router in the proposed scenario.

The *Burst timeout* has been set to 0.01 seconds and the *Burst size* has been configured to 70,000 bytes. Furthermore, 5 Fiber Delay Lines (FDLs) of 0.1 ms have been placed in each core node to get better link utilisation for a given burst drop probability. Finally, the TCP NewReno variant has been used for the simulation as this is the most widely used in Internet [17]. Furthermore, this variant has proved to give the best performance on lossy links [2].

5.2 TCP Performance Under 802.11b and OBS Networks

This section presents the results obtained in the simulation. Fig.2 displays a three dimensional plot of the overall throughput with respect to the packet drop probabilities in the Wireless LAN and OBS networks. The dropping probabilities in the OBS network have been obtained by configuring different batch sizes of the self-similar sources to obtain different background traffics between each edge router. It can be seen how the throughput decreases exponentially as the WLAN dropping probability grows. More interestingly, the graph shows that OBS has a minor influence on the overall throughput in a range of dropping probability values (e.g. [0..0.01]) where WLAN by itself brings the performance down to a 50%.

Several cuts on the three dimensional figure have been made and are represented in Fig.3 on logarithmic scale. The plots include the confidence intervals for a confidence of 95%. The comparison of plots A and B gives evidence of the different response to losses in each subnetwork. Plot B shows the way the WLAN determines the maximum overall throughput that can be reached for a wider range of OBS drop probabilities, always modulated by the state of the OBS network.

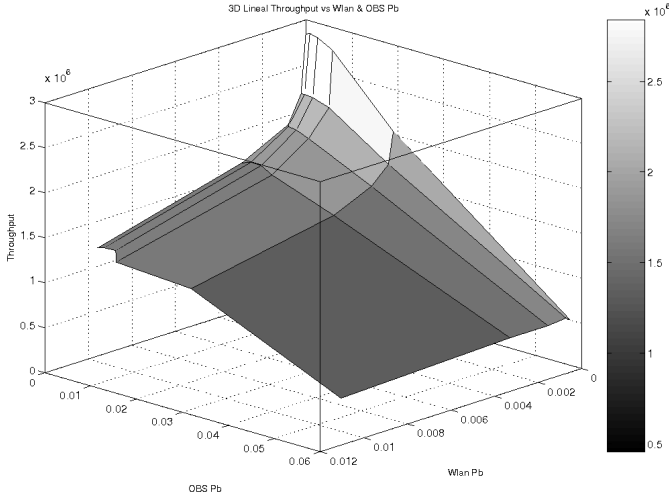


Fig. 2. 3D Mean Throughput Graphic

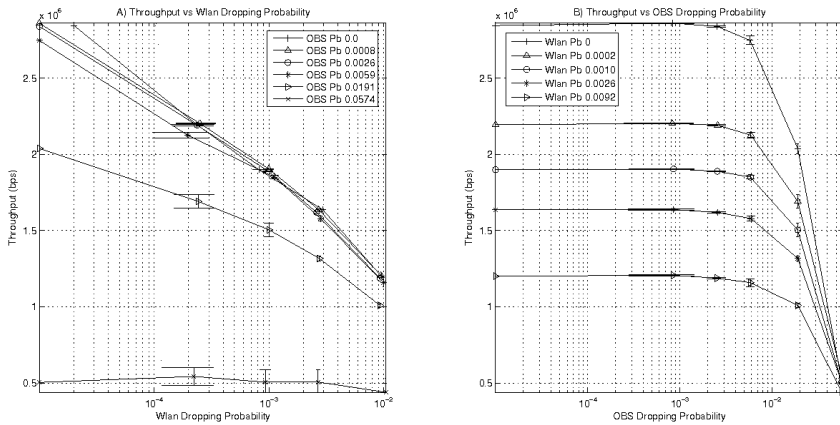


Fig. 3. Mean Throughput Graphic

At first sight, the expected behaviour would be a monotonously decreasing curve with the OBS packet dropping probability. However, when the OBS dropping probability increases, the throughput does not decrease in a significant way before the OBS packet dropping probability is over a given value. Fig.3 shows that this performance-fall value is $5 \cdot 10^{-3}$ in our sample case. This makes WLAN loss dominate over OBS loss for a long range of loss probabilities. RTT explains this behaviour

Fig.4 illustrates the way the RTT grows as the WLAN packet dropping probability does, but this does not happen with the OBS packet dropping probability. This is due to the fact that the OBS dropping probability increases with background traffic

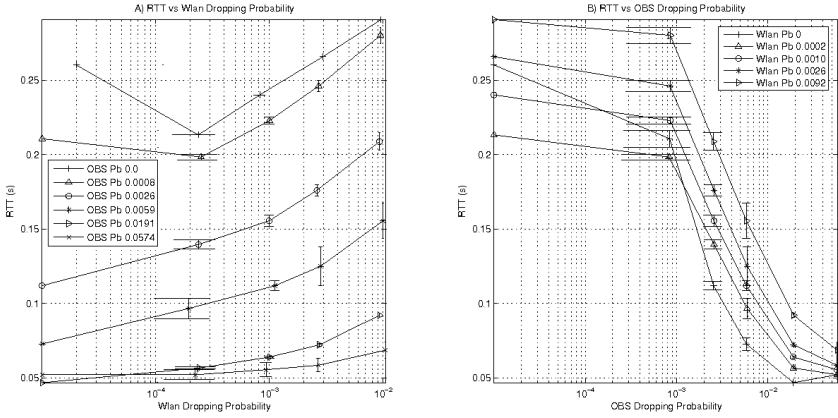


Fig. 4. Mean RTT Graphic in the OBS scenario

load. Consequently, the maximum burst size is reached earlier on average and the RTT decreases. Dropped packets cause a reduction of the Window of TCP, but in OBS this effect is somewhat compensated by a reduction on RTT. Considering that in TCP $Throughput \approx WindowSize / RTT$, the reduction of the RTT softens the effect of the decrement of Window Size and the performance is sustained until the OBS dropping probability is larger than $5 \cdot 10^{-3}$.

6 Conclusions

The presented simulation-based study gives a first understanding of the factors that rule TCP behaviour in a scenario with an OBS core and 802.11 access network. The results show that the throughput achieved, whose maximum is determined by the wireless segment bottleneck, depends heavily on the WLAN dropping probability. The throughput decay is made heavier by the RTT increase caused by retransmissions. The packet loss probability of OBS could approximately be added to this value, given the low drop probabilities considered and the fact that both losses can be assumed independent.

The simulation confirms this fact and provides additional information on how the throughput behaves for the same drop probability at each segment: in a wide range of values, the throughput is dominated by WLAN loss. In the case of WLAN, the RTT grows because of Layer-2 retransmissions, whereas in OBS high drop probabilities are due to high loads, and in these circumstances the RTT shrinks because bursts are sent as soon as the maximum burst size is reached at the edge routers.

As a general consequence, it can be stated that access and core networks can be designed in a largely independent way as long as TCP performance is concerned; the configuration parameters can be selected separately. Furthermore, since TCP is more sensitive to WLAN loss than to OBS loss, more emphasis should be set on improving communication reliability in the radio access than on the core, provided that the OBS network is designed to work far from the throughput-fall point for the carried traffic.

Acknowledgements

This work has been partly supported by the EU under the IST e-Photon/One+ project (FP6-IST-027497) and by the Spanish CAPITAL project (TEC2004-05622-C04-03).

References

1. Chen, Y., Qiao, C., Yu, X.: Optical burst switching: a new area in optical networking research. *Network*, IEEE 18(3), 16–23 (2004)
2. Kumar, A.: Comparative performance analysis of versions of TCP in a local network with a lossy link. *IEEE/ACM Transactions on Networking* 6(4), 485–498 (1998)
3. Xylomenos, G., Polyzos, G.C.: TCP and UDP performance over a wireless LAN. *INFOCOM 2*, 439–446 (1999)
4. Postel, J.: Transmission Control Protocol. RFC 793 (Standard) (September 1981)
5. Allman, M., Paxson, V., Stevens, W.: TCP Congestion Control. RFC 2581 (April 1999)
6. Allman, M., Floyd, S., Partridge, C.: Increasing TCP's Initial Window. RFC 3390 (October 2002)
7. Banchs, A., Perez, X.: Providing throughput guarantees in IEEE 802.11 wireless LAN. In: *IEEE Wireless Communications and Networking Conference, 2002 (WCNC 2002)*, March 2002, pp. 130–138. IEEE Computer Society Press, Los Alamitos (2002)
8. Banchs, A., Perez, X.: Assured and expedited forwarding extensions for IEEE 802.11 Wireless LAN. In: *Quality of Service. Tenth IEEE International Workshop*, May 2002, vol. 237–246, pp. 237–246. IEEE Computer Society Press, Los Alamitos (2002)
9. Spatscheck, O., Hansen, J.S., Hartman, J.H., Peterson, L.L.: Optimizing TCP forwarder performance. *IEEE/ACM Transactions on Networking* 8(2), 146–157 (2000)
10. Bhandarkar, S., Sadry, N., Reedy, A., Vaidya, N.: TCP-DCR: A novel protocol for tolerating wireless channel errors. *IEEE Transactions on Mobile Computing* 4(5) (October 2005)
11. Xiong, Y., Vandenhoute, M., Cankaya, H.C.: Control architecture in optical burst-switched WDM networks. *Selected Areas in Communications, IEEE Journal on* 18(10), 1838–1851 (2000)
12. Gowda, S., Shenai, R., Sivalingam, K., Cankaya, H.: Performance evaluation of TCP over optical burst-switched (OBS) wdm networks. In: *Communications, ICC '03, IEEE International Conference*, May 2003, pp. 11–15. IEEE Computer Society Press, Los Alamitos (2003)
13. Yu, X., Liu, C.Q.Y.: TCP implementations and false time out detection in obs networks. *INFOCOM 2004*, 774–784 (March 2004)
14. Ns-2 network simulator: <http://www.isi.edu/nsnam/ns/>
15. Rappaport, T.S.: *Wireless communications, principles and practice*. Prentice-Hall, Englewood Cliffs (1996)
16. Chakraborty, D., Ashir, A., Suganuma, T., Keeni, G.M., Roy, T.K., Shiratori, N.: Self-similar and fractal nature of internet traffic. *Int. J. Netw. Manag.* 14(2), 119–129 (2004)
17. Medina, A., Allman, M., Floyd, S.: Measuring the evolution of transport protocols in the Internet. *SIGCOMM Comput. Commun. Rev.* 35(2), 37–52 (2005)