

On linear prediction of Internet traffic for packet and burst switching networks

D. Morató, J. Aracil, L. A. Díez, M. Izal, E. Magaña

Dept. Automatics and Computer Science

Public University of Navarra

Campus Arrosadía

31006 Pamplona, SPAIN

Tel: +34 948 16 98 54

Fax: +34 948 16 89 24

email: {daniel.morato, javier.aracil, luisangel.diez, mikel.izal, eduardo.magana}@unavarra.es

Abstract—In this paper, we show that prediction algorithms in the least mean square error sense prove better in a burst rather than in a packet switching network. For the latter, further information about the packet arrival distribution within the prediction interval is required. Regarding burst switching, we compare Optical Burst Switching networks with and without linear prediction to conclude that linear prediction provides a significant improvement in end-to-end latency with low bandwidth waste.

Keywords— Internet performance analysis, traffic prediction, long-range dependence, short-range dependence, traffic measurement.

I. INTRODUCTION

Internet traffic presents long range dependence, non-stationarity and multifractal scaling in short timescales [1], [2]. While *a-priori* link dimensioning is difficult to achieve in practice, the traffic correlation can be exploited to provide dynamic bandwidth allocation. To do so, a traffic prediction method is required, so that resources can be provided to the traffic demand beforehand. We distinguish two fundamental issues in traffic prediction: first the prediction method and secondly the prediction interval. For a traffic forecasting to be effective the prediction method is required to be simple, so that the processing time in estimating the future traffic does not pose a heavy burden in the network interconnection elements. On the other hand, the prediction interval refers to how far into the future can the traffic be predicted with confidence. Such prediction interval is lower bounded by the roundtrip time to the bandwidth allocator and is defined in this paper as τ , being $X(t)$ the random process that gives the number of bytes produced by the traffic source in the interval $[0, t)$ and $\hat{X}(t + \tau)$ an estimate of the process at time $t + \tau$.

The choice of a prediction method is a tradeoff between the prediction interval, prediction error and computational cost. For video and network data traffic, linear prediction methods have been considered in the literature as a simple and effective alternative [3], [4], [5]. Let $Z[n + 1] = X((n + 1)\tau) - X(n\tau) = X(t + \tau) - X(t)$, $n = 0, 1, \dots, N - 1$ with $t = n\tau$. The linear predictor that minimizes the variance of the prediction error $Z[n + 1] - \hat{Z}[n + 1]$ is given by [6, chapter 1]

$$\begin{aligned} \hat{Z}^h[n + 1] &= \hat{X}^h(t + \tau) - X(t) = & (1) \\ &= \alpha_1 Z[n] + \alpha_2 Z[n - 1] + \dots + \alpha_h Z[n - h + 1] = \bar{\alpha}' \mathbf{Z}_n \end{aligned}$$

where h is the predictor length and $\bar{\alpha}$ is a column vector with values α_i ($1 \leq i \leq h$) which are the mean squared linear regression coefficients. Such coefficients can be obtained from the following equation in matrix form

$$\bar{\rho} = \Gamma_k \bar{\alpha} \quad (2)$$

where $\bar{\rho}$ is a vector with values of the autocorrelation function and Γ_k is the correlation matrix [6, chapter 1]. We note that there are a number of algorithms to calculate the predictor in 1 in real time [5], namely taking into account the traffic process history up to time t solely. Thus, the predictor can be used not only for stored video but also for live Internet traffic.

It has been shown [3], [5] that the linear predictor defined in 1 proves adequate for prediction of VBR video [7]. Most interestingly, even though VBR video shows long range dependence, only the short-scale correlations determine the prediction error, with the long-range correlations having a slight incremental value [3]. Such findings are consistent with recent contributions in network performance analysis, reporting that for realistic network scenarios and reasonable target cell loss rates and buffer sizes the correlation structure of the input traffic is important only below a certain scale [8], [9].

In this paper, we further examine the practical suitability of least mean square error predictors in two network scenarios: a link with source rate advertisements (ABR or RCBR-like [10]) and a burst-switching network (optical burst switching [11]). The main contributions of the paper are the following: we show that while in a pure packet switching scenario a traffic predictor with *no prediction error* provides significant queueing delay, the same traffic predictor in a burst switching scenario provides a queueing delay which is close to the theoretical minimum. For the latter, we define the range of values of roundtrip time and burst packetization delay which provide

end-to-end delay gains with the estimator defined in 1. We conclude that linear prediction is suitable to a case demanding *coarse* prediction (packets per interval) only. However, the high variability of Internet traffic in short timescales may require the use of a traffic estimator that takes into account the packet arrival distribution within the interval.

The paper is organized as follows: in section II we define the network scenario and methodology. Section III is devoted to results of both packet and burst switching scenarios and section IV presents the discussion. Finally, section V presents the conclusions that can be drawn from this paper and future work.

II. NETWORK SCENARIO AND METHODOLOGY

Our network scenario is depicted in fig. 1, which shows a generic network setup of an input source equipped with the traffic estimator in 1 and a rate shaper, which is adjusted according to the traffic prediction and queueing backlog at the beginning of the prediction interval. Thus, the input traffic is served with a piecewise constant rate function. On the other hand, such drain rate is reported to a bandwidth allocation unit in the network. In an ATM network, sources can use the Resource Management (RM) cell or, in the Integrated Services Internet, the RSVP PATH message in order to report rate information, as proposed in [10]. In an OBS network, the burst reservation message provides a report of the incoming burst length in a tell-and-go fashion along the path from source to destination. We assume that the network resource comprises a tight control loop between the source and allocator, i.e. the roundtrip time to the allocator is negligible in comparison to the length of the prediction interval. Furthermore, we will assume that the bandwidth allocator provides a bandwidth which is equal to the source advertised rate since our aim is to study the prediction performance and not the scheduling algorithm efficiency in sharing bandwidth among a number of sources. At time $t = n\tau$ the traffic source provides an estimate of the future traffic $Z[n+1]$ in the prediction interval $n+1$ of length τ , together with the queueing backlog $Q[n]$, which is known at the beginning of interval $n+1$ and does not need to be estimated. As a result, the traffic source provides an estimate of the unfinished work at the beginning of the prediction interval. We note that due to the prediction error a queue builds up at the source. An event-driven simulator is developed in C language, consisting of a trace-driven traffic generator and the shaper shown in fig. 1. Incoming packets are marked with the arrival timestamp and the number of packets in queue is updated upon arrival or departure of a packet. At packet departure times the waiting time per packet is obtained and the average waiting time is evaluated at the end of each simulation run.

A. Input traffic

The input traffic is obtained from the ATM Permanent Virtual Circuit (PVC) that links Public University of Navarra (ap-

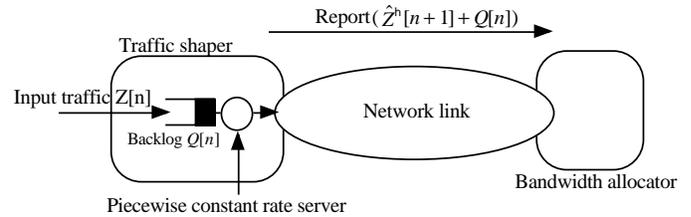


Fig. 1. Network model

prox. 1,500 hosts) to the core router of the Spanish academic network (*RedIris*¹) in Madrid. In order not to introduce artificial correlations in the input traffic we select the measurement intervals such that the utilization factor is always below 0.9 for more than 98% of the time. Note that short-term correlations due to traffic bursts at the link peak rate provide optimistic estimates and the generality of the results becomes jeopardized. Besides, since we wish to study the predictor performance as a function of the multiplexing level we choose to take a set of traffic traces of finite duration N , which are interleaved to produce the different levels of multiplex. By doing so, we mimic the case of an statistical multiplex of several independent sources. Each source is modeled by a trace, noting that our set of traces represents a set of uncorrelated instances of traffic from our 1,500 hosts. In order for this approximation to be valid we carefully check the traces cross-correlation to assess that the traces are indeed independent one another. This point is specially critical with long range dependence traffic, since dependence may be present at long time scales. As a consequence of this sanity check we obtain 32 independent traces of around 10,000 packets each.

Finally, the regression coefficients (α_i) in 1 are obtained per trace *beforehand* and, then, the corresponding simulation run is performed. In a real case, such coefficients are obtained from a training trace and then applied to live traffic, possibly requiring updates as time elapses. Thus, our coefficients represent a best case since they are obtained from the very same trace with which the simulator is fed. In doing so, we focus on the prediction performance of a best case linear predictor and we intentionally leave the problem of adaptation of regression coefficients for future research.

B. Comparison to other traces

Fig. 2 shows the estimated autocorrelation of Internet traffic and the ubiquitous *Star Wars* trace, used in [3], [10]. Regarding the former, we show the mean autocorrelation together with the confidence intervals resulting from calculation of the autocorrelation function of 32 independent traces. Interestingly, the results show that the short-term correlations in the video trace are significantly larger. The reduced variability of video traffic at short timescales may be due to the fact that within the same

¹<http://www.rediris.es>

scene video traffic is smoother than the Internet counterpart at the scene timescale.

On the other hand, we note that the short timescale features of the VBR video are radically different from those of Internet traffic, since the latter presents multifractal scaling and the former shows a traffic arrival process which is determined by the GOP pattern and video server scheduling algorithms. Beyond a certain timescale, the counting process of number of bytes per interval is asymptotically second-order self-similar for both kinds of traffic and therefore Internet and video traces present similar statistical features at large timescales. Indeed, fig. 2 shows that the autocorrelation function of video traffic is close to the Internet traffic counterpart at large lags. However, at short timescales, the packet arrival process is different. In conclusion, our case study represents a trace-driven analysis which differs from previously analyzed cases and accurately portrays a real scenario of a large number of users accessing the Internet.

III. RESULTS

In this section we analyze packet and burst switching networks with sources incorporating the following traffic predictors:

- *Best mean square error linear predictor*: This is the predictor in 1, which provides a minimum on the square error and has been adopted in other papers [3], [5], [4]. While this is the best *linear* predictor we note that *non-linear* predictors, for instance wavelet-based [3], which are tailored to the specific case of long-range dependent processes, can provide further improvements in prediction error.
- *Ideal predictor* (no estimation error): This is a reference predictor which provides the exact number of bytes coming in the next prediction interval. Clearly, the predictor is infeasible but represents a reference model in order to study to which extent network performance is affected if we further reduce the prediction error.

The linear predictor length h is set to 5 samples $Z[n]$, while the prediction interval takes a range of values in the

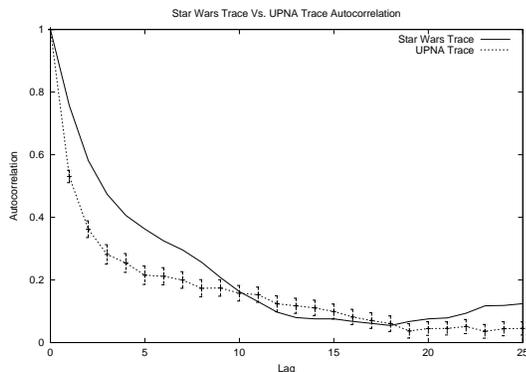


Fig. 2. Autocorrelation for video and data traffic

different simulation runs. From [3] we note that larger prediction lengths do not improve the prediction performance significantly. We calculate the mean square error $e[h] = E[(Z[n+1] - \hat{Z}^h[n+1])^2]$ and the relative improvement $|e[h+1] - e[h]| / e[1]$ which is obtained by making h larger. The results are shown in fig. 3, one curve per trace in our set of 32 traces. We observe that beyond a certain prediction length the predictor performance remains nearly the same. Thus, we choose a value of h equal to 5 in the remaining of the paper.

A. Packet switching

Dynamic bandwidth allocation methods have attracted considerable research attention due to the numerous drawbacks of *a priori* bandwidth allocation for the highly variable Internet and video traffic. We consider the class of piecewise constant dynamic bandwidth allocation methods, such as RCBR [10], which allow bandwidth renegotiation every τ seconds. Assuming that an estimate of the traffic $\hat{Z}^h[n+1] = \hat{X}^h(t + \tau) - X(t)$, $t = n\tau$ is available at time t the bandwidth allocator provides a total bandwidth of $(\hat{Z}^h[n+1] + Q[n]) / \tau$, $Q[n]$ being the queueing backlog at time t . We consider the linear predictor in 1 and, on the other hand, an ideal predictor with no error, namely, a predictor that makes $\hat{Z}^h[n+1]$ equal to $Z[n+1]$. The average waiting time (queueing plus transmission) is shown in fig. 4, as a function of τ , for a single source (approx. 2 Mbps) and a multiplex of 10 (approx. 20 Mbps) and 20 sources (approx. 40 Mbps).

A significant performance gain is obtained with multiplexing level and shorter prediction intervals, consistent with the theory developed in [4]. However, there is a lower bound for the queueing delay *even if the estimator is ideal*. Such lower bound is due to the distribution of packet arrivals within the prediction interval and decreases with the multiplexing level.

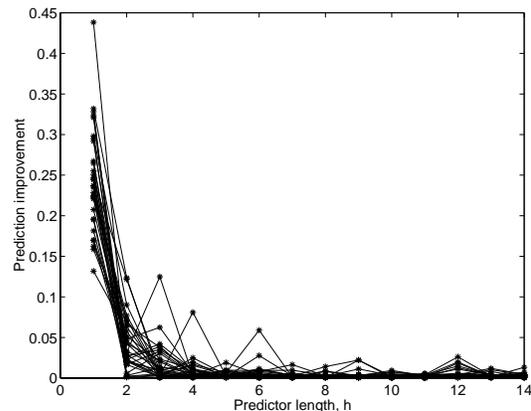


Fig. 3. Relative prediction error

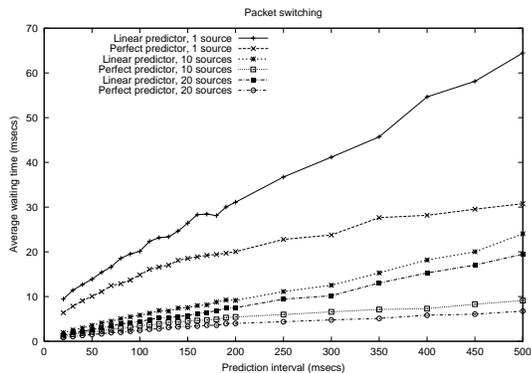


Fig. 4. Average waiting time (packet switching)

B. Burst switching

A number of network transfer modes are based on the principle of burst switching. Such transfer modes provide users with “bursts” so that a chunk of data from the source is transmitted in a single burst. Optical Burst Switching (OBS) [11], for example, performs grouping of several IP packets in a single burst, which can be handled more efficiently than separate packets in the optical network timescale. Thus, OBS provides “coarse packet switching” service in the optical network, namely a transfer mode which is halfway between circuit switching and pure packet switching. First, a resource reservation message is sent along the path from source to destination, and resources are reserved “on the fly” for the incoming burst. Then, the data burst follows. As a result, OBS does not incur the overhead of circuit setup but still resource reservation is performed on a per burst fashion, thus providing enhanced capabilities for QoS discrimination beyond pure packet switching. However, we note that there is a packetization delay, due to the burst assembly time, which is inherent to burst switching, regardless of the prediction. Packets accumulate in a buffer until a timeout expires, and, at that time, the burst is assembled and released to the network [11]. We propose that upon arrival of the first packet in a burst at time $t = n\tau$ a timeout of τ seconds, the prediction interval length, is started. At time t a signalling message with a resource reservation request of $\hat{Z}[n+1] + Q[n]$ is issued by the source and at time $t + \tau$ a data burst of $\hat{Z}[n+1] + Q[n]$ bytes is delivered to the network (see fig. 7(left)). If $Z[n+1] > \hat{Z}[n+1]$ then $Z[n+1] - \hat{Z}[n+1]$ bytes are held in queue. If otherwise $Z[n+1] < \hat{Z}[n+1]$ the burst carries $\hat{Z}[n+1] - Z[n+1]$ padding bytes.

Fig. 5 (top) shows the packet waiting time for the linear predictor and an ideal predictor with no error for a single source, 10 and 20 sources, versus the timeout τ (or prediction interval). For clarity, we also include fig. 5 (bottom), which shows queuing delay only (packetization delay is subtracted from the total delay). Fig. 5 (bottom) also shows the analytical approximation in 6, that will be detailed later. An ideal pre-

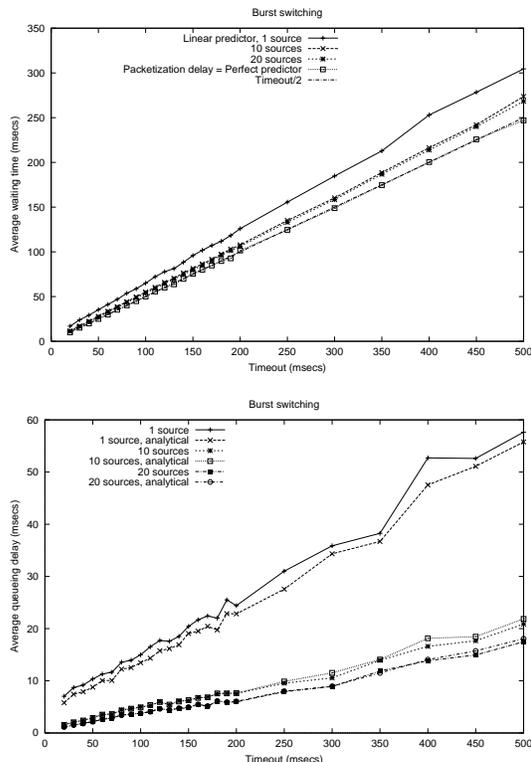


Fig. 5. Average waiting time (OBS): Total (top) and queuing delay only (bottom)

dictor provides optimal performance since the waiting time is equal to the packetization delay, approximately $\tau/2$. We note that *only the number of packets matters and not the particular distribution of packet arrivals within the prediction interval*. On the other hand, a linear predictor provides a queuing delay which is close to the lower bound of packetization delay, specially for small timeout values. In next section we analyze OBS performance with and without prediction. The results show significant performance gains if linear prediction is used.

IV. DISCUSSION

In order to explain the results in the previous section we first show in fig. 6 the relative underestimation error

$$p = \frac{\sum_{i=1}^N Q[i]}{\sum_{j=1}^N Z[j]} \quad (3)$$

as a function of τ , for a single source and a multiplex of 10 and 20 sources. The variable p in 3 represents the percentage of packets which are left for transmission in next interval. In this section, we relate the waiting time in the system with such underestimation error.

From fig. 6 we note that the prediction error decreases with the number of sources, due to the traffic smoothing provided by statistical multiplexing of independent sources. Regarding the

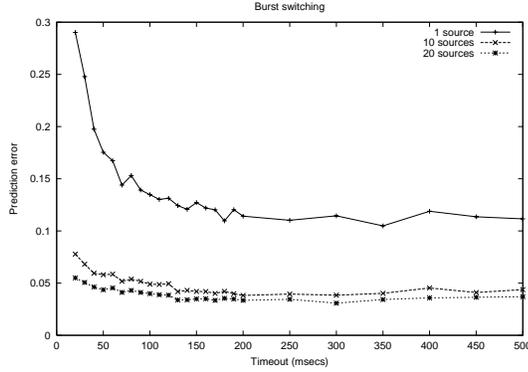


Fig. 6. Underestimation error

packet switching case, we note that even though the prediction performance improves with the multiplex level there is still a lower bound for the queuing delay. The gap between linear and ideal predictor in fig. 4 is due to prediction error and may be shortened with the use of more accurate (non-linear) predictors that further decrease prediction error. Nevertheless, the lower bound set by the ideal predictor is due to the distribution of packet arrivals within the prediction interval. In short timescales, Internet traffic is known to have multifractal scaling [2] and possibly a non-Gaussian marginal distribution with high variability [12]. Such variability affects queuing performance even if we can determine the number of bytes in the prediction interval with no error. For simplicity, consider that packets are constant length and that exactly c packets arrive in interval $[t, t + \tau)$, so that we assign a total bandwidth of c/τ packets per second with an ideal predictor. In a worst case, all packets arrive at the beginning of the interval, namely at time t , and the waiting time W_i for packet i , with $i = 1 \dots n$ is given by

$$W_i = iT \quad (4)$$

where T is the service time of a packet. The average waiting time can be easily derived as follows

$$W = \sum_{i=1}^c \frac{1}{c} W_i = \frac{(c+1)T}{2} \approx \frac{cT}{2} = \frac{\tau}{2} \quad (5)$$

and we note that the average waiting time is equal to half the prediction interval. On the contrary, the best case is obtained when packet i arrives right after packet $i - 1$ has been transmitted and the waiting time is equal to zero. Therefore, the waiting time in a prediction interval is lower bounded by zero and upper bounded by $\tau/2$ in packet switching networks with ideal prediction.

Regarding the burst switching case the average queuing time can be derived taking into account that packets arriving in the interval $[t, t + \tau)$ which are not transmitted in the burst

departing at time $t + \tau$ will be transmitted in the burst departing at time $t + 2\tau$. On the other hand, the packetization delay is approximately equal to $\tau/2$. Such approximation has been validated by simulation, as shown in fig. 5 (top). Thus, the average waiting is given by

$$W = \frac{\tau}{2} + p\tau \quad (6)$$

The curve corresponding to equation 6 (subtracting the packetization delay $\tau/2$) is also shown in fig. 5 (bottom), showing good agreement between theoretical and simulation results. We note that any improvement in the prediction error p implies a linear decrease of queuing delay. Furthermore, an ideal predictor with $p = 0$ provides a queuing delay which is equal to the packetization delay only. Clearly, the distribution of packet arrivals within the prediction interval does not matter in this case. Thus, since linear predictors perform reasonably well for the burst switching case, we propose the use of linear predictors in optical burst switching networks as shown in fig. 7 (right).

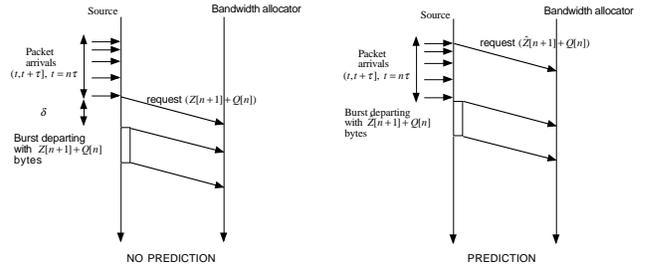


Fig. 7. OBS with and without prediction

Fig. 7 shows a resource reservation protocol for burst switching with and without our linear prediction proposal. While the protocol model is rather simple we note that the analysis performed in this section is not restricted to a particular OBS variant but portrays a generic case of OBS protocol. For the latter, a signalling message is released by the source *once the burst has already been assembled*. An offset time δ between the release of such message and burst transmission is required so that the switches along the path reserve resources for the incoming burst. Thus, the queuing delay is given by:

$$W = \frac{\tau}{2} + \delta \quad (7)$$

If linear prediction is adopted the signalling message can be released *before* the burst has been assembled and $\delta = 0$, as shown in fig. 7 (right). Assuming that $\delta = RTT/2$ where RTT is the roundtrip time to the destination we may calculate the waiting times with and without prediction (6 and 7) versus the timeout values for several RTT values. We subtract the packetization delay from 6 and 7 and show the results in fig. 8. The waiting time in OBS with no prediction, subtracting the packetization delay, can be approximated by $\delta = RTT/2$ (see fig.

5). We note that significant performance improvements can be obtained with the use of a simple linear predictor, which are even more significant as the multiplex of sources is increased, consistent with the availability of bandwidth brought by WDM networks.

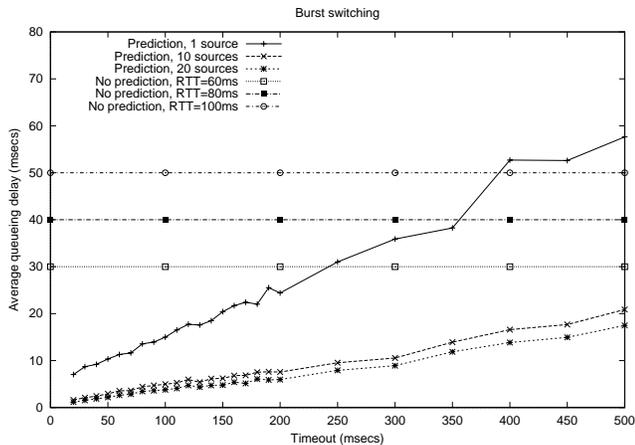


Fig. 8. Comparison of OBS with/without prediction

Finally, fig. 9 shows the bandwidth waste, as percentage of padding bytes per burst when $\hat{Z}[n] > Z[n]$. We note that the waste is also within reasonable values, taking into account the significant performance improvements introduced by the use of linear predictors. For example, if $RTT=60$ ms and the timeout is set to 50 ms, waste is below 5% for a multiplexing level larger than 10 sources. Nevertheless, the average waiting time without prediction is equal to 55 ms and the end-to-end delay (waiting time+ $RTT/2$) is 85 ms whereas the use of a linear predictor reduces the waiting time to 30 ms and the end-to-end delay to 60 ms. By choosing an adequate value of τ the end-to-end delay can be reduced to a value which is close to the minimum (RTT +packetization delay), if a linear predictor is used. The choice of a timeout τ is conditioned by the optical technology used in the network and is normally lower bounded to a minimum value [11], since optical network interconnection elements cannot handle small bursts. Nevertheless, we note from fig. 8 that even for high timeout values the waiting time savings are significant.

V. CONCLUSIONS AND FUTURE WORK

In this paper we have analyzed the suitability of simple linear predictors for traffic prediction in packet and burst switching networks. While packet switching performance is limited by the distribution of packet arrivals within the prediction interval, thus requiring a more precise estimate, we note that burst switching, which only requires a quantitative estimate of packets per interval, performs significantly better in comparison to a case with no prediction.

We are now concerned with determining to which extent

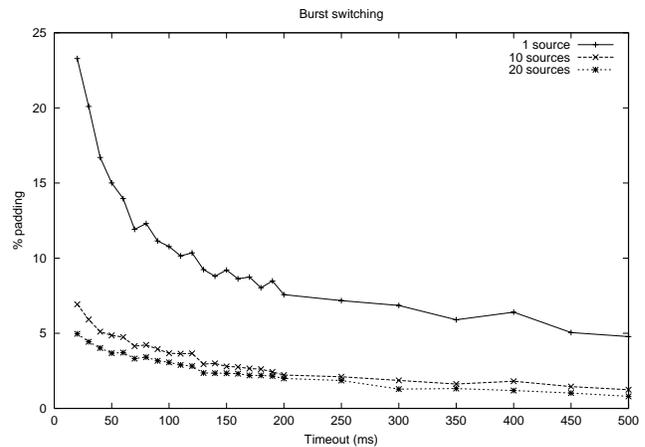


Fig. 9. Percentage of padding bytes per burst

knowledge of the *packet arrivals distribution* within the prediction interval can further improve performance in the packet switching case, since, as we have shown in this paper, knowledge of the number of bytes per interval does not suffice, specially at low multiplexing levels. This is the subject of our present and future research.

REFERENCES

- [1] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of Ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1):1–15, February 1994.
- [2] A. Erramilli, O. Narayan, A. Neidhart, and I. Sanjeev. Performance impacts of multi-scaling in wide area TCP traffic. In *IEEE INFOCOM 00*, Tel Aviv, Israel, 2000.
- [3] Sven A. N. Ostring and Harsha Sirisena. The influence of long-range dependence on traffic prediction. In *Proceedings of ICC 2001*, 2001.
- [4] Aimin Sang and San-qi Li. A predictability analysis of network traffic. In *Proceedings of Infocom 2000*, 2000.
- [5] A. Adas. Supporting real time VBR video using dynamic reservation based on linear prediction. In *IEEE INFOCOM 96*, San Francisco, CA, 1996.
- [6] William W. S. Wei. *Time series analysis. Univariate and multivariate methods*. Addison-Wesley Publishing Company, 1989.
- [7] M.W. Garrett and W. Willinger. Analysis, modeling and generation of self-similar VBR video traffic. In *Proceedings of the 1994 ACM SIGCOMM*, September 1994.
- [8] Daniel P. Heyman and T. V. Lakshman. What are the implications of long-range dependence for VBR-video traffic engineering? *IEEE/ACM Transactions on Networking*, 4(3):301–317, June 1996.
- [9] M. Grossglauser and J.C. Bolot. On the relevance of long-range dependence in network traffic. *IEEE/ACM Transactions on Networking*, 7(5):629–640, October 1999.
- [10] M. Grossglauser, S. Keshav, and D. Tse. RCBP: a simple and efficient service for multiple time-scale traffic. *IEEE/ACM Transactions on Networking*, 5(6):741–755, 1997.
- [11] C. Qiao and M. Yoo. Optical burst switching (OBS) - A new paradigm for an optical Internet. *Journal of High-Speed Networks*, 8(1), 1999.
- [12] J. Aracil and D. Morató. Characterizing Internet load as non-regular multiplex of TCP streams. In *Proceedings of IEEE ICCCN 00*, Las Vegas, NV, October 2000.