

Design and Evaluation of Feedback Consolidation for ABR Point-to-Multipoint Connections in ATM Networks

Sonia Fahmy, Raj Jain, Rohit Goyal, Bobby Vandalore, and Shivkumar Kalyanaraman

The Ohio State University

2015 Neil Ave, DL 395, Columbus, OH 43210-1277, USA

Raj Jain is now at Washington University in Saint Louis, jain@cse.wustl.edu <http://www.cse.wustl.edu/~jain/>

Abstract: The available bit rate (ABR) service is proposed to transport data traffic in asynchronous transfer mode (ATM) networks. ABR is unique because the network switches can indicate to the sources the rates at which they should be transmitting, thus avoiding congestion and efficiently utilizing network resources. A number of algorithms have been developed for extending ABR flow control algorithms for point-to-multipoint connections. In this case, feedback consolidation is required at the branch points to avoid overwhelming the sender with feedback. This paper discusses various design options and implementation alternatives for consolidation algorithms, and proposes a number of novel algorithms. The performance of the proposed algorithms and the original algorithms is compared under a variety of conditions. Results indicate that the algorithms we propose eliminate the consolidation noise (caused if the feedback is returned before all branches respond), while exhibiting a fast transient response.

Keywords: ATM networks, ABR service category, traffic management, congestion control, multipoint communication, feedback consolidation

1 Introduction

Multipoint communication is the exchange of information among multiple parties. Examples of multipoint applications include audio and video conferencing applications, video on demand, distance learning, tele-metering, distributed games, server and replicated database synchronization, advertising, searching and data distribution applications. Multipoint capabilities are essential for ATM networks to efficiently support many applications, including IP multicasting and overlay applications (see figure 1).

The ABR service category attempts to provide possibly non-zero minimum rate guarantees, achieve fairness, and minimize cell loss for data (non real-time) traffic by periodically indicating to ABR sources the rate at which they should be transmitting. ABR flow control requires the sources to send at the rate specified by the network in feedback (resource management) cells. For point-to-multipoint connections, feedback consolidation at the branch points becomes necessary. The operation of feedback consolidation can be explained by figure 2. The consolidation operation avoids the feedback implosion problem,




Figure 1: Multicasting in the Internet

where the number of backward resource management (BRM) cells received by the source is proportional to the number of leaves in the multicast tree. In addition, the allowed rate of the source should not fluctuate due to the varying feedback received from different leaves.




Figure 2: Point-to-multipoint ABR connections

In *point-to-point* ABR connections, the source transmits at the minimum rate that can be supported by all the switches on the path from the source to the destination [6]. The natural extension of this strategy for point-to-multipoint connections is controlling the source to the minimum rate that can be supported by the switches on the paths from the source to *all of the leaves* in the multicast tree, as shown in figure 3. The minimum rate is the technique most compatible with typical data requirements: no data should be lost, and the network can take whatever time needed for data delivery. If the ABR destinations *can* tolerate a certain amount of cell loss, alternative techniques can be used [8].




Figure 3: Taking the minimum in point-to-multipoint ABR connections

A number of consolidation algorithms have been proposed in [1, 2, 9, 8, 10, 11, 12, 13, 14, 15, 16, 17, 18]. Several design and implementation considerations come into play when developing a consolidation algorithm. The oscillations and transient response of the algorithm are important. The algorithm must also be scalable to very large multicast trees. The implemen-

tation complexity, feedback delay, and the overhead of the backward RM cells should not increase with the increase of the number of levels or leaves of the multicast tree.

In this paper, we propose a set of consolidation algorithms that aim at providing a fast transient response, while eliminating consolidation noise. We examine the performance of the proposed algorithms, and compare it to the previous ones in complexity, transient response, consolidation noise, and scalability. The remainder of the paper is organized as follows. The next two sections provide an overview of the ABR flow control mechanism, and a summary of the previous work on point-to-multipoint ABR flow control. A discussion of the various design and implementation issues involved is then presented, followed by a description of the specific underlying congestion avoidance scheme employed. An explanation and pseudocode of the previously proposed consolidation algorithms, as well as the new ones we propose, is presented next. All the algorithms are then simulated and analyzed under a variety of configurations. The paper concludes with a discussion of the tradeoffs among the algorithms.

2 ABR Flow Control

The available bit rate (ABR) service for data traffic in ATM networks periodically indicates to sources the rate at which they should be transmitting. The switches monitor their load and compute the available bandwidth, dividing it fairly among active flows. The feedback from the switches to the sources is sent in resource management (RM) cells which are generated by the sources and turned around by the destinations (see figure 4).




Figure 4: Resource management cells in an ATM network

The RM cells contain the source current cell rate (CCR), in addition to fields that can be used by the switches to provide feedback to the sources. These fields are: explicit rate (ER), the congestion indication (CI) flag, and no increase (NI) flag. The ER field indicates the rate that the network can support at this particular instant. Initially, the ER field is set to a value no greater than the peak cell rate (PCR), and the CI and NI flags are clear. Each switch on the path *reduces* the ER field to the maximum rate it can support, and sets CI or NI if necessary [6].

A component c_j is said to be *downstream* of another component c_i in a certain connection if c_j is on the path from c_i to the destination. In this case, c_i is said to be *upstream* of c_j . RM cells flowing from the source to the destination are called forward RM cells (FRMs) while those returning from the destination to the source are called backward RM cells (BRMs). When a source receives a BRM, it computes its allowed cell rate (ACR) using its current ACR, the CI and NI flags, and the ER field of the RM cell.

3 Related Work

A simple point-to-multipoint ABR algorithm was proposed in [14]. In this algorithm, a register MER (minimum explicit rate), maintains the minimum feedback indicated by the BRM cells received from the branches. Whenever an FRM cell is received, it is multicast to all branches, and a BRM is returned with the MER value as the explicit rate. MER is then reset.

This algorithm suffers from the “consolidation noise” problem when a BRM generated by a branch point does not consolidate feedback from all tree branches [7]. In fact, if a BRM generated by the branch point does not accumulate feedback from *any* branch, the feedback can be given as the peak cell rate (if that branch point itself is not overloaded). In [12, 15] some solutions to this problem are proposed.

To reduce the complexity of the scheme, Ren and Siu [12] propose to forward one of the BRM cells returned by the leaves, instead of turning around the FRM cells of the source. Another alternative would be to pass back the BRM cell only when BRM cells from *all* branches have been received after the last feedback. This idea is also used in [2], but the BRM cell that is allowed to pass back to the source is the last BRM cell to be received with a certain sequence number. These approaches suffer from a slow transient response.

Other schemes, such as those in [1, 9, 10, 17, 18] have been recently developed. Some of these schemes use timers at the branch points, and the BRM cells are sent by the branch points when the timer expires. Timers, however, are expensive to implement in network elements. Others such as [9] store all values returned by the BRM cells from the different branches and update the values as BRM cells are received. The minimum has to be taken every time a BRM cell is to be sent. This can be an expensive operation. Thus, the main problem with new schemes is the complexity of implementation.

4 Design Issues

As previously mentioned, there are several ways to implement the feedback consolidation algorithm at branch points. Each method offers a tradeoff in complexity, scalability, overhead, transient response, and consolidation noise. The tradeoffs can be summarized as follows:

[A] Which component generates the BRM cells (i.e., turns around the FRM cells)? Should the branch point, or should the destination, perform this operation?

[B] What is the condition to trigger sending a BRM at the branch point? Should the branch point wait for feedback from all the branches before passing the BRM cell upstream? Although this eliminates the consolidation noise, it incurs additional complexity, and increases the transient response of the scheme (especially after idle or low rate periods).

[C] Does the scheme scale well? How can the ratio of FRM cells generated by the source to BRM cells returned to the source be controlled? Will the feedback delay grow with the number of branches? For example, if the algorithm waits for an FRM cell to be received before sending feedback, the delay might increase with the number of levels of the multicast tree (see figure 5).

[D] How does the branch point operate when it is also a switch (queuing point)? (see figure 6). The coupling of the switch and branch point functions must be considered. When should the actual rate computation algorithm be invoked?




Figure 5: Multiple branch points and levels




Figure 6: Branch points, switches and their coupling

[E] How are non-responsive branches handled? If the consolidation scheme waits for feedback from *all* the branches before sending a BRM to the source, an algorithm must be developed to determine when a branch becomes non-responsive and handle this case (see figure 7). A timeout or outstanding RM cell count mechanism must be implemented. We believe that an outstanding RM cell count (similar to the mechanism used with source end system rule 6 in [6]) is *superior* to the timeout mechanism used in most algorithms, as the timeout value needs to be constantly re-estimated. Such mechanisms will be the subject of a future study.

[F] How is accounting performed at the branch point? Consolidation algorithms use registers to store values such as the minimum rate given by branches in the current iteration, and flags to indicate whether an RM cell has been received since the last one was sent. Some values, if stored per output port, need an efficient data structure. For example, if the rate returned by each branch is stored (instead of only maintaining the minimum value for this connection for this round), a heap is necessary to enable the minimum operation to be rapidly performed.

5 The ERICA Algorithm

The ERICA algorithm is used in our simulations to calculate the explicit rate (ER) feedback in RM cells based on the load at each port. In this section, we only present the basic features of the algorithm. For a more complete explanation of the algorithm, refer to [3]. The point-to-multipoint algorithms are presented in the next two sections.




Figure 7: Non-responsive branch

ERICA aims at fair and efficient allocation of the available bandwidth to all contending sources. The ERICA scheme periodically monitors the load on each link and determines a load factor, z , the available capacity, and the number of currently active connections (VCs). The load factor, z , is an indicator of the congestion level of the link. The optimal operating point is at an overload value of one. The load factor is calculated as the ratio of the measured input rate at the port to the target capacity of the output link:

$$z \leftarrow \text{ABR Input Rate} / \text{ABR Capacity}$$

where:

$$\text{ABR Capacity} \leftarrow \text{Target Utilization} \times \text{Link Bandwidth} - \text{VBR Usage} - \text{CBR Usage}.$$

Target utilization is a parameter set to a fraction (close to, but less than 100%). Alternatively, it can be a function of the queuing delay.

The switch calculates the quantity:

$$\text{VC-Share} \leftarrow \text{CCR} / z$$

If all VCs change their rate to their VCShare values then, in the next cycle, the switch would experience unit overload ($z = 1$).

The fair share of each VC, FairShare , is also computed as follows:

$$\text{FairShare} \leftarrow \text{ABR Capacity} / \text{Number of Active VCs}$$

A combination of the two quantities FairShare and VCShare is used to rapidly reach optimal operation as follows:

IF $z \leq 1.1$ THEN

ER Calculated \leftarrow Max (FairShare, VCShare, Maximum ER given in previous interval)

ELSE ER Calculated \leftarrow Max (FairShare, VCShare)

Several enhancements to this algorithm avoid transient overloads, and take the queuing delay into consideration when assessing the available capacity. Averaging the measured quantities further improves the performance. These enhancements are described in [3].

6 Consolidation Algorithms

This section describes some previously proposed consolidation algorithms, while the next section proposes new algorithms. In the algorithms presented, ERICA (as explained in the previous section) is employed immediately before sending a BRM on the link. This ensures that the most recent feedback information is sent.

In this and the next section, we describe the first algorithm, and then describe the modifications to algorithm 1 resulting in algorithm 2, and so on, until we arrive at our proposed algorithm, algorithm 7. Thus the description of the algorithms is incremental. The modifications are usually italicized.

6.1 Algorithm 1

This algorithm is a modified version of the algorithm in [14]. The main idea of the algorithm is that BRM cells are returned from the branch point when FRM cells are received, and the BRM cells contain the minimum of the values indicated by

the BRM cells received from the branches after the last BRM cell was sent. FRM cells are duplicated and multicast to all branches at the branch point.

A register, MER, and two flags, MCI and MNI, are maintained for each multipoint VC. MER stores the minimum of the explicit rate (ER) values, MCI is the logical OR of the congestion indication (CI) values, and MNI is the logical OR of the no increase (NI) values indicated in the BRM cells which were received after the last BRM cell was sent. MER is initialized to the peak cell rate, while CI and NI are initialized to zero. Three temporary variables: MXER, MXCI, and MXNI are also used when an FRM cell is received (their values do not persist across invocations of the algorithm). They store the ER, CI and NI from the FRM cell. The algorithm operates as follows.

Upon the receipt of an FRM cell:

1. Multicast FRM cell to all participating branches
2. Let $MXER = ER$ from FRM cell, $MXCI = CI$ from FRM cell, $MXNI = NI$ from FRM cell
3. Return a BRM with $ER = MER$, $CI = MCI$, $NI = MNI$ to the source
4. Let $MER = MXER$, $MCI = MXCI$, $MNI = MXNI$

Upon the receipt of a BRM cell:

1. Let $MER = \min(MER, ER \text{ from BRM cell})$, $MCI = MCI \text{ OR } CI$ from BRM cell, $MNI = MNI \text{ OR } NI$ from BRM cell
2. Discard the BRM cell

When a BRM is to be scheduled:

Let $ER = \min(ER, ER \text{ calculated by rate allocation algorithm for all branches})$

6.2 Algorithm 2

This algorithm is a modified version of the second algorithm in [12]. The only change from Algorithm 1 (as described above) is ensuring that at least one BRM cell has been received from a branch before turning around an FRM. For this purpose, a boolean flag, *AtLeastOneBRM* (initially zero), is set to true when a BRM cell is received from a branch, and reset when a BRM is sent by the branch point (see the italicized portion). As before, MER, MCI, MNI, and here, *AtLeastOneBRM*, are stored for each multipoint VC, and MXER, MXCI, MXNI are temporary variables.

Upon the receipt of an FRM cell:

1. Multicast FRM cell to all participating branches
2. *IF AtLeastOneBRM THEN*
 - A. Let $MXER = ER$ from FRM cell, $MXCI = CI$ from FRM cell, $MXNI = NI$ from FRM cell
 - B. Return a BRM with $ER = MER$, $CI = MCI$, $NI = MNI$ to the source
 - C. Let $MER = MXER$, $MCI = MXCI$, $MNI = MXNI$
 - D. *Let AtLeastOneBRM = 0*

Upon the receipt of a BRM cell:

1. *Let AtLeastOneBRM = 1*
2. Let $MER = \min(MER, ER \text{ from BRM cell})$, $MCI = MCI \text{ OR } CI$ from BRM cell, $MNI = MNI \text{ OR } NI$ from BRM cell
3. Discard the BRM cell

When a BRM is to be scheduled:

Let $ER = \min (ER, ER \text{ calculated by rate allocation algorithm for all branches})$

6.3 Algorithm 3

The main idea here is that the branch point does not turn around the FRMs, but the BRM that is received from a branch immediately after an FRM has been received by the branch point is passed back to the source, carrying the minimum values. A boolean flag, *AtLeastOneFRM*, indicates that an FRM cell has been received by the branch point after the last BRM cell was passed to the source. Again, *MER*, *MCI*, *MNI*, and *AtLeastOneFRM* are stored per multipoint VC. This is a modified version of the third algorithm in [12].

Upon the receipt of an FRM cell:

1. Multicast FRM cell to all participating branches
2. *Let AtLeastOneFRM = 1*

Upon the receipt of a BRM cell:

1. Let $MER = \min (MER, ER \text{ from BRM cell})$, $MCI = MCI \text{ OR } CI \text{ from BRM cell}$, $MNI = MNI \text{ OR } NI \text{ from BRM cell}$
2. *IF AtLeastOneFRM THEN*
 - A. Pass the BRM with $ER = MER$, $CI = MCI$, $NI = MNI$ to the source
 - B. Let $MER = PCR$, $MCI = 0$, $MNI = 0$
 - C. *Let AtLeastOneFRM = 0*
- ELSE Discard the BRM cell

When a BRM is to be scheduled:

Let $ER = \min (ER, ER \text{ calculated by rate allocation algorithm for all branches})$

6.4 Algorithm 4

A variation of this algorithm was presented in [12] as algorithm 4, and another variation using sequence numbers in RM cells was proposed in [2]. The main idea here is that a BRM is passed to the source only when BRM cells have been received from all branches. To count the number of branches from which BRM cells were received at the branch point (after the last BRM cell was passed by the branch point), a counter, *NumberOfBRMsReceived* is incremented the first time a BRM cell is received from each branch (*NumberOfBRMsReceived* is initialized to zero). As before, the *MER*, *MCI*, *MNI*, and *NumberOfBRMsReceived* registers are maintained per multipoint VC. The value of the *NumberOfBRMsReceived* counter is compared to the value of another counter, *NumberOfBranches*, every time a BRM cell is received by the branch point. If the value of *NumberOfBRMsReceived* is equal to *NumberOfBranches*, the BRM cell is passed back to the source, carrying the values of the *MER*, *MCI* and *MNI* registers. [*NumberOfBranches* stores the number of branches of the point-to-multipoint VC at this branch point. It is also stored for each VC, and initialized during connection setup. In addition, if leaf initiated join is allowed (as in UNI 4.0), *NumberOfBranches* must be updated every time a new branch is added to a branch point.]

A flag, *BRMReceived*, is needed for each branch to indicate whether a BRM cell has been received from this particular branch, after the last BRM cell was passed. The flag is stored for each output port and not for each VC, since it is needed

for each branch. Note that a mechanism must be implemented to ensure that BRM cell flow is not stopped in the case of non-responsive branches. Timeouts or RM cell counters can be used for that purpose. This will be the subject of a future study.

Upon the receipt of an FRM cell:

Multicast FRM cell to all participating branches

Upon the receipt of a BRM cell from branch i :

1. *IF NOT BRMReceived _{i} THEN*
 - A. *Let BRMReceived _{i} = 1*
 - B. *Let NumberOfBRMsReceived = NumberOfBRMsReceived + 1*
 2. Let MER = min (MER, ER from BRM cell), MCI = MCI OR CI from BRM cell, MNI = MNI OR NI from BRM cell
 3. *IF NumberOfBRMsReceived is equal to NumberOfBranches THEN*
 - A. Pass the BRM with ER = MER, CI = MCI, NI = MNI to the source
 - B. Let MER = PCR, MCI = 0, MNI = 0
 - C. *Let NumberOfBRMsReceived = 0*
 - D. *Let BRMReceived = 0 FOR all branches*
- ELSE Discard the BRM cell

When a BRM is to be scheduled:

Let ER = min (ER, ER calculated by rate allocation algorithm for all branches)

7 New Algorithms

The main problem with algorithm 4 described in the previous section is its slow transient response. Even when excessive overload has been detected, the algorithm has to wait for feedback from (possibly distant) leaves before indicating the overload information to the source. By that time, the source might have transmitted a large number of cells (which would be dropped due to buffer overflows), resulting in performance degradation. This situation is especially problematic when the source has been idle for some time, and then suddenly sends a burst, so there are no RM cells initially in the network.

The main idea behind the algorithms presented next is that the slow transient response problem should be avoided when a *severe overload* situation has been detected. In this case, there is no need to wait for feedback from all the branches, and the overload should be immediately indicated to the source. In cases of underload indication from a branch, it is better to wait for feedback from all branches, since other branches may be overloaded. This is somewhat similar to the idea behind the backward explicit congestion notification (BECN) cells sent by the switches.

Overload is detected when the feedback to be indicated is *much less* than the last feedback returned by the branch point. The “much less” condition is tested using a multiplicative factor, Threshold. The Threshold value can range from zero to one. A Threshold value of one means that overload is detected when the feedback to be given is less than the current ER (even when it is 99% of the last ER value given); a threshold value of zero means that overload is not detected except when the new rate to be indicated is zero, which, in effect, disables the fast overload indication feature.

An alternative method would be to compare the feedback to be indicated to the *current cell rate (CCR)* or ACR of the VC. Although this may be better because it accounts for upstream bottlenecks, and prevents the transmission of unnecessary BRM cells in such cases, the CCR information may be stale due to the delay from the source to the branch point (it may also be much larger when the source becomes idle or becomes a low rate source after the last FRM was sent), and a large number of BRMs may be sent in such cases. The last feedback indicated by the branch point is a more current value. The minimum of the CCR and last feedback given can be used in the comparison, but this involves some additional complexity, and may slow down the overload response when the CCR happens to have been small, but is currently large.

Note that when a BRM cell is returned due to overload detection before feedback has been received from all branches, the counters and the register values are not reset.

7.1 Fast Overload Indication (Algorithm 5)

In this algorithm, the LastER register maintains the last explicit rate value returned by the branch point (LastER is initialized to the initial cell rate (ICR) of the connection). LastER is stored per multipoint VC, and is compared with the value of MER in step 5 below.

Two temporary variables: SendBRM and Reset are used. SendBRM is set only if a BRM cell is to be passed to the source by the branch point. Reset is false only if a BRM cell is being used to indicate overload conditions, and hence the register values should not be reset. FRMminusBRM is only used for accounting purposes, and will not exist in a real implementation.

Upon the receipt of an FRM cell:

1. Multicast FRM cell to all participating branches
- (* 2. Let $FRMminusBRM = FRMminusBRM + 1$ *)

Upon the receipt of a BRM cell from branch i :

1. Let $SendBRM = 0$
2. Let $Reset = 1$
3. IF NOT $BRMReceived_i$ THEN
 - A. Let $BRMReceived_i = 1$
 - B. Let $NumberOfBRMsReceived = NumberOfBRMsReceived + 1$
4. Let $MER = \min(MER, ER \text{ from BRM cell})$, $MCI = MCI \text{ OR } CI \text{ from BRM cell}$, $MNI = MNI \text{ OR } NI \text{ from BRM cell}$
5. IF $MER < (Threshold \times LastER)$ THEN (* overload is detected *)
 - A. IF $NumberOfBRMsReceived < NumberOfBranches$ THEN
 1. Let $Reset = 0$
 - B. Let $SendBRM = 1$
- ELSE IF $NumberOfBRMsReceived$ is equal to $NumberOfBranches$ THEN
 - A. Let $SendBRM = 1$
6. IF $SendBRM$ THEN
 - A. Pass the BRM with $ER = MER$, $CI = MCI$, $NI = MNI$ to the source
 - B. IF $Reset$ THEN

1. Let $MER = PCR, MCI = 0, MNI = 0$
 2. Let $NumberOfBRMsReceived = 0$
 3. Let $BRMReceived = 0$ FOR all branches
 - (* C. Let $FRMminusBRM = FRMminusBRM - 1$ *)
- ELSE Discard the BRM cell

When a BRM is to be scheduled:

1. Let $ER = \min (ER, ER \text{ calculated by rate allocation algorithm for all branches})$
2. Let $LastER = ER$

7.2 RM Ratio Control (Algorithm 6)

The previous algorithm may increase the BRM cell overhead, since the ratio of source-generated FRM cells to BRM cells received by the source can exceed one. To avoid this problem, we introduce the register *SkipIncrease* which is maintained for each multipoint VC (and initialized to zero). *SkipIncrease* is used to control the RM cell ratio. *SkipIncrease* is incremented whenever a BRM cell is sent before feedback from all the branches has been received. When feedback from all leaves indicates underload, and the value of the *SkipIncrease* register is more than zero, this particular feedback can be *ignored* and *SkipIncrease* decremented. Note that the value of the *SkipIncrease* counter will not increase to large values, since the rate allocation algorithm (such as ERICA) arrives at the optimal allocation within a few iterations, and the explicit rates computed *cannot continue decreasing indefinitely*. Our analysis and simulations have shown that the counter never exceeds small values and quickly stabilizes at zero. A maximum value can also be enforced by the algorithm.

Upon the receipt of an FRM cell:

1. Multicast FRM cell to all participating branches
- (* 2. Let $FRMminusBRM = FRMminusBRM + 1$ *)

Upon the receipt of a BRM cell from branch i :

1. Let $SendBRM = 0$
 2. Let $Reset = 1$
 3. IF NOT $BRMReceived_i$ THEN
 - A. Let $BRMReceived_i = 1$
 - B. Let $NumberOfBRMsReceived = NumberOfBRMsReceived + 1$
 4. Let $MER = \min (MER, ER \text{ from BRM cell}), MCI = MCI \text{ OR } CI \text{ from BRM cell}, MNI = MNI \text{ OR } NI \text{ from BRM cell}$
 5. IF $MER \geq LastER$ AND $SkipIncrease > 0$ AND $NumberOfBRMsReceived$ is equal to $NumberOfBranches$ THEN
 - A. Let $SkipIncrease = SkipIncrease - 1$
 - B. Let $NumberOfBRMsReceived = 0$
 - C. Let $BRMReceived = 0$ FOR all branches
- ELSE IF $MER < (\text{Threshold} \times LastER)$ THEN
- A. IF $NumberOfBRMsReceived < NumberOfBranches$ THEN
 1. Let $SkipIncrease = SkipIncrease + 1$
 2. Let $Reset = 0$

B. Let SendBRM = 1

ELSE IF NumberOfBRMsReceived is equal to NumberOfBranches THEN

A. Let SendBRM = 1

6. IF SendBRM THEN

A. Pass the BRM with ER = MER, CI = MCI, NI = MNI to the source

B. IF Reset THEN

1. Let MER = PCR, MCI = 0, MNI = 0
2. Let NumberOfBRMsReceived = 0
3. Let BRMReceived = 0 FOR all branches

(* C. Let FRMminusBRM = FRMminusBRM - 1 *)

ELSE Discard the BRM cell

When a BRM is to be scheduled:

1. Let ER = min (ER, ER calculated by rate allocation algorithm for all branches)
2. Let LastER = ER

7.3 Immediate Rate Computation (Algorithm 7)

The last two algorithms can offer very fast congestion relief when an overload is detected in a branch of the multicast tree. They do not, however, account for the potential overload situation at the branch point itself: if the branch point is a switch (queuing point), the ERICA algorithm is only performed when the BRM cell is about to be scheduled on the link. In cases when the branch point is itself a switch and queuing point, the immediate rate calculation option invokes ERICA whenever a BRM is received, and not just when a BRM is being sent. Hence overload at the branch point can be detected and indicated according to the fast overload indication option as previously described. Doing this, however, may involve some additional complexity.

The algorithm presented next is the same as *Algorithm 6* in the previous subsection, except for the addition of the ERICA invocation (italicized below).

Upon the receipt of an FRM cell:

1. Multicast FRM cell to all participating branches
- (* 2. Let FRMminusBRM = FRMminusBRM + 1 *)

Upon the receipt of a BRM cell from branch i :

1. Let SendBRM = 0
2. Let Reset = 1
3. IF NOT BRMReceived _{i} THEN
 - A. Let BRMReceived _{i} = 1
 - B. Let NumberOfBRMsReceived = NumberOfBRMsReceived + 1
4. Let MER = min (MER, ER from BRM cell), MCI = MCI OR CI from BRM cell, MNI = MNI OR NI from BRM cell
5. *Let MER = min (MER, minimum ER calculated by rate allocation algorithm for all branches)*

6. IF $MER \geq LastER$ AND $SkipIncrease > 0$ AND $NumberOfBRMsReceived$ is equal to $NumberOfBranches$ THEN

- A. Let $SkipIncrease = SkipIncrease - 1$
- B. Let $NumberOfBRMsReceived = 0$
- C. Let $BRMReceived = 0$ FOR all branches

ELSE IF $MER < (Threshold \times LastER)$ THEN

- A. IF $NumberOfBRMsReceived < NumberOfBranches$ THEN
 - 1. Let $SkipIncrease = SkipIncrease + 1$
 - 2. Let $Reset = 0$
- B. Let $SendBRM = 1$

ELSE IF $NumberOfBRMsReceived$ is equal to $NumberOfBranches$ THEN

- A. Let $SendBRM = 1$

7. IF $SendBRM$ THEN

- A. Pass the BRM with $ER = MER$, $CI = MCI$, $NI = MNI$ to the source
- B. IF $Reset$ THEN
 - 1. Let $MER = PCR$, $MCI = 0$, $MNI = 0$
 - 2. Let $NumberOfBRMsReceived = 0$
 - 3. Let $BRMReceived = 0$ FOR all branches
 - (* C. Let $FRMminusBRM = FRMminusBRM - 1$ *)

ELSE Discard the BRM cell

When a BRM is to be scheduled:

- 1. Let $ER = \min(ER, ER \text{ calculated by rate allocation algorithm for all branches})$
- 2. Let $LastER = ER$

8 Algorithm Summary

Table 1 summarizes the description of the seven algorithms given in the previous two sections, according to the design issues discussed in section 4. As seen in the table, there are a number of design and implementation options. Each of these options affects the performance of the algorithm. The remainder of the paper is devoted to the performance analysis and performance comparison of the algorithms.

9 Performance Analysis

This section provides a performance comparison among all the presented consolidation algorithms, in a variety of configurations with bursty and non-bursty traffic, with and without variable bit rate (VBR) background, and with various link lengths, bottleneck locations, and number of leaves. A number of other configurations was also tested (see [5, 4] for some of the configurations), but only a sample of the results is shown here. In particular, configurations with a large number of leaves at varying distances in the multicast tree were simulated, and the results were consistent with those we present here.

9.1 Parameter Settings

Throughout our experiments, the following parameter values are used:

- [1] All links have a bandwidth of 155.52 Mbps (149.76 Mbps when SONET overhead is accounted for).
- [2] All point-to-multipoint traffic flows from the root to the leaves of the tree. No traffic flows from the leaves to the root, except for RM cells. The same applies for point-to-point connections.
- [3] All sources are deterministic, i.e., their start/stop times and their transmission rates are known. The bursty traffic sources send data in bursts. VBR sources are on/off sources, where the on and off times are 20 ms.
- [4] The source parameter rate increase factor (RIF) is set to one, to allow immediate use of the full explicit rate indicated in the returning RM cells at the source. Initial cell rate (ICR) is also set to a high value (almost peak cell rate). These factors are set to such high values to simulate a *worst case* load situation.
- [5] The source parameter transient buffer exposure (TBE) is set to large values to prevent rate decreases due to the triggering of the source open-loop congestion control mechanism. This was done to isolate the rate reductions due to the switch congestion control from the rate reductions due to source end system rule six.
- [6] All other ABR parameters are set to their default values as specified in [6].
- [7] The switch target utilization parameter was set at 90%. The switch measurement interval was set to the minimum of the time to receive 100 cells and 1 ms.
- [8] The Threshold parameter used in Algorithms 5 to 7 was set to 0.95. This large value was used to illustrate the effect of the fast overload indication, RM ratio control, and immediate rate computation features of the algorithms. A lower value is recommended to be used in practice.

9.2 Simulation Results




Figure 8: WAN parking lot configuration with bursty, persistent and VBR connections

This section discusses the performance of the seven consolidation algorithms by comparing them in a set of configurations. Two graphs are plotted for each configuration: the allowed cell rate for all the ABR sources, and the queue lengths for *overloaded* switches.

Figures 9 through 15 illustrate the performance of the seven different algorithms in a situation where there is both variable capacity and variable demand. These situations offer the toughest challenge for rate allocation algorithms [4]. The configuration simulated is shown in figure 8 (all links are 1000 km). The source indicated by *W* is a bursty source, *S* is a persistent (infinite) source, while *V* is a VBR source. The VBR connection (V1 to dV1) is point-to-point, while the 2 ABR connections (bursty W1 sending to dW1, dW2, dW3; persistent S1 sending to dS1, dS2, dS3) are point-to-multipoint connections.

The ACR graphs of the ABR sources for algorithms 1, 2, and 3 indicate fluctuations and inaccurate (around 140 Mbps)

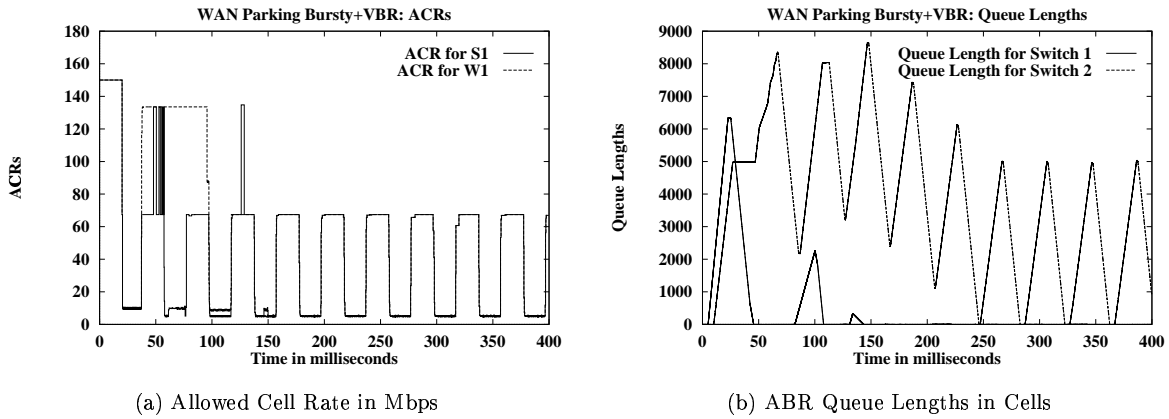


Figure 9: Results for WAN parking lot configuration with bursty, persistent and VBR connections [Algorithm 1]

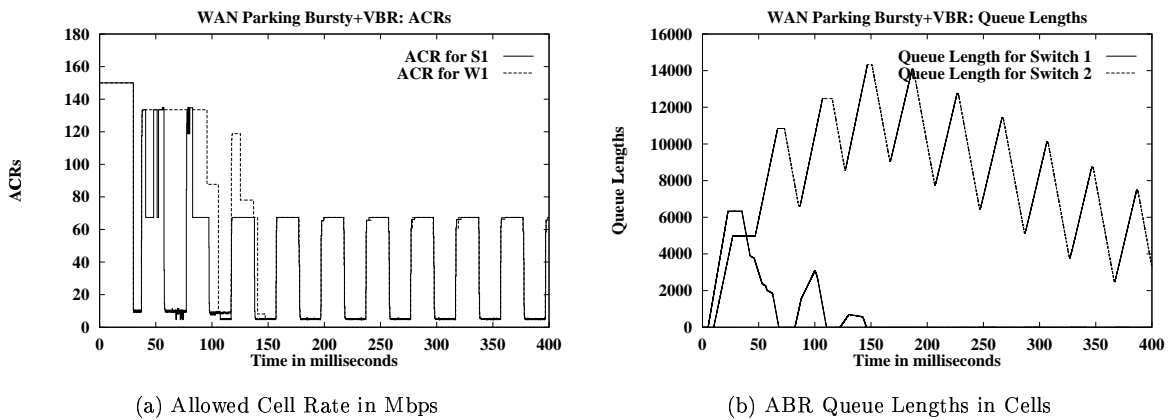



Figure 10: Results for WAN parking lot configuration with bursty, persistent and VBR connections [Algorithm 2]

feedback given in the initial 150 ms. This results in large queues (>5000 cells with every VBR burst) at Switch 2 on the port going to Switch 3 (Switch 3 has negligible queues and is therefore not shown here). The high initial cell rate (ICR) and rate increase factor (RIF) [6] values are the reason for the unusually large initial ABR queues seen for all algorithms.


Algorithm 4 gives more accurate feedback, but the first correct feedback is given after around 50 ms, which results in initially large queues (since ICR is large). Algorithms 5 and 6 produce identical results to algorithm 4, since the bottleneck link is attached to the branch point. Algorithm 7, on the other hand, exhibits a very fast transient response, and gives relatively accurate feedback to both sources. The initial queues caused by high ICR, as well as the queues with every VBR burst are much smaller. Hence, it offers the best performance since it combines the benefits of algorithm 4 with a fast transient response.

The chain configuration, illustrated in figure 16 consists of a point-to-multipoint connection (S1 to dS1, dS2 and dS3) where one of the links on the route to the farthest leaf is the bottleneck link (shared by the point-to-point connection SA to dSA). Also the link lengths increase by an order of magnitude in each of the last two hops (all links from the end systems to the switches are 50 km).

As seen in figures 17 through 23, this configuration is an ideal configuration for illustrating the consolidation noise problem.




(a) Allowed Cell Rate in Mbps




(b) ABR Queue Lengths in Cells

Figure 11: Results for WAN parking lot configuration with bursty, persistent and VBR connections [Algorithm 3]




(a) Allowed Cell Rate in Mbps




(b) ABR Queue Lengths in Cells

Figure 12: Results for WAN parking lot configuration with bursty, persistent and VBR connections [Algorithm 4]




(a) Allowed Cell Rate in Mbps




(b) ABR Queue Lengths in Cells

Figure 13: Results for WAN parking lot configuration with bursty, persistent and VBR connections [Algorithm 5]




(a) Allowed Cell Rate in Mbps




(b) ABR Queue Lengths in Cells

Figure 14: Results for WAN parking lot configuration with bursty, persistent and VBR connections [Algorithm 6]



(a) Allowed Cell Rate in Mbps



(b) ABR Queue Lengths in Cells

Figure 15: Results for WAN parking lot configuration with bursty, persistent and VBR connections [Algorithm 7]




Figure 16: Chain configuration




Figure 17: Results for a Chain configuration [Algorithm 1]

The problem is severe for algorithms 1, 2 and 3 (see figures 17 through 19), and results in rate oscillations, instability, unbounded queues, and unfairness against source SA. Switch 3 is the bottleneck in this configuration, as the link connecting Switch 3 to Switch 4 is the bottleneck link. The queues at Switch 1 and Switch 2 are negligible and hence are not visible on the graph (they are close to zero). The rate of SA remains at half of the bandwidth, while the rate of S1 continues to oscillate around a mean of about 103 Mbps. Although using a scheme such as ERICA+ leads to stability and bounded queues in this case, the persistent rate oscillations result in unacceptable performance and unfairness (*the problem can be mitigated by using small RIF values, but this slows down rate increases*). Algorithms 4, 5 and 6 (figures 20 through 22) avoid the noise completely, but suffer from a slow transient response. The rate of the source S1 only drops after around 60 ms, and by that

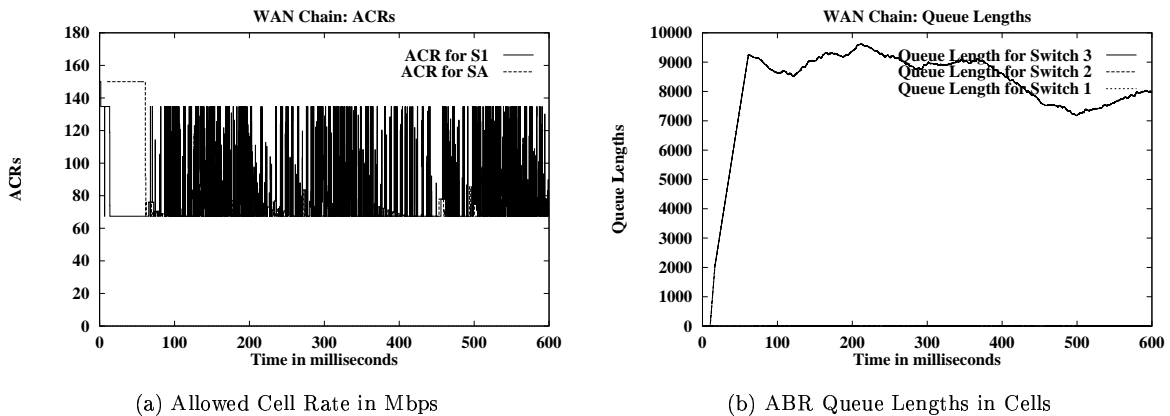
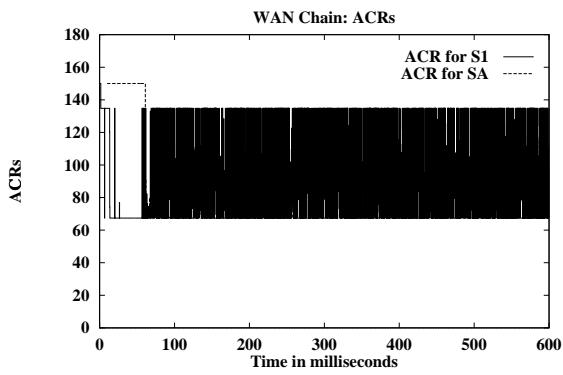



Figure 18: Results for a Chain configuration [Algorithm 2]




(a) Allowed Cell Rate in Mbps




(b) ABR Queue Lengths in Cells

Figure 19: Results for a Chain configuration [Algorithm 3]



(a) Allowed Cell Rate in Mbps



(b) ABR Queue Lengths in Cells

Figure 20: Results for a Chain configuration [Algorithm 4]




Figure 21: Results for a Chain configuration [Algorithm 5]

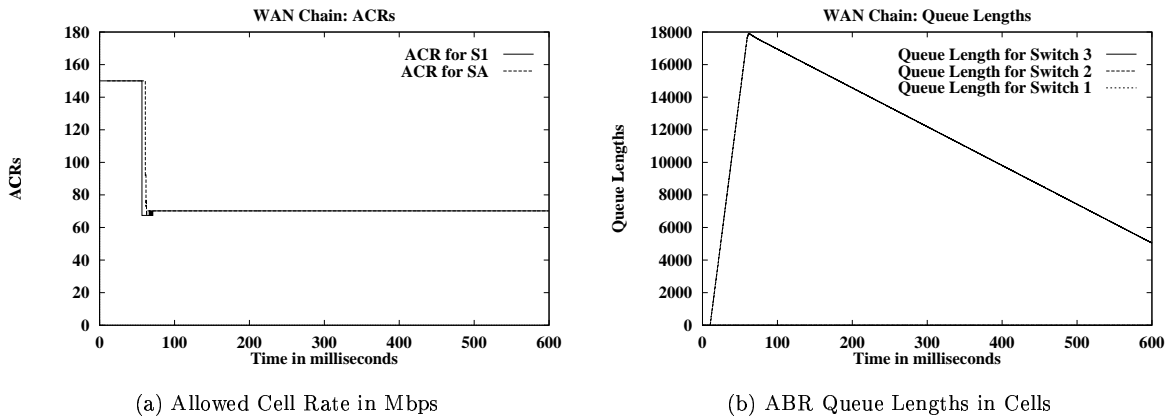


Figure 22: Results for a Chain configuration [Algorithm 6]

time, large queues have built up at the switches. Algorithm 7 yields optimal performance in this case, since the rate of the source S1 immediately drops to its optimal value, as soon as the overload is detected.

Observe that algorithms 5 and 6 also yield near optimal performance (like algorithm 7) if the destination dS3 was further than dS1, as in the configuration in figure 24 (all links from the end systems to the switches are 50 km except for dS3 which is 8000 km away from switch 2). This result is illustrated in figures 25 through 28 which compare the results for algorithms 4, 5, 6 and 7. Here, the chain configuration is modified such that the bottleneck link is closer to the branch point at switch Sw2 than another leaf, namely dS3.

In this case, as seen in figure 25, algorithm 4 wastes a long time waiting for feedback from dS3, while it has already received the bottleneck feedback from Switch 3. Algorithms 5, 6, and 7 send the feedback as soon as the overload situation is indicated by the BRM cell coming from switch Switch 3, and do not needlessly wait for the BRM from dS3. Hence, the 3 new algorithms perform near optimally since the rate of the source S1 goes to the optimal value after only around 20 ms for algorithms 5 and 6, and less than 10 ms for algorithm 7. The maximum queue length is also much smaller than for algorithm 4 (> 16000 cells): for algorithms 5 and 6, it is around 7000 cells, and for algorithm 7, it is less than 3500 cells.

We have observed a similar, but more pronounced, behavior when we simulated configurations with a larger number of leaves




Figure 23: Results for a Chain configuration [Algorithm 7]




Figure 24: Modified chain configuration

at varying distances and at varying levels of the multicast tree. The situation was much worse in those cases with algorithms 1, 2, and 3, which had much more severe noise problems. Algorithm 4 had an extremely slow transient response, while algorithms 5, 6, and especially 7 quickly reached the optimal values, and the queues at the switches were small.

10 Comparison of the Algorithms

This section summarizes the conclusions from the performance comparison of the algorithms. All the algorithms preserve the fairness and efficiency of the point-to-point congestion avoidance algorithm employed. We compare the space and time


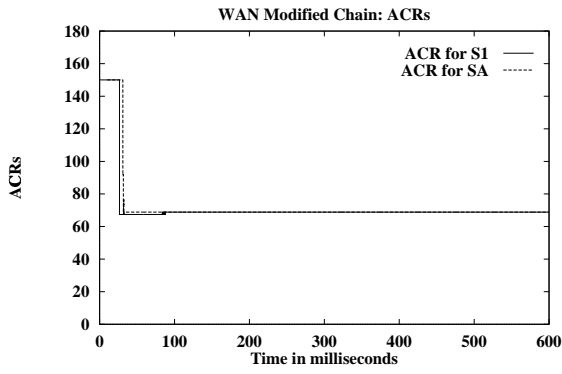



Figure 25: Results for a chain modified configuration [Algorithm 4]




(a) Allowed Cell Rate in Mbps




(b) ABR Queue Lengths in Cells

Figure 26: Results for a chain modified configuration [Algorithm 5]




(a) Allowed Cell Rate in Mbps




(b) ABR Queue Lengths in Cells

Figure 27: Results for a modified chain configuration [Algorithm 6]



(a) Allowed Cell Rate in Mbps



(b) ABR Queue Lengths in Cells

Figure 28: Results for a modified chain configuration [Algorithm 7]

complexity, transient response, consolidation noise, algorithm overhead and scalability, and discuss the interoperability of various algorithms.

10.1 Implementation Complexity

In algorithms 1 and 2, the main source of complexity is that the branch point has to turn around the RM cells. This is somewhat similar to the Virtual Source/Virtual Destination (VS/VD) concept. Most studies argue that turning around RM cells has a high implementation complexity.

Algorithm 3 is definitely the simplest algorithm to implement, since it does not turn around RM cells, and it keeps minimal per-VC accounting information. Algorithm 4 is more complex since it has to maintain the number of branches and the number of branches from which BRMs have been received, and compare those numbers. In addition, it has to maintain a bit for each output port to denote whether a BRM cell has been received from this branch, and some responsiveness detection algorithm-related values.

Algorithms 5 and 6 are slightly more complex since they may also store the last ER sent by the branch point. Alternatively, they can use the CCR of the source, which is already stored and used by most congestion avoidance algorithms (it is used in the ERICA algorithm which we have employed in this study). Hence, the additional complexity over algorithm 4 mainly stems from the comparison of the MER value to the last ER sent or the CCR value, and maintaining the SkipIncrease counter. This additional comparison and integer register do not incur much overhead.

Algorithm 7 is more complex than algorithms 5 and 6, since it invokes the ERICA algorithm for all the branches whenever a BRM cell is received by the branch point, and not only when a BRM cell is to be sent.

10.2 Transient Response

Algorithm 1 exhibits a very fast transient response. Algorithms 2 and 3 also have a reasonable transient response, since, even if there are no RM cells in the network, the feedback is quickly returned on the first BRM arrival.

Algorithm 4 has a slow transient response, since it waits for feedback from all the leaves before sending BRMs. This is especially severe in cases when there are few or no RM cells already in the network, such as during startup periods and for bursty sources. Therefore feedback can be delayed up to a function of the longest round trip times of the leaves. Algorithms 5, 6 and 7 tackle this problem for overload situations. The transient response of the schemes is very fast when an overload is detected downstream (for algorithms 5 and 6), or at this branch and downstream (for algorithm 7). In such cases, the transient response of the scheme is reasonably fast, and potential cell loss and retransmissions are alleviated.

10.3 Consolidation Noise

Algorithms 1, 2, and 3 suffer from severe consolidation noise problems. In particular, algorithms 1 and 3 suffer from unacceptable consolidation noise in some cases, especially with large RIF values (recall figures 17 and 19). Algorithm 2 somewhat alleviates these problems, since BRMs are not sent if no feedback has been received from any of the downstream components. However, it still exhibits considerable noise.

Algorithms 4, 5, 6, and 7 eliminate this problem by waiting for feedback from all branches. Although algorithms 5, 6, and 7 do not wait for feedback from all leaves in cases of overload, this *does not* introduce noise, since the RM cells that are sent faster than the usual cells carry overload information, which would have been conveyed by the next minimum value anyway.

10.4 Scalability Issues

Algorithms should be scalable in the sense that their overhead and feedback delay should not grow with the increase in the number of branch points or levels of the multicast tree.

10.4.1 RM cell overhead

The number of FRM cells generated by the source and the number of BRM cells received by the source should be approximately the same. Algorithm 1 generates a BRM cell at the branch point for every FRM cell it receives, thereby guaranteeing that the BRM to FRM ratio remains one. Algorithms 2 and 3 maintain a BRM to FRM ratio of less than or equal to 1 as follows. Algorithm 2 generates a BRM for an FRM only if a BRM has been received from a leaf since the last BRM was sent by the branch point. Algorithm 3 allows a BRM to pass to the source only if an FRM cell has been received by the branch point after the last BRM cell was forwarded. Therefore both algorithms maintain a ratio that is less than or equal to one (actually, it is strictly less than one for algorithm 2, since the first FRM cell will never be turned).

Algorithm 4 also maintains a ratio of less than or equal to one, since one BRM cell is returned when BRM cells have been received from all branches. Algorithm 5 does not guarantee that the ratio remains at 1, since RM cells carrying overload indication are allowed to quickly return to the source. Algorithms 6 and 7 fix this problem by maintaining a counter that is incremented for every extra RM cell passed, and then decremented (and the BRM cell discarded) in cases of RM cells carrying underload information, if the counter exceeds zero. Hence, over the long run, the ratio is maintained at one. The counter cannot increase indefinitely, since the rates cannot decrease indefinitely, but a maximum value can be enforced. In all cases we have examined, the counter value was always small, because ERICA quickly converged.

10.4.2 Delay sensitivity to the maximum number of branch points on a path (levels of the tree)

Algorithm 1 waits for an FRM cell to arrive before it can send the feedback information it has consolidated from the BRM cells. This has to be done at every branch point, leading to a delay that may increase with the number of levels of the multicast tree. Algorithm 2 suffers from the same drawback, since the algorithm also sends a BRM cell at the branch point when an FRM cell is received.

Algorithm 3 is less sensitive to the number of levels of the multicast tree. The BRM cell is passed to the source only if an FRM cell has been received since the last BRM cell was sent by the branch point. However, it is passed without additional delay.

Algorithms 4, 5, 6, and 7 can be sensitive to the multicast tree levels since BRM cells from all branches are consolidated at every branch point. However, the delay (the time between the transmission of the FRM cell at the source until the source receives the corresponding BRM cell) is mainly dependent on the round trip times from the source to the leaves at that particular time. The round trip times to the leaves can vary with time, dependent on the queuing delay of the switches on

the path of the multicast tree. More than one leaf can affect that delay since BRM cells arrive asynchronously at the branch points.

10.5 Interoperability Issues

The various consolidation algorithms should be able to interoperate with each other if no one algorithm is standardized. Although it seems that all the algorithms can interoperate smoothly with each other, the performance of a network with different algorithms at the different branch points, and point-to-multipoint VCs that branch at several branch points with different algorithms, needs further study. This is especially important because the consolidation algorithm is not standardized. This will be one of the areas of our future research work.

11 Conclusions

Table 2 shows a summary of the results of the comparison between the consolidation algorithms. Note that the main drawback of each algorithm is indicated in bold face. In terms of complexity, algorithm 3 is clearly the simplest. Algorithms 1 and 2 turn around RMs, which is an expensive operation. Algorithm 4 introduces additional complexity to algorithm 3, since it maintains per-branch variables and performs comparisons. Algorithm 5 introduces slightly more complexity to 4; algorithm 6 introduces simple additions to 5; and algorithm 7 introduces some more to 6, but most of the increments are of little complexity.

The transient response of algorithm 1 is fast, but can be erroneous. Algorithms 2 and 3 offer medium response, while algorithm 4 is clearly slow. Algorithms 5, 6, and especially 7, have a fast response when overload is detected. Consolidation noise is a problem with algorithms 1, 2, and 3, especially 1 and 3. The other algorithms overcome this problem.

As for RM cell overhead, the ratio of BRM cells received by the source to FRM cells sent by the source is maintained at unity by algorithm 1. It is less than one for algorithm 2 (at least the first FRM is not returned), and is less than or equal to one for algorithms 3 and 4. Algorithm 5 introduces additional BRM cells in case of overload, while algorithms 6 and 7 ensure the ratio is one over the long run (*lim* in the table means the limit as time goes to infinity).

Finally, the sensitivity of algorithms 1 and 2 to the number of branch points and the levels of the multicast tree is high due to the additional delay waiting for an FRM cell at each branch point, and the additional BRM cells that are turned around in the network at each level. Algorithms 3 to 7 (especially algorithm 3) are somewhat less sensitive to this.

The comparison indicates that algorithms 1 and 2 suffer from complexity and noise problems. Algorithm 3 is good, except for the consolidation noise problem which leads to unacceptable performance in some cases as seen in figure 19. Algorithm 4 provides reasonable performance, but has a slow transient response, which is overcome by the algorithms we proposed (5, 6 and 7). Algorithm 4 and the new algorithms are slightly more complex than algorithm 3, but this can be well worth the performance benefits gained, especially with algorithm 7. Algorithm 7 avoids congestion, while eliminating the consolidation noise problem.

12 Acknowledgments

This research was sponsored in part by Rome Laboratory/C3BC Contract # F30602-96-C-0156.

References

- [1] D. Cavendish, S. Mascolo, and M. Gerla. Rate based congestion control for multicast ABR traffic. In *Proceedings of IEEE GLOBECOM*, volume 2, pages 1114–1118, November 1996.
- [2] Y-Z Cho and M-Y Lee. An efficient rate-based algorithm for point-to-multipoint ABR service. In *Proceedings of IEEE GLOBECOM*, volume 2, pages 790–795, November 1997.
- [3] R. Jain et al. ERICA switch algorithm: A complete description. ATM Forum 96-1172, August 1996.
- [4] S. Fahmy et al. Performance analysis of ABR point-to-multipoint connections for bursty and non-bursty traffic with and without VBR background. ATM Forum 97-0422, April 1997.
- [5] S. Fahmy, R. Jain, R. Goyal, B. Vandalore, S. Kalyanaraman, S. Kota, and P. Samudra. Feedback consolidation algorithms for ABR point-to-multipoint connections. In *Proceedings of IEEE INFOCOM*, volume 3, pages 1004–1013, March 1998.
- [6] The ATM Forum. The ATM forum traffic management specification version 4.0. <ftp://ftp.atmforum.com/pub/approved-specs/af-tm-0056.000.ps>, April 1996.
- [7] D. Hunt. Open issues for ABR point-to-multipoint connections. ATM Forum 95-1034, August 1995.
- [8] T. Jiang, M. Ammar, and E. Zegura. Inter-receiver fairness: a novel performance measure for multicast ABR sessions. In *Proceedings of the ACM SIGMETRICS*, June 1998.
- [9] T. Jiang, E. W. Zegura, and M. Ammar. Improved consolidation algorithms for point-to-multipoint ABR service. In *Proceedings of IEEE Workshop on ATM*, May 1998.
- [10] W. Melody Moh. On multicasting ABR protocols for wireless ATM networks. In *Proceedings of the International Conference on Network Protocols (ICNP) '97, Atlanta*, pages 24–31, 1997.
- [11] W. Molody Moh, Yin Chen, and B. Niizawa. Branch-point algorithms for multicasting ATM ABR protocols. In *Proceedings of IEEE Workshop on ATM*, May 1998.
- [12] W. Ren, K-Y Siu, and H. Suzuki. On the performance of congestion control algorithms for multicast ABR service in ATM. In *Proceedings of IEEE ATM Workshop, San Francisco*, August 1996.
- [13] W. Ren, K-Y Siu, H. Suzuki, and M. Shinohara. Multipoint-to-multipoint ABR service in ATM. *Computer Networks and ISDN Systems*, 1998. To appear.
- [14] L. Roberts. Rate based algorithm for point to multipoint ABR service. ATM Forum 94-0772R1, November 1994.

- [15] K-Y Siu and H-Y Tzeng. Congestion control for multicast service in ATM networks. In *Proceedings of the IEEE GLOBECOM*, volume 1, pages 310–314, 1995.
- [16] H-Y Tzeng and K-Y Siu. On max-min fair congestion control for multicast ABR service in ATM. *IEEE Journal on Selected Areas in Communications*, 15(3), April 1997.
- [17] H. Wang and M. Schwartz. Comparison of multicast flow control algorithms over combined wireless/wired networks. In *Proceedings of the 3rd International Workshop on Mobile Multimedia Communications, Princeton*, September 1996. URL: <http://www.ctr.columbia.edu/~whyu/res.html>.
- [18] H. Wang and M. Schwartz. Performance analysis of multicast flow control algorithms over combined wired/wireless networks. In *Proceedings of IEEE INFOCOM, Japan*, volume 3, pages 1038–1045, April 1997.

All our papers and ATM Forum contributions are available through <http://www.cis.ohio-state.edu/~jain/>

13 Vitae

Sonia Fahmy received her MS degree in Computer Science in 1996 from the Ohio State University, where she is currently a PhD candidate. Her main research interests are in the areas of multipoint communication, traffic management, and performance analysis. She is the author of several papers and ATM Forum contributions. She is a student member of the ACM, the IEEE, and the IEEE Communications and Computer societies.

Professor **Raj Jain** is an active member of ATM Forum Traffic Management group and has influenced its direction considerably. He is a Fellow of IEEE, a fellow of ACM, and serves on the editorial boards of Computer Networks and ISDN Systems, Computer Communications, and the Journal of High Speed Networks. He is the author of two popular books: “FDDI Handbook: High Speed Networking using Fiber and Other Media” published by Addison-Wesley and “The Art of Computer Systems Performance Analysis” published by Wiley. His publications and ATM Forum contributions and can be found at <http://www.cis.ohio-state.edu/~jain/>

Rohit Goyal is a PhD. candidate with the Department of Computer and Information Science at the Ohio State University, Columbus. He received his BS in Computer Science from Denison University, Granville. His work has been published in several papers and ATM forum contributions. His other interests include Distributed Systems, Artificial Intelligence, and Performance Analysis.

Bobby Vandalore received his B.Tech degree in 1993 from Indian Institute of Technology, Madras and MS degree in 1995 from the Ohio State University, both in Computer Science. He is currently a PhD student at the Ohio State University. His main research interests are in the areas of multimedia communications, traffic management, and performance analysis. He is the author of several papers and ATM Forum contributions. He is a student member of the ACM and IEEE-CS.

Shivkumar Kalyanaraman is an Assistant Professor at the Department of Electrical, Computer and Systems Engineering at Rensselaer Polytechnic Institute in Troy, NY. He received his B.Tech from the Indian Institute of Technology, Madras, India in 1993, followed by M.S. and Ph.D. degrees in Computer and Information Science at the Ohio State University in 1994 and 1997 respectively. His research interests include traffic management, multicasting, Internet pricing, and performance

Table 1: Summary of consolidation algorithms

Algorithm	1	2	3	4	5	6	7
FRM turn-around	Branch point	Branch point	Destination	Destination	Destination	Destination	Destination
Wait for all BRMs	No	No	No	Yes	Yes for underload	Yes for underload	Yes for underload
Condition to trigger a BRM from branch point	FRM received	FRM received and a new BRM	BRM received and a new FRM	a new FRM and all BRMs	a new FRM and either all BRMs or overload	a new FRM and either all BRMs or overload, provided no extra	a new FRM and either all BRMs or overload (including current), provided no extra
RM ratio control	through FRM	through FRM	through BRM	through BRM	through BRM in some cases	explicit	explicit
Handling non-responsive branches	unnecessary	unnecessary	unnecessary	necessary	necessary for increase	necessary for increase	necessary for increase
Interacts with switch	No	No	No	No	No	No	Yes
Per-branch accounting	None	None	None	1 bit per branch	1 bit per branch	1 bit per branch	1 bit per branch

analysis of distributed systems. He is a co-inventor in two patents, and has co-authored several papers and ATM forum contributions. He is a member of IEEE-CS and ACM. Internet: <http://www.ecse.rpi.edu/Homepages/shivkuma>

Table 2: Comparison of consolidation algorithms

Algorithm	1	2	3	4	5	6	7
Complexity	High	High	Low	Medium	Medium	Medium	>Medium
Transient Response	Fast	Medium	Medium	Slow	Fast for overload	Fast for overload	Very fast for overload
Noise	High	Medium	High	Low	Low	Low	Low
BRM:FRM at Root	1	<1	≤ 1	≤ 1	may be >1	<i>lim</i> =1	<i>lim</i> =1
Delay sensitivity	High	High	Low	Medium	Medium	Medium	Medium