Link Capacity Allocation and Network Control by Filtered Input Rate in High Speed Networks

San-qi Li, Song Chong, Chia-lin Hwang and Xirong Zhao
Department of Electrical and Computer Engineering
The University of Texas at Austin, TX 78712

Abstract

[1] *We study link capacity allocation for a finite buffer system to transmit multimedia traffic. The queueing process is simulated with real video traffic. Two key concepts are explored in this preliminary study. First, input traffic in the low frequency band stays intact as it travels through a finite-buffer zero-loss system. Second, the link capacity requirement at each node is essentially captured by its low frequency input traffic (filtered at a properly selected cut-off frequency). Hence, one may overlook the queueing process at each node for network-wide traffic flow in the low frequency band. We propose a simple, effective method for link capacity allocation and network control using on-line observation of traffic flow in the low frequency band. Our experimental study points out a possible new direction for control of high speed networks.*

1 Introduction

Link capacity allocation in high speed networks for multimedia services represents one of the most challenging tasks in advanced telecommunication research [1]-[3]. The central objective of link capacity allocation is two-fold: to take advantage of statistical multiplexing for transmission efficiency, and to avoid nodal congestions caused by the arrival of unpredicted bursty traffic. Recently there has been great concern about the effectiveness of statistical multiplexing, primarily due to the high uncertainty of traffic arrivals and the stringent loss/delay constraint on service qualities. The most conservative allocation scheme is to remove statistical multiplexing by allocating the link capacity to each traffic stream according to its peak input rate, which is certainly not cost-effective. On the other hand, the effectiveness of statistical multiplexing via buffering to improve transmission efficiency is found to be somewhat limited. The two most salient features of multimedia traffic in high speed networks are: strong correlation and high burstiness. In stochastic modeling, these features are captured by second- and higher-order input statistics, such as power spectrum, bispectrum and trispectrum in the frequency domain. The study in [4], [5] indicates that queueing performance is much more dependent on second-order input statistics, i.e., the power spectrum, than any higher-order ones. Further, input power in the low frequency band has a dominant impact on queueing performance, whereas input power in the high frequency band can be neglected to a large extent. In practice, a large amount of input power for multimedia traffic in high speed networks is expected in the low frequency band [6]-[8]. As a result, the improvement of transmission efficiency by finite buffering is substantially limited.

One basic element of the traffic descriptor in ATM protocol is the *peak input rate*. It is used by the usage parameter control (UPC) function to regulate certain traffic streams. By definition, a peak input rate is specified by the maximum number of cell arrivals within a given time interval T. In the frequency domain, this is equivalent to implementing a specific low pass filtering function to the original cell arrival process at the cut-off frequency $\omega_c = 2\pi/T$. Its maximum output is called the filtered peak input rate. The key question here is how to select ω_c for the definition of the filtered peak input rate. In one extreme, we choose $\omega_c = 0$ such that the peak input rate is equal to the steady state average input rate. In the other extreme, we may take $\omega_c = 2\pi\mu$, where μ^{-1} is one slot link transmission time per cell, such that the peak input rate corresponds to the maximum number of cell arrivals in a slot. Assume that the link capacity is simply assigned by its peak input rate. If the peak input rate is defined at $\omega_c = 0$, buffer size and queueing delay must be infinite to prevent cell loss. If it is defined at $\omega_c = 2\pi\mu$, buffer size and queueing delay will be zero since statistical multiplexing is virtually removed. Obviously, the peak input rate is strongly dependent on ω_c. The smaller the ω_c is, the less the peak input rate will be. The notion of peak input rate is therefore incomplete without a clear definition of ω_c, whose value is strongly associated with buffer size, link capacity and input traffic characteristics, as revealed by our study.

This paper focuses on the link capacity allocation for a finite-buffer system to transmit a sequence of coded video signals, subject to zero cell loss. We choose to use computer simulation for the following two reasons. First, there is no analytic input model available that can adequately describe multimedia sources. Second, it is always analytically difficult to characterize a filtered input process, which is the focal point in this research. The words *traffic* and *signal* are used interchangeably since our approach takes the combination of queue analysis and signal processing techniques.

Two new concepts will be explored in this experimental study for traffic measurement and link capacity allocation in the frequency domain. First, one can view a finite-buffer system as a transparent pipe for the input signal in the low frequency band, under the condition of negligible cell loss by buffer blocking. In other words, the original characteristics of the input signal in the low frequency band stay intact as it travels at each node and so on throughout the entire network. One may therefore overlook the queueing process at each node for network-wide traffic flow in the low frequency band. Second, the minimum link capacity required by input traffic at each node with a finite buffer capacity subject to zero cell loss is essentially captured by the low frequency input (filtered at a properly selected cut-off frequency). Both static and dynamic link capacity allocations will be examined. For static allocation one can simply assign the link capacity by the filtered peak input rate for a given time interval. For dynamic allocation we adaptively change the link capacity using on-line observation of the filtered input rate. In practice, one can implement dynamic allocation at the network level due to the slow time variation of the filtered input rate in the low frequency band.

Divide the input signal into high and low frequency components. Our investigation shows that the high frequency com-

[1]The research reported here was supported by NSF under grant NCR-9015757 and Texas Advanced Research Program under grant TARP-129.

0-7803-0917-0/93$03.00 © 1993 IEEE

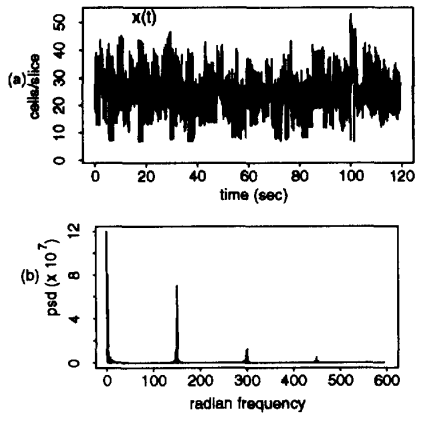

Figure 1: A typical 2-minute video sequence (a) input rate measured in slice time units, (b) power spectrum.

ponent can be absorbed by the finite buffer capacity at each node; the low frequency component will be effectively transmitted through the dynamic link capacity allocation at the network level. Such an input spectral decomposition naturally divides traffic control into two levels: local congestion control of the high frequency component on the link level and global load-balance control of the low frequency component at the network level. We will investigate a new network-wide control scheme based on observation of traffic flow in the low frequency band.

The paper is organized as follows. In Section 2 we study input/output signals in the low frequency band and allocate link capacity by the peak input rate filtered at a properly selected cut-off frequency. Section 3 shows the significant improvement of performance using the dynamic link capacity allocation. The work is extended in Section 4 to a simple, effective network-wide traffic control scheme. The paper is concluded in Section 5.

2 Static Link Capacity Allocation

So far there have been no analytical models that can adequately describe multimedia traffic in high speed networks. For that reason we choose a sequence of data coded from two hours of the movie "Star Wars" to represent our input signal. Its original bit stream was generated by a simulated coder using DCT and Huffman coding without motion compensation, which is available from Bellcore in public domain [9]. The data were recorded in bytes in every 1.4 msec slice time unit, where the bytes are converted into cells, with each cell consisting of 44 bytes of video information plus 9 bytes of ATM protocol overhead. There are 16 slices per frame and 24 frames per second. Fig. 1a shows a typical 2-minute video sequence measured in slice time units. Apparently, the input signal appears to be highly nonstationary. The maximum number of cell arrivals in slice is 52.9, which also represents the peak input rate measured in slice time units; the average is 24.3 cells per slice, which is equivalent to 7.36Mbps. Also plotted in Fig. 1b is the corresponding input power spectral function. Two key observations can be made about the video power spectrum. First, the spikes that appear at the radian frequency $2\pi \times 24$ and its harmonics represent a strong frame correlation. Second, the rest of the video input power is located mainly in a very low frequency band, typically less than 20 radians, capturing the strong scene correlations.

Our objective here is to find the minimum link capacity, μ, in a finite-buffer system for transmission of the above video input signal, subject to no cell loss. When zero buffer space is provided, the link capacity must be assigned by the peak input

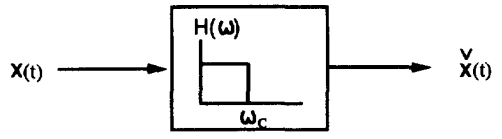

Figure 2: A low-pass filter

rate measured in time units of a cell transmission slot, which is unfortunately not available from the original data collection where the basic time unit chosen for the measurement is one slice. Each slice on average is at least 24 slots, depending on the link capacity. In other words, the original micro dynamics of the video signal at the slot level was filtered out by the slice-moving average operation during data collection. As a result, if the link capacity is assigned by the peak input rate measured in slice units, a small amount of buffer capacity, such as several tens of cells, must be assigned to smooth out the high frequency variations of the input signal within slice. In this case, the corresponding link utilization factor ρ for the above 2-minute video sequence is given by 0.46, which yields $\mu = 52.9$ cells per slice. Notice in Fig. 1a that the peak input occurrence at the time around 100 secends causes a significant reduction of ρ. Without buffering, ρ will be further reduced.

Now, consider a queueing system with buffer capacity of $K = 250$ in cell units. Note that while the input data were recorded in slice units, the queueing process must be evolved in slot units. Converting the input time unit of slice into slot, we consider the worst scenario: all the cells which are generated in each slice are assumed to arrive at the beginning of the first slot in that slice. The simulation study shows that the minimum link capacity at $K = 250$, subject to zero loss, is given by $\mu = 41$ cells per slice for the above 2-minute video sequence, which corresponds to $\rho = 0.59$. By comparing with the zero buffer case, we save at least 22.5 percent of the link capacity by using a buffer capacity of 250 cells. The exact saving can be found if we know the peak input rate measured in slot units.

The key question is how to identify the minimum link capacity required by the input signal in a finite-buffer zero-loss system. Our study indicates that such a minimum link capacity is essentially captured by the inherent peak input rate filtered at a properly selected cut-off frequency ω_c. Apply the input signal, denoted by $x(t)$, to a low pass filter in Fig. 2. The filtered input signal is represented by $\check{x}(t)$. Denote the filter transfer function by $h(t)$ and its Fourier transform by $H(\omega)$. An ideal low-pass filter is defined by

$$|H(\omega)| = \begin{cases} 1 & \text{if } \omega \leq \omega_c \\ 0 & \text{if } \omega > \omega_c \end{cases} \quad (1)$$

In practice, a variety of low-pass filters can be selected. Here we choose an FIR digital low-pass filter with Kaiser window design, which is available in the MATLAB software. The FIR filter is chosen simply because it has much less delay than the IIR filter. The Kaiser window design is used for the sharp discontinuity at ω_c. The time unit in the digital filtering process is equal to one slice interval of the video. Consider that most input power for video, excluding the frame correlations, is located in a very low frequency band, such as $\omega < 20$ in Fig. 1. The corresponding cut-off frequency ω_c for the video filtering, as found below, can be as low as less than 10. In order to filter the video signal in such a low frequency band using slice units, one needs to design the filter with an extremely sharp discontinuity at ω_c, otherwise even taking $\omega_c = 0$ will practically be insufficient for the required low-pass filtering function. For this reason we have designed the above FIR filter with the order of $N = 2000$. One can also

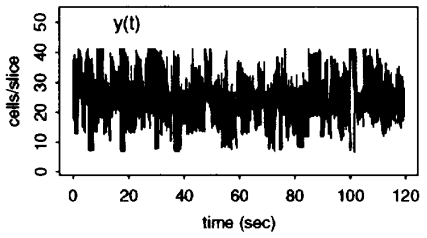

Figure 3: Output video signal

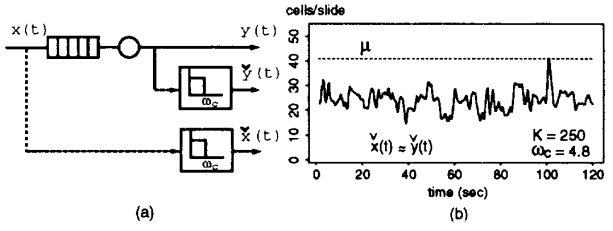

Figure 4: (a) Filter operation, (b) filtered input and output

use multiple cascade filters to achieve a similar effect. Since our focus here is not on filter design, throughout the paper we will use the same FIR filter with Kaiser window design at $N = 2000$ for all filtering purposes.

Static allocation assigns link capacity by the peak input rate filtered at a properly selected ω_c, i.e.,

$$\mu = \max_t \check{x}(t) \tag{2}$$

for a given filter function $H(\omega)$. For instance, applying the 2-minute video signal to the above designed FIR filter, we find $\omega_c = 4.8$ at which $\max_t \check{x}(t) = 41$ cells per slot, which is exactly equal to the minimum link capacity μ at $K = 250$. Before showing other examples, let us explain why the link capacity requirement is essentially captured by the filtered peak input rate. As mentioned earlier, a finite-buffer system without high loss-rate can be viewed as a transparent pipe for transmission of input signal in the low frequency band. To help us understand this new concept, let us compare both input and output signals of the above queueing system at $K = 250$. Denote the output signal by $y(t)$. Fig. 3 shows the output signal $y(t)$ in response to the input signal $x(t)$ in Fig. 1a at $K = 250$ without cell loss. Clearly, $y(t)$ can never exceed the assigned minimum link capacity $\mu = 41$ cells per slice. Let us also examine both input and output signals filtered by the same function $H(\omega)$ at $\omega_c = 4.8$, denoted by $\check{x}(t)$ and $\check{y}(t)$ in Fig. 4a. As shown in Fig. 4b, virtually no difference exists between the two filtered signals, i.e.,

$$\check{x}(t) \approx \check{y}(t)$$

Note that, when a system is designed by $K = 250$ to transmit the 2-minute video sequence at an average input rate equal to 7.36Mbps, it is impossible to hold any arriving cells for more than 14.4 msec (which is even less than one frame interval). Hence, the low frequency input power must stay intact at the output under the zero-loss assumption.

Note that the minimum link capacity in the above case is essentially captured by the peak input rate, which occurs at the time around 100sec as shown in Fig. 1. It could be argued that the reason for $\check{x}(t) \approx \check{y}(t)$ is because of $x(t) \approx y(t)$ in the rest of the period, during which the buffer is lightly loaded. To clarify this argument, let us re-examine the above system based on the

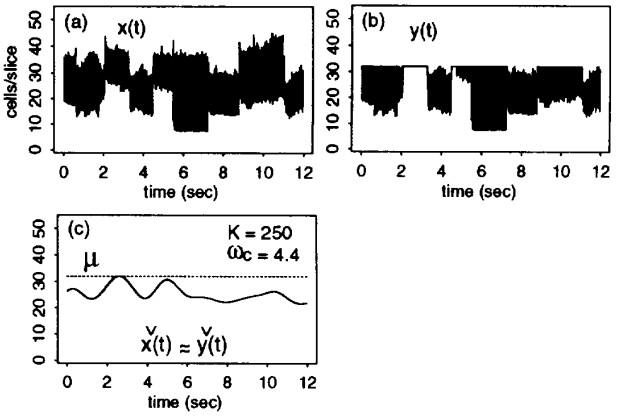

Figure 5: The first 12-second sub-sequence video signal: (a) input, (b) output, (c) filtered input and output.

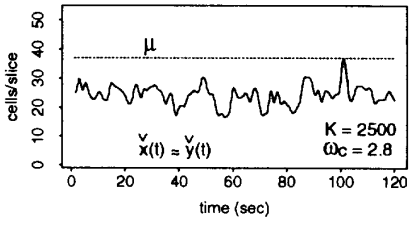

Figure 6: Output video signal filtered at $\omega_c = 2.8$ with buffer capacity $K = 2500$

first 12-second sub-sequence of the 2-minute signal. The average input rate of this 12-second subperiod is 25.6 cells per slice and its peak is equal to 45.2. Without cell loss at $K = 250$, the maximum utilization factor ρ is found to be 0.80. Link capacity is assigned by $\mu = \max_t \check{x}(t)$ at $\omega_c = 4.4$, which is equal to 32 cells per slice. Plotted in Fig. 5a,b are the original $x(t)$ and $y(t)$, both of which are now significantly different. Nevertheless, there is virtually no difference between the two filtered ones, as shown in Fig. 5c at $\omega_c = 4.4$.

Two key concepts are explored in this study. First, the nature of the input signals in the low frequency band basically remains unchanged as they travel through the entire network, under the condition of negligible cell loss at each node. Notice that the condition of negligible cell loss is inevitable for a "well behaved" network to provide quality services. Second, the minimum link capacity required by the input signal at each node subject to zero-cell loss is well captured by its input signal behavior filtered at a properly selected cut-off frequency.

Such a cut-off frequency, of course, is strongly dependent on buffer capacity. Ideally, when the buffer capacity is infinite, ω_c can be zero and the peak input rate is equal to its average (i.e., the DC component), so the minimum link capacity can simply be assigned by the average input rate. In the case of the 2-minute video signal, if buffer size K is increased from 250 to 2500, the corresponding minimum link capacity will be reduced from 41 to 37, and so the link utilization factor ρ is increased from 0.59 to 0.66. Similar results of $\check{x}(t)$ and $\check{y}(t)$ are plotted in Fig. 6 as ω_c is reduced from 4.8 to 2.8, at which $\mu = \max_t \check{x}(t)$.

The cut-off frequency also depends on input signals and the link transmission rate. So far we have only considered a single video source for input. In practice, multiple sources are likely

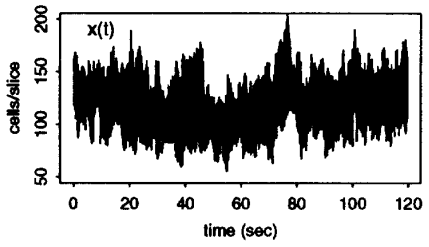

Figure 7: Aggregate input of five video sources.

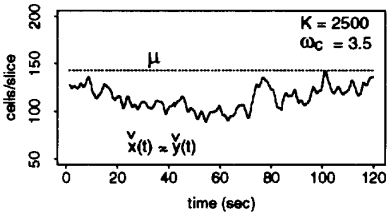

Figure 8: Filtered input and output signals of five video sources at $\omega_c = 3.5$ with buffer capacity 2500

to be statistically multiplexed on a high speed link. The pooling effect is expected to improve queueing performance and link utilization. Let us design a static link capacity allocation for the transmission of five video sources. Each source is represented by a 2-minute video sequence, taken from the same movie "Star Wars" but coded at a different time interval. Fig. 7 shows the aggregate input signal, with its peak rate equal to 204.5 cells in slice unit. The average input rate is equal to 114.2 cells per slice. If link capacity is assigned by the peak input rate in slice, we get $\rho = 0.56$. When the buffer size is $K = 250$, the corresponding ρ increases to 0.72 without causing any cell loss. Again, one obtains $\check{x}(t) = \check{y}(t)$ and $\mu = \max_t \check{x}(t)$ at $\omega_c = 153$. In comparison to the single source case, the aggregate signal should be much smoothed. As a result, the system performance at $K = 250$ is improved from $\rho = 0.59$ of the single source to $\rho = 0.72$ of the multiplexing. Note that the corresponding cut-off frequency ω_c is increased from 4.8 to 153. This is because the queueing performance is dependent on the normalized frequency $\frac{\omega_c}{2\pi\mu}$, instead of ω_c defined in radians [5]. In other words, the larger the link capacity μ is, the higher the cut-off radian frequency ω_c will be for the same queueing and loss performance. The video power spectrum in Fig. 1b has a virtually empty bandwidth ranging from 30 to 140, which essentially divides the power distribution by the scene and frame correlations. That is why we have seen a substantial increase of ω_c from 4.8 to 153 for the static link capacity allocation of five video sources. For 24 frames per second, the fundamental frequency for the frame correlation is equal to $2\pi \times 24$ radians, which is slightly less than 153 (or $2\pi \times 24.3$). In other words, while the video frame correlation has no impact on μ at $K = 250$ for the single video source, it has an impact on the aggregate five video sources due to the increased link capacity. When the buffer capacity K is increased from 250 to 2500, one can furthermore get $\check{x}(t) = \check{y}(t)$ and $\mu = \max_t \check{x}(t)$ at $\omega_c = 3.5$, as shown in Fig. 8. As a result, ρ is up from 0.72 to 0.79, while ω_c is down from 135 to 3.5, which removes the frame correlation effect on μ at $K = 2500$ for the aggregate input.

As discussed in Section 1, the cut-off frequency must be in the range of $0 \leq \omega_c \leq 2\pi\mu$. For the ideal low pass filter in (1), we have $\check{x}(t) = E[x(t)]$ at $\omega_c = 0$ and $\check{x}(t) = x(t)$ at $\omega_c = 2\pi\mu$. Note that only the DC component will be left by the ideal filtering

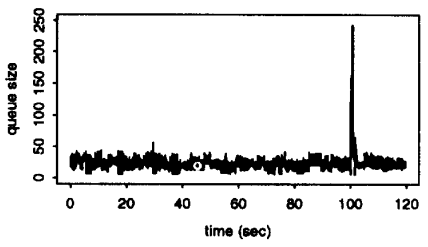

Figure 9: Queueing process with single video source at $K = 250$

at $\omega_c = 0$, and the original $x(t)$ is measured in the transmission slot unit $\frac{1}{\mu}$. Consider two extreme cases of static link capacity allocation: infinite buffering and zero buffering. With infinite buffering, we must have $\mu = E[x(t)]$, or $\mu = \max_t \check{x}(t)$ at $\omega_c = 0$, for the minimum link capacity. With zero buffering, we need $\mu = \max_t x(t)$, which is equivalent to $\mu = \max_t \check{x}(t)$ at $\omega_c = 2\pi\mu$, to avoid cell loss. Generally with a finite buffer capacity, we need to properly select ω_c for the static link capacity allocation by $\mu = \max_t \check{x}(t)$.

Since the filtered peak input rate is measured in a long period, one cannot achieve a high link utilization by the static allocation for bursty traffic. In other words, link capacity is always assigned by the worst scenario of the filtered input for a given time interval. As a result, the queue stays low most of the time and occasionally reaches the maximum at the time occurrence of $\max_t \check{x}(t)$. This phenomenon is well explained in Fig. 9, which shows the queueing process to transmit a single video source at $K = 250$, with respect to the input and output signals in Figs. 1 and 4.

3 Dynamic Link Capacity Allocation

Static allocation assigns link capacity by the filtered peak input rate. In practice, such a peak input rate is difficult to predict. As we have seen, the filtered peak input rate can be substantially lower than the non-filtered one, depending on the cut-off frequency which is a function of buffer size, link capacity and input characteristics. Moreover, such a peak input rate changes dynamically with the number of calls being supported on the link; the holding time of each call for multimedia services can differ by several orders of magnitude. Therefore, it is impossible to effectively implement a policy of static link capacity allocation in a real network environment. The only valid approach is to dynamically allocate link capacity using the on-line observation of the filtered input signal $\check{x}(t)$.

In principle, $\check{x}(t)$ is defined in the low frequency band and so its time variation will be slow and predictable. For instance, we have found $\omega_c < 5$ radians in the above study of the filtered video signal when $K \geq 250$. If the time variation is defined by $\frac{2\pi}{\omega}$ at each frequency ω, most of the corresponding time variations in $\check{x}(t)$ will be in the range of a few seconds or longer. It is therefore possible to allocate the link capacity adaptively by the present $\check{x}(t)$. Such a dynamic allocation can be implemented either locally at each node with adaptive link capacity assignment among individual virtual circuits, or globally at the network level with a dynamic traffic routing scheme. The study in this section shows a significant improvement in performance by implementing such a dynamic link capacity allocation policy.

For simplicity we consider an ideal situation where link capacity can be instantaneously assigned by $\check{x}(t)$, as described in Fig. 10. In other words, the link capacity will be defined as a function of $\check{x}(t)$. Denote the link capacity by $\mu(t)$ for dynamic allocation. We further assume $\mu(t) = c\check{x}(t)$, where c is a control parameter. From $E[\mu(t)] = cE[\check{x}(t)]$ in steady state, one obtains

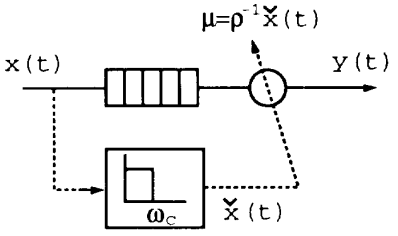

Figure 10: Dynamic link capacity allocation by filtered input rate.

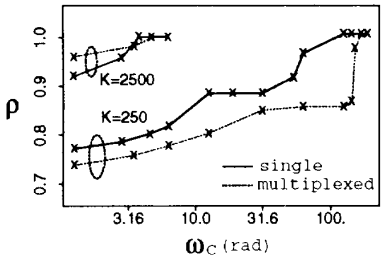

Figure 11: Trade-off between ω_c and ρ for dynamic link capacity allocation at $K = 250$ and 2500 subject to zero-loss

$c = \rho^{-1}$, which leads to

$$\mu(t) = \rho^{-1}\check{x}(t) \tag{3}$$

Since link capacity is synchronously adapted by the filtered input signal, the impact of the low frequency input $\hat{x}(t)$ on the queueing process has now been removed. Note that we must have $\rho < 1$, because some extra link capacity is required for the finite buffer to effectively smooth out the high frequency input.

Let us divide the input signal into high and low frequency inputs, denoted by $\hat{x}(t)$ and $\check{x}(t)$, respectively. Ideally, one can write

$$x(t) = \hat{x}(t) + \check{x}(t)$$

On the one hand, a finite buffer capacity effectively improves system performance with respect to the high frequency input $\hat{x}(t)$. On the other hand, allocating dynamic link capacity effectively improves system performance with respect to the low frequency input $\check{x}(t)$. In practice, $\check{x}(t)$ must stay in a well-founded low frequency band to avoid frequent adaptation of the link capacity.

Let us examine the trade-off between ω_c and ρ. The cut-off frequency ω_c is tuned for the division of $\hat{x}(t)$ and $\check{x}(t)$; the utilization factor ρ is adjusted for the extra link capacity required to resolve the high frequency input $\hat{x}(t)$. The smaller the ω_c, the slower the variation of $\check{x}(t)$ will be, so the link capacity is less frequently adapted as desired in practice. As a result, more input power remains in the high-frequency band for $\hat{x}(t)$, which must be resolved by finite buffering. Hence, in a fixed buffer capacity system, reducing ω_c has the effect of decreasing ρ for an extra link capacity to avoid cell loss. The trade-off between ω_c and ρ in a finite-buffer zero-loss system essentially represents the trade-off between the time adaptability of the link dynamic allocation scheme and the overall average link capacity requirement. The more frequently the link capacity can be adapted, the less the average link capacity will be required. Such a trade-off is found in Fig. 11 for the dynamic link capacity allocation of the video signals used in the previous section. Buffer capacity is fixed at $K = 250$ and 2500, respectively. ρ represents the maximum throughput subject to zero-loss. In contrast with the static allocation in Section 2, the link throughput is significantly improved by dynamic allocation, especially when a large ω_c is selected. For instance, the maximum throughput for the single video source at $K = 250$ is only 0.59 by static allocation, but it can be as high as 0.99 by dynamic allocation. Even if one wants to limit $\omega_c < 10$ for the time adaptability of the dynamic allocation, the link throughput can still be as high as 0.85. In fact, when the buffer size is 2500, the link can be virtually fully loaded by dynamic allocation under the condition of $\omega_c < 10$. One may notice a significant jump at $\omega_c = 2\pi \times 24$ for the multiplexing of five video sources at $K = 250$. This is because the frame correlation is moved from $\hat{x}(t)$ to $\check{x}(t)$ once $\omega_c > 2\pi \times 24$, which significantly improves throughput performance. Also displayed in Fig. 12 is the queueing behavior for the dynamic link capacity allocation of the single video source at $K = 250$. In comparison with the queueing behavior by the static allocation in Fig. 9, the queue length now becomes much more evenly distributed in the time domain. As in static allocation, both input and output signals in the low frequency band must be identical, as shown in Fig. 12 at $\omega_c = 4.8$.

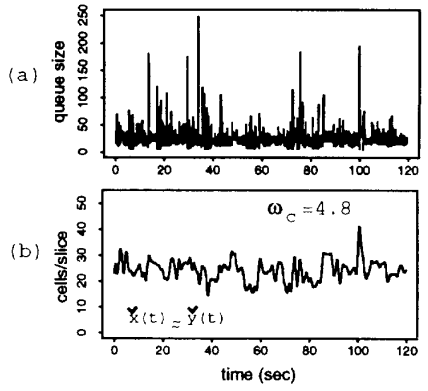

Figure 12: (a) Queueing behavior with dynamic allocation, (b) input and output signals in low frequency band at $\omega_c = 4.8$.

4 Network Control

The above study on a single link can generally be applied network-wide. Essentially, the input signal on each link, denoted by $x_i(t)$ on the i-th link, is decomposed into $\hat{x}_i(t)$ and $\check{x}_i(t)$ at a properly selected cut-off frequency ω_{ci}, $\forall i$. Each link at its entry point is associated with a finite buffer system. While the high frequency input $\hat{x}_i(t)$ is effectively delivered by finite buffering, the low frequency input $\check{x}_i(t)$ has to be transmitted with the dynamic link capacity allocation. Such a decomposition naturally divides traffic control into two levels: local congestion control of each individual $\hat{x}_i(t)$ at the link level and global load-balance control of the overall $\{\check{x}_i(t)\}$ at the network level. As the study in [11], [12] indicates, local congestion control at each individual link, implemented with cell selective discarding and traffic shaping techniques, is effective to the high frequency input only. In terms of low frequency input transmission, one can view the finite buffer system as a transparent pipe with identical input and output.

Here we point out a new direction for network-wide traffic control. In principle, traffic flow in the low frequency band on each link can be approximated by a linear system. This is described in Fig. 13, where traffic on the three input links is denoted by $x_i(t)$, $x_k(t)$ and $x_l(t)$ respectively, and output traffic on the j-th link is denoted by $x_j(t)$. Assuming negligible cell

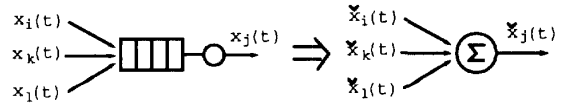

Figure 13: Linear system approximation for traffic flow on each link in low frequency band.

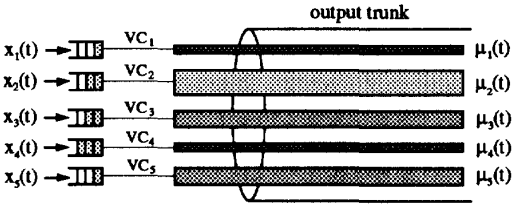

Figure 14: Dynamic link capacity allocation among five video sources by filtering.

loss at the buffer, one can generally write

$$\check{x}_j(t) = \check{x}_i(t) + \check{x}_k(t) + \check{x}_l(t) \tag{4}$$

in the low frequency band. In other words, the queueing process at each node can be overlooked at the network level since it affects only the high frequency traffic flow. Take an extreme case at $\omega_c = 0$ with the ideal low pass filter in (1). Each filtered input consists of a DC component, i.e., $\check{x}_m(t) = E[\check{x}_m(t)]$ for $m = i, k, l$. From (4) we get the average output equal to the sum of the average inputs, which is true for any finite-buffer zero-loss system. Hence, the above proposed frequency decomposition approach to a certain extent can be viewed as an extension of the DC component at $\omega_c = 0$ to a more general low frequency band at some $\omega_c > 0$.

For network-wide traffic analysis, denote the independent original input traffic from the entry node j to the departing node k by $z_{jk}(t)$. Define the capacity of link i by μ_i, which is assumed to be fixed. We say $\mu_i \in (j,k)$ if the traffic $z_{jk}(t)$ is transmitted on link i. One can therefore write

$$\check{x}_i(t) = \underbrace{\sum_j \sum_k}_{\mu_i \in (j,k)} \check{z}_{jk}(t) \tag{5}$$

for the aggregate low frequency input traffic on link i. Given the original input traffic $\{z_{jk}(t)\}$, link capacity $\{\mu_i\}$ and network topology, one can design a dynamic traffic flow control by

$$\max_{\{\check{x}_i(t)\}} \sum_i [\mu_i - \check{x}_i(t)] \quad s.t. \;\; \mu_i \geq \check{x}_i(t) + \Delta_i, \;\; \forall i \tag{6}$$

where Δ_i represents the minimum amount of extra capacity, as mentioned in Section 3, reserved to transmit the high frequency traffic by finite buffering at each node. All real-time traffic $\{\check{x}_i(t)\}$ is monitored by on-line filter operation. Whenever the constraints $\mu_i \geq \check{x}_i(t) + \Delta_i$, $\forall i$, cannot be met, some of the original input traffic at its entry point, $z_{jk}(t)$, must be selectively blocked or regulated. As done in Section 3, one may replace $\check{x}_i(t) + \Delta_i$ with $\check{x}_i(t)/\rho_i$, where ρ_i is the utilization factor on link i.

For connection-oriented services, one may implement the above flow control by dynamic routing. That is, the route of each individual call is changeable in response to network traffic variation in the low frequency band. Similarly for call admission control, when a call request for a certain service class is made, the user is asked to provide the network manager with an estimation of its filtered peak input rate. If such an estimation is not available by the user, some peak value can always be assigned, based on existing statistics of the service class. The estimation, denoted by $\max_t \check{z}_{nm}(t)$ for instance, is added to the existing traffic flow on each link $\{\check{x}_i(t)\}$ with respect to a set of possible routes. The optimal route is determined by

$$\max_{\mu_i \in (n,m)} \sum_i [\mu_i - \check{x}_i(t) - \max_t \check{z}_{nm}(t) \mathbf{1}(\mu_i \in (n,m))] \tag{7}$$

$$s.t. \quad \mu_i \geq [\check{x}_i(t) + \max_t \check{z}_{nm}(t) \mathbf{1}(\mu_i \in (n,m)]/\rho_i, \quad \forall i$$

where $\mathbf{1}(\mu_i \in (n,m))$ is an indicator function equal to one if the new call $\check{z}_{nm}(t)$ is on link i, or zero otherwise. If the above constraints on μ_i's cannot be met, either the new call is rejected, or the routes of existing calls must be changed to re-adjust the existing traffic flow $\{\check{x}_i(t)\}$ in order to accept the new call. Note that each time a decision on the call admission is made, only the filtered peak input rate of the new call, $\max_t z_{nm}(t)$, needs to be estimated in advance; all $\check{x}_i(t)$'s are available from on-line observation. Optimization can be implemented in either distributed or centralized form. If no pre-knowledge about the source is provided, the estimation of $\max_t z_{nm}(t)$ can be conservative since an individual call should not have a significant impact on overall traffic flow. Furthermore, once the call is accepted, such an estimation will have no impact on future call decisions.

The above flow model is an abstract version of our proposed network control scheme. More studies are required for the design of ω_{ci} and ρ_i. A distinct advantage of this network control scheme is its effectiveness to improve system performance, as we found in Sections 2 and 3. Another advantage is its simplicity to operate by overlooking the queueing process at each node. Filter operation is implemented using advanced digital signal processing chips for on-line low frequency traffic measurement on each link.

In our view, the key problem in developing a network-wide traffic control scheme is to find an appropriate cost criterion on each link. The complexity of the involved network-wide queueing analysis so far has severely hindered us from developing simple, effective control techniques for high-speed multimedia service networks. The most significance of this study is to indicate that such a cost criterion on each link is essentially captured by the low frequency input traffic characteristics. Hence, not only is $\{\check{x}_i(t)\}$ easily measured by the filter operation, but also the whole network flow control is formulated as a linear system. Note that our approach does not require the input traffic $\{z_{jk}(t)\}$ to be stationary. Furthermore, global traffic flow control at the network level is well-separated from individual local congestion control on the link level, which is a desired feature in layered protocol architecture design.

Let us take a simple example of transmitting five video signals on a commonly shared transmission link of capacity μ. Each source is associated with a separate buffer and transmitted on a virtual channel, as shown in Fig. 14. The network manager adaptively divides the link capacity into five virtual channels, i.e., $\mu = \sum_{i=1}^{5} \mu_i(t)$, based on the observation of $\{\check{x}_i(t)\}$. The information of $\{\check{x}_i(t)\}$ is either collected from each individual source, or extracted directly by the network manager. As in Section 2, here we take each of the five 2-minute video signals to represent video sources. The overall average input traffic is 114.2 cells per slice. In order to provide equal services, we assume the virtual channel capacity of each source is evenly allocated by

$$\mu_j(t) = \frac{\check{x}_j(t)\mu}{\sum_{i=1}^{5} \check{x}_i(t)}. \tag{8}$$

The above traffic flow was studied in Section 2 at a central-

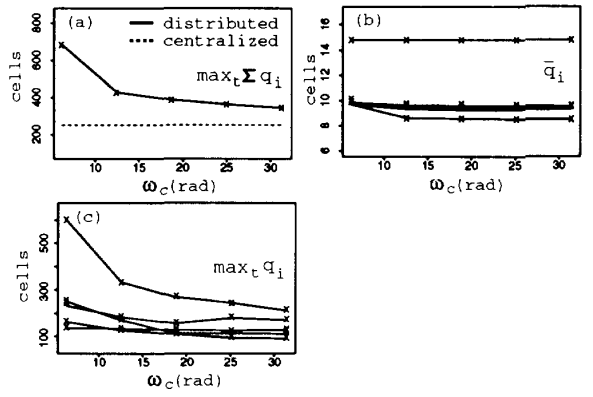

Figure 15: Queueing performance (a) the maximum queue accumulation in five sources, (b) the average queue at each source, (c) the maximum queue at each source.

ized statistic multiplexer with a single commonly shared buffer, where no transmission capacity is wasted unless the commonly shared buffer is empty. The study in Section 2 indicates that the minimum link capacity μ, subject to no cell loss, is assigned by $\rho = 0.72$ for the centralized system at $K = 250$. To compare the effectiveness of the distributed control with that of the centralized one, here we also take $\rho = 0.72$. Denote the i-th queueing process by $q_i(t)$ and its maximum by $\max_t q_i(t)$. The overall maximum is defined by $\max_t \sum_i q_i(t)$. In the centralized case, we found $\max_t \sum_i q_i(t) = 250$ at $\rho = 0.72$. Displayed in Fig. 15a is the performance of $\max_t \sum_i q_i(t)$ in the distributed case as a function of cut-off frequency ω_c. Clearly, the higher the ω_c is, the more frequently the link capacity allocation is adapted to the individual traffic, and so the less the buffer capacity is needed to prevent cell loss. The difference in $\max_t \sum_i q_i(t)$ between the centralized and distributed cases represents the cost of the distributed allocation scheme, which is insignificant especially when ω_c is high. Considering the low frequency range of ω_c we have chosen here, it is practically feasible to implement such a distributed dynamic link capacity allocation scheme among the individual virtual channels. To examine the performance fairness to each individual source, also plotted in Fig. 15b,c are the mean and maximum of individual queues. As one can see, the distributed allocation scheme provides a fair solution to all video sources.

Two key questions remain to be answered regarding the implementation of the above proposed network control scheme. One is how to quantitatively identify ω_{ci} and ρ_i, $\forall i$, which are dependent on buffer capacity, link transmission rate and traffic characteristics at each individual node. The other is how frequently the network routes are practically adaptable to traffic time variation in the low frequency band. Both issues are currently under our investigation. The first one is directly related to queue length and loss rate performance of a finite buffer system in response to the input spectral functions [4, 5, 10]. The second issue will be studied by simulation and analysis as in most data networks.

5 Conclusion

This paper has addressed the issue of link capacity allocation for transmission of multimedia traffic in a finite buffer system. Two concepts have been explored. First, input traffic in the low frequency band stays intact as it travels through a finite-buffer zero-loss system. Hence, one can overlook the queueing process at each node for network-wide traffic flow in the low frequency band. Secondly, the link capacity requirement at each node is essentially captured by its low frequency input signal filtered at a properly selected cut-off frequency. We have proposed a simple, effective method for link capacity allocation and network control using on-line observation of low frequency traffic. The experimental study here points out a possible new direction for control of high-speed networks.

References

[1] G. Gallassi, G. Rigolio and L. Fratta, "ATM: Bandwidth Assignment and Bandwidth Enforcement Policies," Proc. of *Globecom'89*, pp. 1788-1793.

[2] R. Guerin, H. Ahmadi and M. Naghshineh, "Equivalent Capacity and Its Application to Bandwidth Allocation in High-Speed Networks," *IEEE J. Select. Areas Commu.*, Vol. SAC-9, Sept. 1991, pp. 968-981.

[3] M. Decina, T. Tohiatti, P. Vaccari and L. Verri, "Bandwidth Allocation and Selective Discarding for Variable Bit Rate Video and Bursty Data Calls in ATM Networks," Proc. of *IEEE Infocom'91*, pp. 1386-1393.

[4] S.Q. Li and C.L. Hwang, "Queue Response to Input Correlation Functions: Discrete Spectral Analysis," Proc. of *IEEE Infocom'92*, pp. 382-394 (the paper received the Infocom'92 Conference Paper Award).

[5] S.Q. Li and C.L. Hwang, "Queue Response to Input Correlation Functions: Continuous Spectral Analysis," presented at the 7th IEEE Computer Communication Workshop in Hilton Head Island, SC, Oct. 1992 (also submitted to IEEE/ACM Trans. on Networking, July 1992).

[6] P.T. Brady, "A Statistical Analysis of On-Off Patterns in 16 Conversation," *Bell Syst. Tech. J.* Jan. 1968, pp. 73-91.

[7] S.-S. Huang, "Modeling and Analysis for Packet Video," *Proc. of IEEE Globecom'89*, Dallas, Texas, Nov. 1989, pp. 25.2.1-5.

[8] A. Mukherjee, "On the Dynamics and Significance of Low Frequency Components of Internet Load," Distributed Systems Lab Report, MS-CIS-92-83, Computer and Information Science Department, University of Pennsylvania, Dec. 1992.

[9] M.W. Garrett and M. Vetterli, "Congestion Control Strategies for Packet Video," presented in the *Fourth International Workshop on Packet Video*, Kyoto, Japan, August 1991.

[10] H.D. Sheng and S.Q. Li, "Second Order Effect of Binary Sources on Characteristics of Queue and Loss Rate," Proc. of *IEEE Infocom'93*, April 1993, pp. 18-27.

[11] H.D.Sheng and S.Q.Li "Spectral Analysis of Packet Loss Rate at Statistic Multiplexer for Multimedia Services," Proc. of IEEE Globecom'92, Dec. 1992, pp. 1524-1532.

[12] S.Q.Li and S. Chong, "Fundamental Limits of Input Rate Control in High Speed Network," Proc. of *IEEE Infocom'93*, April 1993, pp. 662-671.