MITSUBISHI ELECTRIC RESEARCH LABORATORIES http://www.merl.com

# On ATM Support for Distributed Real-Time Applications

Chia Shen

TR96-01a April 1996

#### Abstract

In this paper, we consider the problem of applying ATM technology in designing distributed industrial plant control applications with integrated real-time communication requirements. In particular, we first examine (1) the design space using current ATM service categories and their corresponding traffic models, and (2) the mathematical and practical implications of the traffic models. We then present a taxonomy of the algorithms available to implement the service categories that can provide performance guarantees. Through the examination, we shed light on the gap between what is usually assumed in research and what the actual ATM network will provide. We then show how real-time communication can be mapped onto ATM service categories, and demonstrate the limitations of ATM services. Finally, we describe schemes which can be employed to overcome the limitations and some open research issues.

Proceedings of IEEE Real-Time Technology and Applications Symposium, June 10-12 1996

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Copyright © Mitsubishi Electric Research Laboratories, Inc., 1996 201 Broadway, Cambridge, Massachusetts 02139



## On ATM Support for Distributed Real-Time Applications

Chia Shen Mitsubishi Electric Research Labs Cambridge Research Center 201 Broadway, Cambridge Mass 02139 shen@merl.com

#### Abstract

In this paper, we consider the problem of applying ATM technology in designing distributed industrial plant control applications with integrated real-time communication requirements. In particular, we first examine (1) the design space using current ATM service categories and their corresponding traffic models, and (2) the mathematical and practical implications of the traffic models. We then present a taxonomy of the algorithms available to implement the service categories that can provide performance guarantees. Through the examination, we shed light on the gap between what is usually assumed in research and what the actual ATM network will provide. We then show how real-time communication can be mapped onto ATM service categories, and demonstrate the limitations of ATM services. Finally, we describe the concept of User Level Multiplexing as a means to overcome the limitations.

#### 1. Introduction

In this paper, we consider the problem of applying ATM technology in designing distributed industrial plant control applications with integrated realtime communication requirements [12]. In such a realtime application environment, many types of communication, including plant monitoring data, operator command and control, video, images and audio, must be accommodated. Traditionally, such distributed plant control systems have been built mostly with proprietary network components. Real-time constraints are met with physical isolation of sub-networks. For example, in order to guarantee the timing constraints of plant control, the human operators and their workstations do not reside on the same physical network as the network in the controlled plant sites. Similarly, video transmission from industrial TV cameras in the plant

monitoring sites is through yet another separate physical network. Consequently, these current plant control networks are inflexible, expensive and difficult to modify.

We envision that the four key technical advantages of ATM, i.e., high bandwidth, end-to-end QoS, virtual connections and flexible topology, will enable us to build future distributed plant control systems on open network technology with integrated communication services. This should lead to cost reduction, more flexibility and more dynamics in our applications. Since ATM is still in its infancy, there exist many challenges before we can actually deploy ATM for such real-time applications. Specifically, to provide guaranteed service for real-time applications in an integrated environment, one challenge is the design of an ATM network architecture that actually provides the endto-end QoS (Quality of Service) guarantees, and a second challenge is the design of middleware algorithms and protocols that (1) expose and explore ATM features, and (2) effectively map applications onto ATM services. This paper explores the relationship between these challenges. In particular, we analyze and compare the research results as well as industrial standards at the ATM network architecture level, and propose solutions at the middleware level.

There has been a lot of research done at the ATM network level. Research at this level is concerned with traffic specification models, CAC (Connection Admission Control) algorithms, and cell scheduling algorithms that can provide deterministic or statistical real-time guarantees for video, audio and data traffic types [5, 6, 8, 9, 10, 15, 18, 21, 20, 25]. Network level research also addresses new protocols and service models, such as the work presented in [7, 22, 24]. Research at the middleware level has just begun. Here research is needed to close the gap between what real-time and multimedia application designer needs and what the ATM network actually provides. Research at

the middleware level deals with problems of QoS mapping, user level multiplexing, transport protocols, connection management, and adaptive QoS management [1, 16]. Middleware level research also needs to address the design of QoS based API that can provide real-time programming services and interface to the applications.

Effective research at the middleware level requires a thorough understanding of what is provided by the communication services in an ATM network. Moreover, to facilitate the design of efficient middleware algorithms and protocols, it is necessary to thoroughly understand the underlying mechanisms which implement ATM communication services.

In this paper, we first examine (1) the design space defined by the current ATM service categories and their corresponding traffic models, and (2) the mathematical and practical implications of the traffic models. Then we present a taxonomy of the algorithms available to implement the service categories that can provide performance guarantees. Through out this examination, we shed light on the gap between what is usually assumed in research and what the actual ATM network will provide. We then show how real-time communication can be mapped onto ATM CBR and rt-VBR service categories. We demonstrate that real-time periodic communication is not necessarily best implemented with CBR, although some of the previous research [9, 15] has in essense assumed the CBR traffic model as the hard real-time communication traffic model. We then show the limitations of ATM services and how User Level Multiplexing can be used to overcome the limitations.  $^{1}$ .

The rest of the paper is organized as follows. In Section 2, we describe the design space in ATM and analyze traffic models used for CBR and rt-VBR service categories. In Section 3, we present a coherent taxonomy of the traffic scheduling algorithms available for the implementation of real-time ATM services. Section 4 illustrates how real-time communication and the corresponding QoS requirements can be mapped onto ATM. We demonstrate the limitations of ATM in Section 5 and describe User Level Multiplexing in Section 6. Section 7 concludes the paper.

### 2. Design Space in ATM

When we map application traffic onto ATM, our design space is defined by the service categories that ATM promises to provide. In figure 1, we illustrate this concept. There are five service categories currently defined by the ATM Forum Traffic Management Subworking Group [2], namely CBR (Constant Bit Rate), rt-VBR (real-time Variable Bit Rate), nrt-VBR (non-real-time Variable Bit Rate), ABR (Available Bit Rate), and UBR (Unspecified Bit Rate). Traffic control will be applied to meet QoS requirements of applications using these services. We classify the traffic control schemes for each of the service categories into two types open loop and closed loop control. In particular, open *loop* control is imposed for CBR, rt-VBR and nrt-VBR, where the source QoS requirement with respect to delay bound is guaranteed only if the source does not exceed its specified sending rate. ABR service is controlled via *closed loop* feedback mechanisms in which the network periodically tells the source how much traffic a source can generate according to the network load in order to avoid undesirable amount of cell loss. No delay bound can be guaranteed for ABR service. UBR service is not subject to any types of network traffic control, therefore no QoS is provided to UBR traffic either.

Given the ATM service categories, a distributed real-time system designer needs to intellegently map the actual application requirements onto these services. In the middle box in Figure 1, we list some of the important issues such a designer must take into consideration. These issues include the guarantee ratio for hard real-time connections, the level of network utilization and the degree of multiplexing that can be obtained among real-time connections, as well as among realtime and non-real-time connections. We will return to address some of these issues in sections 4 and 5.

Open-loop control is particularly suited for real-time communication. By mapping real-time communication onto either CBR or rt-VBR services, performance guarantees can be obtained from the network with respect to delay bounds, and loss ratio. Thus, in order to fully understand the usefulness and the limitations of ATM for real-time applications, we must first understand what exactly CBR and VBR services can provide us with by examining and analyzing their properties. In the following, we discuss (1) the traffic model used for specifying the characteristics of CBR and rt-VBR traffic, (2) the mathematical and practical implications of the traffic model in Section 2.1, and (3) the cell scheduling disciplines and connection admission control (CAC) algorithms that can be employed by the network to implement CBR and rt-VBR services in Section 3.

<sup>&</sup>lt;sup>1</sup>In this paper, we use the term User Level Multiplexing to mean multiplexing at any level of the system above the ATM cell multiplexing level, i.e., it can be at application level, at the operating system or protocol stack, or at the network interface driver level.



Figure 1. Design space defined by ATM service categories.

#### 2.1. Traffic Models of CBR and rt-VBR

The traffic and QoS parameters used to specify CBR and rt-VBR traffic are summarized in Table 1. In the table, PCR, SCR, and MBS are traffic parameters to characterize a traffic source, while CTD and CDV are QoS parameters of a connection. Both CTD and CDVare end-to-end delay requirements measured from the exit of the source to the entry of the destination of a connection. CTD should be the sum of the fixed propagation delay on a path plus the worst case CDV. According to the current ATM Forum Traffic Management Specification Version 4.0, a CBR connection is characterized by its PCR. A rt-VBR connection also specifies SCR and MBS in addition to PCR. CBR is designed to support real-time connections that requires a static amount of bandwidth and tightly constrianed cell delay variation, while rt-VBR is for real-time connections which require transmission rate that may vary with time or is bursty. The QoS parameters for both CBR and rt-VBR connections include maximum CTD and CDV.

The traffic models for CBR and rt-VBR are both deterministic models. Traffic specification with a single parameter of PCR for CBR traffic is straightforward. rt-VBR is more complicated. The traffic generated at the source by a rt-VBR connection is characterized by three traffic parameters (PCR, SCR, MBS). Intuitively, rt-VBR traffic is a stream of ATM cells such that

- the long term average cell rate is SCR,
- 1/PCR is the minimum inter-cell arrival time, and
- during any time interval of length MBS \* 1/SCR, the maximum number of cells that can arrive with an inter-cell arrival time of 1/PCR is MBS.

Let's analyze the properties of traffic generated from a rt-VBR connection  $VC_i$ . Given the three traffic parameters of  $VC_i$ , i.e.,  $PCR_i$ ,  $SCR_i$  and  $MBS_i$ , at the traffic source we can derive (1) the traffic contraint function [3], (2) the minimum interval between maximum bursts of cells, and (3) the maximum earliness of a cell.

The traffic constraint function f(t) at the first switch is

$$f(t) = MBS_i + SCR_i * t$$

That is, let N be the maximum number of cells generated by a conforming rt-VBR source during any time interval t, then  $N \leq f(t)$ . The minimum time interval I between maximum size bursts of cells is

$$I = MBS * (1/SCR - 1/PCR)$$

The maximum earliness  $\tau$  of a cell is

$$\tau = (MBS - 1) * (1/SCR - 1/PCR)$$

In an ATM network, cells from different connections are multiplexed at each switch. Even if we have a complete deterministic characterization of the traffic at the

Description	
Peak Cell Rate: The cell rate that a source can never exceed.	
Sustainable Cell Rate: The long term average cell rate.	
Maximum Burst Size	
End-to-end Cell Delay Variation induced by buffering and cell scheduling.	
End-to-end Cell Transfer Delay from the exit of a source to the entry of a destination.	

Table 1. CBR and VBR traffic and QoS specification parameters

source, this cell multiplexing along the path of a connection will result in traffic distortion. With respect to VBR traffic, when traffic rate restoration is not assumed to be done at each switch, we can derive the traffic arrival specification for the switches along the way as follows: Let  $d_i^j$  be the delay at the *j*th switch for  $VC_i$ . If we ignore the link propagation delays for now, then the three parameter traffic arrival specification  $A_j$  for a rt-VBR traffic of  $VC_i$  at switch *j* for 1 < jwould be

$$A_j = (PCR^l, SCR_i, MBS_i + SCR_i * \sum_{k=1}^{k=j-1} d_i^k)$$

subject to  $\sum_{k=1}^{k=n} d_i^k \leq CDV_i$ , where *n* is the total number of switches along the path of connection  $VC_i$ , and  $PCR^l$  is the link speed. The constraint  $\sum_{k=1}^{k=n} d_i^k \leq CDV_i$  came from the fact that if the network admits the rt-VBR connection  $VC_i$ , then the QoS parameters  $CDV_i$  and  $CTD_i$  must have been guaranteed. This implies that the traffic distortion inside the network cannot be unbounded in practice.

The above functions are useful in calculate the buffer requirements, and to generate conforming traffic for a connection. Note that one desirable property of rt-VBR service is that it allows bursts of cells, rather than only evenly distributed cells as in the CBR service. Moreover, these bursts of cells from a connection need not be smoothed out into a traffic with rate SCR at the entrance to the network as assumed by some of the previous research [9]. Therefore, we can model realtime messages with bursts of cells.

We would like to mention briefly the practical advantages of such simple deterministic traffic model used for CBR and rt-VBR — This deterministic traffic model ensures that only a fixed number of leaky buckets is needed by the virtual scheduler for cell transmission at a source (e.g., in the network interface card for a workstation) and by the traffic policer at the UNI (User Network Interface) or NNI (Network to Network Interface) [3]. Since a leaky bucket can be implemented with one counter and a single timer [17], the traffic model is thus simple enough to enable the network traffic to be monitored and controlled in real-time at ATM data rates. In this section, we have presented the design space in ATM for real-time applications through analysis of the traffic model for CBR and rt-VBR. We next describe a taxonomy of the cell scheduling algorithms that can be used to implement these ATM services.

## 3. Cell Scheduling Disciplines

With the traffic model and the corresponding specifications described in the previous section, the network allocates resources in terms of buffer space, cell transmission rate and delay bounds to a connection<sup>2</sup>. This resource allocation is translated into specific cell scheduling algorithms at cell transmission time and specific CAC algorithms at connection establishment time both at the ATM switches along a communcation path, as well as at the source and destination system. Since CAC algorithms entirely depend on the cell scheduling algorithm, and traffic and QoS specification, we focus on cell scheduling disciplines only in this section.

There are in general three classes of cell scheduling algorithms, i.e., First Come First Serve (FCFS), Static Priority (SP), and Dynamic Priority (DP). These three classes of algorithms differ in their ability to provide delay bounds and the complexity of their implementation.

#### First Come First Serve (FCFS)

Even though it is a general belief that FCFS is not a real-time scheduling algorithm, it has been shown in [15, 20] that deterministic delay bounds can be obtained taking into consideration of the peak cell rates of all the connections. The attractive feature of FCFS is its implementation simplicity. However, there are a number of obvious drawbacks of FCFS that make this scheduling discipline unsuitable for real-time systems with integrated communication requirements:

• Since the maximum delay in a FCFS scheduler is the same for all connections, it can only provide the same worst case delay bound to all connections. This is not good enough in an integrated en-

<sup>&</sup>lt;sup>2</sup>Since we are concerned with hard real-time applications in this work, we always require lossless transmission. Thus we do not explicitly deal with loss ratio as a QoS parameter in this paper.

vironment where real-time connections may have diverse QoS requirements. Using FCFS will in essence force us to guarantee all real-time connections with the delay bounds of the connection with the most stringent QoS requirement.

• It does not have the firewall property [10] i.e., the property that the guarantee to a connection is not affected by the behavior of other connections sharing the same scheduler. Thus, using FCFS,connections are not protected from misbehaving sources. This is a very important requirement in our distributed plant control application.

#### Static Priority (SP)

The class of Static Priority (SP) algorithms include any packet scheduling algorithms that assign the priority of a connection at connection set up time, and thus all cells belonging to the same connection will have the same priority thereafter upon arrival. Examples of SP algorithms include Rate Monotonic [13], a rate controlled version of Rate Monotonic called TCRM [9], and RCSP [19]. Algorithms in the SP class have the following properties:

- Only a fixed number of delay bounds can be provided.
- All connections with the same priority have the same delay bound these connections are serviced in FCFS order. Because of this limitation, the firewall property can only be provided to connections belonging to different priorities.
- To implement SP algorithms, a network switch need only search or sort by the number of priority levels, not by the number of cells in the queue. Thus, if the number of priority levels is sufficiently small, this search can be implemented very fast (e.g., by hashing, table lookup, electronic hardware, etc.). However, if the number of priority levels is large, say 64K or more, then finding the first static priority level with a non-empty queue is computationally just as complex as sorting a queue of cells with that many entries.

#### Dynamic Priority (DP)

Dynamic Priority (DP) is theoretically the most powerful class of packet scheduling algorithms among the three classes presented in this section in the sense that (1) these algorithms can guarantee each connection an individual delay bound, (2) the firewall property holds, and (3) they can fully utilize the link bandwidth. Algorithms in this class include variations of the well-known EDF (Earliest Deadline First) algorithm [4, 21, 25], Virtual Clock [23], and PGPS (Packet-by-Packet Generalized Processor Sharing) [14] <sup>3</sup>. Although Virtual Clock and PGPS are rate-based <sup>4</sup> scheduling algorithms, lately it has been shown that they can also provide delay guarantees to connections [5, 18].

However, not all DP algorithms are equal in (1) their efficiency to utilize the network bandwidth while providing inidividual delay bounds, (2) their CAC calculation, and their implementation complexity. We elaborate on these points briefly here.

• Efficiency:

Rate-based algorithms, such as Virtual Clock and PGPS, provide end-to-end delay guarantee at the cost of coupling reserved rates and delay bounds. This coupling is inherent since these algorithms schedule cells/packets entirely based on the reserved cell rate of a connection. Thus to obtain a smaller delay bound, a connection is forced to ask the network to reserve a high rate, even though it may not need that rate. EDF-based algorithms do not have this direct coupling between delay bound and reserved rate for a connection. As it turns out, this rate-delay coupling is the fundemental difference between rate-based and EDF-based algorithms since it decides the complexity of both CAC calculation and implementation of these algorithms as explained below.

• CAC calculation:

To avoid the rate-delay coupling, EDF-based algorithms need exact CAC calculations (i.e., a schedulability test) that depend on the traffic patterns of all flows sharing the scheduler. CAC calculation for rate-based algorithms is simple, it only needs to check that the aggregated rate allocation will not exceed the portion of the server capacity reserved for real-time traffic.

• Implementation:

All DP algorithms require (1) to calculate a cell's priority upon cell arrival, and (2) to sort/search backlogged cells with priority values. Here we assume the most powerful and general EDF model, i.e., it is work-conserving without rate control at the input to a switch. Since all DP algorithms

 $<sup>^{3}\</sup>mathrm{These}$  are only example algorithms. It is by no means an exhaustive list.

 $<sup>^4</sup>$ Note here that the term rate-based should not be confused with the rate based feedback flow control scheme developed for ABR traffic.

need a search/sort on cell priorities, their implementation complexity differ in the priority calculation.

To compare the complexity of the dynamic priority calculation of EDF-based and rate-based algorithms, we observe the following:

— In the calculation of the priority (e.g., the virtual clock value in the Virtual Clock algorithm) for each arriving packet/cell, rate-based algorithms only needs the previous priority value (which may depend on the arrival time of the previous cell) and an algorithm dependant constant as described in [5]. This calculation is simple and can be done efficiently with one read and one addition.

— In the calculation of the deadline value for each packet/cell arrival, EDF-based algorithms depend not only on the previous cell arrival time, but also on the position of the cell in a real-time message/packet. To see this, let us examine the two ways that the deadline of a cell can be calculated:

- (1) A cell can carry the deadline value with it. This requires the source system to assign this value and keep track of which cells belonging to the same message (counting). But this may require using fields either in the ATM cell header or the payload, thus may be infeasible to implement in a real network.
- (2) The switches can calculate this deadline value. This reqires that the switch either counts the number of cells — cells of different messages may become 'lumped' together enroute due to traffic distortion, or recognizes the end-of-message cell.

Actually, the simplicity of the priority calculation for rate-based algorithms is due to their inherent coupling of the guaranteed rate with the delay bounds. Similarly, the complexity of the general deadline value calculation results from the decoupling of reserved rate and delay bounds guaranteed.

In the above, we have shed light on the intrinsic properties of cell scheduling algorithms that can provide deterministic delay guarantees. This is summarized in Table 2. The remaining question is which scheme(s) is (are) the right choice(s) for CBR and rt-VBR services.

From the discussions in this section, one can see that cell scheduling algorithms in all three classes can be used for CBR and rt-VBR service categories since they can all provide delay bounds to satisfy QoS requirements, as long as CAC is applied. However, for the sake of network efficiency and for the diversity of multimedia and real-time application requirements, FCFS is certainly not a good choice for future ATM networks. SP is already feasible in today's ATM switches. As far as EDF-based algorithm is concerned, although powerful, determining the position of cells within a message or the end of messages may require the cooperation of the SAR (cell assembly and disassembly) process at a ATM network interface at an end system and the switches inside the network. This may be difficult unless all network components in the network agree to do so <sup>5</sup>.

On the other hand, rate-based dynamic priority algorithms are very attractive. In addition to supporting individual priorities and the firewall property, another advantage of rate-based algorithms is that there exist methods to determine end-to-end delay bounds in heterogeneous networks. In [5], a class of Guaranteed Rate (GR) scheduling algorithms was defined, and a method for determining an upper bound on the end-to-end delays for a network of switches each of which employs a scheduling algorithm in the GR class was developed. It was shown that many rate-based algorithms, such as Virtual Clock and PGPS, belong to GR. This means, even when switches in the network employ different GR scheduling algorithms, delay bounds still can be guaranteed. This is particularly attractive since we cannot require all the network vendors to use the same scheduling algorithm. Moreover, there is evidence that it will be feasible to search a queue of cells electronically in the near future [11], therefore we postulate that ratebased dynamic priority algorithms will be employed for future CBR and rt-VBR traffic.

In this section, we have presented a taxonomy of the scheduling algorithms that can be used to implement CBR and rt-VBR service categories to provide deterministic guarantees. With the ATM design space and its implementation in mind, we next present how to map real-time traffic onto ATM services, and why CBR is not necessarily the best or feasible implementation choice.

<sup>&</sup>lt;sup>5</sup>If AAL5 packets can be assumed to be used by the applications, then the End of Message (EOM) bit can be used since EOM is already used for Early Packet Discard algorithms by some ATM switches. However, we must then deal with the case of when the cell carrying the EOM bit is lost or corrupted.

	Number of delay bounds	Firewall Property	Implementation
FCFS	1	No	No special support
SP	Fixed	Per priority	Per priority queueing
DP	Arbitrary	Per connection	Per cell priority calculation
			Per cell sorting/searching

Table 2. Properties of cell scheduling algorithms

## 4. Mapping Real-Time Traffic onto ATM Services

Real-time applications usually model their time constrained communication with a 3-parametered realtime connection, (P, M, D), where P is the period of message generation, M is the maximum message size, and D is the deadline by which the entire message (not individual cells) must be delivered to the destination. Since we consider ATM networks specifically in this paper, we let M be the maximum number of ATM cells (53 bytes/cell) in a message.

To map such a real-time connection onto an ATM virtual connection (VC), at first glance, it may seem obvious that one should use CBR service due to the periodicity of the model as follows:

$$PCR = M/P$$

However, this CBR mapping with PCR = M/P is only viable for cases where  $D \ge P$ . In distributed industrial plant control applications, it is not uncommon to have real-time communication requirements with D < P. For example, operators often need to periodically issue commands to controlled sites. The period between each command issuing can be much larger than the deadline by which the commands should arrive at the controlled sites and be executed.

If D < P, we have two choices. One is to map onto a CBR VC with PCR = M/D. However, a better choice in terms of network bandwidth usage efficiency is to map the real-time connection onto a rt-VBR VC. Given a real-time connection with parameters P = 8 \*ctt, M = 4, and D = M \* ctt + 2ctt, where ctt = one cell transmission time, the following is a valid mapping to a rt-VBR VC (remember that a rt-VBR is charaterized by three traffic parameters and two QoS parameters):

• PCR = linkspeed, SCR = 0.5 \* linkspeed, MBS = 4, CTD = 2ctt, CDV = 2ctt.

To see that this is a valid mapping, we can use the formula for I from Section 2.1 to calculate the minimum interval between maximum bursts of cells, which will confirm that the real-time connection can indeed

transmit periodic messages of size M at rate P while still conforming to the rt-VBR specification.

In this section, we have given examples to illustrate how real-time communication requirements might be mapped onto ATM services. Even though CBR is a very simple communication service to use, certain realtime communication requirements need to be mapped onto the more complicated rt-VBR service instead. Next we examine the performance limitations that may arise in doing such mapping.

## 5. Limitations of ATM for Real-Time Communication

ATM is very different from traditional shared media networks because it is based on cell multiplexing. In shared media networks, such as FDDI, once a real-time connection obtains the access right to the network, it can transmit at the link speed until its allocated network resource is used up, or it is explicitly preempted from the network. However, communication in ATM is forced into cell multiplexing among connections sharing the same link or the same switch output based on their *rate* specifications, as well as their *CTD* and *CDV* QoS requirements.

The level of multiplexing by a network depends on how the network allocates resources with respect to meeting the delay bounds required by the applications. Since EDF-based scheduling algorithms are not likely to be widely adopted, if at all, by ATM switch vendors in the near future, the network would have to translate the delay bound requirements into rate or bandwidth reservations. As we have discussed in Section 3, rate based scheduling couples delay bounds with rates reserved. This coupling leads to fundamental limitations in network performance with respect to the number of bursty real-time connections that can be admitted <sup>6</sup>.

To guarantee a given delay bound, the network needs to ensure that even when the traffic source generates cells at its peak rate, the desired end-to-end delay bound expressed as CTD and CDV should still be met.

<sup>&</sup>lt;sup>6</sup>Here we use the term rate-based to mean any cell scheduling algorithms other than EDF-based algorithms because, in essence, all the other algorithms provide delay guarantees by allocating or reserving a certain rate for a real-time connection.

Thus, rate allocation inherently depends on the value of M/CTD. However, in reality, the real-time connection does not need a rate of r = M/CTD all the time. This over commitment of the network bandwidth can result in very low admission rate for real-time connections. To see this point, using the same example as illustrated in the last section, this real-time connection will need all the bandwidth half of the time. If we allocate a rate of  $r = \frac{M}{M*ett+2ett}$  to this VC, the network can hardly guarantee any other real-time connections.

To design distributed real-time applications using ATM as the underlying network technology, we next describe user level multiplexing as part of middleware protocols to overcome the above discussed network limitations.

#### 6. User Level Multiplexing

Intuitively, as the name implies, we can map multiple real-time traffic into one rt-VBR or one CBR. For example, if we know that the aggregated traffic of two real-time channels will not exceed some rt-VBR traffic specification, and each one of them cannot efficiently use one rt-VBR individually, we can do User Level Multiplexing. User Level Multiplexing can be applied locally at individual nodes, and globally at the distributed system.

In the following, we give an example to show how User Level Multiplexing can be done on a node. Suppose we have two real-time connections  $VC_i$  and  $VC_j$  to the same destination with parameters ( $P_i =$  $8ctt, M_i = 4, D_i = (M + 2)ctt$ ), and ( $P_j = 10ctt, M_j =$  $4, D_j = P_j$ ). (These parameter values are somewhat extreme values for simplicity.) Suppose these are the only two real-time connections and they will only go through one ATM switch. If we map them onto individual rt-VBR VCs, we obtain:

 $VC_i$ :  $PCR_i = linkspeed$ ,  $SCR_i = 0.5linkspeed$ ,  $MBS_i = 4$ ,  $CTD_i = 2ctt$ ,  $CDV_i = 2ctt$ .

 $VC_j$ :  $PCR_j = linkspeed$ ,  $SCR_j = 0.4linkspeed$ ,  $MBS_j = 4$ ,  $CTD_j = 6ctt$ ,  $CDV_j = 6ctt$ .

Using peak rate allocation, they will not both be admitted into the network. However, we can multiplex them at the source end system into one rt-VBR channel  $VC_k$  with parameters  $PCR_k = linkspeed$ ,  $SCR_i = 0.9*$ linkspeed,  $MBS_k = 8$ ,  $CTD_k = 2ctt$ , and  $CDV_k = 2ctt$ . At the user level, the end system ensures that the message from  $VC_i$  is always transmitted before the message from  $VC_i$ .

We observe the following from the above example using User Level Multiplexing:

• In doing User Level Multiplexing, we are actually using EDF-based scheduling at the traffic source to overcome the rate-based limitations in the network. And in doing so, we achieve better control over the allocation of the network resources.

A full presentation of our User Level Multiplexing schemes for both the local case and the distributed case is out of the scope of this paper. We would like to mention that in a distributed industrial plant control system, we have more control over the application processes, thus cooperative User Level Multiplexing is a feasible choice.

#### 7. Conclusion and Future Work

This paper is the first effort in examining the ATM technology for the application of distributed industrial plant control applications. Our industrial plant control application operates in an environment in which many types of communication traffic exist, including video from industrial TV cameras, audio from plant monitors, numerical data from plant sensors, operator control commands, as well as images from plant database. Due to space limitations, we only addressed issues for hard real-time data in this paper. As part of our on-going research, we are integrating the above mentioned different types of multimedia traffic onto ATM networks.

Recently, other research has also addressed the shortcomings of the ATM VBR service with respect to its suitability for compressed video transmission. Most of this work advocates either QoS renegotiation [7, 22], or separate burst scheduling [10]. Neither of these schemes are applicable for the type of hard realtime communication dealt with in this paper.

For future research, we are continuing to address the issues of evaluating User Level Multiplexing algorithms for both local and distributed applications, and analyzing their performance and implementation implications. Besides dealing with mapping application QoS onto network services, we are also designing programming models and protocols in our middleware to support client-server type of communication in a realtime network on ATM.

Finally, it is our hope that the analysis presented in this paper can serve the distributed real-time application designers who are considering to employ ATM technology as guidelines in their system and application design.

#### 8. Acknowledgement

The author wishes to thank John Howard for helpful technical discussions during this research. The author also wishes to thank Hugh Lauer in particular for his careful review of the paper and the many excellent suggestions he made to strengthen some of the technical arguments and presentations of the paper.

### References

- [1] V. Bansal, R. Siracusa, J. Hearn, G. Ramamurthy, and D. Raychaudhurt. Adaptive QoS-Based API for ATM Networking. In Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.
- [2] A. F. T. Committee. Traffic Management Specification 4.0, January 1996.
- [3] R. Cruz. A Calculus for Network Delay, Part I: Network Elements in Isolation and Part II: Network Analysis. *IEEE Transactions on Information Theory*, 37(1):114-141, January 1991.
- [4] D. Ferrari and D. Verma. A Scheme for Real-Time Channel Establishment in Wide-Area Networks. *IEEE Journal on Selected Areas in Communications*, 8(3):368-379, April 1990.
- [5] P. Goyal, S. Lam, and H. Vin. Determining End-to-End Delay Bounds in Heterogeneous Networks. In Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.
- [6] P. Goyal, H. Vin, C. Shen, and P. Shenoy. A Reliable, Adaptive Network Protocol for Video Transport. to appear in Proceedings of the 15th IEEE INFOCOM, March 1996.
- [7] M. Grossglauser, S. Keshav, and D. Tse. The Case Against Variable Bite Rate Service. In Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.
- [8] E. W. Knightly, D. E. Wrege, J. Liebeherr, and H. Zhang. Fundamental Limits and Tradeoffs of Providing Deterministic Guarantees to VBR Video Traffic. In *Proceedings of ACM SIGMETRICS*, Ottawa, Canada 1995.
- [9] S.-K. Kweon and K. G. Shin. Traffic-Controlled Rate-Monotonic Priority Scheduling of ATM Cells. to appear in Proceedings of the 15th IEEE INFOCOM, March 1996.
- [10] S. Lam and G. Xie. Burst Scheduling Networks: Flow Specification and Performance Guarantees. In Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.
- [11] H. Lauer, A. Ghosh, and C. Shen. A General Purpose Queue Architecture for an ATM Switch. In proceedings of First Annual Conference on Telecommunica-

tions R&D in Massachusetts, volume 6, pages 17-22, October 1994.

- [12] H. Lauer, C. Shen, R. Osborne, J. Howard, Q. Zheng, M. Takegaki, H. Shimakawa, and I. Mizunuma. Digital Audio and Video in Industrial Systems . In Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.
- [13] C. L. Liu and J. W. Layland. Scheduling Algorithms for multiprogramming in a Hard-Real-Time Environment. Journal of ACM, 20(1), 1973.
- [14] A. Parekh. A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks. PhD thesis, MIT, 1992.
- [15] A. Raha, S. Kamat, and W. Zhao. Admission Control for Hard Real-Time Connections in ATM LANs. to appear in Proceedings of the 15th IEEE INFOCOM, March 1996.
- [16] J. Salehi, Z. Zhang, J. Kurose, and T. Towsley. Supporting Stored Video: Reducing Rate Variability and End-to-End Resource Requirements through Optimal Smoothing. To appear in Proceedings of ACM SIG-METRICS, 1996.
- [17] J. Turner. New Directions in Communications (or Which Way to the Information Age? IEEE Communications Magazine, 25(8):8-15, October 1986.
- [18] G. Xie and S. Lam. Delay Guarantee of Virtual Clock Server . IEEE/ACM Transactions on Networking, 3(6), December 1995.
- [19] H. Zhang and D. Ferrari. Rate-Controlled Static-Priority Queueing. In Proceedings of IEEE INFO-COM, 1993.
- [20] H. Zhang and D. Ferrari. Improving Utilization for Deterministic Service In Multimedia Communication . In Proceedings of International Conference on Multimedia Computing and Systems, May 1994.
- [21] H. Zhang and S. Keshav. Comparison of Rate-Based Service disciplines. In *Proceedings of ACM SIG-COMM*, Sept. 1991.
- [22] H. Zhang and E. Kightly. A New Approach to Support Delay-Sensitive VBR Video in Packet-Switched Networks. In Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995.
- [23] L. Zhang. VirtualClock: A New Traffic Control Algorithm for Packet Switching Networks. In *Proceedings* of ACM SIGCOMM, pages 19–29, August 1990.
- [24] Z. Zhang, Z. Liu, J. Kurose, and D. Towsley. Call Admission Control Schemes under the Generalized Processor Sharing Scheduling. To appear in Proceedings of Telecommunication Systems, May 1996.
- [25] Q. Zheng, K. Shin, and C. Shen. Real-Time Communication in ATM Network. In Proceedings of 19th Annual Local Computer Network Conference, pages 156–165, Minneapolis, Minnesota, October 1994.