

Buffer Size Requirements for Delay Sensitive Traffic Considering Discrete Effects and Service-Latency in ATM Switches

Steven Wright¹ and Yannis Viniotis²

¹ BellSouth Science & Technology, 41G70 BSC, 675 West Peachtree St NE,
Atlanta, GA 30375 USA

Steven.wright@snt.bellsouth.com

² Dept. ECE, North Carolina State University, Raleigh , NC 27695-7911,USA.
candice@eos.ncsu.edu

Abstract. Various approaches to buffer size and management for output buffering in ATM switches supporting delay sensitive traffic are reviewed. Discrete worst case arrival and service functions are presented. Using this format, bounds are developed for buffer size under zero cell loss for leaky bucket constrained sources. Tight bounds are developed for the case of discrete arrival functions with fluid servers and fluid arrival functions with discrete servers. A bound on the buffer size is also proposed for the case of discrete arrival and service process. While this bound is not exact, the maximum gain that could be achieved by a tighter bound is bounded. In some cases it is possible to reduce the buffer size requirements through over allocation of link bandwidth. Feasibility conditions for this scenario are developed.

1 Introduction

ATM network support [13] for delay-sensitive QoS requirements requires ATM switches to support the QoS guarantees. Low service-latency schedulers, e.g. RRR [4] provide a mechanism for supporting delay QoS guarantees if sufficient switch resources - including buffer space and bandwidth - are available. The focus of this paper is the buffer requirements rather than the scheduler design. The buffer size constraints switch performance (e.g. Connection Admission Control) while representing a significant fraction of interface costs for wire speed interfaces. This paper develops expressions for the buffer size requirements (see equations (1),(11),(12)and (13)) that can be used in buffer optimization problems for specific schedulers, particularly simple Latency-Rate (\mathcal{LR}) schedulers (see [10]) derived from Weighted Round Robin. These expressions are used to explore potential buffer reductions. In contrast to other buffer studies (see [6],[8],[9]), we provide more precise formulations for the buffer requirements due to discrete effects in the arrival and service processes as well as considering service-latency. The paper also provides quantified examples with a specific WRR scheduler to achieve reductions in buffer size requirements. We assume a simple output buffering arrangement such as that used for Burstiness-Class based Queuing (B-CBQ, see [14]), where several queues (α,β,γ) are serviced by an \mathcal{LR} scheduler which provides each queue with a guaranteed

bandwidth and (potentially) delay bound guarantees. The buffer size formulations in this paper focus on the buffer requirements for an individual queue. Multiple Pt-Pt or MPt-Pt ATM connections are allocated to a queue (i.e. connections $VC_{\alpha,1..n}$ are directed to queue α). \mathcal{LR} Delay bound guarantees require source traffic that is leaky bucket (σ, ρ) constrained with a worst case burst of size σ , and arrival rate ρ . In this paper, we assume worst-case aggregate leaky-bucket-constrained arrival function at a queue, i.e., periodic bursts of size σ . (see e.g. [3]). The buffer is cleared with a period T given by the ratio σ/ρ . We are particularly interested in the case where the buffer requirements can be reduced from the maximum burst size (σ). Buffer size requirements are typically considered at switch design time, however for some switch designs buffers may be reallocated between queues while the switch is operational. For such situations, the computational complexity of estimating buffer requirements is an important issue. The paper is organized as follows. Section 2 places this work in the context of prior work. The buffer size, assuming fluid models for arrivals and service, and also considering service-latency, is considered in section 3. The buffer size requirements under discrete arrival and/or service functions (again considering service- latency) are discussed in section 4. Section 5 explores numerically the feasibility and magnitude of buffer size reductions possible by rate over-allocation for WRR schedulers. Conclusions are provided in section 6.

2 Prior Work

Previous work had focussed largely on ideal fluid arrival and service processes (e.g. WFQ) or had simply assumed that the allocated rate matched in the requested rate exactly, rather than considering some potential for over allocation of bandwidth in order to minimize buffer occupancy. The absolute delay bounds both for the GPS schedulers and the more generic version for \mathcal{LR} schedulers [10] rely on a maximum buffer size and a guaranteed service rate in order to produce the delay bound. The tradeoff of additional guaranteed rate for reduced buffer requirements was identified by [3] in the context of work on equivalent bandwidth capacity. They developed an approach to buffer allocation for ATM networks which reduced the two resource allocation problems (buffer & bandwidth) to a single resource allocation (“effective bandwidth”) problem. [8] built on the work of [3] and separated the buffer/bandwidth tradeoff into two independent resource allocation problems. While work on scheduler design (e.g.[4],[10]) has identified several issues related to service-latency, the impact on the buffer-bandwidth tradeoff has not been explicitly considered. The theory related to fluid models of arrival and service curves has recently been extended into an elegant network calculus (e.g., [1], [2], [7]). While these approaches typically are used to derive end-to-end network delay bounds, they can also be used to provide assertions regarding buffering requirements. In the realm of equipment design, the service curves, and network calculus, are an intermediate form and equipment optimizations must be recast in terms of scheduler design parameters. In practice, ATM switches must deal with discrete data units (cells) and consider the effects of service-latency on buffer requirements. Previous work on discrete buffer sizing for

WRR has been largely empirical (e.g. [9], [6], [5]) and addressed towards loss based QoS parameters, rather than delay-sensitive QoS requirements.

3 Buffer Size with Fluid Arrival and Fluid Service with Service-Latency

The main result for the buffer size (b) in the case of fluid arrivals and service processes is given in equation (1). Three cases must be considered:

- if the service-latency (L) is such that service starts after arrivals have reached the maximum value (which occurs at $t=\tau_\sigma$), then the supremum occurs at $t=\tau_\sigma$ and $b=\sigma$, and no reduction in buffer requirements is possible.
- if the arrival rate (ρ_a) is less than (or equal to) the service rate (ρ_s), then the supremum occurs at $t=L$, and $b=\rho_a L$.
- if the arrival rate is greater than the service rate, then the supremum occurs at $t=\tau_\sigma$ and $b=\sigma-\rho_s(\tau_\sigma-L)$.

$$b \geq \begin{cases} \sigma & \text{if } \{L \geq \tau_\sigma\} \\ \sigma - \rho_s(\tau_\sigma - L) & \text{if } \{L < \tau_\sigma\} \text{ and if } \{\rho_a > \rho_s\} \\ \rho_a L & \text{if } \{L < \tau_\sigma\} \text{ and if } \{\rho_a \leq \rho_s\} \end{cases} \quad (1)$$

An example of the buffer space requirements is illustrated in Fig. 1 using equation (1). In order to reduce the buffer requirements below σ , the service-latency must be less than σ/ρ_a , (10msec in this example). Simply increasing the service rate may not reduce the buffer requirements below σ . Even for the cases where $L < 10\text{msec}$, increasing the service rate beyond ρ_a provides no additional benefit which is consistent with the results of [8]. If the service rate, ρ_s , is less than the peak arrival rate, ρ_a , then only smaller reductions from σ are possible. While the potential reduction of 10 cells for the connection in this example may not seem significant, we recall that there may be several thousand connections on an ATM interface at OC-3 or higher rates. Also, the burst size, used in this example is very small for VBR traffic.




Fig. 1 Buffer Space Requirements for Fluid Arrivals and Fluid Service with Service-Latency. The parameters for the example are: $\sigma = 10$ Cells; $\rho = 250$ Cells/Sec; $\rho_a = 1000$ Cells/Sec; $500 \leq \rho_s \leq 5000$ Cells/Sec; $0 \leq L \leq 40$ msec.

4 Effect of Discrete Arrivals and Discrete Service on Buffer Size

The fluid model ignores discrete effects in the arrival and service functions. Significant buffer reductions are still possible, even after allowing for discrete effects in the arrival and service functions. In contrast to the fluid arrival functions typically based on two (σ, ρ) or three (σ, ρ, ρ_a) leaky bucket parameters, we use a worst-case leaky-bucket discrete arrival function based on four parameters— $(\sigma, \rho, \rho_a, k)$; where k is the step size in the same data units as σ (see Fig. 4). Similarly, we need to move from the two-parameter (L, ρ_s) fluid service function of section 0 to a worst-case discrete service function based on three parameters (L, ρ_s, m) where m is the service step size in the same data units as σ (see Fig. 2). The main result of this section is the formulation for the discrete arrival function in equation (4) and for the discrete service function in equation (8). Also presented are a pair of fluid functions that bound each of these discrete functions. These are equations (5) and (6) for the bounds on the discrete arrival function and equations (9) and (10) for the bounds on the discrete service function.

4.1 Arrival Functions

The incoming arrivals are not eligible for service until the time, τ_a , given by equation (2). Equation (3) illustrates the time, τ_σ , at which the maximum arrival burst, σ , is reached. The worst-case discrete arrival function is then defined by equation (4). We can consider the discrete arrival function as bounded by two fluid functions: $A_{\max}(t)$ (refer Equation (5)) and $A_{\min}(t)$ (refer equation (6)). The maximum error in the bounding functions is one step, k . As $k \rightarrow 0$, the discrete model is less relevant and $A(t)$ converges towards the $A_{\min}(t)$ function.

$$\tau_a = \frac{k}{\rho_a} \quad (2)$$

$$\tau_\sigma = \tau_a \left\lceil \frac{\sigma - k}{k} \right\rceil \text{ for } \sigma \geq k. \quad (3)$$

$$A(t) = \begin{cases} 0 & t \leq 0 \\ k \left(1 + \left\lfloor \frac{t}{\tau_a} \right\rfloor \right) & 0 \leq t < \tau_\sigma \\ \sigma & \tau_\sigma \leq t \leq T \end{cases} \quad (4)$$

$$A_{\max}(t) = \begin{cases} 0 & t \leq 0 \\ k + \rho_a t & 0 \leq t < \tau_\sigma \\ \sigma & \tau_\sigma \leq t \leq T \end{cases} \quad (5)$$

$$A_{\min}(t) = \begin{cases} 0 & t \leq 0 \\ \rho_a t & 0 \leq t < \tau_\sigma \\ \sigma & \tau_\sigma \leq t \leq T \end{cases} \quad (6)$$

4.2 Service Functions

The generalized worst case discrete service function is shown in Fig. 2, where service proceeds in discrete steps up until the maximum buffer occupancy has been served. Many link schedulers offer rate guarantees over a longer time-scale than one cell transmission time, and service some quanta ($m > 1$ cell) of data at a step [11]. The period, τ_s , associated with a service step is given by equation (7). We chose to separate the effects and leave L to reflect service-latency associated with the start of the service function beyond the scheduling rate granularity. This formulation reflects that buffers are freed at the end of the service time-scale. A discrete version of the service function is then given by equation (8). The worst case discrete service function is bounded by two fluid service functions: $S_{max}(t)$ (refer equation (9)) and $S_{min}(t)$ (refer equation (10)). The maximum difference between in the bounds is m . As $m \rightarrow 0$, $\tau_s \rightarrow 0$ and the discrete model becomes less relevant as it converges towards $S_{max}(t)$.

$$\tau_s = \frac{m}{\rho_s} \quad (7)$$

$$S(t) = \begin{cases} 0 & 0 \leq t \leq L \\ \min(A(t-L), m \left\lfloor \frac{t-L}{\tau_s} \right\rfloor) & L \leq t \leq T \end{cases} \quad (8)$$

$$S_{max}(t) = \begin{cases} 0 & 0 \leq t \leq L \\ \min(A(t-L), \rho_s(t-L)) & L \leq t \leq T \end{cases} \quad (9)$$

$$S_{min}(t) = \begin{cases} 0 & 0 \leq t \leq L + \tau_s \\ \min(A(t-L-\tau_s), \rho_s(t-L-\tau_s)) & L + \tau_s \leq t \leq T \end{cases} \quad (10)$$

4.3 Fluid Arrivals and Discrete Service

Consider the fluid arrival process and discrete service process with service-latency in Fig. 2(a). Using the discrete service function may provide an improved result (a potential reduction up to m) if we can evaluate the supremum. The main result is presented in equation (11) (refer to [12] for the proof). In brief, the intuition on locating the supremum from the fluid case is extended by whether the supremum occurs at the first step in the service function after $t=L$, i.e. $t=\tau_2$, or at the last step in the service function prior to $t=\tau_\sigma$ i.e. $t=\tau_5$. An example of the per-connection buffer requirements for the case of $m=3$ is shown in Fig. 2(b), where the supremum occurs at τ_2 (~6msec). The arrival function used was a fluid arrival with parameters $\{\sigma = 10$ Cells; $\rho = 250$ Cells/Sec; $\rho_a = 1,000$ Cells/Sec $\}$.

$$b \geq \begin{cases} \sigma & \text{if } \{\tau_2 \geq \tau_\sigma\} \\ \sigma - m \left\lfloor \frac{\tau_\sigma - L}{\tau_s} \right\rfloor & \text{if } \{\tau_2 < \tau_\sigma\} \text{ and if } \{\rho_a > \rho_s\} \text{ and if } m \leq \sigma - \rho_a \tau_s \\ \rho_a \tau_s - m \left\lfloor \frac{\tau_\sigma - L}{\tau_s} \right\rfloor + m & \text{if } \{\tau_2 < \tau_\sigma\} \text{ and if } \{\rho_a > \rho_s\} \text{ and if } m > \sigma - \rho_a \tau_s \\ \tau_2 \rho_a & \text{if } \{\tau_2 < \tau_\sigma\} \text{ and if } \{\rho_a \leq \rho_s\} \end{cases} \quad (11)$$

At low latencies, when the service rate exceeds the peak arrival rate, the buffer requirements can be significantly reduced. With $\rho_s = \rho_a = 1000 \text{ cells/sec}$ in Fig. 3, the buffer requirements are reduced to 3-8 cells depending on the latency of 0-5mS. This is a reduction of 20-70% from a buffer size based on the peak burst size, $\sigma (=10 \text{ cells}$ in this example). When the service rate is below the peak arrival rate, the effects of increasing service-latency are impacted by the service step ($m=3$) as illustrated by the series of plateaus in Fig. 3. These are not seen in the fluid arrival and service model shown in the example of Fig. 1.

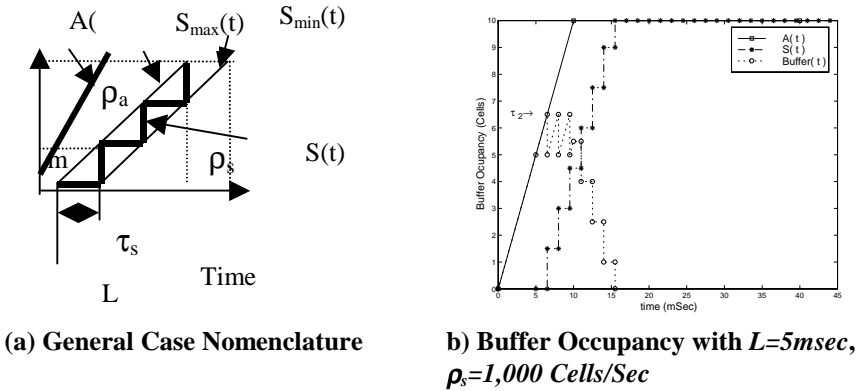


Fig. 2 Fluid Arrivals and Discrete Service




Fig. 3 Buffer Size as a Function of L and ρ_s for Worst-Case Fluid Arrivals and Discrete Service with service-latency


4.4 Discrete Arrivals and Fluid Service

The main result in this section is equation (12) (refer to [12] for the proof), which presents a formula for the worst case buffer requirements when there is a discrete arrival function and a fluid service function. In brief, the intuition on locating the supremum from the fluid case is extended by whether the supremum occurs at the first step in the arrival function after $t=L$, i.e. $t=\tau_1$, or at the last step in the arrival function prior to $t=\tau_\sigma$, i.e., $t=\tau_4$. The formulation is made more complex by consideration of

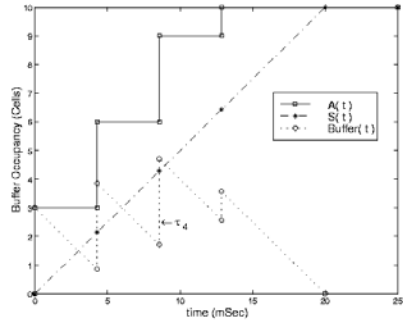
the arrival step size, k , in relation to the maximum burst size σ . Fig. 4 (a) illustrates the general case. In Fig. 4(b), a discrete arrival function is considered that has the following characteristics: $\sigma = 10$ Cells; $\rho = 400$ Cells/Sec; $\rho_a = 7$ Cells/Sec; $k = 3$ Cells. In Fig. 4(b) time is shown on the x-axis, and $A(t)$ is shown as a solid line. $S(t)$ is shown as alternating dash-dot line, and the maximum buffer occupancy is shown as a dotted line. In Fig. 4(b) the last step is less than k . The maximum buffer occupancy occurs at τ_4 (~ 9 msec).

$$b \geq \begin{cases} \sigma & \text{if } \{L \geq \tau_\sigma\} & \text{or if } \{k \geq \sigma\} \\ \sigma - (\tau_\sigma - L)\rho_s & \text{if } c1 = \text{TRUE} & \text{and if } A(\tau_4) \geq \sigma - \tau_a\rho \\ A(\tau_4) - \rho_s(\tau_\sigma - L - \tau_a) & \text{if } c1 = \text{TRUE} & \text{and if } A(\tau_4) < \sigma - \tau_a\rho \\ k \left\lceil \frac{L}{\tau_a} \right\rceil & \text{if } c2 = \text{TRUE} & \text{and if } \left\{ k \leq \frac{\tau_1 - L}{\rho_s} \right\} & \text{and if } \tau_1 \neq \tau_\sigma \\ A(\tau_1) - S(\tau_1) & \text{if } c2 = \text{TRUE} & \text{and if } \left\{ k > \frac{\tau_1 - L}{\rho_s} \right\} & \text{and if } \tau_1 \neq \tau_\sigma \\ \max(A(L), (A(\tau_1) - S(\tau_1))) & \text{if } c2 = \text{TRUE} & \text{and if } \tau_1 = \tau_\sigma \end{cases} \quad (12)$$

where $c1 = \{L < \tau_\sigma\} \text{ AND } \{\rho_a > \rho_s\}$
 $c2 = \{L < \tau_\sigma\} \text{ AND } \{\rho_a \leq \rho_s\}$



(a) $\tau_a < L < \tau_\sigma, \rho_s > \rho_a$



(b) Buffer Occupancy with $L=0$ msec, $\rho_s=500$ cells/s The discrete arrival function used was: $\sigma = 10$ Cells; $\rho = 250$ Cells/S; $\rho_a = 1000$ Cells/S; $k = 3$ Cells.

Fig. 4 Discrete Arrivals and Fluid Service

Fig. 5 illustrates the buffer size behavior as a function of L and ρ_s . Note that the minimum buffer size in Fig. 5 is not zero, but the $k(=3$ cells) on the vertical axis. When $\rho_s > \rho_a$, the buffer size shows discrete steps due to the discrete arrivals. When $\rho_s < \rho_a$, the buffer size is a continuous function. Although this example uses a small value

of σ (~10 cells), for connections with large values of σ , the potential reduction in buffer requirements could be significant. This numerical example is also considering the effects on one connection. Aggregating the result over the number of connections supported at the interface provides for a potentially large reduction in buffer requirements. For example, if we can operate with a low latency of ~1msec and a $\rho_s = 3,000$ Cells/sec, a 70% reduction in buffer requirements can still be achieved even after considering discrete arrival effects.




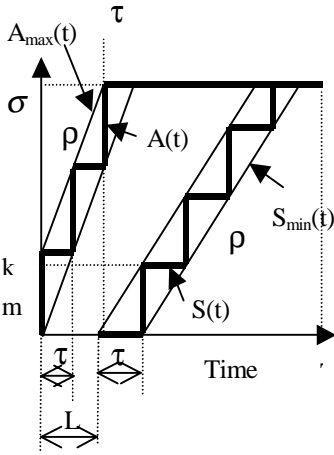
Fig. 5 Buffer Occupancy Supremum as function of L and ρ_s

4.5 Combined Discrete Arrival and Departure Effects


The effect of combining discrete arrivals and service is shown in Fig. 6(a). The intuition followed in the fluid cases to locate the supremum is not sufficient here. In Fig. 6(b), rather than showing the evolution of discrete $A(t)$ and $S(t)$ functions (as in e.g. Fig. 2(b)) we chose to show only the detail of the buffer occupancy for an example where the supremum does not occur at the location intuitively expected. Consider the case where the periods (τ_a, τ_s) of the arrival and service curves are not equal. As the phase of the two discrete functions changes, eventually a point is reached where one discrete function steps twice between two steps in the other discrete function. This results in local maxima or minima in the buffer occupancy. This invalidates the previous intuition about where the supremum occurs.

Where both k and m are small with respect to σ , the bounding model considering discrete effects is appropriate. The fluid model bounding the discrete arrival function was derived as equation (5). The worst-case fluid model bounding the discrete service function was derived as equation (10). Combining these two leads to equation (13) (refer to [12] for the proof). If only one of k, m was small, then a bounding model could be derived based on the material in the previous sections. The supremum can be evaluated by exhaustively computing the buffer occupancy at each discrete step in the arrival or service functions. If neither k nor m is small, then there will be few steps making the numerical solution of the buffer supremum faster.

$$b \geq \begin{cases} \sigma & \text{if } \{(L + \tau_s) \geq \tau_\sigma\} \\ \sigma - (\tau_\sigma - L - \tau_s)\rho_s & \text{if } \{(L + \tau_s) < \tau_\sigma\} \text{ and } \{\rho_a > \rho_s\} \\ k + (L + \tau_s)\rho_a & \text{if } \{(L + \tau_s) < \tau_\sigma\} \text{ and } \{\rho_a \leq \rho_s\} \end{cases} \quad (13)$$



(a) General Nomenclature



(b) Buffer Occupancy Behavior Detail

Fig. 6 Combining Discrete Arrivals and Service

Fig. 7 illustrates the bound on buffer size required after considering the discrete effects using equation (13). This is very similar in shape to Fig. 1, with the minimum possible buffer size now being k (3 Cells in this example) rather than zero as in Fig. 1. The position of the knee where buffer reduction below σ occurs is also moved slightly (to $L \sim 8$ msec) due to the additional service-latency introduced by discrete service effects. Despite the discrete effects, significant buffer reductions are still possible where the service-latency can be reduced below this knee. As in Fig. 1, when the latency is below the knee point, increasing the service rate to match the peak arrival rate provides potential buffer reductions, increasing beyond the peak arrival rate does not provide any additional reduction in buffer size. Although reduction in the number of cells in this example is small, the potential reduction is significant for bursty connections with large values of σ , and when considered aggregated over the potentially large numbers of connections on an interface.




Fig. 7 Bound on Buffer Space Considering Discrete Effects with $k=3, m=2$

5 Buffer Size Reduction in WRR

In this section, we provide a numerical illustration of the feasibility and scale of buffer size reduction for a specific WRR scheduler. We assume a node with n identical input lines of speed C Cells/Sec. The maximum number of simultaneous arrivals is then one per link, i.e. $k=n$ (for a MPt-Pt connection). The peak arrival rate is $\rho_a = kC$, and the maximum aggregate arrivals are still constrained to σ . τ_a is given by Equation (14). The maximum number of arrivals in the worst-case-burst have been received by the time given in equation (3). Consider a scheduler that offers a rate guarantee of some fraction of the link bandwidth (C) to a class i . This class is given a integer weight ϕ_i . A simple interpretation is that the integer weights, ϕ_i , correspond to the number of units of service to be provided in one round by the WRR server, where the units of the weights are the units of service (e.g. bits, cells, packets). We assume the step size for the discrete service is $m = \phi_i$. Then ρ_s is given by equation (15), and τ_s is given by equation (16)

$$\tau_a = \frac{k}{\rho_a} = \frac{1}{C} \quad (14)$$

$$\rho_s = \frac{\phi_i}{\sum_{\forall j} \phi_j} C \quad (15)$$

$$\tau_s = \frac{m}{\rho_s} = \frac{\sum_{\forall j} \phi_j}{C} \quad (16)$$

The worst-case service-latency (L) is the service-latency due to the discrete arrival processes. i.e. $L = \tau_a$. The buffer space requirements (assuming k and m are small with respect to σ) corresponding to this can then be derived by substituting in equation (13).

5.1 Feasible Region for Reduced Buffer Size

We assume k is a positive integer, and for all non-trivial cases $C > \rho_s$ and hence $\rho_a \geq \rho_s$. From equation (13), we have the constraint on service starting before τ_σ . This can be developed into a constraint on the maximum frame size as shown in equation (17). This constraint is illustrated in Fig. 8 where frame sizes that permit buffer reduction are in the region below the surface. The frame size constraint is linear in σ , but has an inflexion point as a function of k . In order to have a reasonably large frame size to accommodate connections with large σ , we need to keep k small, e.g. $k < 10$.

$$\text{Maximum frame size} = \sum_{\forall j} \phi_j < \frac{\sigma_i - k}{k} - 1 \quad (17)$$

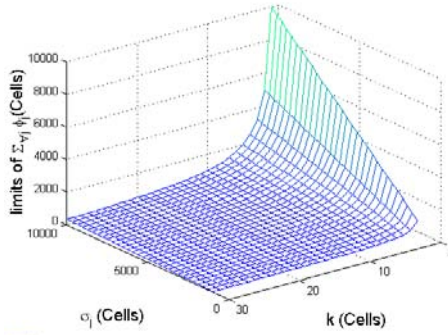



Fig. 8 maximum frame size ($\sum_{\forall j} \phi_j$) as a function of k and σ

5.2 Reduced Buffer Size


Equation (13) can be reduced in this region of interest as equation (18).

$$b_i \geq \sigma_i - \left(\frac{\sigma_i - k}{k} - 1 - \sum_{\forall j} \phi_j \right) \frac{\phi_i}{\sum_{\forall j} \phi_j} \tag{18}$$

Fig. 9 illustrates the behavior of equation (18). To avoid worst case buffer overruns, the buffer should be provisioned above this surface. Significant reductions in buffer size are only possible for small values of k e.g. systems with small fan-in. Although there is some sensitivity to the burst size, σ_i , at large values of the burst size σ_i , even a small percentage reduction can represent a large amount of memory. This example shows buffer reductions of up to 10% from the burst size σ_i , but further reductions are also possible by considering the effects of the weights used.



(a) as function of k and σ for the case of $\phi_i = 10, \sum_{\forall j} \phi_j = 100$.



(b) as functions of the connection weights for the case of $k=10$ Cells and $\sigma_i=10,000$ Cells

Fig. 9 Buffer Size Requirement per Inequality (18)

Buffer size reductions from σ_i are clearly possible. However this appears to only be significant when the proportional weight allocated to the over-allocating connection is close to 1 and the maximum frame size ($\sum_{\forall j} \phi_j$) is small.

6 Buffer Sizing Conclusions

In this paper, more precise (than fluid model) formulations for buffer size requirements have been developed to consider the effects of worst-case discrete arrival and discrete service, as well as service latency. When both m and k are small compared to σ the formulation for the combined case of discrete arrivals and discrete service (see equation (13)) provides a suitably tractable bound on the worst-case. The worst-case formulations for discrete arrivals and fluid service (see equation (12)) or fluid arrivals and discrete service (see equation (11)) are exact solutions. These formulations should be used for buffer size calculations when the discrete nature of either the arrival or service function must be considered (i.e. when one of m or k is not small compared to σ). When both m and k are not small compared to σ then the evolution of the buffer occupancy needs to be evaluated in more detail. The computational complexity of determining the supremum of the buffer occupancy reduces to evaluating the buffer occupancy at each step in the arrival and service functions. With both m and k large, there will be few steps between the onset of service and the completion of the arrival burst and an exhaustive evaluation is feasible in a reasonable time. The numerical application of equation (13) to the selection of weights for a discrete WRR scheduler has been provided as an illustration. In this example, buffer size reduction below σ is possible, (up to 10% in Fig. 9) but the gains are most significant when the sum of all the weights is small, the burst size, σ is large, and the fan-in, k , is small. In these cases, even a small percentage reduction (say 10%) in the buffer requirements can result in a significant reduction in the number of cells that must be stored. For network elements requiring high-speed memory buffers, the cost savings can be significant.

References

1. Agrawal,R., R.Cruz, C.Okino, R.Rajan, "Performance Bounds for Flow Control Protocols", IEEE/ACM Transactions on Networking, Vol.7,No.3, June 1999.
2. Chang,C-S, "On Deterministic Traffic Regulation and Service Guarantees: A Systematic Approach by Filtering", IEEE Transactions on Information Theory, Vol. 44, No. 3, May 1998, pp1097-1110.
3. Elwalid, A., D. Mitra, R.H. Wentworth, "A New Approach for Allocating Buffers and Bandwidth to Heterogeneous Regulated Traffic in an ATM Node", IEEE Journal on Selected Areas in Communications, Vol 13, No. 6, August 1995, pp 1115-1127.
4. Garg, R., X.Chen, "RRR: Recursive Round Robin Scheduler", IEEE Globecom'98, paper S94-4, November 1998.

5. Guerin,R., S.Kamat, V.Peris, R.Rajan, "Scalable QoS Provision through Buffer Management", Computer Communications Review, Vol. 28, No. 4, 1998
6. Katevenis, M., S. Sidiropoulos, C. Courcoubetis, "Weighted Round Robin cell Multiplexing in a General Purpose ATM Switch Chip", IEEE Journal on Selected Areas in Communications, Vol. 9, No. 8, October 1991.
7. LeBoudec, J-Y, "Application of Network Calculus to Guaranteed Service Networks", IEEE Transactions on Information Theory, Vol. 44, No. 3, May 1998, pp1087-1096
8. LoPresti, F., Z-L.Zhang, J.Kurose, D.Towsley, "Source Timescale and Optimal Buffer/Bandwidth Tradeoff for Regulated Traffic in an ATM Node", Proceedings of IEEE INFOCOM'97 .
9. Sasaki,G., "Input Buffer Requirements for Round Robin Polling Systems", Performance Evaluation, Vol. 18, 1993, pp 237-261.
10. Stiliadis, D., A. Varma, "Latency-Rate Servers: A General Model for Analysis of Traffic Scheduling Algorithms", IEEE Infocom'96, San Francisco, March 1996, pp 111-119.
11. Wright, S., Y. Viniotis, "Simulation Study of Maximal Delays in Various Link Scheduling Algorithms", IEEE Globecom'98, paper S117-7, November 1998.
12. Wright, S., "Delay Bounds and Connection Admission Control for Burstiness-Class Queuing", Ph.D. Thesis, North Carolina State University, 1999.
13. Wright, S., Y. Viniotis, "ATM Network Procedures Supporting Delay QoS", IEEE/IEICE ATM Workshop 99, Kochi City, Kochi, Japan May 24-27, 1999.
14. Wright, S., Y.Viniotis, "Burstiness-Class Based Queueing in ATM Networks Supporting Delay QoS Bounds", IEEE Infocom 2000, Tel Aviv, Israel, March 2000.