Issues on performance evaluation of switching nodes with shallow buffer under time-correlated traffic

Moisés R. N. Ribeiro[1] and Mike J. O'Mahony[2]

[1]LABTEL, Dept. de Eng. Elétrica – Universidade Federal do Espírito Santo (UFES)
Caixa Postal 01-9011 – 29060-970– Vitória – ES – Brasil

[2]Centre for Network Research – Electronic and System Dept. (ESE)
University of Essex, CO4 –2SQ - Colchester - UK

moises@ele.ufes.br, mikej@essex.ac.uk

Abstract. *A model suitable for describing time-correlated arrivals and traffic forwarding at time-slotted output-buffered switching nodes is thoroughly investigated regarding marginal distribution, correlation structure, load spatial partition, and traffic aggregation. Straightforward, though exact, analytical tools are developed and validated against numerical simulations. Results show that realistic outcomes can still be obtained from appropriate low-complexity analytical models since strictly correlated forwarding and hot-spot do not significantly influence most important traffic features which the investigated model is able to represent. Moreover, long-range dependence may not be an issue when evaluating performance of nodes with limited buffer space (e.g. photonic packet switching).*

1. Introduction

Traffic models hold an important role in packet-switched networks performance evaluation. They are fundamental in producing trustworthy representation for the dynamics of packet arrivals, forwarding, and buffering within switching nodes. Time-correlation among packets is a very demanding feature to be dealt with by buffers; mainly if low packet loss probability is required. As a result, many authors argue against traffic model with Short-Range Dependence (SRD) as it seems an obsolete practice in face of the fact that modern traffic sources actually possess long-range dependence (LRD) (self-similarity) in their correlation structure [8]. However, the way packets are distributed across output ports, the statistics of this convergence of incoming packets over outlets, and the buffer limited memory for past events are issues usually neglected by such critics. This paper investigates forwarding characteristics of a simple analytical model able to represent SRD and shows that it may even evaluate node performance under more strict conditions than its self-similar counter part.

1.1. Model complexity issues

The primary function of switching nodes is to sort out incoming packets onto output ports according to packets' header destination field. This process is here called forwarding. Packet arrival patterns strongly influence node performance. Fig.1(a) illustrates bursty arrivals, hot-spot, and temporal/spatial correlation in time-slotted nodes. The speed-up factor is four, i.e. it is an internally non-blocking switch and

therefore even four packets can be simultaneously forwarded to a given outlet. The numbers inside the packets represent the output destination port. Assuming that links serving outlets of the node shown in Fig.1(a) are able to forward just one packet per output in a time-slot, forwarded arrivals beyond this limit have to be either discarded or treated by a technique of contention resolution (e.g. output buffering). For instance, there is contention for the packets arriving in the sixth time-slot on inputs 1 and 4. In addition, notice that packets on inlet 4 arrive in a cluster of three packets all destined to output 3 while the burst on input 1 has packets headed for different outlets. Time-correlation may be present in case packets are arriving in such clusters on a regular basis. This correlation may last for few time-slots or even go through very long periods of time. The latter is the cause of the so-called self-similarity (or fractal) behaviour observed in Ethernet [8] and multimedia traffic [4][13]. Temporal correlation has a major influence on degrading buffer efficiency as a means for contention resolution. By receiving such bursty arrivals, it is likely that buffers will soon be completely taken resulting in some packet being lost.

Not only temporal effects are present in a switching node, traffic spatial distribution is also an aspect to be considered. A given port (or a group of them) may receive, on average, more packets than the remainder ports. In the illustrative example given in Fig. 1(a), traffic loads on inlets 1 and 4 are higher than on inputs ports 2 and 3 while outlets 3 and 4 are destinations for most of the arrivals. These traffic imbalances (also known as hot-spots) certainly pose more complex challenges to the mechanisms of contention resolution but this may also bring the opportunity to do asymmetrical node design as discussed in [1]. A considerable spatial (negative) correlation effect takes place at output of nodes with hot-spots as the "population" of packets to be forwarded in a time-slot is finite. The more packets go to a given output the fewer are destined to the remnants.

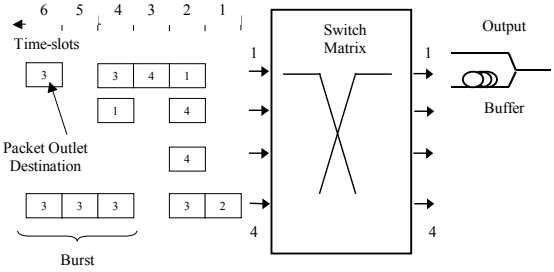

(a)

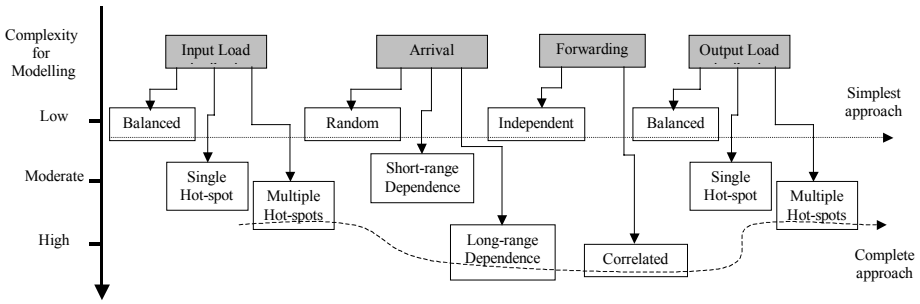

(b)

Fig. 1 -. Traffic features. (a) Arrival patterns. (b) Model complexity.

Figure 1(b) presents a classification system in which a vertical scale indicates complexity for modelling traffic features and packet forwarding. It is intuitive that the more complex is the model, the closest to the real behaviour it becomes. Therefore, bound approaches can be devised, named here simplest and complete. For the former, one may develop a model to evaluate node performance considering the incoming traffic as totally uncorrelated (random) and packets evenly distributed among outlets (independent forwarding). On the other extreme, there is a complete model that would include long-range temporal correlation in the incoming traffic representation. This kind of traffic often produces bursts that are composed of packets going to the same output. Therefore, correlated forwarding must be implemented in order to model those bursts crossing switching nodes accordingly. Independent forwarding, in this case, would alter traffic correlation structure resulting in unrealistic outcomes for performance evaluation. Finally, this complete model should be also able to represent the load imbalances and consequent spatial correlation that is present in real networks.

The analytical effort that must be put into the development of the complete model can be prohibitive while the simplest model might produce results of little practical interest. Nonetheless, this dilemma may be avoided when one considers that there is a whole range of options in between these two extreme models that could well produce reasonable results. A solution balancing performance and model complexity might be achieved by choosing, beforehand, only relevant aspects for implementing a model. For instance, for a node equipped with shallow buffers, which is often the case for photonic nodes where buffering is implemented with fibre-delay lines [1][12], the description of incoming traffic with LRD is not relevant given that buffers cannot realise such long correlation due to the lack memory concerning events happened long ago. Moreover, how contention resolution is implemented also counts. Having an incoming correlated and non-balanced traffic (such as the burst on inlet 1 in Fig. 1(a)) is important for input-buffered nodes but it is a completely irrelevant feature for output-buffered architectures, unless packets are to be forwarded to the same output (e.g. burst on input 4 in Fig. 1(a)). In conclusion, a fairly simple analytical model accounting for SRD with correlated forwarding and balanced loads (input and output) may produce acceptable traffic characterisation for output-buffered nodes. This paper evaluates if this option for simplicity may jeopardise accuracy in representing traffic relevant features for assessing buffering efficiency in dealing with switch external blocking (contention).

Another important aspect regarding performance evaluation is whether to use analytical or numerical methods for assessing node functioning. Analytical models bring forth elegant solutions but they are time-consuming to develop. Furthermore, assumptions and simplifications have to be made in order to keep its complexity within acceptable levels. Numerical models, on the other hand, can significantly reduce analytical effort besides allowing a more realistic representation of tasks performed within a switching node. However, its downside lies in the prohibitively long time taken to assess performance indicators such as low packet loss probability levels (e.g. 10^{-7} and below) required from nodes sitting in the of core transport networks. This paper makes use of both analytical and numerical techniques to investigate traffic features of the model under analysis. Numerical models are applied to validate the proposed analytical tools but they are also used to go beyond the limits imposed by analytical complexity, so that inaccuracies arising from assumptions and simplifications can be established.

1.2 Related work and contributions

The study presented in this paper is focused on source representation and traffic forwarding developed in [5]. Regardless of being a widely used model, little is known about traffic features produced by this method and the relevance of various factors that compose traffic statistics reaching output buffers. Investigations into traffic superposition of correlated sources can also be found in [7][10][14] and buffer performance under such traffic is carried out in [5][9][12]. The validity of this approach needs to be assessed when traffic crossing the switch is actually LRD instead of SRD. The role played by both marginal distribution and correlation structure on buffer performance is yet to be clarified. The remainder of this paper is organized as follows. The traffic model itself is described in Section 2 alongside with proposed methodologies to analysing its marginal distribution, correlation structure, and traffic aggregation features. Section 3 compares traffic of the investigated model against SRD numerical models and self-similar features. Section 4 addresses how traffic reaching output buffers behave against incoming traffic characteristics and node size in order to support interpretation of buffer performance outcomes. Furthermore, an analysis for buffers fed with SRD and LRD traffic is also provided in Section 4 to demonstrate the roles played by memory depth, marginal distribution, and correlation length in buffer performance assessment. Finally, the conclusions are drawn in Section 5.

2. Model description and assessment of its traffic features

A time-slotted switching node with internally non-blocking space matrix with nN inputs and N outputs is analysed. This feature imposes a speed-up factor of nN for the switch fabric. The aim is to find out traffic characteristics at a given outlet, namely marginal distribution and correlation structure for the stochastic process representing arrival at output buffers.

2.1 Traffic source

The model considers input traffic per port as a two-state system, more specifically High and Low for the representation used in Fig. 2(a). Sources are independent and identically distributed (i.i.d.) and they are either producing a continuous stream of packets while in High state or no traffic during Low state. The transition probabilities R_{HL} and R_{LH} stand for High to Low and Low to High respectively. Each source spends, on average, $1/R_{LH}$ on Low state and $1/R_{HL}$ time-slots on High state. The latter is also the mean burst length β (burstiness) produced by each source. Applying local balance boundary to the source model shown in Fig. 2(a) one finds (1)

$$\Pr ob[state = High] = R_{LH} / (R_{LH} + R_{HL}) \quad (1)$$

2.2 Forwarding

Balanced input and output loads ($0 \leq \rho \leq 1$ on average per input) are assumed. The forwarding process considers that, in the long term, the incoming traffic from nN sources is equally distributed across the N output ports. However, an ingenious way is utilised in [5] to address forwarding via a hybrid representation that could be fitted between independent and correlated approaches presented in Fig. 1(b) as each burst is forwarded in a correlated way to a given output. Consequently, one only needs to keep

track of sources addressed to the tagged output. As packets are exclusively released in the High state, Prob[state=High]=ρ/N and from (1) the probability of transition from Low to High can be calculated as stated in (2).

$$R_{LH} = \left(\frac{\rho}{N} R_{HL}\right) \Big/ \left(1 - \frac{\rho}{N}\right) \tag{2}$$

The evolution in time for the number of sources in High state (and addressed to the tagged output) is modelled by a Markov chain, illustrated in Fig. 2(a), in which any transition among the S_i ($0 \leq i \leq nN$) states is allowed. For each source, the transition probabilities are taken as geometric distributed with mean R_{HL} and R_{LH} for sources leaving state H and state L respectively. As a result, the number of input sources in state H and addressed to the outlet under observation in a given time-slot is a binomial distribution, where i represents the current state while j stands for the number of sources in High state one time-slot ahead. An auxiliary variable z is brought in to represent all possible combinations regarding transitions of individual sources, finally allowing the calculation of transition probability from i to j ($i,j \in \{0,1,2..,nN\}$) in a compact expression as shown in (3).

$$Q_{i,j} = \Pr ob[S(T+1) = S_j \setminus S(T) = S_i] = \sum_{z=\max(0,x-y)}^{\min(x,n\cdot N-y)} \Delta_L(z,i) \cdot \Delta_H(z-i+j, nN-i) \tag{3}$$

$$\Delta_L(u,v) = \binom{v}{u} R_{HL}^{u} \left(1 - R_{HL}\right)^{v-u} \quad ; \quad \Delta_H(u,v) = \binom{v}{u} R_{LH}^{u} \left(1 - R_{LH}\right)^{v-u}$$

The resulting arrival pattern at a given output, represented here by a discrete stochastic process A(t), is the summation of individual contributions in each time-slot as illustrated in Fig. 2(b) for three inputs that have packets addressed for the tagged outlet. The marginal distribution and two random variables A_τ and $A_{\tau+k}$, from temporal samples of A(t) taken at $t=\tau$ and $t=\tau+k$ respectively are also shown.

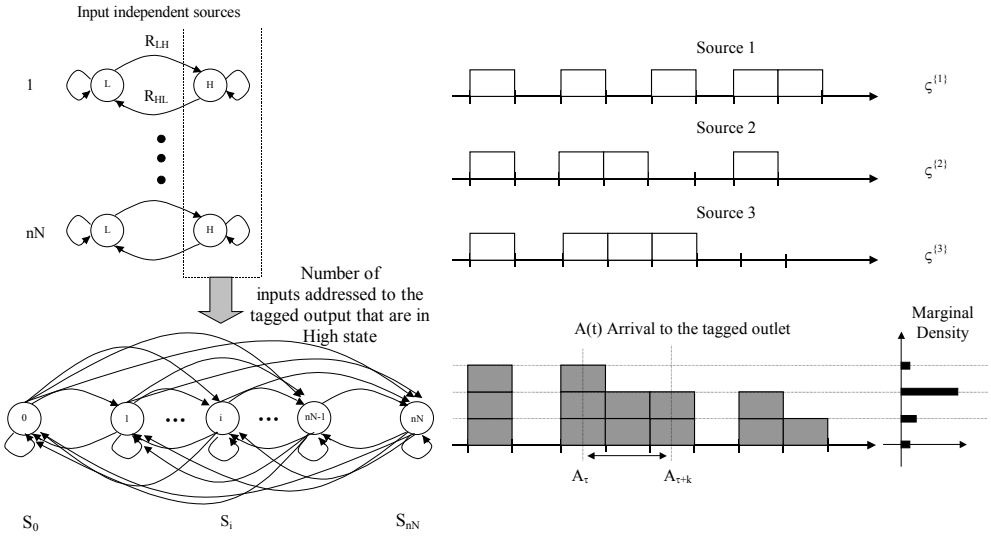

(a) (b)

Fig. 2 – (a) Source representation. (b) Illustration for traffic headed for an outlet.

2.3 Marginal distribution

Two methods for obtaining the marginal distribution for traffic reaching an output buffer are discussed, namely, steady-state solution for the Markov chain representing state of input sources, and convolution of individual contributions within a time-slot.

2.3.1 Analysis via Markov chain

The steady-state solution for the evolution in time of number of sources in High state can then be found by solving the equation system presented in (4). The underlined S and Q are the state and transition probabilities, in vector and matrix form respectively; e is a unitary column vector, i.e. $e=[1\ 1\ \ldots\ 1]^T$, with nN+1 elements while $*$ stands for matrices product.

$$\begin{cases} \underline{S} = \underline{S} * \underline{\underline{Q}} \\ \underline{S} * \underline{e} = 1 \end{cases} \tag{4}$$

Once packets are released with probability 1 for sources at High state, vector S may be seen as the marginal distribution for the discrete-valued stochastic process A(t), i.e. $S_a=p_A(a) = \text{Prob}[A=a]$, $a \in \{0,1,2..,nN\}$.

2.3.2 Analysis via Central Limit Theorem (CLT)

Provided that the traffic addressed to a given output is simply a summation of independent random variables $\varpi_{\{i\}}$, $\varpi \in \{0,1\}$, which are temporal samples from $\varsigma^{\{i\}}(t)$ coming from $i \in \{1,2,\ldots,nN\}$ inputs within a time-slot as shown in Fig. 2(b), one is able to find out the marginal distribution at the outlet under analysis by convoluting ($\otimes$) the individual contributions from each input as stated in (5)

$$p_A(a) = p_\varsigma^{\{1\}}(\varpi) \otimes p_\varsigma^{\{2\}}(\varpi) \otimes \cdots \otimes p_\varsigma^{\{nN\}}(\varpi) \tag{5}$$

For the assumptions used in Section 2.2, (6) gives the probabilities of sending either one or no packet to the tagged output in a time slot.

$$p_\varsigma(0) = \text{Prob}[\varpi = 0] = \left(1 - \frac{\rho}{N}\right); \qquad p_\varsigma(1) = \text{Prob}[\varpi = 1] = \frac{\rho}{N} \tag{6}$$

For CLT under balanced load, one may easily find the well-known binomial probability density function in (7) by convoluting nN i.i.d. sources.

$$p_A(a) = \binom{nN}{a} \left(\frac{\rho}{N}\right)^a \left(1 - \frac{\rho}{N}\right)^{nN-a} \tag{7}$$

This result may provide a proof for the memoryless nature of the aggregated traffic marginal distribution. In other words, despite the presence of time-correlation in each source that composes the total arrival, Eq. (4) should produce the same results as Eq. (7). The latter expression is clearly memoryless. Eq. (5) allows forwarding with hot-spots to be analysed while (4) is valid for i.i.d. sources only. Before reaching Gaussian shapes for $nN \to \infty$, one should expect outcomes from (5) to go through intermediate profiles, e.g. Poisson-like, for a limited number of input sources.

2.4 Correlation structure

It is of great interest the correlation structure present in such aggregation of time-correlated yet independent sources. An exact and simple method is developed here to perform this investigation. In addition, the concepts of correlation analysis of traffic traces and self-similarity are discussed in order to introduce a proposal to obtain variance-plot charts from analytical models.

2.4.1 Correlation coefficient function

Provided that A(t) is a stationary process, the two samples that are k time-slots apart, $k \in \{1, 2, 3, \ldots, \infty\}$, A_0 and A_k, have the same marginal distribution $p_A(a)$ but they may not be independent. In order to evaluate the correlation coefficient between them, the joint probability density $p_{A_o, A_k}(a_0, a_k)$ is needed. Fig. 3 shows a diagram that illustrates the relationship between time samples. The arrows represent transition probabilities calculated in (3).

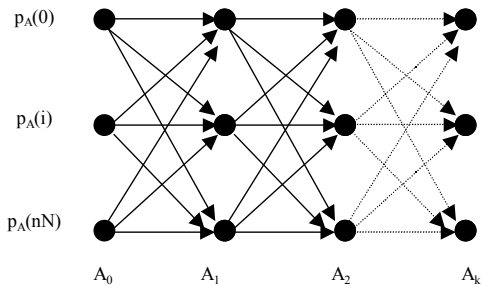

Fig. 3. - Correlation coefficient analytical calculations.

One can see from Fig. 3 that the joint density is obtained by multiplying $p_A(a)$ by the conditional probabilities of transition (given by Q) for each step forward in time. If Q is stationary this may be represented as in (8) where (•) stands for scalar product and pp is a square matrix composed nN+1 repetitions of $p_A(a)$, as shown in (9).

$$\underline{\underline{p_{A_1,A_k}}} = \underline{\underline{pp}} . [\underline{\underline{Q}}(t_1 \setminus t_0) * \cdots * \underline{\underline{Q}}(t_k \setminus t_{k-1})] = \underline{\underline{pp}} . [\underline{\underline{Q}}]^k \tag{8}$$

$$\underline{\underline{pp}} = \begin{bmatrix} \leftarrow \underline{p_A} \rightarrow \\ \vdots \\ \leftarrow \underline{p_A} \rightarrow \end{bmatrix} \tag{9}$$

Finally, the correlation coefficient r(k) for A(t) can be found via covariance calculation, as stated by (10), where E[.] represents the average over the ensemble.

$$r_A(k) = \frac{E[A(t_0, t_k)] - E[A]^2}{E[A^2] - E[A]^2} \tag{10}$$

2.4.2 Long-range dependence and (second-order) self-similarity

A wide-sense stationary process X(t) is long-range dependent (LRD) if its correlation coefficient $r_{ss}(k)$ is nonsumable (i.e. $\sum_k |r_{ss}(k)| \to \infty$), meaning that samples from a traffic trace are still related to each other no matter how far apart they are taken. An exact (second order) self-similar process is fully characterised by the Hurst factor H (0.5<H<1) as shown in (11).

$$r_{ss}(k) = \begin{cases} \frac{1}{2}[(k+1)^{2H} - 2k^{2H} + (k-1)^{2H}] & 0 < k < \infty \\ 1 & k = 0 \end{cases} \quad (11)$$

By increasing the observation window (time-scale) in which a stochastic process is analysed, one may see the effects of long-range dependence manifesting its graphical fractal-like behaviour as illustrated in [8][15]. This phenomenon is due to the slow reduction of variance as the process is viewed in coarser scales. A convenient way to see the resilience of traffic variance and to estimate the Hurst factor is by constructing a new stochastic process of non-overlapping blocks of X(t), each of them gathering m samples as shown in (12), where m∈ {1, 2, … }. This is somehow equivalent to observing X(t) in a time scale that is m coarser than the used by the original trace. Notice that $\lim_{m\to\infty} X^{\{m\}} = \mu$ (convergence to the mean) consequently $\sigma^2_{X^{\{m\}}} \to 0$ as m goes towards infinity. It is of great interest how fast this convergence is.

$$X_t^{\{m\}} = \frac{1}{m} \sum_{i=tm-m+1}^{tm} X_i \quad (12)$$

Plotting the variance of $X^{\{m\}}$ (normalised by the variance of X(t)) against the aggregation factor m in log-log scale, one can use the slope (S) of this graph for Hurst parameter estimation, where $\hat{H} = (S+2)/2$. It is worth looking at the two extreme values for H. For H→0.5 the variance of $X^{\{m\}}$ reduces at the same rate m increases, which actually means that the stochastic process is completely uncorrelated. This lack of correlation can be confirmed by the fact that in (11) where $r_{ss}(k)=0$ for H=0.5 and K>1, meaning that no self-similarity will be present as the time scale of observation is increased. In this case, variance is quickly smoothed by aggregation. On the other hand, for H→1 the variance stays virtually the same regardless of the aggregation factor m. This is explained by the autocorrelation $r_{ss}(k)=1$ for ∀ k. In summary, the variance of aggregated traffic depends upon the correlation structure of the original samples. Therefore, long-range dependence implies (second order) self-similarity. In other words, if samples are strongly related to each other the aggregation will result in a process very similar to the original one (at least the variance).

Correlation coefficient can be estimated using (13) for traces with M samples, where $\hat{\mu}_X$ and $\hat{\sigma}^2_X$ are temporal mean and (unbiased) variance estimation.

$$\hat{r}(k) = \frac{\sum_{i=1}^{M-k} [X(t_i) - \hat{\mu}_X][X(t_i - t_k) - \hat{\mu}_X]}{(M-k)\hat{\sigma}^2_X} \quad (13)$$

$$\hat{\mu}_X = \frac{1}{M}\sum_{i=1}^{M} X(t_i) \qquad \hat{\sigma}_X^2 = \frac{1}{M-1}\sum_{i=1}^{M}[X(t_i) - \hat{\mu}_X]^2$$

2.4.3 Traffic aggregation analysis for variance-plots

From the interpretation of variance-plots in Section 2.4.2, one is able to devise an analytical method to obtain such charts. An aggregated version of A(t) directly depends upon the correlation structure of its original samples. Another way to see the aggregation process needed for variance-plots is as a summation of m correlated random variables, which is then multiplied by 1/m in order to obtain an averaged value over m samples. As a result, the analytically evaluated variance of $A^{[m]}$, represented here by $Var[A^{[m]}]$, is the variance of this gathering of m samples multiplied by $1/m^2$ as seen in (14).

$$Var[A^{[m]}] = \frac{1}{m^2}\left\{\sum_{u=1}^{m}\sum_{v=1}^{m} E[A_u, A_v] - E[A]^2\right\} \qquad (14)$$

In order to produce the variance plot required for the estimation of H, the variance obtained in (14) must be normalised by the initial variance (m=1). Assuming that the stochastic process is wide-sense stationary, one obtains the normalised variance, NV(m), as in (15).

$$NV(m) = \frac{Var[A^{[m]}]}{Var[A^{[1]}]} = \frac{1}{m^2}\sum_{u=1}^{m}\sum_{v=1}^{m} r_A(|u-v|) = \frac{1}{m^2}\left[m + 2\sum_{u=1}^{m-1}(m-u)\, r_A(u)\right] \qquad (15)$$

Using the limits for correlation coefficient in (15), convergence to 1/m and 1 takes place, as it would be expected, for uncorrelated and fully correlated traffic respectively.

3. Model evaluation

Traffic reaching output buffers according to the model presented in Section 2.1 and 2.2 (called here basic model with hybrid forwarding) are checked against numerical simulation. Comparisons are performed within a framework used to analyse self-similar traffic, developed in Section 2.4, in order to provide a clear judgement to what extend the analytical model under analysis may neglect important traffic features. An ad hoc numerical simulator generates traces with 10^5 samples for a node with 4x16 input ports and 16 outlets. The basic purpose here is to validate the analytical model and to provide some insight into the traffic features where analytical investigation would be rather complex. The offered load (ρ), from each 16x4 independent sources, is 0.8 and it is evenly distributed across the 16 output ports (except in Section 3.3 where output hot-spot is investigated). A mean burst length (β) with 16 time-slots was chosen for the results presented in this Section.

3.1 Basic model with hybrid forwarding

For this case the offered load is actually 0.8/16, as seen in (7), due to the forwarding approach utilised. The marginal distribution is shown in Fig. 4(a) for both analytical and numerical models. It is also presented a Poisson density function with the same mean found in the trace clearly showing that the model generates Poisson-like densities for this node size and offered load. As it might be expected, both methods for evaluating

the marginal distribution presented in Section 2.3 produce the same result. This is the evidence that the number of active sources gathered over a outlet is itself a memoryless process.

Temporal correlation properties are exposed by autocorrelation coefficient function and aggregation analysis, shown in Fig. 4(b) and (c) respectively. Fig. 4(b) shows the agreement between the autocorrelation obtained through the analytical procedure described in (8)-(10) and numerical evaluation of traffic traces in (13). One can also observe that the correlation values are negligible for lags over 60 time-slots. On the other hand, Fig. 4(c) shows that the variance for the aggregated version of this traffic remains virtually unchanged (H=0.96) until the observation period exceeds the mean burst length (16 time-slots). After this point, a steep reduction in variance against m takes place as the trace becomes uncorrelated. And this is also the explanation for the good matching found to a line in parallel with the curve for H=0.5. The best fitting line for the whole range studied (m=1000) would produce an estimated Hurst factor H=0.82. However, it is easy to conclude from Fig. 4(b) that this is a meaningless result since traffic correlation present in this trace is SRD. Nonetheless, it must be said that the trace has stronger correlation properties (and so has the analytical model) than an equivalent self-similar trace with H=0.82 for lags below 29 time-slots. As far as correlation goes, this result means that nodes with shallow buffers may be under more severe circumstances with this trace than with its self-similar counterpart with H=0.82.

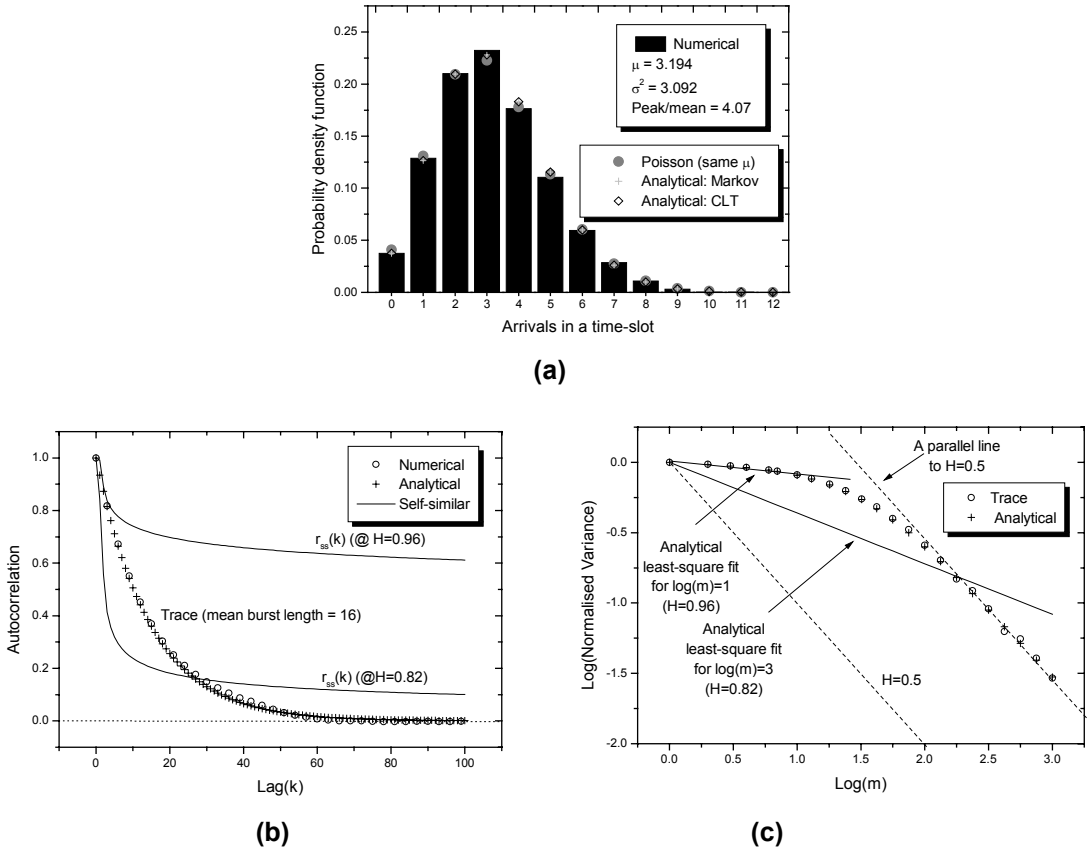

Fig. 4. – Model. (a) Marginal distribution. (b) Autocorrelation. (c)Variance-plot.

3.2 Strictly-correlated forwarding

The point to be addressed here is whether the simplified approach towards packet forwarding may undermine the validity of results obtained analytically. A numerical model with strictly correlated forwarding is developed, i.e. the whole burst is always headed for a given outlet. Traffic features are then compared with the forwarding hybrid approach discussed in Section 3.2. Statistical study of a trace considering balanced load with correlated forwarding is presented in Fig. 5. For the sake of illustration, a fragment of the trace from a given incoming source is presented in Fig. 5(a) where packets are shown according to the output they are headed for versus time-slot intervals. The off period of this source is also shown. In this particular fragment, arrivals destined to output 6, for example, consist of a 63 time-slot burst followed by a single packet while no traffic is generated for outlets 1 and 3 during this period of observation.

As far as statistical traffic features considered in this paper are concerned, no significant difference is noticed between the hybrid model and the strictly correlated forwarding as one can see by comparing marginal densities, aggregation and autocorrelation functions from Fig. 4 and Fig. 5. Further experiments were performed for different offered load, node size and burstiness factor and both models kept producing very similar results. Nonetheless, it is noteworthy that convergence to zero in Fig. 5(c) is slower for long lags than in Fig. 4(b), which is also reflected in the mismatch observed in Fig. 5(d) between aggregations factor gathering more than 300 time-slots and the parallel line to H=0.5. This may be a slightly sign of LRD or even non-stationarity in the traffic trace.

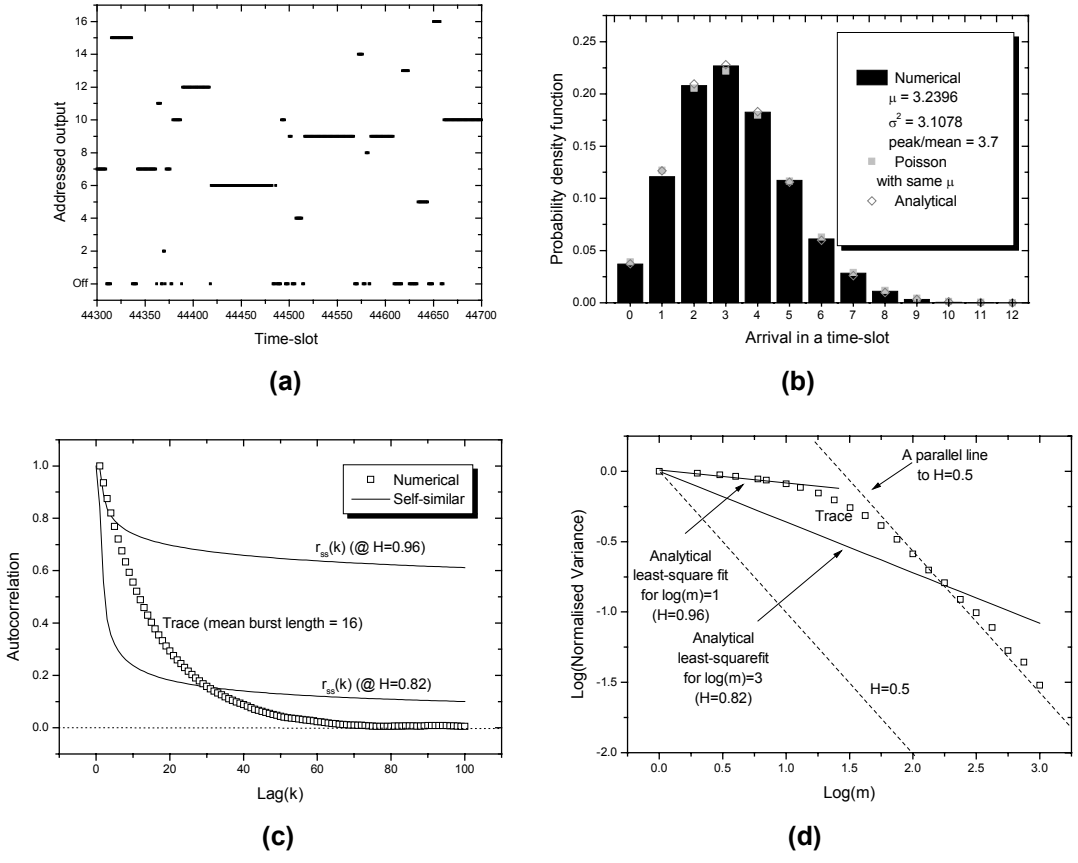

Fig. 5. Strictly correlated forwarding. (a) Trace fragment (b) Marginal distribution. (c)Autocorrelation function, and (d) Variance-plot.

3.3 Strictly correlated forwarding plus single hot-spot

One more feature is included on top of strictly correlated routing. A single traffic hot-spot takes place at output port five. In Fig. 6(a) the spatial distribution of packets and the numerical result obtained are compared. Outlet five was set to be five times more likely to receive packets. However, this implies that other ports will receive fewer packets, which indicates (negative) space correlation. For the balanced load case, each output would receive 6.25 % of the packets but here this number is reduces to 5% due to the fact that 25% of incoming packets are directed to the hot-spot. This result can be seen in Fig. 6(a). The marginal distribution at the hot-spot along with CLT analytical result are shown in Fig. 6(b). There is also a Gaussian curve using the mean and variance found in the trace generated by simulation. No alteration is observed for correlation structure, and consequently for aggregation function as well, compared to results shown in Fig. 4(b) and Fig. 5(c) such results are omitted in Fig. 6.

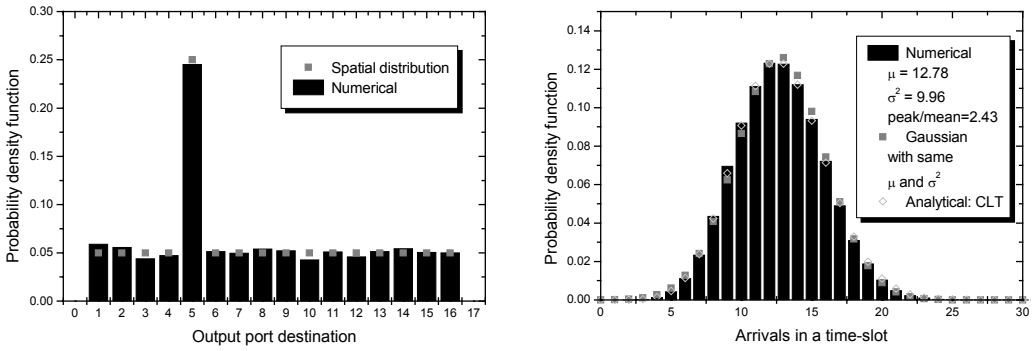

Fig. 6. Hot-spot. (a) Spatial distribution and (b) Marginal distribution

It is noticed, however, that the marginal distribution is affected (expectedly converge to a Gaussian-like shape) by concentration of arrival over a given outlet. This paper does not investigate performance implications of traffic imbalances. Nevertheless, one can see that any performance degradation will be due to marginal distribution changes as correlation span appears to reach a saturation level, for the range of interest, after the aggregation of many sources is performed. Studies concerning ATM switches under uncorrelated traffic and hot-spot can be found in [16] while [9] provides analytical approaches for speed-up limitations and correlated traffic. Photonic node design taking into account traffic imbalances due to routing is numerically investigated in [1].

4. Discussion

Once the model under analysis has proved to be accurate enough in representing traffic features, issues related to performance assessment are now discussed in this section, allowing some important conclusion to be drawn concerning node performance evaluation.

4.1 Asymptotical behaviour for correlation span and marginal distribution

Before assessing buffer performance, one should carefully examine the features of the traffic reaching the buffer in order to better interpret outcomes of such experiments and avoid performing unnecessary trials. The results presented in Fig. 7(a) are for

summation of the correlation coefficients (from k=1 to k= 100) versus node size for different burstiness and offered load. This method is used here to obtain very concise results for correlation behaviour against node and traffic characteristics. In order to produce traffic with a given mean burst length, the only option left for the model is to place burst far apart in order to comply with the offered load that has been chosen. One should remember that correlation structure has also to do with the off period of a traffic source [2] [3]. As a result, it is seen in Fig. 7(a) that highly correlated traffic is produced when simulations using light traffic are performed. In addition, the correlation obtained in this case becomes less sensitive to node size. The higher the traffic burstiness, the wider is the difference between curves for light and heavy traffic. Another important outcome is that the asymmetry (n) between input and output ports has virtually no impact on the results shown in Fig. 7(a). One should notice that results converge to the respective summation for the correlation coefficient from a single on-off source [2] with geometrically distributed transition coefficients for RLH→0 shown by dotted lines. This may mean the aggregation of incoming independent sources over switch outlets causes little increase in burstiness.

A clear indication provided by Fig. 7(a) is that little difference is to be expected from assessments of output-buffered nodes with more than 16 ports, as far as influence of correlation on performance is concerned. Although it is evident that this traffic model only produces SRD, the correlation length is unbounded as Fig. 7(a) demonstrates when β is increased. At this point, it is important to gain an insight into the influences of LRD on performance of nodes with short memory (limited buffer space). For instance, take a bufferless node. In such case no information is kept about past events, therefore time-correlated arrivals should cause absolutely no penalty to node performance. The relationship between buffer depth and correlation span has been studied, from different viewpoints, in [4] and [13]. Both approaches, however, have found a linear connection between correlation span and buffer depth. In other words, there is just a relevant correlation length (correlation horizon in [4] and critical time scale in [13]) that should be present in the traffic model in order to properly assess performance of a node with a given buffer depth. Therefore, the inclusion of long-range dependence may only have minor importance in the context of the present study. Additional evidences to support this argument can be found in results presented in [6] and [15]. Fig. 4(b) and (c) show that the level of SRD match (or even exceed) correlation levels from (exact) self-similar traffic sources within the range that might be realised by photonic switching nodes as feasible optical buffer depths are usually well below 50 time-slots [1].

Marginal distribution plays an important role in influencing buffer performance. This traffic feature for the model under analysis is expected to converge to a Poisson-like distribution (and to Gaussian-like for large means). Comparisons for random variables distributions are generally performed via Quantile-Quantile plots [11] or simply by displaying cumulative distributions together. Nevertheless, these charts only provide a qualitative, and therefore subjective, appreciation of convergence. A simple summation of squared difference is proposed here instead. Results presented in Fig. 7 are insensitive to burstiness variations. For nodes larger than 8x8, the marginal density quickly converges to Poisson regardless of node asymmetry factor (n). Although the results presented in Fig. 7 is only for ρ=0.8, the lower the offered load, the faster is this convergence.

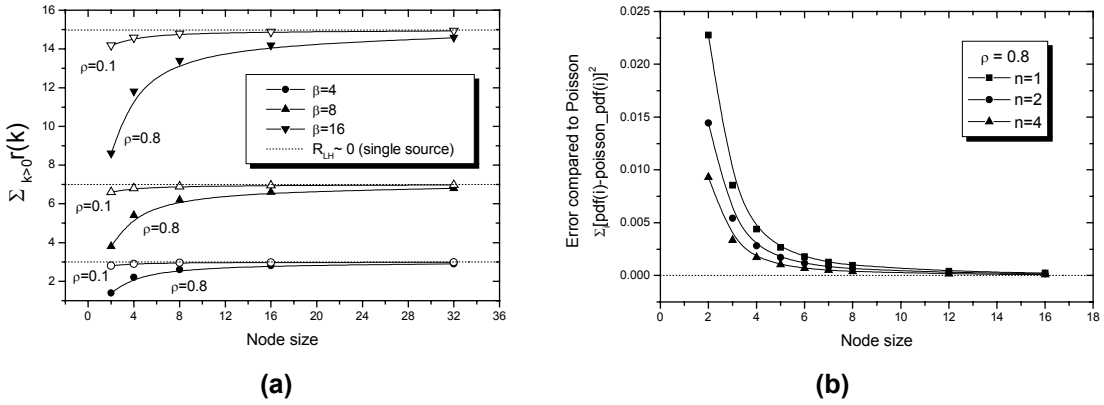

Fig. 7. Asymptotical behaviour vs. node size. (a) Summation over 100 time-slots for correlation coefficient. (b) Comparison with Poisson for ρ=0.8.

In summary, performance assessment for a given buffer depth are not expected to vary much nodes for nodes with more than 16x16 ports, as seen in [12], since both correlation and marginal distribution remain practically unchanged for such node size onwards.

4.2 Influence of SRD, LRD, and marginal distribution on buffer performance

Packet loss probability will be assessed for buffers fed with the traffic model under analysis and with traces bearing LRD.

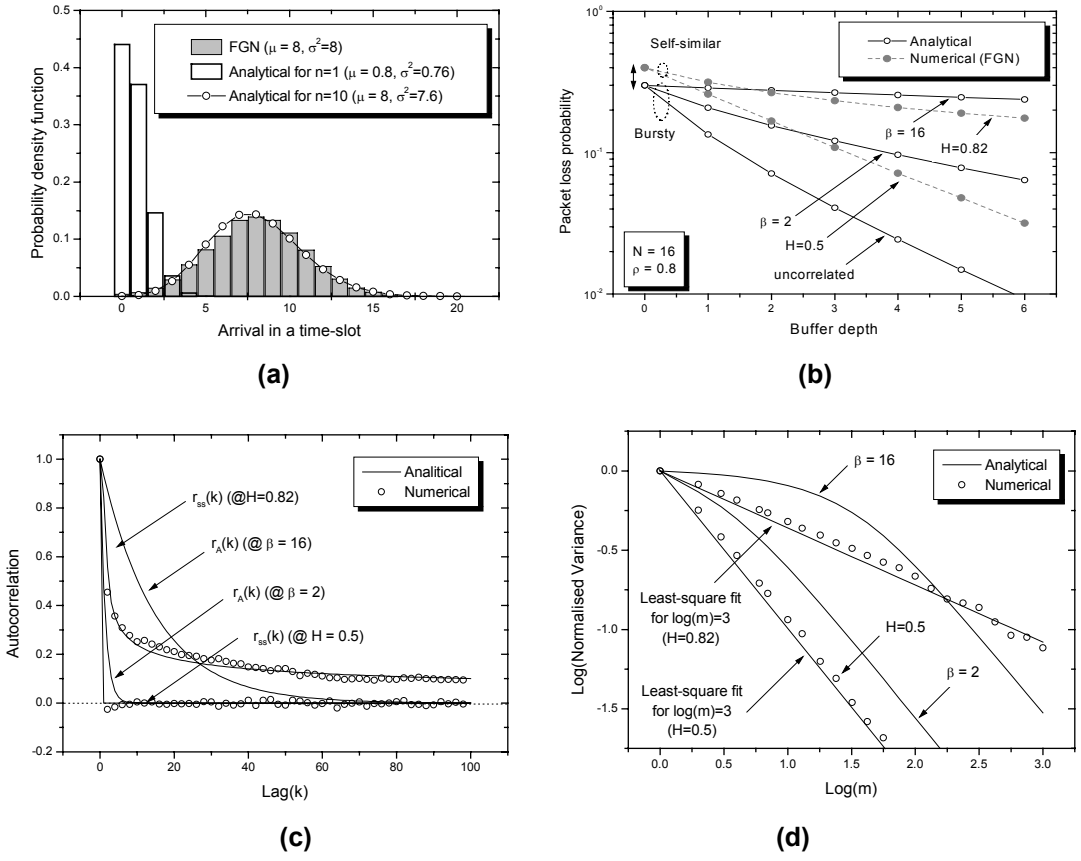

Fig. 8 - Comparison with self-similar traffic performance. (a) Marginal distribution. (b) Packet loss probability. (c) Autocorrelation. (d) Variance-plot.

Analytical buffer modelling proposed in [12] is applied in this Section while an exact (second order) self-similar trace, derived from Fractional Gaussian Noise (FGN) [11], is used in an ad hoc numerical simulator. Although both cases use N=16 and $\rho = 0.8$, slightly different marginal distributions are intentionally used in order to highlight its influence over buffer performance. The numerical model produces Gaussian marginal distributions and it was set to generate a traffic trace with $\mu=8$ and $\sigma^2=8$. If the analytical model was evaluated with n=1, it would produce $\mu=0.8$ and $\sigma^2=0.76$. For the sake of convenience it is assumed, however, that an asymmetry factor n=10 is used, mean and variance become ten times higher and, therefore, the curve takes a Gaussian-like shape. These marginal distributions can be seen in Fig. 8(a). As a result, a small difference would still remains between variances of numerical and analytical traffic models. Outlets of the node under analysis are assumed to be ten times faster than its inlets. As a result, each buffer actually receives an offered load of 0.8. Packet loss probability against buffer depth is presented in Fig. 8(b). The effect of different variances is best seen in Fig. 8(b) for bufferless nodes (i.e. at buffer depth=0). The model with LRD predicts a worse performance at this point than the analytical model bearing SRD. This certainly is due to the marginal distribution characteristics since no time-correlation can be realised in the absence of memory. In order to isolate the effects of marginal distribution and time-correlation, performance offset caused by marginal distribution should be removed. One may still use Fig. 8(b) but self-similar curves have to shifted down in order to find the best match for uncorrelated traffic ($\beta \rightarrow 1$) and H=0.5. At these points similar loss probability outcomes are expected from both models (provided they have the same incoming traffic). As far as shallow buffers are concerned, it is clear that the self-similar model with H=0.82 predicts packet loss probabilities as SRD analytical model with $\beta=2$. The variance-plot in Fig. 8(d) for $\beta=2$ matches self-similar trace with H=0.82 for small traffic aggregations ($m<10^{0.5}$) corroborating this finding for the range of buffer depth studied. It is important to draw attention to the fact that packet loss probability for the SRD model with $\beta=16$ is far less affected by buffer deepening than LRD model with H=0.82. This is because the analytical model possesses higher correlation levels within lags comparable with the buffer depth, as seen in Fig. 8(c). Note that FGN model generates exact (i.e. it is not asymptotic) second order self-similarity traffic as can be checked in Fig. 8(c) as traces fit results from (11) for $\forall$ k; and so does it for $\forall$ m in Fig. 8(d).

5. Conclusions

A comprehensive study into the traffic features of a simple analytical SRD analytical model with hybrid (correlated-random) forwarding was performed. Statistical tools were developed, within an analytical framework, allowing direct comparisons between the investigated model and the so-called self-similar traffic. Numerical simulations were employed for further investigations into features arising from strictly correlated forwarding, hot-spot, and buffer performance under self-similar traffic. Although it is intuitive that the basic analytical model could appropriately represent first order traffic features, it is surprising to find out very little discrepancies in marginal distribution and correlation structure when compared with models that account for correlated forwarding and hot-spots. This is possibly due to the negligible effect of aggregation compared to the individual source behaviour. It is important to highlight that, even with moderate burtiness, the analytical model may evaluate node performance under more strict

conditions than by using the so-called self-similar model over swallow buffers. In case buffer depths are well beyond the burst mean length, it must be expected a perform prediction for LRD traffic worse than the one found using SRD models. But remember that the outcome also depends on marginal distributions. Further investigation is needed to generalise the model studied for network performance assessment (end-to-end) taking into account issues such as network topology and routing algorithms.

References

[1] Castañón G., et al. "Asymmetric WDM all-optical packet switched routers", In Proc. OFC'00, paper WD-4, 2000.

[2] Garroppo R. G., S. Giordano, M. Isopi, and M. Pagano, "On the implications of the off periods distribution in two-state traffic models" IEEE Communications Letters, Vol. 3, No. 7, pp. 220-222, July 1999

[3] Garroppo R. G., S. Giordano, M. Pagano, and G. Procissi, "On the relevance of correlation dependencies in on/off characterization of broadband traffic " In proc. ICC'2000, Vol. 2, pp. 811-815, 2000

[4] Grossglauser M., and J-C Bolot, "On the relevance of long-range dependence in network traffic", IEEE Trans. on Networking, Vol. 7, No. 5, pp. 629-640, Oct. 1999.

[5] Hou T-C., and A. K.Wong, "Queueing Analysis for ATM Switching of Mixed Continuous-Bit-Rate and Bursty Traffic", In Proc. INFOCOM'90, Vol..2, pp. 660-667, 1990.

[6] Huebner F., D. Liu and J. M. Fernandez, "Queueing Performance Comparison of Traffic Models for Internet Traffic". In Proc. of GLOBECOM '98, pp 471–476, 1998.

[7] Ide I., "Superposition of interrupted poisson processes and its application to packetized voice multiplexers " in Proc. ITC'89(ITC-12) pp1399-1405, 1989.

[8] Leland W. E., et al. "On the self-similar nature of Ethernet traffic (extended version), IEEE/ACM Trans. on Networking, Vol. 2, No. 1, pp. 1-15, Feb. 1994.

[9] Makhamreh I.I., N. D. Georganas and D. McDonald, "Analysis of an output-buffered ATM switch with speed-up constrains under correlated and imbalanced bursty traffic" IEE Proc. Communications Vol. 142, No. 2, pp. 61-66, April 1995.

[10] Noros I., J. W. Roberts, A. Simonian, and T. Virtamo "The superposition of variable bit rate in ATM multiplexer", IEEE Journal of Selec. Areas in Comm.; Vol. 9, No. 3, pp. 378-387, April 1991

[11] Paxson V., "Fast approximation of self-similar network traffic" Computer Communications Review, Vol. 27, N. 5, pp. 5-18, October 1997

[12] Ribeiro M. R. N., and M. J. O'Mahony, "On the benefits of WDM and wavelength conversion in photonic switching", Revista da Sociedade Brasileira de Telecomunicações, Vol. 15, No. 1, pp. 1-9, June 2000.

[13] Ryu B. K., and A. Elwalid, "The importance of long-range dependence of VBR video traffic in ATM traffic engineering: myths and realities", In Proc. SIGCOMM'96, pp. 3-14, 1996.

[14] Sohraby K., "On the theory of general on-off sources with applications in high-speed networks", In Proc. INFOCOM'93, Vol. 2, pp. 401-410, 1993.

[15] Tarongí, D., D. Rodellar, J. M. Torner, J. Solé-Pareta, S. Borgione, "Traffic characterization using optical based packet switches with Poisson and fractal traffic sources", In Proc. ONDM' 01,. 2001

[16] Yoon H., M. T. Liu, and K. Y. Lee, "The knockout switch under nonuniform traffic", In Proc. of GLOBECOM'88, Vol. III, pp. 1628-1634, 1988.