Feedback Control Mechanisms for Real-Time Multipoint Video Services*

Brett J. Vickers†, Meejeong Lee‡ and Tatsuya Suda†

† Dept. of Information & Computer Science
University of California
Irvine, CA 92697-3425

‡ Dept. of Computer Science
Ewha Womans University
Seoul, Korea

To appear in the IEEE Journal on Selected Areas of Communications

Abstract

While existing research shows that reactive congestion control mechanisms are capable of providing high video quality and channel utilization for point-to-point real-time video, there has been relatively little study of the reactive congestion control of point-to-multipoint video, especially in ATM networks. Problems complicating the provision of multipoint, feedback-based real-time video service include (1) implosion of feedback returning to the source as the number of multicast destinations increases, and (2) variance in the amount of available bandwidth on different branches in the multipoint connection.

In this paper, a new service architecture is proposed for real-time multicast video, and two multipoint feedback mechanisms to support this service are introduced and studied. The mechanisms support a minimum bandwidth guarantee and the best effort support of video traffic exceeding the minimum rate. They both rely on adaptive, multi-layered coding at the video source and closed-loop feedback from the network in order to control both the high and low priority video generation rates of the video encoder. Simulation results show that the studied feedback mechanisms provide, at the minimum, a quality of video comparable to a CBR connection reserving the same amount of bandwidth. When unutilized network bandwidth becomes available, the mechanisms are capable of exploiting it to dynamically improve video quality beyond the minimum guaranteed level.

1 Introduction

Most of the early approaches to supporting the transmission of real-time video in ATM networks relied on traditional, preventive congestion control mechanisms (i.e., call admission and usage parameter control). Some services relying on preventive congestion control, like Constant Bit Rate (CBR) service, are simple to support but place a ceiling on the amount of bandwidth that can be used to support the connection, and thus, they sacrifice multiplexing gain. Furthermore, in the context of video, CBR renders the video output process constant by dynamically adjusting the quantization, or coarseness, of the video sequence, thereby causing the video quality to fluctuate, sometimes severely [1]. Other preventive congestion control mechanisms, such as those used to support Variable Bit Rate (VBR) service, require the accurate prediction of complex traffic descriptors during call admission. This is necessary to provide a statistical multiplexing gain and support the required quality of service. However, due to real-time video's often unpredictable nature, traffic descriptors are difficult, if not impossible in some cases, to predict accurately during call admission. Therefore VBR service for real-time video is often not feasible without complex online measurement and/or bandwidth renegotiation mechanisms.

Due to the problems preventive congestion control techniques have with supporting certain types of network traffic, there has recently been a shift toward the use of a combination of preventive and reactive (i.e., feedback-based) congestion control schemes. This shift occurred first in the data realm, with Available Bit Rate (ABR) service being perhaps the best example [2], but it has also occurred to some extent in the video realm, particularly with regard to live, real-time video. The notion of using feedback from the network to control the video generation rate has been studied by several researchers [3, 4, 5]. In previous work, the authors of this paper have shown that this type of rate-based feedback control mechanism reduces the cell loss in a network, allows graceful degradation in the perceptual image quality during periods of congestion, and simplifies call admission parameter estimation [6]. They also studied the interaction between reactively controlled video and ABR data, with results showing that reactively controlled video is more stable than reactively controlled ABR data.

While a number of researchers have studied feedback-based reactive congestion control mechanisms for video as mentioned above, most of the existing work examines only point-to-point video transport. Few have looked at the problem from a multipoint perspective. Although some have argued that reactive congestion control in the point-to-point case scales simply to the point-to-multipoint case,

*This research is supported in part by grants from the National Science Foundation, the University of California MICRO program, the Pacific Bell CalREN program, Hughes Aircraft, Nippon Steel Information and Communication Systems Inc. (ENICOM), Hitachi Ltd., Hitachi America, Tokyo Electric Power Company, and Nippon Telegraph and Telephone Corporation (NTT).

there are at least two reasons to believe otherwise. First, there is the problem of *feedback implosion*. As the number of destinations a source communicates with increases, so too does the volume of feedback returning to the source. The amount of feedback must therefore be reduced in an efficient and systematic manner. Second, there is the problem of *available bandwidth variation*, where different branches of a multipoint connection have differing amounts of available bandwidth on them. It is not sensible to reduce the output rate of the video source, and thereby the video's quality, to accommodate lack of available bandwidth on a single branch. This would have the effect of penalizing the perceived video quality at all destinations. Rather, the loss-tolerant nature of video should be taken into account in order to increase throughput across non-congested links while gracefully degrading video across congested links.

Unless a reactive congestion control mechanism for real-time video takes the above two issues into account, it will not adequately scale to a large number of destinations. This paper proposes and investigates a reactive congestion control architecture that utilizes an adaptive, multi-layered encoding technique and efficient multicast feedback mechanisms to provide a scalable real-time video multicast service for ATM networks. The proposed mechanisms provide a minimum guarantee of bandwidth but also allow the user to exceed this amount when the network has unutilized capacity. They require only two simple traffic descriptors at call admission time: the minimum guaranteed cell rate (MCR) and the peak cell rate (PCR).

In order to solve the problem of feedback implosion, two feedback mechanisms are suggested. In the first mechanism, called the feedback polling mechanism, the source polls each destination one at a time for congestion feedback, thereby reducing the amount of feedback returning from destinations. In the second mechanism, called the feedback coalescence mechanism, every destination responds to probe cells, but switches on the return paths merge feedback cells to reduce the amount of feedback arriving at the source.

In order to ameliorate the problem of available bandwidth variation, the proposed mechanisms utilize adaptive, multi-layered encoding combined with feedback-based rate control. The adaptive multi-layered encoding divides the real-time video stream into high and low priority streams, and the feedback mechanisms control the output rates of each of these streams to account for the congestion state of the network. The high priority cell rate is adjusted to approximate the amount of bandwidth available on the multipoint connection's most congested branch, while the low priority cell rate is adjusted to make use of other branches' unutilized bandwidth. Congested paths are allowed to drop low priority cells, thereby gracefully degrading video quality on paths where bandwidth is scarce, while noncongested paths may transmit both the high and low priority cells, achieving high quality video on paths where bandwidth is plentiful.

Other solutions to the problems of feedback implosion and available bandwidth variation have been proposed. These include probabilistic feedback [7, 8] and hierarchical acknowledgements [9]. Probabilistic feedback ameliorates feedback implosion by having each destination return a feedback signal with a probability less than or equal to one. In the context of multicast video, however, probabilistic schemes do not allow for the rapid detection of the most congested branch of a multipoint connection, and hence the high priority video rate cannot be quickly targeted according to the most congested path. Hierarchical acknowledgement prevents feedback implosion by designating selected destinations as special feedback consolidating nodes. Its primary drawback, however, is that the designated destinations may become performance bottlenecks, especially in the case of high speed video sessions where a significant amount of feedback is necessary for stability.

An existing approach to dealing with available bandwidth variation is the receiver-driven layered multicast mechanism [10]. The source generates a fixed number of video layers, and destinations "subscribe" to as many layers as they have the bandwidth to receive. One of the drawbacks of receiver-driven schemes is that they limit the destinations to choosing among the layers the source is willing to provide. The source does not adjust the amount of video it generates according to the current state of the network. The feedback mechanisms described in this paper, on the other hand, allow the source to gauge the network state and adapt the video generation rates for each video layer accordingly.

The remainder of this paper is organized as follows. Section 2 presents an overview of the MPEG video encoding process. In section 3, the proposed encoding and feedback mechanisms are described in detail. Section 4 discusses the results from an extensive simulation study in which the proposed mechanisms are analyzed. Finally, concluding remarks are given in section 5.

2 Brief Review of MPEG Standards and Related Concepts

The MPEG video encoding standards and related concepts are referred to throughout the paper and are therefore briefly reviewed in this section.

The ISO Motion Picture Experts Group (MPEG) has designed both video and audio compression algorithms for storage and transmission of full motion video and associated audio [11]. MPEG breaks video data into a number of hierarchical layers. At the highest layer is the video sequence, which is composed of a number of pictures (or frames). A frame is divided into a number of *slices*, and each slice contains some number of *macroblocks* of size 16×16 pixels. The slice size may be as small as one macroblock or as large as an entire picture. However, it must remain constant over the entire video sequence. A macroblock is then broken into a number of 8×8 blocks. Typically, the first four of these blocks cover the area of the 16×16 macroblock and describe macroblock's luminance (Y) values, and the remaining two blocks cover the chrominance, or color, of the entire 16×16 macroblock. The first chrominance block (C_r) describes one color component of the macroblock, while the second chrominance block (C_b) describes another. Because the chrominance blocks are only 8×8 pixels in size, they contain the appropriate color component of every other pixel in both the horizontal and vertical directions. The triple (Y, C_r, C_b) is the digital equivalent of the (Y,U,V) triple used for encoding NTSC video, and can easily be transformed to and from a red-green-blue (RGB) triple.

MPEG uses both intraframe coding and interframe cod-

ing techniques to achieve a high degree of compression. Intraframe coding techniques are used to exploit the spatial redundancy between adjacent pixels of a picture. The two-dimensional *Discrete Cosine Transformation* (DCT) [11] is applied independently to each block, resulting in an 8×8 block of DCT coefficients. The information representing the large spatial extent of the image is mapped into the low frequency coefficients, while information representing the finer detail is mapped into high frequency coefficients. These coefficients are then digitally quantized. Following quantization, the two dimensional DCT coefficients are rearranged (zig-zag scanned) into one dimension and sequenced with the lowest frequencies first and the highest frequencies last. With coarser quantization, the higher frequency coefficients often become zero since there is usually little variation in the image at high frequencies. This rearrangement and quantization increases the number of consecutive zero coefficients so that significant compression can be achieved by using a run-length coding technique. In run-length coding, the rearranged and quantized DCT coefficients are represented as a sequence of coefficient run-level pairs, denoted (*run*, *level*), where *run* is the number of zeros before the next non-zero coefficient, and *level* is the magnitude of the next non-zero coefficient. A unique variable length code is assigned to every possible run-level pair. In the MPEG standards, shorter variable length codes are assigned to run-level pairs that tend to occur more frequently in video coded images, thereby resulting in further compression.

In addition to intraframe coding, MPEG uses interframe coding to reduce temporal redundancy in video sequences. It does this by encoding the movements of macroblocks from frame to frame rather than re-encoding the same (or similar) macroblock repeatedly. The encoding algorithm determines if a macroblock has moved from one frame to the next and, if so, locates the direction in which the macroblock has moved. It then calculates a motion vector, which is encoded with the image. Any difference between the original macroblock and the matched area, i.e., the prediction error, is also encoded. In MPEG, this motion compensation may be bidirectional. In bidirectional temporal prediction, also called motion compensated interpolation, a macroblock may be predicted from a matching macroblock in the past reference picture (forward prediction), the future reference picture (backward prediction), or an average of the two (interpolation). For predicted macroblocks, only motion vectors and prediction error are encoded.

In MPEG video terminology, an I frame refers to a picture that is intraframe coded for every macroblock, a P frame is a picture whose macroblocks are either intraframe coded or predictively interframe coded in the forward direction, and a B frame is a picture whose macroblocks are either intraframe coded or bidirectionally interframe coded. B frames fall between I and/or P frames, and the frame sequence is encoder-dependent. For instance, a typical repeating frame generation sequence is IBBPBB. For smooth playback, the P or I frames are typically transmitted before their associated B frames.

The MPEG-2 standard enhances the MPEG encoding mechanisms described above by providing scalability extensions. MPEG-2's scalability features allow for the segmentation of video data into high and low priority video streams. These multiple streams, or *layers*,[1] are encoded in such a way that some are more tolerant to loss than others. Those layers most tolerant to loss are marked as low priority, while those that cannot withstand any loss are marked as high priority.

Two important MPEG-2 scalability extensions include data partitioning and SNR (signal-to-noise ratio) scalability. *Data partitioning* is a relatively simple method for dividing an encoded video bit stream into two layers with different priorities. As described earlier in this section, the quantized and rearranged DCT coefficients are represented as a sequence of run-level pairs. To split this stream of run-level pairs into two, a priority breakpoint is introduced. This priority breakpoint is an integer specifying the number of run-level pairs per block to place into the high priority stream. All remaining run-level pairs are placed into the low priority stream. *SNR scalability*, like data partitioning, defines a mechanism for the encoding of two layers of video, where the two layers provide different levels of picture quality for an encoded video sequence. Video signals are encoded into two bit streams called the base layer and enhancement layer. These layers correspond to the high priority layer and the low priority layer, respectively. In the base layer, the DCT coefficients are coarsely quantized to provide a minimum image quality. In the enhancement layer, the residual DCT coefficients resulting from the coarse quantization of the base layer are then requantized using a second, finer quantizer. In this paper, we consider only data partitioning due to its simplicity.

3 Rate Control and Encoding Mechanisms

3.1 Rate Control Feedback Mechanisms

Closed loop feedback is utilized to control both the overall cell rate and the high priority cell rate of the video source. The source periodically transmits a forward feedback cell (or, in ATM terminology, a resource management cell). When the destination receives the forward feedback cell, it returns a backward feedback cell to the source. Switches or destinations may mark feedback cells to indicate congestion on the source-to-destination path. In this paper, we assume a closed loop feedback mechanism that uses binary congestion indication. In other words, feedback is looped between the source and the destinations, and a single bit is used to express the congestion status of the network.

Two feedback mechanisms, both using closed loop feedback and binary congestion indication, are studied. The first mechanism, called *feedback polling*, uses feedback cells to selectively probe source-to-destination paths for congestion. This mechanism is simple and requires no modification to existing ATM switches. The second mechanism, called *feedback coalescence*, probes all source-to-destination paths simultaneously but requires switches to merge feedback cells on their return paths. Although this mechanism requires a feedback cell-merging functionality in ATM switches, it is more scalable than feedback polling. Both mechanisms mark forward and backward feedback cells as high priority to prevent their loss.

Both feedback mechanisms provide guaranteed service to all video traffic under the user-specified minimum cell

[1] Priority layers are not to be confused with the hierarchical layers of MPEG encoding.

rate (MCR). However, they also provide best effort service for any video traffic exceeding the minimum cell rate up to the peak cell rate (PCR). Restated, the feedback mechanisms support dynamic allocation of the network's unutilized bandwidth, thereby improving network bandwidth utilization and video quality.

The goal of these two feedback mechanisms is to adjust the high priority cell rate according to the state of the *most* congested branch in the multipoint connection while allowing the overall cell rate to increase and make use of unutilized bandwidth. This goal is chosen for three reasons. First, by setting the high priority cell rate according to the status of the most congested branch in the connection, the loss of high priority cells in the network is minimized. This is important, because high priority cell loss has a much more deleterious effect on video quality than low priority cell loss. Second, by raising the high priority cell rate to the maximum level allowed by the most congested link in the connection, the impact of low priority cell loss on other links is minimized. If there are fewer low priority cells being generated, there are fewer candidates for loss during periods of congestion. Any low priority cell loss that *does* occur is limited to the least important parts of the image. Third, by allowing the overall cell rate to exceed the rate allowed by the most congested branch of the connection, throughput and video quality to those destinations belonging to less congested multicast branches are improved. Thus, video quality is not determined by the "weakest link" in the connection.

3.1.1 The Feedback Polling Mechanism

Feedback polling is a relatively simple rate control mechanism capable of determining the congestion state of a multipoint connection without requiring modification to existing ATM switches.

For each N_{rm} video cells transmitted, the video source transmits a single forward feedback cell. To provide rapid notification of congestion, N_{rm} should be a relatively small integer, but not so small that the overhead of generating and transmitting feedback cells is severe. Like ABR service, we recommend a value of N_{rm} equal to 16, resulting in a 5.9% feedback overhead. With each feedback cell, the source simultaneously probes two destinations for different types of information. The first type of probe is used to locate the most congested link in the multipoint connection and adjust the high priority cell rate accordingly. The second type of probe observes the congestion state of the entire multipoint connection and uses this information to adjust the overall cell rate.

Each destination is assigned a unique identification that the video source uses while polling. When the source generates a forward feedback cell, it indicates the targets for both probes. Because the feedback cells are multicast, every destination participating in the point-to-multipoint video call receives the probe. However, only the two destinations being polled respond with a backward feedback cell. Furthermore, the information contained in the backward feedback cell depends on the type of probe the destination is receiving. Since each forward feedback cell contains two probe targets, transmission of a forward feedback cell results in the return of two backward feedback cells.

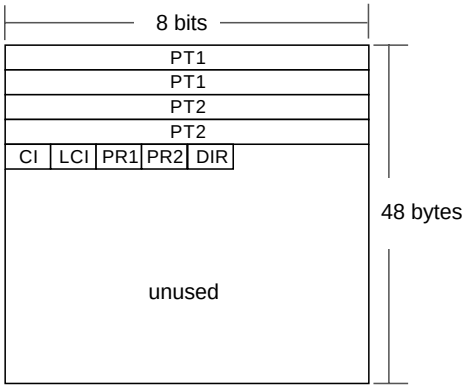

Figure 1: Feedback cell payload for the feedback polling mechanism

Fig. 1 shows the format of the feedback cells. This format is used for both forward and backward feedback cells and is based loosely on the binary congestion indication provisions of the Enhanced Proportional Rate Control Algorithm (EPRCA) devised by the ATM Forum for ABR service [2]. The fields contained in the cell are as follows:

Probe#1 Target (PT1) This 16-bit field contains the identifier of the destination that should respond with a backward feedback cell for a probe of type 1. As mentioned earlier, the first type of probe is used to locate the most congested branch in the multipoint connection.

Probe#2 Target (PT2) This 16-bit field contains the identifier of the destination that should respond with a backward feedback cell for a probe of type 2. The second type of probe is used to gauge the overall congestion state of the multipoint connection.

Congestion Indication (CI) This one-bit field indicates whether the source-to-destination path has experienced congestion and is used only in backward feedback cells. The destination, which examines the EFCI bits of incoming cells, sets this field if forward congestion is indicated. Optionally, switches may also mark this field to indicate congestion on the connection in the source-to-destination direction.

Low-Priority Congestion Indication (LCI) This one-bit field indicates whether the source-to-destination path is in danger of losing a high priority cell. The destination may mark this bit if it perceives itself as not receiving enough low priority cells. The switch does not participate in the marking of this bit. This field is only used in backward feedback cells.

Probe #1 Response (PR1) This one-bit field is set by the destination if it is responding to a probe of the first type.

Probe #2 Response (PR2) This one-bit field is set by the destination if it is responding a probe of the second type.

Feedback Direction (DIR) This one-bit field indicates whether the feedback cell is in the forward (DIR=0) or backward (DIR=1) direction.

The feedback polling mechanism uses two probes, and hence it requires two algorithms — one to locate the most congested branch in the multicast tree and one to observe the overall congestion state of the tree. In the following, we describe these two algorithms.

Algorithm for locating the most congested branch of the multicast tree

The algorithm listed in Fig. 2 locates the most congested branch in the multicast tree and adjusts the video source's high priority cell rate accordingly.

First, in order to explain this algorithm, we introduce the notion of "low priority traffic congestion." Low priority traffic congestion is indicated by a destination when the fraction of low priority cells being received becomes very small, i.e., when the loss of a high priority cell is imminent. A destination sets the low priority traffic congestion indication to true when the fraction of low priority cells it is receiving falls below a given threshold, T_L. For example, consider the case where a switch on the path from the source to the destination is congested and is dropping a large number of low priority cells. Due to the dropping of low priority cells, the destination may receive only a small fraction, say 5%, of its cells as low priority cells. If T_L were set to 0.1 (or 10%), the destination would indicate low priority traffic congestion in the next backward feedback cell, resulting in a high priority cell rate reduction at the source. The effect of the threshold is to predict when the loss of a high priority cell is imminent and to request that the source reduce its high priority cell rate in order to prevent such a loss. The threshold value determines the range of accuracy within which the destination may predict the loss of a high priority cell.

We now discuss the algorithm the video source uses to locate the most congested branch in the connection. The algorithm has two states: the SEARCH state, in which the source actively searches for the most congested branch in the connection, and the PROBE state, in which the source continually probes a congested branch while simultaneously reducing its high priority cell rate. In the SEARCH state, the source probes each destination, one at a time. Whenever a forward feedback cell arrives at a destination, the probe target #1 (PT1) field is examined. If the identifier contained in the PT1 field matches the destination's identification, then the destination examines its low priority traffic congestion status, generates a backward feedback cell, and sets the LCI bit accordingly. When the source receives a backward feedback cell with LCI=1, it enters the PROBE state. While in the PROBE state, the source continually probes the destination that generated the LCI indication. Each time an LCI indication is received from the destination, the source decrements its high priority cell rate by an amount proportional to its current value. The source returns to the SEARCH state only after the destination being probed returns a feedback cell with LCI=0. If during the SEARCH state the source completes an entire cycle of probes without receiving a low priority traffic congestion indication, the source increases its high priority cell rate additively and begins another search cycle.

This algorithm has the effect of reducing the high priority cell rate rapidly whenever a high priority cell is in danger of being lost. Further, it increases the high priority cell rate only after every destination has indicated a lack of low priority traffic congestion. It is therefore a conservative algorithm. The algorithm's goal is to produce a high priority cell rate that approximates the available bandwidth on the multipoint connection's most congested link.

Algorithm for observing the overall congestion state of the multicast tree

In addition to controlling the high priority cell rate, the feedback polling algorithm also controls the overall cell rate. It does this through a second probe algorithm that operates independently from the first algorithm. This algorithm is listed in Fig. 3. It observes the overall congestion state of the multipoint connection and adjusts the overall cell rate accordingly.

First, we explain the behavior of the switch. Switches are assumed to determine congestion according to a simple two-threshold scheme: when an output port's buffer occupancy exceeds an upper threshold, the switch considers the buffer congested until its occupancy drops below a lower threshold. When a buffer is congested, the switch marks the EFCI (Explicit Forward Congestion Indication) bit in the header of the cell being served to indicate congestion to the destination. Upon receiving the cell, the destination examines the EFCI bit and stores the congestion state of the forward path. When the destination is probed, it generates a backward feedback cell and marks the feedback cell's Congestion Indication (CI) bit if the last cell the destination received had its EFCI bit marked.

Switches on the path back to the source may also set the congestion indication (CI) bit of the backward feedback cell. This can be done if switches implement a "look-back" congestion detection functionality, where the switch, upon serving a backward feedback cell from the destination to the source, observes the congestion status of the buffer in the source-to-destination direction [2]. If the buffer has become congested, the switch marks the backward feedback cell's congestion indication field. This functionality improves the responsiveness of network feedback by shortening the amount of time necessary for negative congestion feedback to arrive. It is, however, optional; switches are not required to implement it to achieve a working feedback polling mechanism.

Upon receiving the backward feedback cell, the source stores the cell's congestion indication (CI) bit in a congestion history. The congestion history maintains the last N congestion indications received by the source, where N is the total number of destinations in the multipoint connection. From the congestion history, the source calculates a target overall cell rate (TCR) as follows:

$$\text{TCR} = \text{MCR} + \frac{N_0}{N} \times (\text{PCR} - \text{MCR}) \qquad (1)$$

where N_0 is the number of zeroes in the congestion history, and MCR and PCR are the minimum and peak cell rates of the connection. If the current overall cell rate is less than

Source Pseudo-Code

```
Initialization

  HCR := MCR
  LCI_count := N
  state := SEARCH
  last_probe := 0

While VC_on_line do

  if (receive RM(DIR=backward, PR1=TRUE, PT1, LCI))
    if (state = PROBE)
      if (PT1 = last_probe)
        if (LCI = TRUE)
          HCR := HCR - (HCR * Nrm)/RDF
          HCR := MAX(HCR, MCR)
        else
          LCI_count := N
          state := SEARCH
        endif
      endif
    else if (state = SEARCH)
      if (LCI = TRUE)
        HCR := HCR - (HCR * Nrm)/RDF
        HCR := MAX(HCR, MCR)
        last_probe := PT1
        state := PROBE
      else
        last_probe := (last_probe + 1) mod N
        LCI_count := LCI_count - 1
        if (LCI_count = 0)
          HCR := HCR + AIR*Nrm
          HCR := MIN(HCR, PCR)
          LCI_count := N
        endif
      endif
    endif
  endif

end While
```

Parameters and Variables

N	Number of destinations
MCR	Minimum Cell Rate
PCR	Peak Cell Rate
RDF	Rate Decrease Factor
AIR	Additive Increase Rate
Nrm	One RM cell for every Nrm data cells
HCR	Allowed High Priority Cell Rate
LCI	Low-Priority Traffic Congestion Indicator
DIR	The feedback cell direction
PT1	Feedback cell's probe target ID
PR1	Probe #1 response indicator
state	Algorithm state: PROBE or SEARCH
last_probe	ID of last target probed
LCI_count	the number of destinations left in the search cycle
LT	Lower Threshold (0 < LT < 1)

Destination Pseudo-Code

```
While VC_on_line do

  if (receive_data_cell)
    count := count + 1
    if data_cell is low priority
      lo_count := lo_count + 1
    endif
  endif

  if (receive RM(DIR=forward, PT1=my_ID))
    if (lo_cell_count < cell_count * LT)
      send RM(DIR:=backward, LCI:=TRUE,
              PT1:=ID, PR1:=TRUE)
    else
      send RM(DIR:=backward, LCI:=FALSE,
              PT1:=ID, PR1:=TRUE)
    endif
    lo_cell_count := cell_count := 0
  endif

end While
```

Figure 2: Feedback polling mechanism: algorithm for locating the most congested branch

the target cell rate, it is increased additively. If, on the other hand, the current overall cell rate is greater than the target cell rate, it is decreased proportionally.

This heuristic has the effect of controlling the overall cell rate such that some paths are overloaded with more low priority cells than they can handle, but others are provided with enough video traffic to make use of unutilized bandwidth. The target cell rate is determined by the degree of congestion on all the source-to-destination paths. At one extreme, when all source-to-destination paths are congested, the target rate equals the minimum cell rate. At the other extreme, when all source-to-destination paths are uncongested, the target rate equals the peak cell rate. For congestion levels falling between these two extremes, the target cell rate is set at an intermediate level corresponding linearly to the degree of congestion in the multicast connection.

It should be noted that Eq. (1) is not the only heuristic that could have been used to calculate the target rate. It is, however, a logical one, because it linearly maps the target rate to the degree of congestion in the connection. The loss of a single video cell in the network results in the loss of an entire block or slice of video data. Hence, dropping video cells in the network has a more deleterious effect on video quality than a graceful reduction in the encoding quantization at the source. It is therefore desirable to reduce the overall cell rate as the number of congested paths grows, rather than simply allowing more paths to drop cells. A target cell rate that decreases linearly with the number of congested paths achieves this desirable result. See the appendix for a comparison of this heuristic with two other possible TCR heuristics.

3.1.2 The Feedback Coalescence Mechanism

While the feedback polling mechanism has the advantage that it requires no special modification to ATM switches, its scalability may be compromised as the number of destinations, and thereby the length of the probe cycle, grows. For this reason, we have also studied a more scalable mechanism, called the feedback coalescence mechanism, that captures the state of the entire multicast tree with each backward feedback cell. However, it requires that switches be capable of merging backward feedback cells on the return paths to the sources. While this functionality does not exist in many existing ATM switches, it would be relatively simple to implement. Further, the ATM Forum has proposed feedback consolidation for its ABR service [2], and

Source Pseudo-Code

```
Initialization

  ACR := MCR
  count := next_dest := 0

While VC_on_line do

  if (now >= next_cell_time)
    if (cell_to_send)
      send data_cell with EFCI=0
      if (count = 0)
        send RM(DIR:=forward, PT2:=next_dest,
                CI:=0)
        next_dest := (next_dest + 1) mod N
      endif
      count := (count + 1) mod Nrm
    endif
    next_cell_time := next_cell_time + 1/ACR
  endif

  if (receive RM(DIR=backward, CI))
    update CI_history(CI)
    Nzero := number of zeroes in CI_history
    TCR := MCR + (Nzero/N) * (PCR-MCR)
    if (ACR < TCR)
      ACR := ACR + Nrm*AIR
      ACR := MIN(ACR, PCR)
    else
      ACR := ACR - (ACR*Nrm)/RDF
      ACR := MAX(ACR, MCR)
    endif
  endif

end While
```

Algorithm Parameters and Variables

N	Number of destinations
MCR	Minimum Cell Rate
PCR	Peak Cell Rate
RDF	Rate Decrease Factor
AIR	Additive Increase Rate
Nrm	One RM cell for every Nrm data cells
ACR	Allowed Overall Cell Rate
TCR	Target Cell Rate
CI	Traffic Congestion Indicator
CI_history	A record of the last N CI bits
DIR	The feedback cell's direction
PT2	Feedback cell's probe target ID
PR2	Probe #2 response indicator
next_dest	ID of the next target to probe
count	cell transmission counter
VC_CI	Congestion state of connection

Destination Pseudo-Code

```
While VC_on_line do

  if (receive_data_cell)
    VC_CI = EFCI state of cell
  endif

  if (receive RM(DIR=forward, CI))
    if (VC_CI = 1)
      send RM(DIR:=backward, CI:=TRUE, PR2:=TRUE)
    else
      send RM(DIR:=backward, CI:=FALSE,
              PR2:= TRUE)
    endif
  endif

end While
```

Figure 3: Feedback polling mechanism: algorithm for observing the overall congestion state of the multicast tree

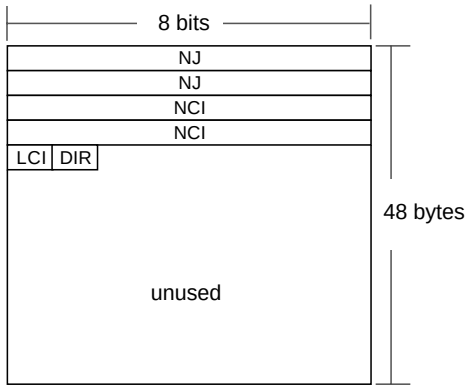

Figure 4: Feedback cell payload for the feedback coalescence mechanism

thus, such functionality in switches may become common in the near future.

As shown in Fig. 4, the feedback coalescence mechanism introduces several new fields to the feedback cell payload. They are as follows:

Number of Joined Cells (NJ) This 16-bit field indicates the number of backward feedback cells whose information has been merged into this cell.

Number of Congestion Indications (NCI) This 16-bit field specifies the number of congestion indications that have been merged into this cell.

Low Priority Traffic Congestion Indicator (LCI) This one-bit field indicates whether any of the destinations whose backward feedback cells have been merged into this cell are experiencing low priority traffic congestion.

Feedback Direction (DIR) This one-bit field indicates whether the feedback cell is in the forward (DIR=0) or backward (DIR=1) direction.

With the exception of the DIR field, these fields are used only in backward feedback cells.

The algorithm for the feedback coalescence mechanism is relatively straightforward and is listed in Fig. 5. As in the feedback polling mechanism, the source multicasts a forward feedback cell to all the destinations for every N_{rm} video cells transmitted. When a destination receives the forward feedback cell, it analyzes the path's congestion state, generates a backward feedback cell, and sends it back to the source. The destination analyzes the congestion state of the path in the same way as it does for the feedback polling mechanism. If the destination considers the source-to-destination path congested (i.e., if the last arriving cell had its EFCI bit set), then it sets the backward feedback cell's NCI field to 1. If the destination observes low priority traffic congestion, it sets the backward feedback cell's LCI bit. Finally, the destination sets the NJ field to 1 to indicate that the cell contains feedback information from only one destination thus far.

Source Pseudo-Code

```
Initialization

  ACR := HCR := MCR
  count := 0

While VC_on_line do

  if (now >= next_cell_time)
    if (cell_to_send)
      send data_cell with EFCI=0
      if (count = 0)
        send RM(DIR:=forward)
      endif
      count := (count + 1) mod Nrm
    endif
    next_cell_time := next_cell_time + 1/ACR
  endif

  if (receive RM(DIR=backward, NJ, NCI, LCI))
    TCR := MCR + (NJ-NCI)/NJ * (PCR-MCR)
    if (ACR > TCR)
      ACR := ACR - (ACR*Nrm) / RDF
      ACR := MAX(ACR, MCR)
    else
      ACR := ACR + AIR*Nrm
      ACR := MIN(ACR, PCR)
    endif
    if(LCI = TRUE)
      HCR := HCR - (HCR*Nrm) / RDF
      HCR := MAX(HCR, ACR)
    else
      HCR := HCR + AIR*Nrm
      HCR := MIN(HCR, PCR)
    endif
  endif

end While
```

Parameters and Variables

MCR	Minimum Cell Rate
PCR	Peak Cell Rate
RDF	Rate Decrease Factor
AIR	Additive Increase Rate
Nrm	One RM cell for every Nrm data cells
ACR	Allowed Overall Cell Rate
HCR	Allowed High Priority Cell Rate
TCR	Target Overall Cell Rate
LCI	Low-Priority Traffic Congestion
NJ	# of feedback cells merged
NCI	# of congestion indications merged
CI	Traffic Congestion Indicator
DIR	The feedback cell's direction
count	cell transmission counter
VC_CI	Congestion state of connection

Destination Pseudo-Code

```
While VC_on_line do

  if (receive_data_cell)
    VC_CI = EFCI state of cell
    cell_count := cell_count + 1
    if data_cell is low priority
      lo_cell_count := lo_cell_count + 1
    endif
  endif

  if (receive RM(DIR=forward))
    if (lo_cell_count < cell_count * LT)
      LCI := TRUE
    else
      LCI := FALSE
    endif
    lo_cell_count := cell_count := 0
    NCI := VC_CI
    send RM(DIR:=backward, NJ:=1, NCI, LCI)
  endif

end While
```

Figure 5: Feedback coalescence mechanism: algorithm

Intermediate switches merge backward feedback cells in the following manner. We assume the switch contains a VC table of some sort from which the switch can determine the number of backward feedback cells it must merge into one cell for a given multipoint connection. Before transmitting a merged backward feedback cell to the next hop, the switch waits for one of two conditions to become true: (1) at least one backward feedback cell has been received from each subtree of the multicast tree, or (2) a feedback merging timer for the connection expires. If more than one backward feedback cell arrives on a single input port before either of these conditions becomes true, the old feedback information from that port is discarded in favor of the information contained in the newly arriving feedback cell. Fig. 6(a) illustrates the first case, where backward feedback cells from all three output ports have arrived before the merging timer for the connection expires. In this case, all three backward feedback cells are successfully merged into a single cell and sent to the next hop. Fig. 6(b) illustrates second case, where no backward feedback cell has arrived from the center port before the timer expires. Furthermore, a more recent feedback cell has arrived from the upper port. In this case, the most recent feedback cells from the top and bottom ports are merged into a single cell and sent to the next hop when the timer expires. The advantage of this simple merging scheme is that it requires no synchronization of backward feedback cells returning from downstream destinations.

The switch merges the information in the backward feedback cells according to the following policy:

$$\begin{aligned} \mathrm{NJ}' &= \sum_{i=1}^{n} \mathrm{NJ}_i \\ \mathrm{NCI}' &= \sum_{i=1}^{n} \mathrm{NCI}_i \\ \mathrm{LCI}' &= 1 \text{ iff } \mathrm{LCI}_i = 1 \text{ for any } i = 1, \ldots, n \end{aligned}$$

where n is the number of backward feedback cells being joined by the switch, primed variables indicate the contents of the outgoing feedback cell, and subscripted variables indicate the contents of the incoming backward feedback cells. Under this merging policy, the NJ field counts the number of feedback cells that have been joined, the NCI field maintains the number of congestion indications that have been joined, and the LCI bit indicates whether any destinations have experienced low priority congestion.

Intermediate switches may also implement the look-back congestion detection functionality. However, instead of marking a single bit in the payload of the feedback cell as the feedback polling mechanism does, the switch sets the incoming feedback cell's NCI field equal to its NJ field if the port from which it arrived has become congested. As in the case of feedback polling, the look-back congestion detection functionality is optional for feedback coalescence.

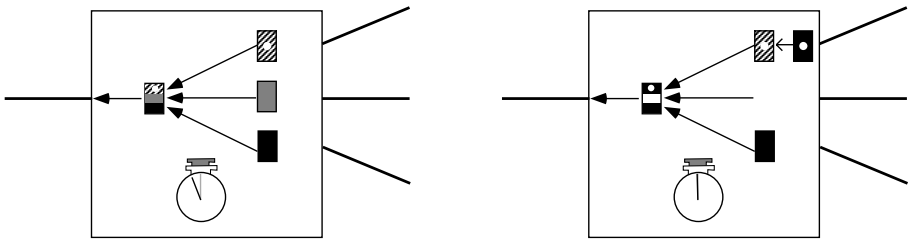

Figure 6: Feedback coalescence. (a) successful merging of feedback cells before timer expires, (b) merging of available feedback cells when timer expires.

Upon receiving the backward feedback cell, the source adjusts its high priority and overall cell rates accordingly. To calculate its high priority cell rate, the source simply examines the cell's LCI bit, and if it is set, decreases the high priority cell rate proportionally. Otherwise, the source increases its high priority cell rate additively. To determine its overall cell rate, the source first calculates a target cell rate based on the fraction of congested source-to-destination paths indicated by the feedback cell. The target cell rate TCR is calculated as follows:

$$\text{TCR} = \text{MCR} + \frac{\text{NJ} - \text{NCI}}{\text{NJ}} \times (\text{PCR} - \text{MCR}) \qquad (2)$$

where MCR and PCR are the minimum and peak cell rates for the connection. This heuristic sets the target cell rate according to the current degree of congestion being experienced by the connection. For example, if 50% of the source-to-destination paths are congested, the target rate is set to the midpoint between the minimum and peak cell rates. The source uses the target rate to decide whether to increase or decrease its current overall cell rate. If the source's current overall cell rate is less than the target cell rate, an additive increase on the rate is performed. If it is larger than the target cell rate, a proportional decrease is performed.

3.2 Adaptive Video Encoding

In order to meet the target high priority and overall cell rates, an adaptive, multi-layered encoding technique is adopted. Control of the overall output rate of the video encoder requires adjustment of the encoder's quantization (or coarseness) parameter. The video buffer is filled by the encoder and served at an output rate equal to the overall cell rate. As the video buffer's occupancy changes, the quantization parameter (Q) of the video encoder is adjusted to prevent buffer overflow and underflow. When the buffer occupancy grows large (usually due to a reduction of the controller's output rate), Q is increased, thereby reducing the total generation rate of video data. This has the effect of reducing the overall video quality, but in a graceful manner. Conversely, if the buffer remains largely empty (usually due to an increase in the rate controller's output rate), Q is decreased, thereby increasing the combined output rate and the overall image quality.

Quantization parameter adjustment achieves control of the total output rate from the video encoder. In addition, data partitioning (as described in section II) is utilized to control the output rates of high and low priority information. As each macroblock is encoded and packed into ATM cells, the priority breakpoint is adjusted to produce the desired high priority cell rate.

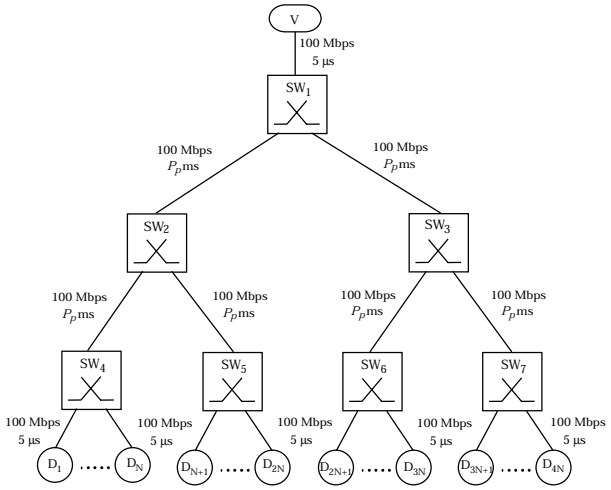

Figure 7: Simulation model

4 Simulation Model and Numerical Results

4.1 Simulation Model

In order to study the interaction between the feedback and encoding mechanisms described in the previous section, the network model shown in Fig. 7 was devised and a number of simulations performed.

Network Model

The simulation model consists of seven ATM switches, SW_1 through SW_7, interconnected in a binary tree topology. See Fig. 7. The simulated video source of the multicast connection is attached to the root switch SW_1, while all destinations participating in the multicast connection are located two switch hops from the source and are equally distributed among the leaf switches, SW_4 through SW_7 (N destinations per switch).

Links between adjacent ATM switches are assigned a propagation delay of P_p, a value which simulations may vary between 10 μs and 10 ms. Propagation delays of 10

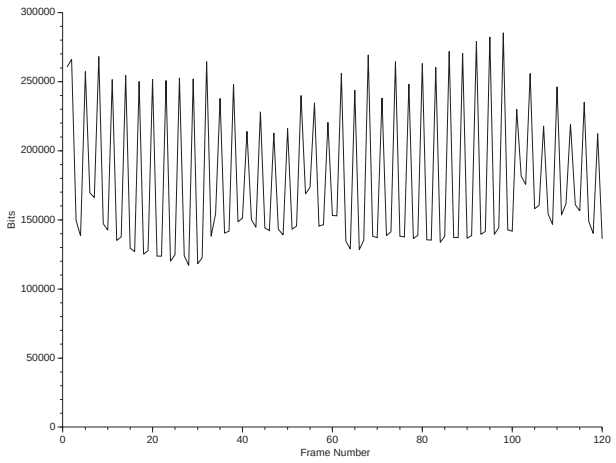

Figure 8: Bits per frame in a sequence of the motion picture *The Return of the Jedi* at MPEG quantization step Q=1

μsec and 10 msec correspond roughly to distances of 2 km and 2000 km, respectively, and these values are chosen to allow for investigation of the proposed video service in typical LAN and WAN environments. The links connecting the source and destinations to switches are assigned propagation delays of 5 μs. Each link's capacity is set to 100 Mbps. No assumptions about the loss of feedback cells have been made; a feedback cell is just as likely as any other high priority cell to be lost by the network.

Video Traffic Characteristics

Thirty seconds (720 frames) of video from the motion picture, *The Return of the Jedi*, are used in all simulations. The video data is stored in raw format and encoded using the rate control and encoding mechanisms described in the previous section. The encoder used in our simulations is a software-based MPEG-1 encoder [12] with modifications to achieve multi-layered data partitioning. Fig. 8 shows the number of bits per frame for the test video sequence when encoded at the optimal quantization level (Q=1) with a repeating frame sequence of IBBPBB. At the standard motion picture frame rate of 24 Hz, this video sequence is capable of generating a mean bit rate of 6.67 Mbps at optimal quantization.

Interfering Traffic

In order to simulate a realistic network environment, interfering traffic is generated and inserted at each switch. Interfering packets are generated according to a Poisson distribution with interarrival rate $\lambda = 6.5 \times 10^{-5}$. Packet lengths are geometrically distributed with an average length of 8 cells, resulting in an average interfering load of $\rho = 0.5$ (or 50 Mbps) on all links in the network.

In some simulations, several links are selected as bottleneck links. Two types of bottleneck traffic are employed: steady or oscillating. Packets for the steady bottleneck traffic are generated according to a Poisson distribution with interarrival rate $\lambda = 3.3 \times 10^{-5}$, and packet lengths are geometrically distributed with an average length of 8 cells. This results in a total load of $\rho = 0.98$ on the selected bottleneck links, leaving approximately 2 Mbps available, on average, for video traffic.

For simulations requiring oscillating interfering traffic on bottleneck links, two alternating Poisson processes with differing interarrival rates are used to generate packets. One process with an interarrival rate $\lambda_1 = 3.3 \times 10^{-5}$ alternates with another whose interarrival rate is $\lambda_2 = 6.5 \times 10^{-5}$. The period of oscillation is set at 600 ms, and packet lengths are again distributed geometrically with an average length of 8 cells. The result of this type of interfering bottleneck traffic is an interfering load that oscillates between $\rho = 0.5$ and $\rho = 0.98$. The primary reason for introducing oscillating interfering traffic is to study the stability and reaction times of feedback mechanisms, particularly as a function of propagation delay.

Source and Switch Buffers

ATM switches are assumed to be non-blocking and output-buffered. Each output port is allocated 200 cells of buffer space, none of which is shared by other ports. Buffers are served on a first-come first-serve basis. The video encoder is also equipped with a transmission buffer of 200 cells. These buffer sizes result in a worst case queueing delay of $\left(\frac{200\text{cells}}{\text{MCR}}\right) + 3 \times \left(\frac{200\text{cells}}{100\text{Mbps}}\right)$ for the given topology. The first term is the worst-case queueing delay imposed by the 200-cell video buffer, and the second term is the queueing delay imposed by buffers in the three switch hops between the source and the destination. For example, with a minimum cell rate of 1500 cps, the worst case queueing delay is 135.8 ms. Switches are assumed capable of performing the look-back congestion detection functionality.

A first-come first-serve service discipline is used primarily because of its simplicity, but also because it provides a worst case performance estimate of the mechanisms under investigation. More sophisticated service disciplines, such as fair queueing [13], earliest date due [14] and virtual clock [15], can undoubtedly provide improved performance and quality of service guarantees.

Rate Control Parameters

In simulations, the video source is given a guaranteed minimum cell rate (MCR) of 1500 cells per second (approx. 0.6 Mbps) and a peak cell rate (PCR) of 20000 cells per second (approx. 8.5 Mbps). These values are chosen to simulate the case in which a relatively small amount of bandwidth is reserved, but the utilized bandwidth is able to grow significantly when more network bandwidth becomes available. One forward feedback cell is generated for every 16 video cells transmitted (N_{rm}=16). A loss threshold of $T_L = 0.1$ is used to determine low priority traffic congestion. No assumptions about the loss of feedback cells have been imposed; a feedback cell is just as likely as any other high priority cell to be lost.

Both algorithms have two rate adjustment parameters, the additive increase rate (AIR) and the rate decrease factor (RDF), that control the sizes of increases and decreases in the high priority and overall rates of transmitted video. The AIR is expressed in units of cells per second and determines the size of rate increases. When a rate increase occurs, the AIR is first multiplied by N_{rm} and then added to the current rate. In the simulations, AIR is set conservatively to 1 cell per second. The second rate adjustment parameter (RDF) is a unitless constant that determines the magnitude of a rate decrease. When a rate decrease occurs, the current rate is divided by the RDF. The result is

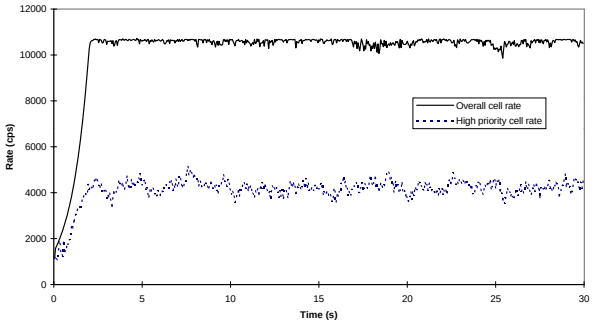

Figure 9: Overall and high priority cell rates for the feedback coalescence mechanism (100 destinations, 50% of paths congested)

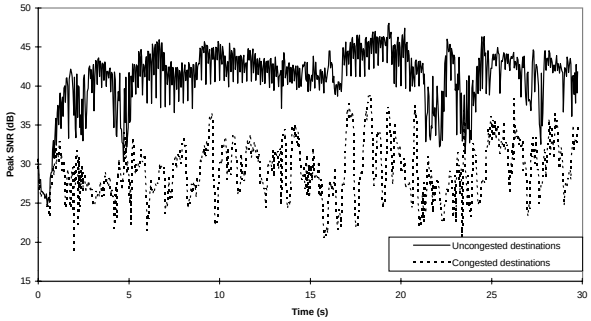

Figure 10: Peak SNR of video received at destinations for the feedback coalescence mechanism (100 destinations, 50% of paths congested)

then multiplied by N_{rm} and subtracted from the current rate, resulting in a rate decrease that is proportional to the current rate. In the presented simulation study, RDF is set to 1024.

4.2 Numerical Results

A series of simulations are run using various numbers of destinations, inter-switch propagation delays, and degrees of network congestion in order to obtain performance measures such as the source's average overall cell rate, its average high priority cell rate, the video cell discarding rate, and the signal-to-noise ratio (SNR) of the video received by destinations.

4.2.1 Transient Behavior

To examine the performance of the feedback mechanisms, we first study the transient behavior of a system using the feedback coalescence mechanism. (The transient behavior of the feedback polling mechanism, because it is qualitatively similar to that for the feedback coalescence mechanism, is not presented here.) For this simulation, we consider a network with 100 destinations (i.e., $N = 25$). Half the source-to-destination paths are congested by an steady interfering Poisson packet load of $\rho = 0.98$ on the link between SW_1 and SW_2. The propagation delay between switches is set at 1 ms.

Fig. 9 illustrates the overall and high priority cell rates achieved during a 30 second simulation of the feedback coalescence mechanism. As the figure shows, the feedback coalescence algorithm is capable of providing stable overall and high priority cell rates, even in the face of bursty Poisson interfering traffic. Note also that the source transmits video at an overall rate exceeding 10000 cells per second, thereby exploiting some of the unutilized bandwidth on uncongested links. The high priority cell rate is controlled by congestion on the multipoint connection's most congested link and therefore remains close to the average available bandwidth on the bottleneck link, namely 2 Mbps (or approximately 5000 cps). However, even the high priority rate exceeds the minimum guaranteed cell rate of 1500 cps. While the multipoint connection using the feedback coalescence mechanism reserves the same amount of network bandwidth as a 1500 cps CBR connection, the adaptive feedback and rate control mechanisms allow the feedback coalescence mechanism to exploit available bandwidth, unlike CBR.

In Fig. 10, a measure of the quality of video received by the destinations is illustrated. One curve plots the peak signal-to-noise ratio (SNR) of the video received by uncongested destinations, while the other plots the peak SNR for video received by a congested destinations. (As a shorthand, "uncongested destinations" refer to destinations terminating uncongested paths, while "congested destinations" refer to destinations terminating congested paths.) The peak SNR is a measure of the difference between the original image and the received image. For a given video frame, the peak SNR is calculated according to the following equation:

$$\text{Peak SNR} = 10 \times \log_{10}\left(\frac{255^2 \times n}{\sum_{i=1}^{n}(P_i^r - P_i^c)^2}\right)$$

where n is the number of pixels in the frame, P_i^r is the 8-bit luminosity value of the i-th pixel in the reference frame, and P_i^c is the 8-bit luminosity value of the i-th pixel in the frame being compared to the reference frame. To obtain the graph in Fig. 10, the following process was followed. First, the video sequence was reconstructed at each destination. Low and high priority cell losses were taken into account, and no error resilience algorithms were performed to conceal the effects of high priority cell loss. Then, the peak signal-to-noise ratio was calculated for each reconstructed frame and plotted versus time. As the figure shows, congested destinations receive a video sequence whose quality is worse than that received by uncongested destinations. This is due primarily to the loss of low priority cells on the bottleneck link, but is also partially due to occasional high priority cell loss. The oscillations in the upper curve are a result of MPEG's varying frame types; B and P frames typically have a lower SNR than do I frames.

Fig. 11 illustrates the cell discarding rate for the transmitted video on the bottleneck link as a function of time. It generally falls between 10% and 35%, with fluctuations being determined largely by the variation in MPEG frame sizes. The losses shown in this figure are almost entirely

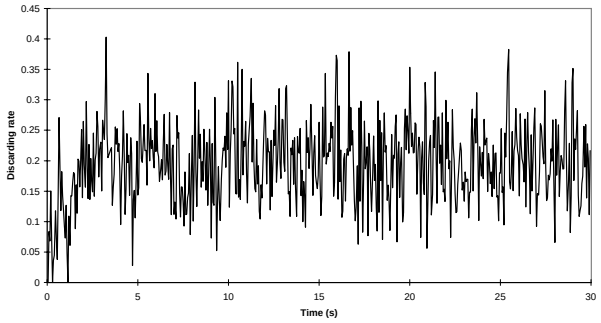

Figure 11: Cell discarding rate for video on the bottleneck link for the feedback coalescence mechanism (100 destinations, 50% of paths congested)

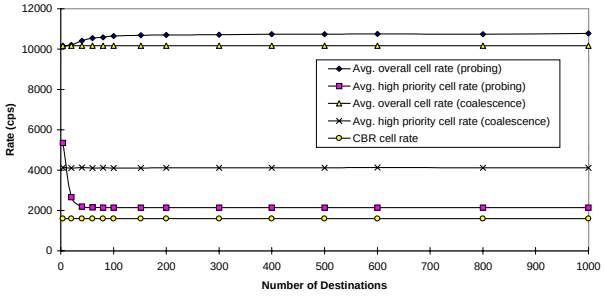

Figure 14: Average video rate vs. destination count for feedback polling and feedback coalescence (50% of paths congested by steady bottleneck load)

low priority cell losses. These losses determine the quality of video that congested destinations receive.

Figs. 12 and 13 show a sample frame from the video sequence as received by uncongested and congested destinations, respectively. A partial blurring effect occurs in the congested destination's image due to the loss of low priority cells, but the video quality is subjectively tolerable, especially when the frames are being displayed at a rate of 24 per second.

4.2.2 Scalability with Destination Count

It is important to understand how the feedback polling and feedback coalescence mechanisms scale with the number of destinations in the multipoint connection. To study this, the number of destinations in the multicast connection is varied while holding all other parameters fixed. Half of the source-to-destination paths are again congested by introducing a steady, high priority interfering Poisson load of $\rho = 0.98$ on the link between SW_1 and SW_2 (see Fig. 7). A switch-to-switch propagation delay (P_p) of 1 ms is used and the following statistics are collected:

- *the average overall cell rate of the video source*: the video source's overall cell transmission rate averaged over the duration of the simulation,
- *the average high priority cell rate of the video source*: the video source's high priority cell transmission rate averaged over the duration of the simulation, and
- *the average peak signal to noise ratio*: the peak signal-to-noise ratio of the video sequences received by the destinations, averaged over the duration of the simulation.

Fig. 14 illustrates the effect of increasing the number of destinations on the average overall and high priority cell rates of the video source for the feedback polling and feedback coalescence mechanisms. For comparison, the transmission rate of a CBR connection reserving the same amount of bandwidth as the two feedback mechanisms is shown. Note that both feedback mechanisms provide a relatively constant average overall cell rate as the number of

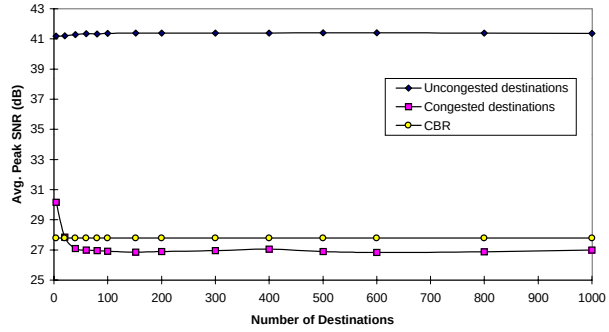

Figure 15: Average peak SNR vs. destination count for feedback polling (50% of paths congested by steady bottleneck load)

destinations in the connection is increased. The feedback coalescence mechanism also provides a stable average high priority cell rate as the number of destinations increases. The feedback polling mechanism's average high priority cell rate, on the other hand, decreases rapidly and is dominated by the minimum cell rate when the number of destinations is greater than 50. This limiting of the high priority cell rate is due to the lengthening of the probe cycle as the number of destinations grows; when the probe cycle is long, it takes a longer amount of time before a high priority rate increase can take place.

Figs. 15 and 16 illustrate the impact of the two mechanisms on the quality of video received by destinations. Fig. 15 shows the feedback polling mechanism's impact on video quality and compares it to the video quality achieved by a CBR connection reserving an equivalent amount of bandwidth. The average peak SNR is used as a metric to gauge the quality of the video received by destinations. To obtain the average peak SNR for a set of destinations, the peak SNR values for each frame in the video sequence are first averaged over the simulation's duration. These averages are then averaged together with those of other destinations in the set. Thus, Fig. 15 gives a measure of the quality of video received by two types of destinations par-

Figure 12: Sample frame of video received by an uncongested destination

Figure 13: Sample frame of video received by a congested destination

ticipating in the feedback polling mechanism: congested and uncongested destinations.

Fig. 15 shows uncongested destinations receive video of high quality. This is because the feedback polling mechanism allows the video source to increase its overall cell rate and make use of unutilized bandwidth on those paths. Congested destinations, on the other hand, receive video of lower quality, due largely to the loss of low priority cells on those paths. The shape of the curve for congested destinations reflects the downward trend in the high priority cell rate seen in Fig. 14. Finally, compare these two curves to the quality of video obtained using a single-layered CBR connection reserving the same amount of network bandwidth as the feedback polling algorithm. While congested destinations receive quality slightly below that of a CBR connection when the total number of destinations is greater than 20, uncongested destinations receive video of a much higher quality than the CBR connection. Thus, when the network has available bandwidth, feedback polling is capable of providing, on average, a higher video quality to all destinations than can CBR.

Fig. 16 shows the quality of video received by congested and uncongested destinations under the feedback coalescence mechanism and compares them to the quality of video that would be received by a CBR connection reserving an equivalent amount of bandwidth. Here, the feedback coalescence curves are both flat due to the mechanism's independence of destination count. Furthermore, for congested destinations, video quality is improved over that of the feedback polling mechanism *and* CBR. For instance, in a multipoint connection with 500 destinations, the feedback coalescence mechanism provides video with an average peak SNR of 29.5 dB to congested destinations, while CBR achieves a value of 27.8 dB and feedback polling achieves a value of 26.9 dB for the same destinations.

Together, Figs. 14, 15 and 16 illustrate the scalability of the feedback polling and feedback coalescence mechanisms. Of the two, the feedback coalescence mechanism is more scalable, since each forward feedback cell causes the return of a backward feedback cell describing the congestion state of the entire multipoint connection. However, the feedback polling mechanism, while not capable of optimizing its high priority cell rate for more than 50 destinations, is still capable of improving on the service CBR offers to uncongested destinations. Thus, both mechanisms are capable of providing better video quality to multiple destinations than CBR when there is unutilized bandwidth

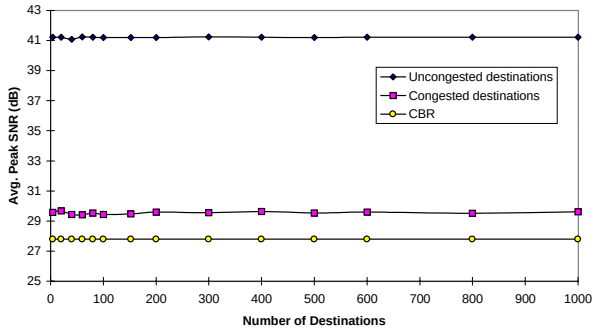

Figure 16: Average peak SNR vs. destination count for feedback coalescence (50% of paths congested by steady bottleneck load)

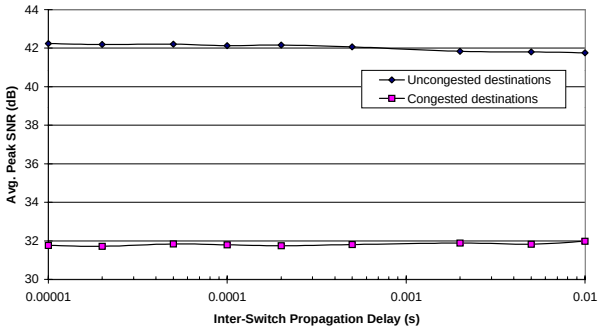

Figure 18: Average peak SNR vs. inter-switch propagation delay for feedback polling (100 destinations, 50% of paths congested by oscillating bottleneck load)

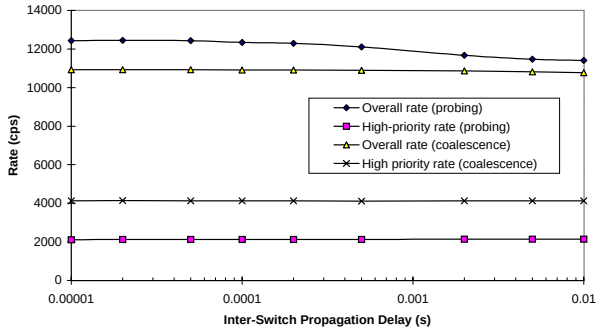

Figure 17: Average video rate vs. inter-switch propagation delay for feedback polling and feedback coalescence (100 destinations, 50% of paths congested by oscillating bottleneck load)

available in the network.

4.2.3 The Impact of Propagation Delay

In the second set of simulations, the impact of propagation delay on the performance of the two feedback mechanisms is assessed by varying the inter-switch propagation delay from as little as 10 μs to as much as 10 ms. The number of destinations is held constant at 100, with 25 destinations per leaf switch, and interfering high priority traffic is introduced on the links from SW_2 to SW_4 and SW_5. The interfering bottleneck traffic is Poisson, but it oscillates between two loads, $\rho = 0.98$ and $\rho = 0.50$, every 300 ms. This oscillation is introduced to test the reaction times and stability of the feedback mechanisms in the face of increasing propagation delays.

Fig. 17 illustrates the average overall and high priority cell rates for each mechanism as a function of the inter-switch propagation delay. The average overall cell rate for the feedback polling mechanism decreases as the propagation delay increases, while the average overall cell rate for the feedback coalescence mechanism declines only slightly. These declines occur due to the increased delay of accurate congestion information to the video source. Each time the interfering load oscillates between its two load levels, there is a lag between its actual occurrence and the notification to the video source of its occurrence. The look-back congestion functionality helps to ameliorate this lag by short-cutting the feedback response when a switch buffer has just become congested. However, it does little to improve the response when the switch buffer has just become uncongested; nearly a full round-trip time is required before notification of the switch's uncongested state can take place. Thus, as the propagation delay increases, the overall cell rate is diminished due to slower feedback response. This effect is less pronounced in the case of the feedback coalescence mechanism, because its feedback cells report the state of the entire multipoint connection. The feedback polling mechanism, on the other hand, must cyclically probe one destination at a time, and thus it may take longer to locate and respond to the conditions on the congested link.

The high priority cell rates for both mechanisms are unaffected by the increase in propagation delay. For the feedback polling mechanism, the flatness is explained by the domination of the high priority cell rate by the minimum cell rate. Even at low propagation delays, a probe cycle with 100 destinations is long enough to keep the average high priority cell rate near the minimum cell rate. For the feedback coalescence mechanism, however, the average high priority cell rate is not limited by the minimum cell rate. Rather, it is controlled by the number of positive and negative low priority traffic congestion indications returning to the source. Since the lengthening of the propagation delay does little to change the ratio of positive and negative LCI indications sent to the source, the average high priority cell rate essentially remains flat as propagation delay increases.

Figs. 18 and 19 illustrate the impact of propagation delay on the video quality received by congested and uncongested destinations. The modest declines in video quality shown in these figures reflect the modest declines exhibited in the overall and high priority cell rates of Fig. 17.

Percentage of Paths Congested	Feedback Polling Mechanism			Feedback Coalescence Mechanism		
	Avg. overall cell rate (cps)	Avg. high priority cell rate (cps)	Avg. peak SNR to all destinations (dB)	Avg. overall cell rate (cps)	Avg. high priority cell rate (cps)	Avg. peak SNR to all destinations (dB)
25%	13381	2079	38.5	13221	4192	42.6
50%	10898	2146	34.2	10176	3320	35.0
75%	7597	2066	30.8	5958	2798	32.5

Table 1: Impact of path congestion on the feedback polling and coalescence mechanisms

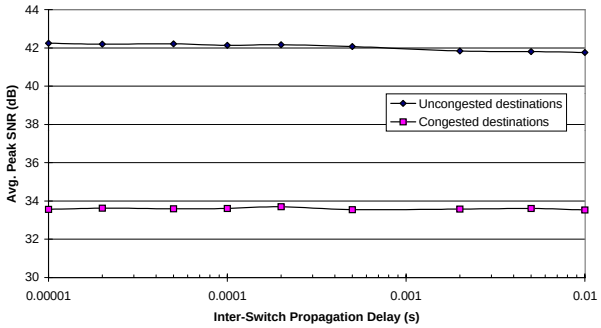

Figure 19: Average peak SNR vs. inter-switch propagation delay for feedback coalescence (100 destinations, 50% of paths congested by oscillating bottleneck load)

4.2.4 The Impact of the Number of Congested Paths

Finally, the performance of the two feedback mechanisms under various degrees of network congestion is studied. Three degrees of congestion are examined: one in which one quarter of the source-to-destination paths are congested, one in which half of them are congested, and one in which three quarters of them are congested. To generate these three different degrees of congestion, a non-oscillating, interfering Poisson packet load of $\rho = 0.98$ is introduced on the following links:

- 25% congestion: the link between SW_2 and SW_4,
- 50% congestion: the links between SW_2 and SW_4; and SW_2 and SW_5,
- 75% congestion: the links between SW_2 and SW_4; SW_2 and SW_5; and SW_3 and SW_6.

The inter-switch propagation delay (P_p) is 1 ms, and the number of destinations is fixed at 100.

Table 1 lists several measures of the mechanisms' performance as a function of the percentage of congested source-to-destination paths. For both the feedback polling and feedback coalescence algorithms, the video sources respond to larger degrees of congestion by reducing their overall cell rates. This is the desired effect. As the network becomes more and more congested, the video source reduces its overall cell rate to avoid overloading more links than necessary. The average high priority cell rate for feedback coalescence also shows a decline as the number of congested paths increase. This decline is due to a greater probability that any one destination will indicate low-priority traffic congestion when there are more congested paths. The average high priority cell rate for feedback polling shows no decline, because, as Fig. 14 showed, feedback polling's high priority cell rate is limited by the minimum cell rate when the total number of destinations is greater than 50.

The average peak SNR values reflect the average video quality received by all 100 destinations. As the number of congested paths increases, the average video quality drops as expected. However, the feedback coalescence mechanism produces an average video quality that is higher than that for the feedback polling mechanism, due largely to its ability to gauge the congestion status of the connection rapidly and to adjust its high priority cell rate appropriately, even for a large number of destinations.

5 Conclusion

We have presented two novel feedback mechanisms for reactively controlling the rate of high and low priority video in a point-to-multipoint congestion. The rates of transmitted video are controlled according to the congestion status of the network, with the high priority cell rate being determined by the available bandwidth on the multipoint connection's most congested link and the overall cell rate being determined by the connection's global congestion status.

In order to investigate the performance of the proposed video service architecture, we performed a series of simulations in which actual video sequences were transmitted through a simulated network. Results from these simulations showed that (1) the feedback polling and feedback coalescence mechanisms are capable of providing better video quality than CBR service when bandwidth is available in the network, (2) both mechanisms can scale to a large number of destinations, with feedback coalescence being particularly immune to the destination total, (3) increases in end-to-end propagation delay reduce the quality of video received by destinations, but not to such a degree that the mechanisms are useless in a wide area environment, and (4) the feedback mechanisms can adapt to varying degrees of congestion in the network, controlling high priority and overall cell rates to improve video quality when possible and gracefully degrade video quality when necessary.

Percentage of Paths Congested	Avg. peak SNR to all destinations (dB)		
	Linear Heuristic	PCR Heuristic	Semilinear Heuristic
25%	39.2	39.3	39.4
50%	35.4	35.7	35.7
75%	33.1	31.6	32.1

Table 2: Impact of path congestion on three TCR heuristics

In future work, we would like to augment the feedback mechanisms with explicit rate control. The current mechanisms rely on binary congestion indications, but with explicit rate control, the destination or the switches may indicate the exact rates at which a source may transmit high and low priority video. We would also like to study the use of the MPEG-2 scalability extension known as SNR scalability. Data partitioning, while simple, has the negative side effect that as the overall rate increases with the high priority rate remaining constant, the quality of high priority video decreases. The use of the SNR scalability extension does not suffer from this effect to the same degree.

References

[1] A.R. Reibman and A.W. Berger. Traffic Description for VBR Video Teleconferencing Over ATM Networks. *IEEE Trans. on Networking*, 3(3):329–339, June 1995.

[2] ATM Forum Traffic Management Specification Version 4.0. *ATM Forum Technical Committee, Traffic Management Working Group*, April 1996.

[3] H. Kanakia, P.P. Mishra, and A. Reibman. An Adaptive Congestion Control Scheme for Real-Time Packet Video Transport. *IEEE/ACM Transactions on Networking*, 3(6):671–682, Dec. 1995.

[4] C.M. Sharon, M. Devetsikiotis, I. Lambadaris, and A.R. Kaye. Rate Control of VBR H.261 Video on Frame Relay Networks. *Proc. of IEEE ICC '95*, pages 1443–1447, 1995.

[5] Y. Omori, T. Suda, and G. Lin. Feedback-Based Congestion Control for VBR Video in ATM Networks. *Proc. of the 6th IEEE Workshop on Packet Video*, 1994.

[6] B. J. Vickers and T. Suda. An ATM Service Architecture for the Transport of Adaptively Encoded Live Video. *Proc. of the IEEE International Conference on Computer Communications and Networks*, Oct. 1996.

[7] R. Yavatkar and L. Manoj. Optimistic Strategies for Large-Scale Dissemination of Multimedia Information. *Proc. of ACM Multimedia*, pages 13–20, August 1993.

[8] J.C. Bolot, T. Turletti, and I. Wakeman. Scalable Feedback Control for Multicast Video Distribution in the Internet. *Proc. of SIGCOMM*, pages 58–67, Aug. 1994.

[9] J. C. Lin and S. Paul. RMTP: A Reliable Multicast Transport Protocol. *Proc. of IEEE INFOCOM '96*, pages 1414–1424, March 1996.

[10] S. McCanne, V. Jacobson, and M. Vetterli. Receiver-driven Layered Multicast. *Proc. of SIGCOMM*, pages 117–130, Aug. 1996.

[11] Generic Coding of Moving Pictures and Associated Audio Information. *Recommendation H.262 of the SC29/WG11 committee (MPEG), draft submitted to ISO-IEC/JTC1 SC29*, March 1995.

[12] A.C. Hung. PVRG-MPEG Codec 1.1 Documentation. *Portable Video Research Group, Stanford University*, 1993.

[13] A. Demers, S. Keshav, and S. Shenker. Analysis and Simulation of a Fair Queueing Algorithm. *Proc. of SIGCOMM*, pages 1–12, Sept. 1989.

[14] D. Ferrari and D. Verma. A Scheme for Real-Time Channel Establishment in Wide-Area Networks. *IEEE J. Select. Areas Commun.*, SAC-8(3):368–379, April 1990.

[15] L. Zhang. Virtual Clock: A New Traffic Control Algorithm for Packet Switching Networks. *Proc. of SIGCOMM*, pages 19–29, Sept. 1990.

6 Appendix

As described in section 3.1.1, the feedback control mechanisms were simulated using the following target cell rate heuristic:

$$\text{TCR} = \text{MCR} + \frac{N_0}{N} \times (\text{PCR} - \text{MCR}) \qquad (3)$$

where N is the total number of destinations in the multipoint connection and N_0 is the number of uncongested destinations. However, this linear heuristic is not the only possibility for calculating the video encoder's target overall cell rate. Other candidate heuristics include:

- **The PCR heuristic**: always set the target cell rate equal to the peak cell rate (i.e., TCR=PCR)
- **The semi-linear heuristic**: the target cell rate remains equal to the peak cell rate until over 50% of the destinations become congested, at which point the TCR begins to fall linearly to the minimum cell rate.

$$\text{TCR} = \begin{cases} \text{PCR} & \text{if } N_0 \leq N/2 \\ \text{MCR} + \frac{N_0}{N/2} \times (\text{PCR} - \text{MCR}) & \text{otherwise} \end{cases} \qquad (4)$$

Using the simulation model of Fig. 7 and the feedback coalescence mechanism, the video qualities achieved using these three heuristics are compared. For this set of simulations, we consider a network with 400 destinations and a propagation delay between switches of 1 ms. The three degrees of path congestion used in section 4.2.4 are then examined for each of the three TCR heuristics. Table 2 summarizes the results. Of the three heuristics simulated, all three perform nearly identically up to the point at which the multipoint connection is 50% congested. However, for a network that is 75% congested, the linear heuristic achieves the best video quality averaged over all destinations, while the PCR heuristic results in the lowest quality. The PCR and semilinear heuristics result in a higher target cell rate at the 75% congestion level, causing more cells to be dropped in the network than the linear heuristic. The dropping of video cells in the network has a more deleterious effect on video quality than a graceful reduction in the encoding quantization at the source.