Dynamic Time Sharing: A New Approach For Congestion Management

Adrian Popescu

University of Karlskrona/Ronneby
Dept. of Telecommunications and Mathematics
S - 371 79 Karlskrona
Sweden
adrian@itm.hk-r.se

Parag Pruthi

Bellcore
445 South Street
Morristown, NJ 07960 - 6438
USA
parag@bellcore.com

Abstract

A new approach for bandwidth allocation and congestion control is reported in this paper, which is of the Rate Controlled admission with Priority Scheduling service type. It is called *Dynamic Time Sharing (DTS)*, because of the dynamic nature of the procedure for resource partitioning to allocate and guarantee a required bandwidth for every traffic class. This approach is based on guaranteeing specific traffic parameters (bandwidth requirements) through a policing unit, and then optimizing the bandwidth assignment within the network for specific parameters of interest (like delay or jitter, and loss). The optimization process is based on the parameters guaranteed by the policing unit. A batch admission policy is used at the edges of the network according to a specific framing strategy to follow the traffic characteristics (e.g., the traffic constraint function) of different traffic classes. On the other hand, another framing (congestion control) strategy is used within the network, which is based on different (delay/loss) requirements of the traffic classes. Proper management of bandwidth and buffer resources is provided in every (switch) node of the network, such as to guarantee the diverse performance of interest.

1 Introduction

To take advantage of statistical multiplexing and to efficiently utilize network resources, issues like efficient bandwidth allocation and congestion control must be solved. These are difficult problems, and proper methods of handling traffic such as to provide (worst case) loss, delay and jitter guarantees for all traffic classes, irrespective of transmission speed, network size and traffic mix, seem generally not yet sufficiently understood [Jai95], [Rob95]. The main difficulties stem from different and diverse problems, such as the uncertainties when modeling the statistical behavior of different traffic sources, the so-called process of clustering (i.e., an increased variability observed in the interarrival times of packets) that takes place in a packet network with conventional FIFO queueing (with the consequences of destroying the original smoothness characteristics of traffic and of network congestion), the diverse mix of traffic types with different characteristics and service requirements, the high speed of transmission in optical fibers, etc.

The new issues that must be dealt with in the design of efficient traffic management and congestion control mechanisms for B-ISDN networks are mainly related to the self-similar nature of bursty (packet) traffic [LTWW94], [BSTW95] as well as the large bandwidth-delay products available in broadband networks. For instance, admission control schemes that do not account for the highly bursty nature of packet traffic will probably be unable to meet QoS objectives when subjected to realistic network traffic. Designs based on traditional traffic assumptions may compromise the potential effectiveness of call admission control approaches like the often used linear CAC (and the related equivalent bandwidth). Furthermore, the concept of QoS has been trivialized in most CAC algorithms to be only the cell loss probability. The introduction of new measures, e.g., delay, delay jitter, and conditional cell loss rates will probably further restrict the set of admissible traffic mixtures and, as a

result, the above mentioned CAC may become even more insufficient.

On the other hand, due to long propagation delays inherent in broadband environments, long idle times of sources as well as requests for highly variable bandwidth connections resulting from a myriad of new services and applications (such as WWW, MBONE, network advertisement) have significantly increased the complexity of the control mechanisms in B-ISDN networks. The fundamental issue of the dominance of propagation delay in the face of very high transport speeds and highly bursty traffic, with diverse characteristics, means that reactive (adaptive feedback) control algorithms, like for instance window-based schemes, face significant challenges [Kle92]. Such mechanisms are performance limited by the round trip delay of the feedback information, which may not match the temporal constraints of many broadband services. Moreover, large amounts of buffer space is required in nodes to minimize the risk of buffer overflow.

Furthermore, several new strategies for congestion control proposed for high-speed networks, such as Virtual Clock, Delay Earliest Due Date (Delay-EDD), Jitter Earliest Due Date (Jitter-EDD) [ZF94], Dynamic Time Windows [MLF92] and Generalized Processor Sharing [PG93], are not very advantageous either. The main problems connected with these mechanisms (so-called *sorted priority queue* mechanisms) stem from the difficulties of providing diverse worst-case guarantees as well as from the need to monitor the traffic on a per-packet or per-connection basis at the switching nodes, a process that requires additional hardware and/or software capabilities.

We believe, in contrast, that solutions where the service controller is based solely on the aggregate traffic of each class (so-called *rate-based servers*) are more advantageous, due to their simplicity and the choice to open up the processing bottleneck. Examples of such mechanisms are Stop-and-Go [Gol91] and Hierarchical Round Robin (HRR) [KKK90]. These mechanisms, which are based on using diverse time-framing strategies, do not need to monitor the individual packets at every switching node. Instead, service disciplines compatible with the FIFO policy are used, where the data clustering is reduced, and this offers advantages in terms of protocol simplifications and performance guarantees.

It is also important to make a distinction between traffic characteristics (e.g., bandwidth requirements) and traffic requirements (e.g., delay and/or jitter requirements in data transport). This is an important observation, since it may have a decisive impact on the solution of choice for traffic management and congestion control. Namely, we believe that, in broadband environments, one should separate the rate adaptation/allocation from delay allocation. This concept is advantageous, since it offers the choice to decouple the interdependencies that exist, at other time-framing mechanisms like Stop-and-Go and HRR, between the frame size at admission and the queueing delay and/or the granularity of bandwidth allocation.

We therefore suggest a new approach for bandwidth allocation and congestion control, which is of the Rate Controlled admission with Priority Scheduling service type. It is called *Dynamic Time Sharing (DTS)*, because of the dynamic nature of the procedure for resource partitioning to allocate and guarantee a required bandwidth for every traffic class. This approach is based on guaranteeing specific traffic parameters (bandwidth requirements) through a policing unit, and then optimizing the bandwidth assignment within the network for specific parameters of interest (like delay or jitter, and loss). The optimization process is based on the parameters guaranteed by the policing unit. A batch admission policy is used at the edges of the network according to a specific framing strategy to follow the traffic characteristics (e.g., the traffic constraint function) of different traffic classes. On the other hand, another framing (congestion control) strategy is used within the network, which is based on different (delay/loss) requirements of the traffic classes. Proper management of bandwidth and buffer resources is provided in every (switch) node of the network, such as to guarantee the diverse performance of interest.

The remainder of this paper is organized as follows. The components of a network model with deterministic services are presented in section 2. The network model used for study is described in section 3. The sections 4 and 5 are devoted to the problem of network dimensioning, with the specific components (admission framing and service framing), such as to provide worst-case performance guarantees. Section 6 concludes the paper.

2 Components of a Network Model with Deterministic Services

Four types of traffic handling capabilities are currently viewed as relevant for ATM applications, namely Constant Bit Rate (CBR), Variable Bit Rate (VBR), Available Bit Rate (ABR) and Unspecified Bit Rate (UBR) [ATM95], [ITU95].

The loss- and delay-sensitive VBR traffic class is perhaps the most challenging for traffic control and congestion management. Some of the main problems are related in this case to the need for a proper resource reservation scheme to allocate network resources for each VBR traffic stream. Preventive control mechanisms are used to provide the contracted QoS. Precise source traffic descriptors are required to capture the burstiness and time correlations of different VBR sources and, at the same time, they should be simple enough to permit for the network to enforce or police. Another important problem is related to the process of clustering and ways to handle with it.

There are two classes of services that can be provided by a network for VBR traffic [Fer90]. These are of the type with *deterministic guarantees* (so-called *hard guarantees*), where diverse QoS parameters (throughput, loss, delay, etc.) are guaranteed even in the worst case, and of the type with *statistical guarantees* (so-called *soft guarantees*), in which probabilistic guarantees are provided by the network and where statistical multiplexing is used to improve the network efficiency in resource utilization (at the price of losses and/or delay/jitter).

Given the importance that (worst-case) loss and delay performance has in the provision of relevant QoS guarantees for VBR traffic (especially video traffic [Oht94]) as well as the drawbacks in using of stochastic traffic models (difficulties in capturing the self-similar traffic behavior, in policing and shaping, in dealing with heterogeneous sources and QoS requirements, etc.), we focus our efforts on network models with deterministic services. Accordingly, deterministic traffic models are considered that model the VBR traffic in terms of a worst-case description. It is also important to notice that such models do not necessarily involve peak rate allocation and low resource utilization for network [Pop94], [Kni96], [WKZL96].

There are three main components that have a decisive impact on the performance of a network model with deterministic services [WKZL96]. These are the analytical model used for the worst-case description of traffic, the admission control procedure as well as the scheduling discipline used at the network multiplexers.

A worst-case representation of a traffic source can be viewed as a model to provide an absolute upper bound on a source's traffic arrival [Cru91], [Gol91], [KZ95], [WKZL96]. This representation is usually valid for the aggregated amount of received data in a specific time interval but it could be valid for the total number of arrivals (packets or cells) in the specific interval as well. The model is parameterized in the sense that a source can efficiently specify its traffic characteristics and the network can use these parameters to provide the requested QoS as well as efficient resource utilization. The absolute upper bound is referred to as a *traffic constraint function* that provides an upper bound on traffic (one parameter) for every time interval of a specific length (another parameter) to capture burstiness properties [WKZL96]. The traffic constraint function is specific for different applications and reflects the properties of the process generating traffic as well as the properties (hardware and/or software) of the generating unit.

In the literature, there have been several models proposed to capture the interval-length dependent behavior of VBR traffic. Among them, the most relevant are the models proposed by Cruz [Cru91], Kurose [Kur92] and Knightly and Zhang (D - BIND model) [KZ95]. It is important to notice that the worst-case parametrization of these models can be specified with the help of a traffic constraint function that bounds the source [WKZL96].

Finally, the selection of particular solutions for scheduling and for admission control involves trade-offs between accuracy and complexity. Some of the most important problems considered in design are related to the process of clustering, provision of delay and/or jitter guarantees, provision of requested services (CO or CL), non-degradation of the service of existing connections, simplicity, etc. Unfortunately, due to the process of clustering, the burstiness of a connection's traffic accumulates along with the path from source to destination and this poses serious problems in the provision of

requested QoS. The existent solutions either reserve more resources to burstier traffic, or the traffic patterns are partially or completely reconstructed at each switch [ZF94]. The drawback in this case is that reconstructing traffic patterns at each switch requires non-work-conserving scheduling disciplines that may increase the average delay and decrease the average throughput.

3 Model Description

The network model used for study has a ring configuration with a number of n stations $\{S_1, \dots, S_i, \dots, S_n\}$, and a number of m substations connected to each station $\{SS_{i1}, SS_{i2}, \dots, SS_{ij}, \dots, SS_{im}\}$, where SS_{ij} represents the substation j connected to station i . It is assumed that every substation SS_{ij} can provide any type of traffic (i.e., subclass of traffic), which is decided according to different performance experiments.

Two parallel, wavelength- or time-separated channels (w_d and w_c), with time synchronization provided among them, are used for the transfer of data and for the corresponding protocol. All stations may transmit and receive on both channels.

I types of traffic subclasses are considered within the network, which are denoted by $\{t_1, \dots, t_I\}$. They are segregated according to their delay sensitivity and bandwidth needs. However, it should be mentioned that the traffic subclasses t_j , defined for service discipline within the network, may not be the same as the traffic subclasses defined at the edges of the network and partitioned on the basis of specific smoothness properties, e.g., the traffic constraint function. This is because of the adopted three-parameter characterization of traffic, called (Λ_i, D_i, T_i) -smooth, which does not contain any explicit information about the delay sensitivity of a specific application that belongs to the traffic class (simplicity reasons). The same class of traffic, with a specific traffic constraint parameter/function, may refer to multiple subclasses of traffic/services, which are segregated according to different delay requirements/sensitivities. This is in fact a consequence of the difference that may exist between the temporal properties of one class of traffic and the (delay) requirements of the same traffic. For simplicity reasons, we consider in the following that the subclasses of traffic, as defined within the network, correspond to those defined at the network edges. Information about delay requirements is implicitly contained in the type/subclass of traffic, as defined at the network edges.

A model that is similar to the Sustainable Cell Rate (SCR) model, defined by the ATM Forum [ATM95], is used to represent the worst-case behavior of the traffic. Time-invariant upper bounds are used in the definition of smoothness properties. A data/packet stream is said to comply with a smoothness constraint if the number of packets, or cells, as well as the amount of incoming data (bits), contained in any time interval are less than some specific upper bounds that depend on the interval length and the characteristics of application.

Total number of packets or cells (Λ_i) and total amount of data (D_i) in a specific time interval (T_i) are considered in determination of smoothness constraints. A data stream is said to be (Λ_i, D_i, T_i) -smooth if, during every interval of length T_i , the total number of packets, or cells, is limited to Λ_i , and the aggregated amount of received data is limited to D_i . No considerations are given to other figures, such as peak rate, average (bit) rate, average (burst) length, average number of packets in T_i , etc.

Having defined the model for worst-case description of traffic, we focus our efforts on the admission control and the scheduling discipline.

The admission control requires that a data stream/packet, which belongs to subclass i , must be (Λ_i, D_i, T_i) -smooth in order that the network provides the requested QoS. When the entering traffic does not comply with this smoothness, corrective actions are taken.

The protocol for bandwidth allocation and congestion control is a mechanism acting at the packet level. It emulates packet switching for I traffic classes, and it determines when a new incoming packet, with a given smoothness and latency requirement, should be given access to w_d , and for how long. This is a distributed control mechanism where every station stores, in a queueing table, the appropriate state of the system and, in specific buffers, the data waiting for access or passing through the

station [Pop94], [PP96]. A HOL non-preemptive model, with a single cyclic server and I queues of type FIFO, is used to model the transmit multiqueue system in the control channel. On the contrary, a HOL preemptive model, with a cyclic server and $(2I + 1)$ queues of type FIFO, is used to model the transmit multiqueue system in the data channel. The allocation of resources in both (w_c and w_d) channels is done, in every station, on the basis of the required delay performance (and loss probability), the smoothness properties of the traffic and the underutilized part of the allocated bandwidth. Worst-case guarantees for delay and loss probability are considered in the partitioning of resources, regardless of the channel speed, network size and traffic mix.

Every arriving packet will cause a reservation cell to be sent onto w_c in the current reservation frame, or one of the subsequent reservation frames. This cell circulates around the whole ring, and comes back to the transmitting station, where it is removed from w_c . It informs all stations, among them the receiving station, about the data message to be transmitted in one of the next frames in w_d . It contains information about the addresses (station, substation and process) of the transmitter and receiver, class of service, the temporal position of the first slot allocated for data transport and the length of message. Upon reception of reservation cells by stations, the corresponding packets (or parts of them) are assigned, by every station, free service slots in w_d , within the specific service framing and following the specific service discipline.

Following reception of one reservation cell from a substation, each station updates the content of its queueing table in the control channel. This table contains information needed to determine the state in the w_d channel, i.e., the occupied service slots. Based on this information, as well as on the information collected from other incoming reservation cells (upstream stations), the control mechanism is able to deduce the assignment of the outgoing free service slots in w_d . It also schedules these free slots to service the data messages waiting in station, according to fairness criteria. The scheduling unit is based on the specific policy used for resource partitioning, i.e., framing strategy, and the fairness is provided, in every station, by servicing different data messages such as the worst-case delay and loss probability are guaranteed for all traffic subclasses, regardless of the network size and traffic mix.

Further details on station modeling and diverse control mechanisms are presented in [Pop94] and in the extended paper [PP96].

4 Admission Framing

There are two distinct framing strategies used for the bandwidth allocation and congestion control, i.e., a packet/message admission policy imposed at the network edges, and another framing that is used for service (at the switching nodes) within the network. The admission policy makes use of a set of specific frames, denoted by F_k^{adms} (where $k = 1$ to I), which are chosen according to smoothness characteristics of different subclasses of traffic t_k . These frames are the same with the parameters T_i used in the three-parameter characterization of traffic (Λ_i, D_i, T_i) -smooth. On the other hand, the service policy makes use of another set of specific frames, denoted by F_k^{serv} (where $k = 1$ to I), which are chosen according to delay requirements in the transport of data.

The purpose of the admission policy is to provide admission of different packets/messages into the network according to their specific bandwidth requirements, i.e., maximum average number of arrivals and maximum average rate, and to preserve these requirements. This policy can be stated as follows: once a specific QoS is contracted between the user (i.e., substation) and the network (i.e., station), for a specific subclass of traffic t_k (where $k = 1$ to I), the arriving packet/message stream belonging to subclass t_k is required to be (Λ_k, D_k, T_k) -smooth in order to be admitted to the network.

The smoothness properties considered here are defined in the strongest sense in order to make them phase-independent. That means maximum average bandwidth requirements are considered, in the definition of frames F_k^{adms} , such as the (Λ_k, D_k, T_k) -smoothness holds in any frame F_k^{adms} , no matter the choice of temporal origin chosen in the definition of frames. Accordingly, specific rates must be allocated in w_c , and in w_d , to serve messages of class k received in the frame with number x (denoted by $F_k^{\text{adms},x}$). These rates are given by

$$r_{d,k}^x = \frac{D_k^x}{F_k^{adms,x}} \quad (1)$$

for the data channel w_d , and

$$r_{c,k}^x = \frac{\Lambda_k^x}{F_k^{adms,x}} \quad (2)$$

for the control channel w_c .

The parameters D_k^x and Λ_k^x represent the amount of data, and requests (for message transport), received from one substation in the time interval with number x , i.e., $F_k^{adms,x}$.

The rate $r_{d,k}^x$ is dimensioned in number of bits/frame, whereas $r_{c,k}^x$ is expressed in number of requests/frame. In a time period of 1 sec, we therefore have rates of

$$R_{d,k} = \sum_{x=1}^{\left\lfloor \frac{1 \text{ sec}}{F_k^{adms}} \right\rfloor} r_{d,k}^x \quad (3)$$

$$R_{c,k} = \sum_{x=1}^{\left\lfloor \frac{1 \text{ sec}}{F_k^{adms}} \right\rfloor} r_{c,k}^x \quad (4)$$

The rate $R_{d,k}$ is expressed in number of bits per second, and the rate $R_{c,k}$ is expressed in number of requests per second. The conditions for rate allocation associated with the admission policy are

$$\sum_{i=1}^n \sum_{j=1}^m \sum_{k=1}^l R_{d,k,i,j} \leq n_d \quad (5)$$

for the data channel w_d , and

$$\sum_{i=1}^n \sum_{j=1}^m \sum_{k=1}^l R_{c,k,i,j} \leq n_c \quad (6)$$

for the control channel w_c .

The parameters $R_{d,k,i,j}$ and $R_{c,k,i,j}$ represent the rates in w_d , and in w_c , associated with substation j connected to station i , and for (generated) traffic of class k . They are dimensioned in number of bits/sec, and number of requests/sec. The parameters n_d and n_c represent the capacity of w_d , and of w_c , channels, dimensioned in number of bits/sec, and number of cells/sec (dedicated only for protocol purposes). The admission policy therefore involves a specific framing strategy used for admission, according to different contracted QoS, as well as a proper dimensioning of the network, such that the (bandwidth) resources in w_d and w_c are not overwhelmed.

5 Service Framing

Having different packets/messages admitted to network, according to the above-mentioned admission policy, the next question to be answered is: how to serve these messages, inside the network, in a fair way? The answer to this question is decided according to the specific parameters of interest (e.g., delay, jitter, or loss) contracted for the specific traffic classes.

A set of (multiple) frames F_k^{serv} is used to provide the transport of data. It should be mentioned

that the frames with the same number x in w_d data channel ($F_{d,k}^{\text{serv},x}$), and in w_c control channel ($F_{c,k}^{\text{serv},x}$), are not simultaneous. Instead, they are delayed, from each other, due to different conditions (e.g., transmission delays) in data and control channels. A frame synchronization mechanism provides, in every station, proper correlation between these frames [Pop94].

The underlying idea for the service framing strategy is to provide "logical containers" within which diverse messages are transported, through the network, with the corresponding performance guarantees. This policy ensures that messages, of class k , admitted for transport, at the source station, during the frame with number x (i.e., $F_{d,k}^{\text{serv},x}$) will be "served" exclusively during this frame, such that they arrive to destination station within the same frame, and not in one of the subsequent frames (Fig. 1).

Fig. 1. Example of Transport Service in w_d

It therefore turns out that the end-to-end (transport) delay of a specific message of class k , traveling the network, from a given source station to a given destination station, is deterministic (in certain limits). This delay has two deterministic components, one related to network configuration (i.e., number of intermediate stations between source and destination, and distance between stations), and the second component related to class of priority (i.e., frame size). There is also one more unspecified component, which is related to the (indefinite) position of the message within the arriving frame to the source station, as well as the statistical properties of the messages of higher priority, which are admitted, for transport, during the same frame. It is reminded that, inside one service frame (in the data channel), an admitted message has preemptive priority over admitted messages of lower priority. It can be also preempted in service by (admitted) messages of higher priorities.

Two conditions must be fulfilled for the rate allocation associated with the service policy:

$$\sum_{i=1}^n \sum_{j=1}^m \sum_{y=1}^k r_{d,y,i,j} \leq C_{d,k} \quad (7)$$

$$\sum_{i=1}^n \sum_{j=1}^m \sum_{y=1}^k r_{c,y,i,j} \leq C_{c,k} \quad (8)$$

The parameters $C_{d,k}$ and $C_{c,k}$ (where $k = 1$ to l) represent the capacity, dimensioned in number of bits per service frame, and in number of cells per service frame, of one service frame of type k (i.e.,

F_k^{serv}) in data channel, and in control channel. The parameters $r_{d,y,ij}$ and $r_{c,y,ij}$ represent the message of class y (where $y = 1$ to k), generated by the substation j connected to station i (and dimensioned in requested number of bits per admission frame), and the request for this message (per admission frame).

It is also assumed that each frame size, used for service, is a multiple of smaller frame sizes, i.e.

$$F_k^{\text{serv}} = \xi_k \cdot F_{k-1}^{\text{serv}} \quad (9)$$

where $k = 1$ to I , and ξ_k is an integer.

Though the sizes of service frames are chosen based on delay requirements of the specific application, they are also related to the corresponding admission frames in a way to provide the conditions requested for worst-case guarantees. For instance, the following constraints must be satisfied by the aggregate service capacities $C_{d,k}$:

$$C_{d,1} \geq n \cdot m \cdot r_{d,1} \quad (10)$$

$$C_{d,2} \geq C_{d,1} \cdot \left\lceil \frac{F_2^{\text{serv}}}{F_1^{\text{adms}}} \right\rceil + n \cdot m \cdot r_{d,2} \quad (11)$$

$$C_{d,3} \geq C_{d,1} \cdot \left\lceil \frac{F_3^{\text{serv}}}{F_1^{\text{adms}}} \right\rceil + C_{d,2} \cdot \left\lceil \frac{F_3^{\text{serv}}}{F_2^{\text{adms}}} \right\rceil + n \cdot m \cdot r_{d,3} \quad (12)$$

Finally:

$$C_{d,k} \geq \sum_{x=1}^{k-1} C_{d,x} \cdot \left\lceil \frac{F_k^{\text{serv}}}{F_x^{\text{adms}}} \right\rceil + n \cdot m \cdot r_{d,k} \quad (13)$$

In a similar way:

$$C_{c,k} \geq \sum_{x=1}^{k-1} C_{c,x} \cdot \left\lceil \frac{F_k^{\text{serv}}}{F_x^{\text{adms}}} \right\rceil + n \cdot m \cdot r_{c,k} \quad (14)$$

The conclusion therefore is that, to provide worst-case performance guarantees, a *DTS* network has a size that is upper bounded (for any k) to:

$$n \cdot m \leq \left(\left(\frac{C_{d,k} - \sum_{x=1}^{k-1} C_{d,x} \cdot \left\lceil \frac{F_k^{\text{serv}}}{F_x^{\text{adms}}} \right\rceil}{r_{d,k}} \right)_{\min} ; \left(\frac{C_{c,k} - \sum_{x=1}^{k-1} C_{c,x} \cdot \left\lceil \frac{F_k^{\text{serv}}}{F_x^{\text{adms}}} \right\rceil}{r_{c,k}} \right)_{\min} \right)_{\min} \quad (15)$$

Some of the main performance metrics of interest are related in this case with the data transfer at a packet level, i.e., end-to-end delay (first order moments), delay jitter (second order moments), throughput and buffer requirements. Detailed theoretical studies are presented in [Pop94], [PP96].

One of the most important observations that have been made in these studies is that different traffic classes have different, and specific, performance limitations. Whereas in the case of traffic classes with higher priority (characterized by high arrival rates and small batch lengths) the limitation is given mainly by the (limited) capacity available in one service frame in the control channel, the situation is completely changed in the case of traffic classes of lower priorities (which are characterized by

low arrival rates and large message lengths). Here it is the (limited) capacity available in one service frame in the data channel that limits the performance. Furthermore, the performance results clearly indicate that dynamic mechanisms for resource partitioning should be also used in the control channel.

At this point, the fundamental differences that exist between this approach, called DTS, and the Stop-and-Go approach [Gol91], need to be pointed out. They refer partly to the conservation of smoothness properties, partly to end-to-end delay.

Concerning the smoothness properties, both mechanisms provide conservation, in certain limits, of the application smoothness properties. However, the smoothness properties are conserved, in the DTS case, at temporal scales that are compatible with the delay requirements of different traffic classes, as given by the service frames. It is not the case at Stop-and-Go, where the smoothness properties follow the common frames used both for admission and service. They get also related, in this case, to the number of intermediate stations between source and destination [Gol91].

Second, the end-to-end delay (variation) is much lower in the DTS case, since it is limited here to only one (service) frame. This is to compare with the Stop-and-Go solution, where the end-to-end delay (deterministic component) has values between once and two times the product between frame size and the number of intermediate stations between source and destination [Gol91].

The main drawback connected with the framing strategy takes different forms in the case of Stop-and-Go and DTS mechanisms. In order to provide performance guarantees, as specified in the contracted QoS, a rate-based service discipline must accept non-work conserving [ZF94]. This means that it might happen that the network/node server is idle, though there are packets waiting for service (at the network edges or inside the nodes), but that are not yet admitted for service. The Stop-and-Go mechanism is typical for such kind of problems. It can be reduced/eliminated through the provision of a (lowest priority) traffic class with no delay guarantees, with best-effort service.

Similarly to the non-work conserving problem, the DTS mechanism could also have problems that take the form of non-fairness. This means that, in order to provide delay guarantees, as contracted in QoS, it might happen that the network/transport server is servicing traffic messages of a specific priority, though there are messages of higher priority admitted into the network and waiting for service. This problem is mainly related to the rate allocation policy chosen for service, frame sizes, and the mechanism used for resource sharing/scheduling in stations. Such problems can be reduced/eliminated through better mechanisms used for resource sharing (for instance, dynamic mechanisms based on partial sharing) as well as a proper rate allocation associated with the service policy.

6 Conclusions

A new approach for bandwidth allocation and congestion control has been suggested to provide bounded delay and jitter requirements for loss-free traffic. We call this *Dynamic Time Sharing (DTS)* and it is of the Rate Controlled admission with Priority Scheduling service type. This approach is based on guaranteeing specific traffic parameters (bandwidth requirements) through a policing unit, and then optimizing the bandwidth assignment, within the network, for specific parameters of interest (like delay or jitter, and loss) regardless of traffic statistics, network size and transmission speed.

Network dimensioning to provide worst-case guarantees has been presented for the case of a ring topology. Two components are involved in this dimensioning, namely the admission framing and the service framing, and the relation between them and the admissible number of traffic sources has been deducted. Furthermore, the fundamental differences that exist between the DTS approach and the Stop-and-Go approach have been pointed out, with the indication of the advantages of DTS.

The DTS approach is applicable to other topologies as well. This is topics for further research. Other areas of interest are related to assessing the applicability of this approach to VBR video traffic, study of the efficiency in utilization of network resources, diverse trade-offs between QoS guarantees and network utilization as well as utilization of better models for worst-case description of traffic.

References

- [ATM95] ATM Forum Traffic Management Specification, Version 4.0, 1995.
- [BSTW95] Beran, J., Sherman, R., Taqqu, M.S. and Willinger, W., "Long-Range Dependence in Variable-Bit Rate Video Traffic," *IEEE Transactions on Communications*, Vol. 43, No. 2/3/4, pp. 1566 - 1579, February/March/April 1995.
- [Cru91] Cruz, R.L., "A Calculus for Network Delay, Part 1: Network Elements in Isolation," *IEEE Transactions on Information Theory*, Vol. 37, No. 1, pp. 114 - 131, January 1991.
- [Fer90] Ferrari, D., "Client Requirements for Real-Time Communication Services," *IEEE Communications Magazine*, pp. 65 - 72, November 1990.
- [Gol91] Golestani, S.J., "A Framing Strategy for Congestion Management," *IEEE Journal on Selected Areas in Communications*, Vol. 9, No. 7, pp. 1,064 - 1,077, September 1991.
- [ITU95] ITU - T Recommendation I.371, Traffic Control and Congestion Control in B-ISDN, 1995.
- [Jai95] Jain, R., "Congestion Control and Traffic Management in ATM Networks: Recent Advances and A Survey," *Computer Networks and ISDN Systems*, 1995.
- [KKK90] Kalmanek, C.R., Kanakia, H. and Keshav, S., "Rate Controlled Servers for Very High-Speed Networks," in Proceedings of *IEEE GLOBECOM'90*, San Diego, California, pp. 12 - 20, December 1990.
- [Kle92] Kleinrock, L., "The Latency/Bandwidth Tradeoff in Gigabit Networks," *IEEE Communications Magazine*, pp. 36 - 40, April 1992.
- [Kni96] Knightly, E.W., "H - BIND: A New Approach to Providing Performance Guarantees to VBR Traffic," in Proceedings of *IEEE INFOCOM'96*, San Francisco, California.
- [Kur92] Kurose, J., "On Computing Per - Session Performance Bounds in High - Speed Multi - Hop Computer Networks," in Proceedings of *ACM SigMetrics'92*, Newport, June 1992.
- [KZ95] Knightly, E. and Zhang, H., "Traffic Characterization and Switch Utilization Using Deterministic Bounding Interval Dependent Traffic Models," in Proceedings of *IEEE INFOCOM'95*, 1995.
- [LTWW94] Leland, W.E., Taqqu, M.S., Willinger, W. and Wilson, D.V., "On the Self-Similar Nature of Ethernet Traffic (Extended Version)," *IEEE/ACM Transactions on Networking*, Vol. 2, No. 1, pp. 1 - 15, February 1994.
- [MLF92] Mukherjee, A., Landweber, L. and Farber, T., "Dynamic Time Windows and Generalized Virtual Clocks: Combined Closed Loop/Open Loop Congestion Control," in Proceedings of *IEEE INFOCOM'92*, Florence, Italy, pp. 322 - 332, May 1992.
- [Oht94] Ohta, N., *Packet Video Modeling and Signal Processing*, Artech House, 1994.
- [PG93] Parekh, A.K. and Gallager, R.G., "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single Node Case," *IEEE/ACM Transactions on Networking*, Vol. 1, pp. 344 - 357, June 1993.
- [Pop94] Popescu, A., "A Parallel Approach to Integrated Multi-Gbit/s Communication over Multi-wavelength Optical Networks," Ph.D. dissertation, Royal Institute of Technology, Stockholm, 1994.
- [PP96] Popescu, A. and Pruthi, P., "Dynamic Time Sharing: A New Approach For Congestion Management (Extended Version)," the 4th IFIP ATM Workshop, Bradford, July 1996.
- [Rob95] Roberts, J.W., "What Traffic Handling Capabilities for the B-ISDN?" in Proceedings of *9th ITC Specialists Seminar*, The Netherlands, November 1995.
- [WKZL96] Wrege, D.E., Knightly, E.W., Zhang, H. and Liebeherr, J., "Deterministic Delay Bounds for VBR Video in Packet - Switching Networks: Fundamental Limits and Practical Tradeoffs," *IEEE/ACM Transactions on Networking*, June 1996.
- [ZF94] Zhang, F. and Ferrari, D., "Rate - Controlled Service Disciplines," *Journal of High - Speed Networks*, 3(4), pp. 389 - 412, 1994.