On the Benefits of Traffic Segregation in WDM Networks

Fernando L. Trazzi Jr.[1] , Moisés R. N. Ribeiro[1] , and Helio Waldman[2]

[1]LABTEL – DEL – CT – UFES – C. P. 01-9011-29060-970 – Vitória – ES – Brasil

[2]DECOM – FEEC – UNICAMP – C. P. 6101-13083-970 – Campinas – SP – Brasil

{fernando,moises}@labtel.ele.ufes.br,
waldman@decom.fee.unicamp.br

Abstract. *This paper investigates whether performance of WDM optical switching may be improved by traffic shaping. Heavy-tailed Pareto traffic transmission split into different wavelengths, according to the holding time, is here proposed as a means of traffic shaping aimed at the minimization of mean delay, jitter, or blocking-time probability in optical burst switching. It is found that burst segregation across n wavelengths (n x $M/P_T/1$) may outperform classical WDM (M/P/n) systems as far as overall residual-service time is concerned. The existence of optimal segregation thresholds is analytically demonstrated. Results for the whole traffic show one order of magnitude reduction in mean delay, while jitter and blocking-time are, respectively, two and three orders of magnitude better than the classic WDM non-segregated approach.*

1. Introduction

The evolution of circuit switched optical networks will use Optical Burst Switching (OBS) [Qiao1999] as the most likely step in the direction of packet switched architectural paradigm. One of the main goals of OBS is to avoid the downsides of circuit switched systems such as lightpath setup/teardown delay, bandwidth over-provisioning, and complex control plane. OBS may provide optical networks with the agility and granularity needed to better match spatial and temporal dynamics of Internet Protocol (IP) traffic. On one hand, highly variable and time-correlated IP traffic [Leland1994] has deep implications for OBS design. The actual impact depends on how the IP packets are gathered to build up optical burst payloads [Necker2003] and how optical buffering is employed (if adopted at all) in core nodes [Buchta2003]. On the other hand, the end-to-end time-delay experienced by IP-based applications will be impacted by OBS network performance. Indeed, the strategies chosen for burst assembly as well as burst buffering are again key issues but now from the point of view of OBS clients. Real-time applications (such as voice and interactive video conference) are extremely sensitive to delay and jitter that are introduced in the packet stream [Giroux1999]. In addition, non-real time transport protocols, such as Transmission Control Protocol (TCP), rely on the first and second order statistics collected from (round-trip time) end-to-end delay to perform flow control [Tanenbaum2003]. Therefore, efforts must be put into the minimization of time-delay features at edge nodes as well as core nodes to, simultaneously, meet efficiency in OBS network design, high throughput to bulk traffic, and Quality of Service (QoS) targets of premium clients.

This paper addresses the residual-service time at queues across the network as it assumes the model “blocked calls delayed” to handle contention. It has a twofold contribution - to bring evidences that residual-service time of intersected elements lies at the root of waiting time statistics and demonstrates that traffic might be segregated, according to holding time, in order to reduce the effects of heavytailness in queues. While the former gives valuable insights into the role of the diverse factor in queuing dynamics, the latter analytically investigates how to take advantage of path diversity allowed by wavelength division multiplexing (WDM) systems. The paper is organized as follows. The context of application is presented in Section 2. The basic modeling is given in Section 3 while metrics for comparing segregated and non-segregated approaches over WDM are presented in Section 4. The segregation policies investigated in this paper are introduced in Section 5. Results and discussion on the their implications on OBS network design are given in Section 6. Finally, Conclusion brings final remarks and future work.

2. Optical Burst Switching and QoS

An illustration for the OBS network architecture is given in Fig. 1. Edge nodes gather traffic from client networks (e.g. IP) and assemble their packets in proper forward equivalence classes (FECs) buffers according to edge node destination. Another task assigned to edge node is to produce signaling for OBS routers (core nodes). These header-like messages should include, among other information, the burst length. They are sent, some time ahead of the burst, through control channels along the path chosen for the burst to reach its destination (another edge node). OBS core nodes are plain optical switching matrices that might include optical delay-lines for dealing with contention. Therefore, connection pattern to be set is the result of signaling received on-line through control units connected to control channels. The burst itself follows the signaling message (after offset-time) without waiting for acknowledgement, i.e. whether the connections were successfully established or not. Note that this eliminates the setup/teardown delay of circuit switching but QoS will be impacted due to the increase the likelihood of blocking in one of core nodes along the route. The investigation into QoS awareness for OBS can be divided into burst assembly methods at edge nodes, resource reservation techniques, and forms for reducing burst blocking probability due to contention in core nodes.

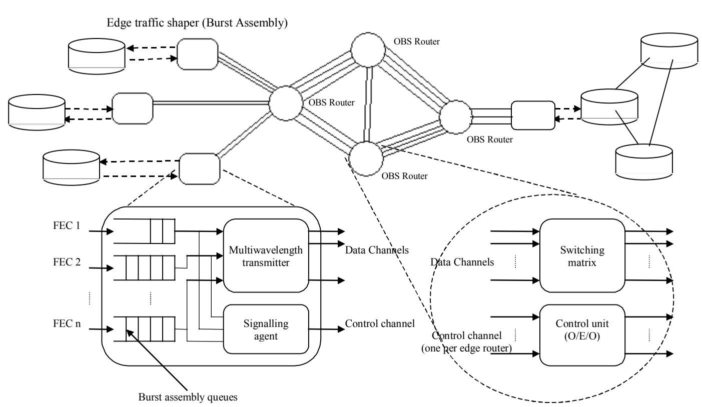

Figure 1. Illustration for Burst switching network [Morato2003].

The choice for the assembly technique is closely related to the kind of traffic being transported. Timer-based algorithms suit time-constrained traffic as a maximum delay for assembling a burst is set; regardless of payload size produced. On the other hand, the transport of bulk traffic has no such constraint. Then bursts are only released when accumulated traffic achieves a minimum payload size. From OBS buffer design standpoint, however, it is shown that self-similarity is transferred from IP incoming traffic to the burst stream in spite of the assembly method employed [Necker2003]. Furthermore, even approaches such as "files-over-bursts", which are designed to reduce control packets overhead, lead to Pareto-distributed burst sizes [Morato2003]. Implications of these results for OBS performance may be realized as highly variable delays experienced by bursts when crossing just few nodes and failed attempts to reduce contention through fiber delay lines. Note that offset time based reservation approach to QoS [Yoo2000] might also be severely impacted by unpredictable skew induced through highly variable delays on bursts.

In summary, traffic statistic emerges as a key issue in QoS for burst switching as much as it is in IP networks. It is argued in this paper that problem is not the heavytailness alone but the even heavier tail statistics of residual-service time of intersected bursts. It is been previously shown in [Waldman 2004] that multiple servers (WDM) may significantly reduce that influence; even bringing first moments of unbounded Pareto distributed bursts down to finite figures through the use of a minimum number of wavelengths. Moreover, the role of limiting the length of bursts was unveiled as an efficient way to improve delay statistics. This paper, in its own right, pinpoints the role of residual-service time in waiting time statistics and takes further the improvements suggesting the use of traffic segregation across wavelengths.

3. Basic Modeling

The main goal here is the comparison between delay statistics in queuing systems M/P/n (Markovian arrival, Pareto service, and n servers) and n of $M/P_T/1$ (Markovian arrival, truncated Pareto, and one server). In others words, the performance of non-segregated approach system is checked against the segragated traffic over the same number of WDM channels. The focus is on residual-service time. Bursts are assumed independently and identically distributed (i.i.d.) events. They arrive following a random variable A, in intervals exponentially distributed given by arrival rate λ, so that $E[A]=1/\lambda$. The holding time is represented by the random variable T. It is assumed $p_T=\text{Prob}\{T\geq t\}$ follows Pareto distribution in (1), where α is the shape factor, t_{min} and t_{max} are, respectively, the minimum and maximum bursts length allowed.

$$p_T(t)=\begin{cases}0, & t<t_{\min}\\ \dfrac{\alpha\, t_{\min}{}^{\alpha}}{t^{\alpha+1}}\left(\dfrac{1}{1-r^{\alpha}}\right), & t_{\min}\leq t<t_{\max}\end{cases} \quad \text{, where } r=\frac{t_{\min}}{t_{\max}} \tag{1}$$

The k^{th} moment for T is, therefore, obtained as stated in (2)-(4):

$$E[T]=\int_{t_{\min}}^{t_{\max}} t p_T(t)dt=\frac{\alpha}{\alpha-1}t_{\min}\left(\frac{1-r^{\alpha-1}}{1-r^{\alpha}}\right) \tag{2}$$

$$E[T^2] = \int_{t_{\min}}^{t_{\max}} t^2 p_T(t)\,dt = \frac{\alpha}{\alpha-2} t_{\min}^2 \left(\frac{1-r^{\alpha-2}}{1-r^{\alpha}}\right) \tag{3}$$

$$E[T^k] = \int_{t_{\min}}^{t_{\max}} t^k p_T(t)\,dt = \frac{\alpha}{\alpha-k} t_{\min}{}^k \left(\frac{1-r^{\alpha-k}}{1-r^{\alpha}}\right) \tag{4}$$

3.1. Residual-service time seen by arriving bursts under FIFO scheduling

Arriving bursts may either gain immediate access to service or intersect another burst being served. In the latter case, there may also be a backlog ahead of them, as illustrated in Fig. 2. The random variable B represents server occupation and, therefore, its events may only assume two values, i.e. $b \in \{0,1\}$, with probabilities (1-ρ) and ρ, respectively, with $\rho = E[T]/E[A] = \min(\lambda E[T], 1)$.

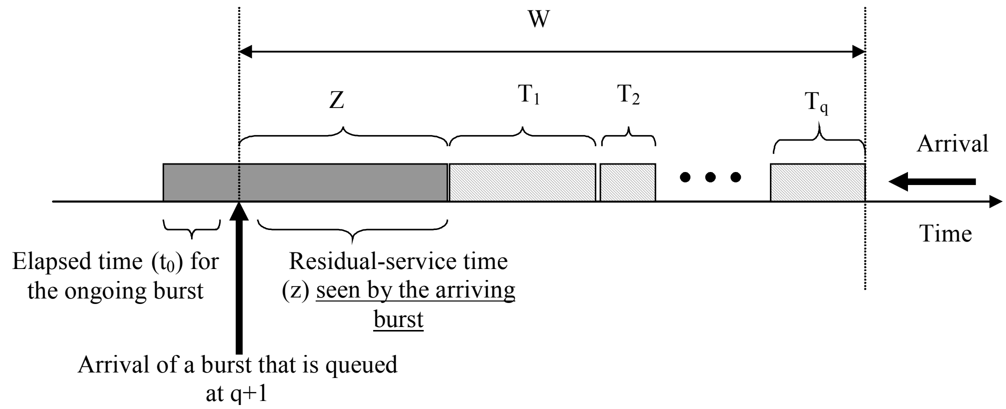

Figure 2. Arrival of a burst under FIFO scheduling.

Contrary to intuition, the intersected burst at the moment of arrival (darker filling in Fig. 2) does not follow T. Actually, longer bursts are more likely to be intersected than shorter ones. In other words, it is the random variable T sampled at random times rather than samples taken randomly from the sample space; that would result in T itself. The residual–service time seen from intersecting burst standpoint in Fig. 2 is represented by random variable Z, with p.d.f. presented in (5) [Waldman 2004].

$$p_Z(z \backslash b=1) = \begin{cases} \dfrac{1}{E[T]}, & z < t_{\min} \\ \dfrac{1}{E[T](1-r^{\alpha})}\left[\left(\dfrac{t_{\min}}{z}\right)^{\alpha} - r^{\alpha}\right], & t_{\min} \le z < t_{\max} \end{cases} . \tag{5}$$

It is clear that $p_Z(z \backslash b=0)=0$ as no residual-service time is seen by a burst arriving at a moment the server is not busy. Thus, $p_Z(z)=p_Z(z \backslash b=1)\rho$. Notice that the k^{th} moment from Z can be related to the $(k+1)^{th}$ moment of T, as it is stated in (6)-(8).

$$E[Z] = \rho \frac{(\alpha-1)}{2(\alpha-2)} t_{\min} \left(\frac{1-r^{\alpha-2}}{1-r^{\alpha-1}}\right) = \frac{\rho E[T^2]}{2E[T]} \tag{6}$$

$$E\left[Z^2\right] = \rho \frac{(\alpha-1)}{3(\alpha-3)} t_{\min}^2 \left(\frac{1-r^{\alpha-3}}{1-r^{\alpha-1}}\right) = \frac{\rho E\left[T^3\right]}{3E[T]} \tag{7}$$

$$E\left[Z^k\right] = \rho \frac{(\alpha-1)}{(k+1)\left[\alpha-(k+1)\right]} t_{\min}{}^k \left(\frac{1-r^{\alpha-(k+1)}}{1-r^{\alpha-1}}\right) = \rho \frac{E\left[T^{k+1}\right]}{(k+1)E[T]} \tag{8}$$

The situation encountered by the arrival is the accumulation of q burst (see Fig. 2), each of which following the distribution of T_i, awaiting service on top of residual-service time. The p.d.f. for waiting time for service (represented by random variable W) could be described as in (9).

$$p_W(w \backslash b=1) = p_Z(z \backslash b=1) + p_{T_1}(t_1 \backslash b=1) + \cdots + p_{T_q}(t_q \backslash b=1) \tag{9}$$

Provided that bursts are i.i.d. with p.d.f. $p_T(t)$, independent to B, and that backlog is unbounded, the k^{th} moment for $p_W(t)$ with $\rho<1$ can be found [Takács1962] [Kleinrock1975]. After arithmetical manipulations, a recursive expression using moments from Z is yielded in (10).

$$E\left[W^k\right] = \frac{1}{(1-\rho)} \sum_{j=1}^{k} \binom{k}{j} E\left[Z^j\right] E\left[W^{(k-j)}\right] \tag{10}$$

It is now evident the importance of residual-service times Z over the moments of W. Consequently, one should expected that efforts to reduce moments of Z will be directly reflected in the statistics for waiting time in queue. For k=1, Eq. (10) turns into the well-known Pollaczek-Khintchine mean delay for M/G/1. It is also noteworthy that (10) implies (11).

$$E\left[Z^k\right] < \infty \Rightarrow E\left[W^k\right] < \infty \text{ for } \rho<1 \tag{11}$$

Indeed, it is shown that Eq.(11) holds even for GI/GI/1 systems, i.e. general-but independent-for both arrival and service statistics [Scheller-Wolf2000]. This may be a strong indication of predominance of service over arrival statistics in queuing stability as far as waiting time goes.

3.2 Multiserver scenario (WDM)

Let Z_1, Z_2, ..., Z_n in Fig. 3 be the residual-service times of the n processes intersected by the new arrival. The residual-service time for one out n busy wavelengths to be released under scheduling policy “first available server”, will then be as in (12).

$$Z = \min\left[Z_1, Z_2, \ldots, Z_n\right] \tag{12}$$

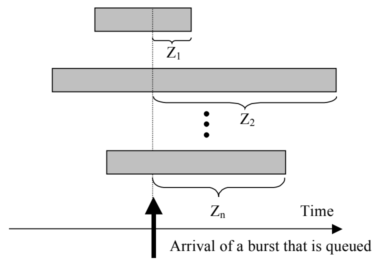

Figure 3. Residual-service time seen by arriving bursts in multiserver scenario.

Let $p_Z^{\{n\}}\left(z \setminus b^{\{1\}}=1,\cdots,b^{\{n\}}=1\right)$ be the p.d.f. for residual-service time under first available server scheduling given that arriving bursts finding all n servers busy [Waldman2004].

$$p_Z^{\{n\}}\left(z \setminus b^{\{1\}}=1,\cdots,b^{\{n\}}=1\right)=\begin{cases}\frac{n}{E[T]}\left(1-\frac{z}{E[T]}\right)^{n-1}, t<t_{\min}\\ n\left(\frac{1}{E[T](1-r^{\alpha})}\right)^{n}\left[\left(\frac{t_{\min}}{z}\right)^{\alpha}-r^{\alpha}\right]\left[z\left(\frac{1}{\alpha-1}\left(\frac{t_{\min}}{z}\right)^{\alpha}-r^{\alpha}\right)-\frac{\alpha}{\alpha-1}t_{\max}r^{\alpha}\right]^{n-1}, t_{\min}\le t<t_{\max}\end{cases} \tag{13}$$

Let $b^{\{i\}}=1$ stand for the following event: i-th wavelength busy. The total load $\rho^{\{n\}}$, which is no longer limited to 1 but to n instead, is evenly distributed across n wavelengths. Provided that any $b^{\{i\}}\neq 1$ leads $p_Z^{\{n\}}(z)$ to zero, then $p_Z^{\{n\}}(z)=p_Z^{\{n\}}(z\setminus b^{\{1\}}=1,\cdots,b^{\{n\}}=1)\left(\rho^{\{n\}}n^{-1}\right)^{n}$ is found. The modeling for queuing with multiple servers is not as simple as in (9) and it is not possible, to the best of our knowledge, to reach closed expressions such as (10) to account for waiting time moments of W in M/G/n. Nevertheless, residual-service time plays a central role in waiting time statistics for single server and multiserver alike. It is shown in [Scheller-Wolf2000] that for GI/GI/n Eq. (11) remains valid to $\rho^{\{n\}}<(n-1)$, as the load per server (ρ) stay bellow 1. Furthermore, under lightly loaded servers, i.e. $\rho^{\{n\}}<1$, finite mean waiting time becomes viable for lower and lower moments of T as n approaches infinity in (14).

$$E\left[T^{(n+1)/n}\right]<\infty \Rightarrow E[W]<\infty \text{ for } \rho^{\{n\}}<1. \tag{14}$$

4. Metrics

The scenarios investigated consider wavelengths at full use, i.e. each server is loaded at $\rho\to 1$. This provides a common framework to compare performance of segregated against non-segregated approaches. Actually, network operators standing point could be well represented by this framework since they seek best use of deployed resources. As a result, the strategy adopted here is to provide comparisons using solely the residual-service time. In doing so, queuing length and the effects of load on waiting time are ruled out as causes for differences in performances of segregated vs. non-segregated systems. This, however, does not mean that such effects are unimportant but only that they do not lie at the root of waiting time moments (see Eq. (10)). In addition, they are not quantifiable for $\rho\to 1$ and an analytical framework encompassing them all, at this stage of investigation, would be rather complex and it would, perhaps, fail to provide conclusive outcomes.

The metrics employed here are overall traffic average delay, delay variance, and blocking time probability. The first two metrics are well known among real-time and interactive services as statistical descriptors for QoS [Giroux1999]. Blocking time probability [Morato2003], in contrast, imposes a hard, i.e. deterministic, deadline to be met by individual bursts, e.g. those bearing critical mission data, instead of usual statistical bounds given by the former two metrics. The size of playout buffers, in real time applications over the Internet, is another illustration for how such deterministic limit is imposed on traffic. Packets experiencing delays above the playout buffer size are discarded. It is important to stress that throughput of bulk traffic is affected by a combination of the metrics. Widespread a non-real time protocol, i.e. TCP, relies on delay statistics to adaptively tune the timeout, which trigger flow control mechanisms [Tanenbaum2003]. As far as burst switching design is concerned, delay experienced by bursts should be kept as close as possible to blocking time so that the length fiber delay lines (used as buffers to deal with contention) are maintained at a minimum [Yoo2000]. The metrics described above are now presented for both (non-segregated) WDM and segregated traffic approaches.

4.1. (Non-segregated) WDM Approach

A pool (with n wavelengths) serves the incoming traffic. Observe in Fig. 4 that $p_T(t)$ is present in all n wavelengths.

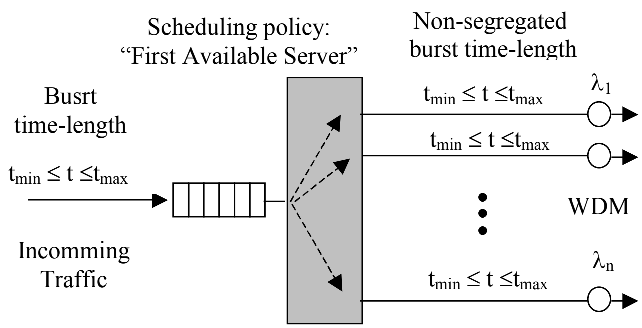

Figure 4. Non-segregated WDM traffic.

The metrics for WDM approach are calculated through straightforward expressions given in (15)-(17), which stand, respectively, for average delay, variance and blocking time probability.

$$\mu_{WDM}(n) = \int_0^{t_{\max}} z.p_z^{\{n\}}(z)dz \tag{15}$$

$$Var_{WDM}(n) = \int_0^{t_{\max}} z^2.p_z^{\{n\}}(z)dz - (\mu_{WDM})^2 \tag{16}$$

$$Pb_{WDM}(n, t_{th}) = \int_{t_{th}}^{t_{\max}} p_z^{\{n\}}(z)dz \tag{17}$$

Note that (17) also depends on the deadline (t_{th}) set and it is, as far as a single-hop and residual-service time alone are concerned, the probability of bursts experiencing delays exceeding t_{th}.

4.2. Segregated Traffic Approach

In contrast to (non-segregated) WDM scheme, each wavelength here bears different slices taken from $p_T(t)$. As it is illustrated in Fig. 5, the incoming traffic goes through a classifier and this element segregates burst according to its length.

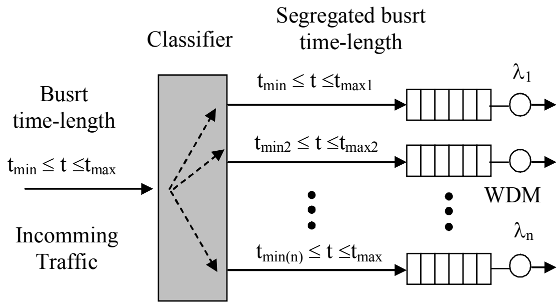

Figure 5. Segregated WDM traffic.

The metrics for segregated traffic, in order to allow direct comparison with (non-segregated) WDM, take the overall delay and its variance as well as the weighted blocking time probability. However, the traffic on each wavelength derives from (1) and the residual-service time from (13) using the set of definitions given in (18).

$$\begin{cases} r_i = t_{\min_i} / t_{\max_i} \\ t_{\min_{i+1}} = t_{\max_i} \\ 0 \le z_i \le t_{\max_i} \\ \mathbf{T}_{\max} = \left(\mathrm{t}_{\max_1}, \mathrm{t}_{\max_2}, \cdots, \mathrm{t}_{\max_{n-1}}\right) \end{cases} \tag{18}$$

The vector containing the maximum holding time for wavelengths 1 to n-1 is denominated $\mathbf{T}_{max}$. The last wavelength receives t_{max} from the incoming traffic itself. The overall mean delay is then stated in (19), which for $\rho \rightarrow 1$ yields (20).

$$\mu_{SGR}\left(n, T_{\max}\right) = \sum_{i=1}^{n} \left[\int_{0}^{t_{\max_i}} z_i . p_{z_i}(z_i) dz_i . \int_{t_{\min_i}}^{t_{\max_i}} p_T(t) dt \right] \tag{19}$$

$$\mu_{SGR}\left(n, T_{\max}\right) = \sum_{i=1}^{n} \left[\frac{(\alpha - 1)}{2(\alpha - 2)} t_{\min_i} \frac{\left(1 - r_i^{\alpha-2}\right)}{\left(1 - r_i^{\alpha-1}\right)} \left(\frac{t_{\min}}{t_{\min_i}} \right)^{\alpha} \frac{\left(1 - r_i^{\alpha}\right)}{\left(1 - r^{\alpha}\right)} \right] \tag{20}$$

The overall variance is presented in (21) and in (22) is its particular case when $\rho \rightarrow 1$.

$$Var_{SGR}\left(n, T_{\max}\right) = \sum_{i=1}^{n} \left[\int_{0}^{t_{\max_i}} z_i^{\ 2} . p_{z_i}(z_i) dz_i . \int_{t_{\min_i}}^{t_{\max_i}} p_T(t) dt - \left(\mu_{SGR}\right)^2 \right] \tag{21}$$

$$Var_{SGR}\left(n, T_{\max}\right) = \sum_{i=1}^{m} \left[\left[\frac{(\alpha - 1)}{3(\alpha - 3)} \cdot t_{\min_i}^2 \cdot \left(\frac{1 - r_i^{\alpha-3}}{1 - r_i^{\alpha-1}} \right) - \left[\frac{(\alpha - 1)}{2(\alpha - 2)} \cdot t_{\min_i} \cdot \left(\frac{1 - r_i^{\alpha-2}}{1 - r_i^{\alpha-1}} \right) \right]^2 \right] \cdot \left(\frac{t_{\min}}{t_{\min_i}} \right)^{\alpha} \cdot \left(\frac{1 - r_i^{\alpha}}{1 - r^{\alpha}} \right) \right] \tag{22}$$

Finally, (23) presents the weighted time blocking probability while (24) brings the expression for the case under investigation in this paper ($\rho \rightarrow 1$).

$$Pb_{SGR}(n, t_{th}, T_{\max}) = \sum_{i=1}^{n} \left[\int_{t_{th}}^{t_{\max_i}} p_{z_i}(z_i) dz_i \cdot \int_{t_{\min_i}}^{t_{\max_i}} p_T(t) dt \right] \tag{23}$$

$$Pb_{SGR}(n, t_{th}, T_{\max}) = \sum_{i=1}^{n} \left[\frac{1}{\alpha\left(1 - r_i^{\alpha-1}\right)} \cdot \left(\frac{t_{\min}}{t_{\min_i}}\right)^{\alpha} \cdot \left[\left(\frac{t_{\min_i}}{t_{th}}\right)^{\alpha-1} + \left(\frac{t_{\min_i}}{t_{th}}\right)^{-1} \cdot r_i^{\alpha} \cdot (\alpha - 1) - \alpha r_i^{\alpha-1} \right] \cdot \left(\frac{1 - r_i^{\alpha}}{1 - r^{\alpha}}\right) \right] \tag{24}$$

5. Burst Length Segregation Policies

Choosing $\mathbf{T}_{\max}$ properly is the crucial point now. Two segregation strategies are investigated, namely, logarithmical and optimal. While the former is based on uniform ratio between $t_{\min}$ and $t_{\max}$ for all wavelengths, the latter provides the upper bound, i.e the best attainable performance for the segregated approach.

5.1. Logarithmical Segregation

In this case the whole range r of the incoming traffic is divided into n segments (r_i is constant) as described in (25).

$$r_i = \sqrt[n]{r} \qquad \text{for i = 1, 2, ... , n.} \tag{25}$$

From (18) $\mathbf{T}_{\max}$ is readily found given r_i.

5.2. Optimal Segregation

It is possible to find optimum $\mathbf{T}_{\max}$ but only for a single metric. Therefore, there is a $T_{\max}^{\{\mu\}}$ for overall mean, a $T_{\max}^{\{\mathrm{var}\}}$ for overall variance, and a $T_{\max}^{\{P_b\}}$ for weighted blocking time probability. Gradient method is here used for minimization of each metric and it can be expressed as in (26) for mean, (27) for variance, and (28) for time blocking probability.

$$\nabla \mu_{SGR}\left(n, T_{\max}^{\{\mu\}}\right) = 0 \tag{26}$$

$$\nabla Var_{SGR}\left(n, T_{\max}^{\{Var\}}\right) = 0 \tag{27}$$

$$\nabla Pb_{SGR}\left(n, T_{\max}^{\{Pb\}}\right) = 0 \tag{28}$$

where $\nabla = \partial / \partial t_{\max_1}, \cdots, \partial / \partial t_{\max_{n-1}}$

6. Results

Results are presented for shaping factor α=1.2 for incoming traffic burst lengths. This shape factor is a representative case for Internet traffic [Crovella1997] and it is assumed that little change is seen after edge nodes build up optical bursts [Necker2003]. Fig. 6 presents results for logarithmical burst segregation against (non-segregated) WDM for the same number n of wavelengths. There are also results from an ad hoc numerical simulator for validation of analytical results. The good agreement between numerical and analytical results is clear in Fig. 6(a) and Fig. 6(b). Nevertheless, results are slightly skewed only for variance in the WDM approach with t_{max} that may reach 10000 times t_{min} (i.e. r =10^{-4}). This difference is basically down to constraints of the numerical simulation in capturing (higher order) statistics of such highly variable traffic using a finite number of events [Gross2002].

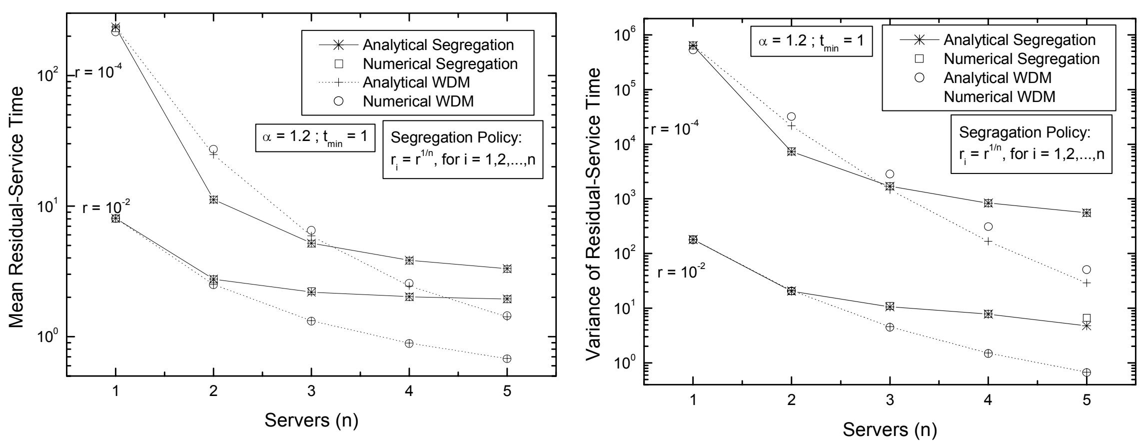

(a) (b)

Figure 6. Residual-service time for log. segregation. (a) Mean and (b) Variance .

The most important outcome from Fig. 6 is the reduction in the overall mean residual delay as well as its variance when using n=2. Nearly ten-fold reduction is seen in variance (Fig. 6(b)) while average delay is decreased to half of its (non-segregated) WDM counterpart with two wavelengths for r=10^{-4}. On the other hand, the WDM approach mostly outperforms systems with segregated traffic elsewhere. Separating the incoming traffic into logarithmically sized limits is just a first approach to the problem. One should actually seek out the maximum improvement by segregating traffic across wavelengths via optimal segregation thresholds.

Results in Fig. 7 give the maximum attainable improvements for the segregated traffic approach against the number of wavelength in use. The best improvement is found at n=2, as in logarithmical segregation, but overall delay and variance are improved for other values of n. For instance, for r=10^{-4} the variance of optimally segregated traffic is still slightly better than the WDM approach for n=3 while with logarithmical splitting in Fig. 6(b) only allow to match WDM performance. It is evident from Fig. 7(a) and Fig. 7(b) that the higher the range of burst variation, the more advantageous is the use of traffic segregation. One order of magnitude delay reduction is found for incoming traffic with r=10^{-6} in Fig. 7(a), while variance in Fig. 7(b) is decreased nearly 100 times compared with WDM for two wavelengths.

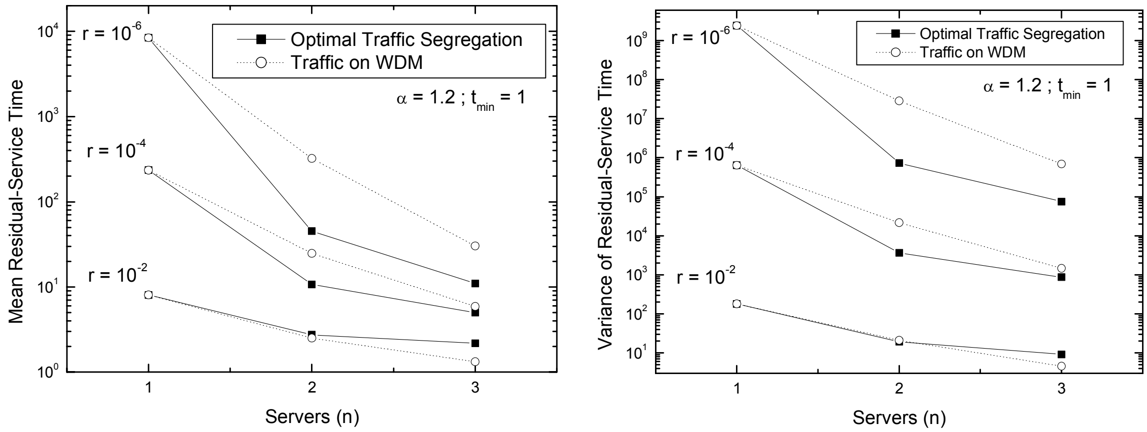

(a) (b)

Figure 7. Residual-service time for optimal segregation. (a) Mean and (b) Variance.

One should bear in mind that $\mathbf{T}_{max}$ allowing the best results for the overall mean delay is diverse from the one that minimizes the variance. Results obtained for logarithmical segregation, on the other hand, uses a single $\mathbf{T}_{max}$. Fig. 8 illustrates the optimum segregation points under n=2, i.e. $\mathbf{T}_{max} = \{t_{max1}\}$, for overall mean and variance. Indeed, quite different optimal points are obtained depending on whether minimization is aimed at mean or variance. For example, with $r = 10^{-6}$ in Fig. 8(a) yields t_{max1}=512.86 as the best choice to reduce mean, which contrast with t_{max1}=5980.5 in Fig. 8(b) for achieving the least possible variance in the overall traffic under the same conditions.

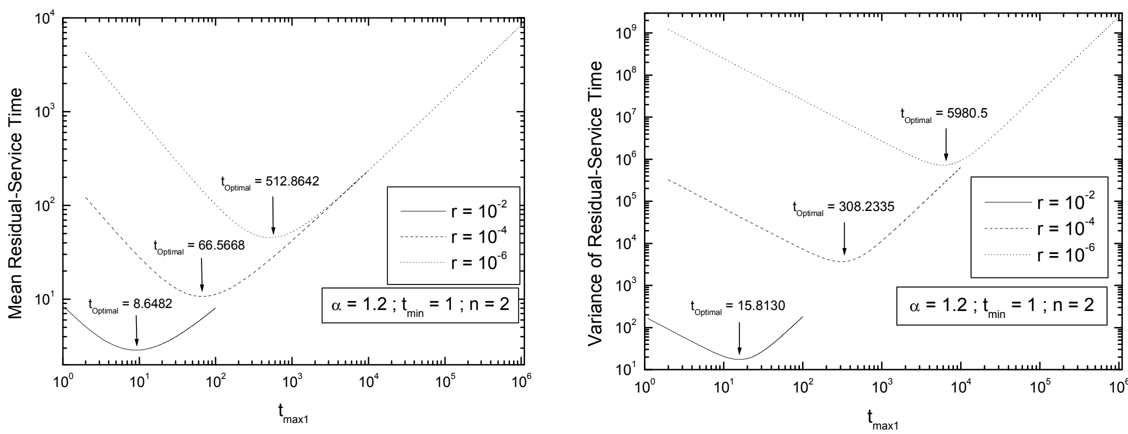

(a) (b)

Figure 8. Residual-service time showing optimal $\mathbf{T}_{max}$. (a) Mean and (b) Variance.

Note that logarithmical segregation with n=2 would use t_{max1} =10 (for $r=10^{-2}$), 100 (for $r=10^{-4}$), and 1000 (for $r=10^{-6}$). This result helps explain the good performance of logarithmical segregation found in Fig. 6. For logarithmical segregation, $\mathbf{T}_{max}$ presented above falls in between optimum values obtained for mean and variance in Fig. 8. More elaborated trade-off solutions for $\mathbf{T}_{max}$, i.e. those properly balancing mean and variance in order to meet particular performance targets of applications using OBS networks, can be found by merging (26) and (27) into (29).

$$\nabla\left[\mu_{SGR}\left(n,T_{\max}\right)+\beta Var_{SGR}\left(n,T_{\max}\right)\right]=0 \tag{29}$$

where β is the parameter gauging the relative importance of the two objectives. Taking values of $\beta\rightarrow 0$ one can represent the predominance of mean while with $\beta\rightarrow\infty$ is used to express that the major concern is the traffic overall variance. As a result, it should be expected the convergence of $\mathbf{T}_{\max} \rightarrow T_{\max}^{\{\mu\}}$ and $\mathbf{T}_{\max} \rightarrow T_{\max}^{\{\mathrm{var}\}}$, respectively, for these extreme values for β.

The last results presented here concern blocking time (P_b) performance. Recall that this metric assesses the probability of bursts arriving after a given timeout (represented by t_{th}). It is then clear that $P_b\rightarrow 0$ as $t_{th}\rightarrow t_{max}$. This means that all bursts meet the deadline since no intersected burst will delay the next burst for more than t_{max}. Provided two wavelengths are available, Fig 9 shows the time blocking probability P_b against t_{th} for (a) $r=10^{-2}$ (b) $r=10^{-4}$ (c) $r=10^{-6}$. In addition, Fig.9(d) presents $t_{max1}=t_{opt}$ found for each value of deadline (t_{th}).

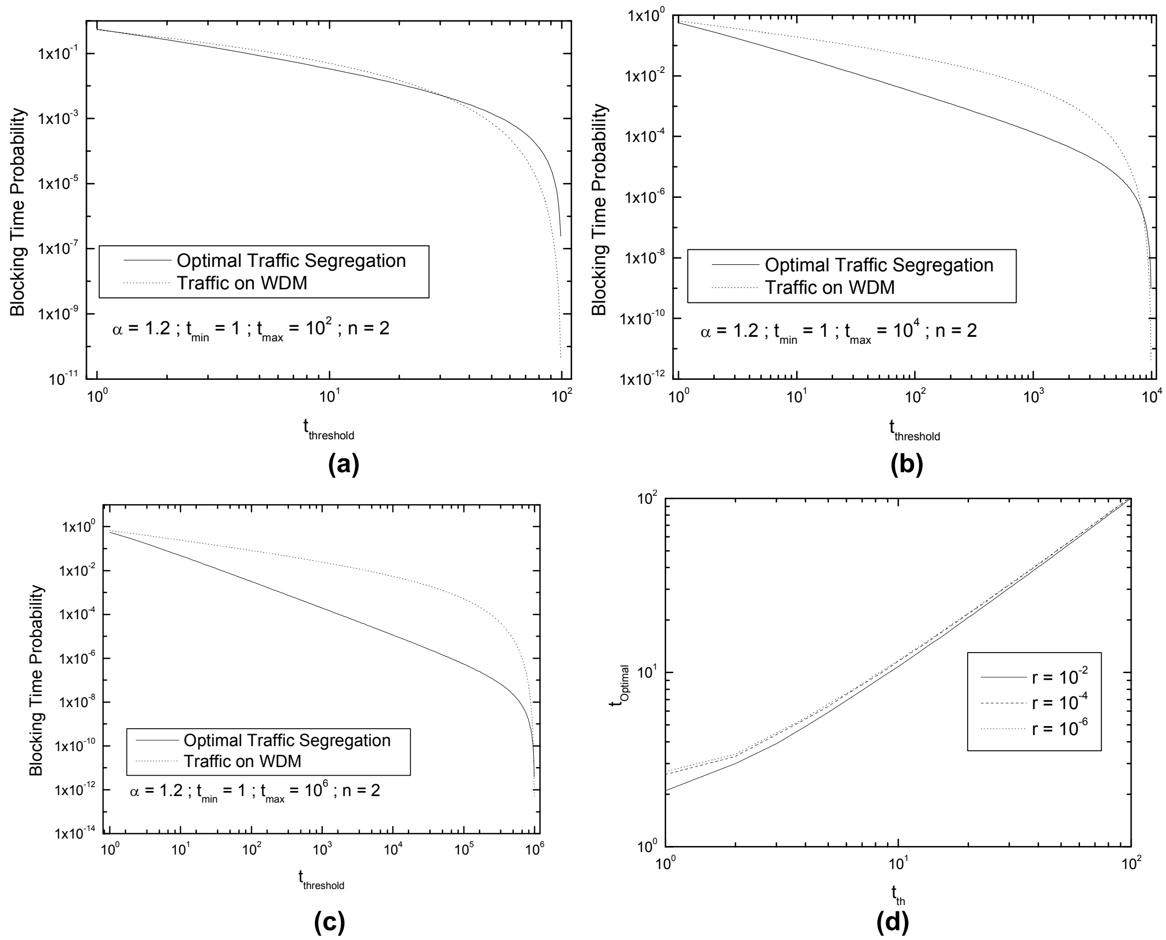

Figure 9. Blocking time probability for optimally segregated traffic with two wavelengths. (a) $r=10^{-2}$. (b) $r=10^{-4}$. (c) $r=10^{-6}$. (d) Optimal segregation against traffic deadline.

On the whole, similar trends for the relationship between improvements brought by traffic segregation and the maximum allowed bursts lengths are once more observed. For instance, while modest results are obtained for r $=10^{-2}$ in Fig. 9(a), it is possible to significantly reduce P_b from 10^{-3} to 10^{-6} when t_{th} is set at 10^5 in Fig. 9(c). Bursts differing in as much as six orders of magnitude, i.e. $r=10^{-6}$, benefit more from traffic segregation than those differing in only one hundred times. However, these results should be interpreted along with curves in Fig. 9(d). The optimal values for t_{max1} actually point out to the fact that $t_{max1} \rightarrow t_{th}$ for deadlines set beyond 10^2. This means that the first wavelength should bear the traffic meeting the deadline while the second carries traffic that is very likely breach it. Consequently, the second wavelength is solely used by traffic that most of it will be discarded at the receiver end. Indeed, the optimal traffic segregation in Fig. 9 perhaps suggests to network designers that the income traffic with deadlines beyond 10^2 should be truncated at t_{th} and carried on a single wavelength as the best engineering solution to the problem of OBS supporting traffic with hard deadlines. Nevertheless, for stringent deadlines (e.g. $t_{th}=10^2$) there is use for traffic segregation with one and two order of magnitude improvements for over all performance, as seen in Fig. 9(b) and (c), respectively.

7. Conclusion

This paper presented an analytical approach (along with numerical validation) aimed at best use of WDM resources. Traffic transported separately, according to its holding time, was compared with highly variable demands placed upon a pool with n wavelengths. Traffic segregation proved to be an advantageous design strategy from the standpoint of reduction of residual-service times. Under realistic traffic assumption ($\alpha=1.2$), benefits start to emerge as soon as the maximum allowed holding time exceeds the shortest one by two orders of magnitude. The larger is the ratio between them, the more prominent are the improvements brought by segregation. Particularly, the best results come into view at n=2. Significant reductions, of one and two orders of magnitude, were found for mean and variance, respectively, with $r=10^{-6}$. Furthermore, time blocking probability, under same conditions, may reach impressive figures around one thousand times better than non-segregated WDM.

This paper highlights the fact that residual-service time is the most important part of the problem of queuing under heavy-tailed service so that enhancements found in delay statistics are likely to prevail the test of realistic traffic load scenarios. Although it managed to fairly contrast two different approaches, one should bear in mind that comparisons were carried out under fully loaded wavelengths. Queuing system are not stable at this point. Improvement figures may change as sound traffic loads are used. For instance, notice that a factor $(\rho^{\{n\}}/n)^n$ multiplies the residual-service time for (non-segregated) WDM approach. This may lessen the differences found once after segregation wavelengths may become unevenly loaded. On the other hand, segregation allows fairness improvements and may accept a convenient incoming traffic shape in order to set wavelength traffic loads accordingly. Therefore, finding out the proper traffic load for stable queues is an issue in its own right and future investigations will tackle that along with the development of new performance metrics.

Acknowledgements

This work was supported by CAPES, CNPq, and Project Giga/CPqD/Finep/Funttel.

References

Buchta, H.; Patzak, E., Saniter, J. and Gauger, C. (2003) "Limits of effective throughput of optical burst switches based on semiconductor optical amplifiers", Optical Fiber Communications Conference (OFC 2003), p. 215 - 217 vol.1.

Chan, W.C., Lu, T.-C and Chen, R.-J (1997) "Pollaczek-Kinchin formula for the M/G/1 queue in discrete time with vacations.", in IEE Proc.-Comput. Digit.Tech, vol. 144, no. 4, pp. 222-226.

Crovella, M.E. and Bestavros, A. (1997) "Self-Similarity in World Wide Web traffic: evidence and possible causes", IEEE/ACM Transactions on Networking, vol. 5, n. 6, pp. 835-846.

Giroux N., S. Ganti, (1999) "Quality of service in ATM networks: state-of-the-art traffic management", Prentice-Hal,l New Jersey.

Gross, D., Shortle, J.F., Fischer, M.J. and Masi, D.M.B. (2002) "Difficulties in Simulating Queues with Pareto Service", in Proc. 2002 Winter Simulation Conference, p. 407-415.

Kleinrock, L. (1975) "Queueing Systems, vol.1: Theory", Wiley.

Leland, W., Taqqu, M., Willinger, W. and Wilson, D.V. (1994) "On the self-similar nature of Ethernet traffic", IEEE/ACM Transactions on Networking, vol. 2, pp. 1-15.

Morato, D., Izal, M., Aracil, J., Magaña, E. and Miqueleiz, J. (2003) "Blocking time analysis of OBS routers with arbitrary burst size distribution", in Proceedings of GLOBECOM 2003, pp. 2488-2492.

Necker, M., Gauger, C. and Bodamer, S, (2003) "Does burst assembly really reduce the self-similarity?", Optical Fiber Communications Conference (OFC 2003), p. 487 – 488.

Qiao, C. and M. Yoo, (1999) "Optical Burst Switching (OBS) - A New Paradigm for an Optical Internet", Journal of High Speed Networks, vol. 8, pp. 69-84, Jan 1999.

Roberts, J.W. (2000) "Engineering for Quality of Service", in Self-Similar Network Traffic and Performance Evaluation, Ch. 1, edited by K. Park and W. Willinger, Wiley.

Scheller-Wolf, A. (2000) "Further Delay Moment results for FIFO Multiserver Queues", Queueing Systems: Theory and Applications Vol. 34 , Issue 1-4, pp. 387-400.

Takács, L. (1962). Introduction to the Theory of Queues, Oxford University Press, NY.

Tanenbaum, A. S. (2003) "Computer Networks" Fourth Edition, Prentice Hall PTR.

Waldman, H., Trazzi Jr.,F.L. and Ribeiro, M.R.N. (2004) "Statistics for Residual Lifetime of Intersected Bursts in QoS-Aware WDM Networks", in IEEE The Third International Workshop on Optical Burst Switching (WOBS 2004), (http://www.broadnets.org/2004/workshop-papers/OBS/WOBS310_Waldman_final.pdf)

Yoo, M., C. Qiao and S. Dixit, (2000), "QoS Performance of optical burst switching in Ip-Over-WDM" IEEE JSAC., vol. 18.no10, pp. 2062-2071, Oct. 2000.