****

ATM Forum Document Number: ATM_Forum/98-0807

****

TITLE: Effect of RM cell interval on ABR feedback: A simulation study using OPNET

****

SOURCE: Rohit Goyal, Sonia Fahmy, Raj Jain, Bobby Vandalore
The Ohio State University,
Department of Computer and Information Science,
ous, OH 43210
ate.edu

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

Shobana Narayanaswamy
MIL3 INC,
3400 International Drive, NW
Washington, DC 20008
Phone: 202-364-4700
snaraya@mil3.com

The presentation of this contribution at the ATM Forum is sponsored by NASA.

****

DISTRIBUTION: ATM Forum Technical Committee
Traffic Management Working Group

****

DATE: July, 1998 (Portland)

****

ABSTRACT: In this contribution, we present an analysis of the effect of changing RM cell intervals on ABR performance. We describe our newly developed ABR model in OPNET. This OPNET ATM model contains enhanced features to support the QoS capabilities of ATM, and a comprehensive ABR feedback model. We describe the various features of the ATM model, and use it for our simulation and analysis.

****

****

Effect of RM cell interval on ABR feedback: A simulation study using OPNET

Rohit Goyal, Sonia Fahmy, Raj Jain, Bobby Vandalore

The Ohio State University,
Department of Computer and Information Science,
2015 Neil Ave, DL 395, Columbus, OH 43210
Phone: 614-688-4482
{goyal,fahmy,jain}@cis.ohio-state.edu

Shobana Narayanaswamy

MIL3 INC,
3400 International Drive, NW
Washington, DC 20008
Phone: 202-364-4700
Snaraya@mil3.com

ABSTRACT

In this contribution, we present an analysis of the effect of changing RM cell intervals on ABR performance. We describe our newly developed ABR model in OPNET. This OPNET ATM model contains enhanced features to support the QoS capabalities of ATM, and a comprehensive ABR feedback model. We describe the various features of the ATM model, and use it for our simulation and analysis.

1 Introduction

The ATM Forum has been investigating how to transport real-time multimedia applications over ATM networks. The SAA (Service Aspects and Applications) group at ATM Forum has approved specifications for transporting the MPEG-2 service [MPEG2] over ATM networks [VOD]. It is well known that bandwidth demands of video can be easily adjusted to meet the available bandwidth. In several situations, it may be cost-effective to adjust video quality to match the available bandwidth. There have been a limited number of studies that addressed the problem of transporting real-time video with feedback control. [LAKSH] shows how ABR explicit rate feedback can be used to transport compressed video. The video sources adapt to the required rate by modifying the quantization value of an MPEG compression algorithm. [KANAK1] discusses transporting packet video adaptively by using binary feedback. [KANAK2] proposes an adaptive congestion control scheme to transport packet video. Distributed feedback control can also be used to achieve fair bandwidth sharing among video sources. Recently, in [DUFF] an algorithm for transporting smoothed compressed video over explicit rate networks is given. In this study, a small number of frames are stored at the source and are used for smoothing the traffic. The rate adaptation is performed by using adaptive video encoding. In [VICKERS], multi-layered video source traffic is transported over a multicast network. The sources adapt to network congestion based on the feedback by adding or dropping video layers. In spite of these limited studies, video over feedback controlled networks is not a fully solved problem and a number of ideas remain to be explored.

ABR sources send an RM cell after every Nrm-1 (usually Nrm = 32) cells. The sources adjust their rate when they receive these backward RM (BRM) cells. At high data rates, a low RM cell interval can result in a high frequency rate variations in the ABR feedback. One of the goals of transporting video over ABR is to minimize the rate variations, which in turn will reduce variations in the quality of service. Users want a constant quality of service in a real-time application such as real-time video. Hence, it is necessary to reduce the rate variations to provide low variations in quality of service.

One way of reducing the ABR rate changes is to send RM cells less frequently, i.e., Nrm should be large, instead of 32. Sending RM cells at end of each video frame is one possible option. Another method to reduce variation is to increase the length of the averaging interval which some switch algorithms, such as the ERICA algorithm, use.

This contribution has two main goals:

1. To present a preliminary study of the impact of varying the Nrm values on ABR performance.

2. To present a new ATM model in OPNET that is used for these experiments.

2 The OPNET Model

OPNET is a modeling and simulation tool [MIL31] that provides an environment for analysis of communication networks. The tool provides a three layer modeling hierarchy. The highest layer, referred to as the network domain, allows the definition of network topologies. The second layer, referred to as the node domain, allows definition of node architectures (data flow within a node). The third layer (process domain) specifies logic or control flow among components in the form of a finite state machine.

3 The OPNET ATM Model Suite

The OPNET ATM model suite (AMS) described in [MIL32] supports many of the characteristics of ATM networks. The model suite provides support for signaling, call setup and tear-down, segmentation and re-assembly of cells, cell transfer, traffic management and buffer management. Standard ATM nodes such as routers, stations, bridges and switches are provided to facilitate building of common topologies used for the design and analysis of ATM networks.

Traffic management within AMS incorporates functions such as call admission control, policing using a continuous-state leaky bucket implementation (GCRA), call-based queuing, priority scheduling and collection of standard statistics such as end-to-end delay and end-to-end delay variation.

Reference Topology

The example network topology used for the design and development of traffic management functions within AMS represents an N-source configuration shown in Figure 1. Source and destination end-systems are connected to a pair of ATM switches that communicate via a bottleneck link.

The node architecture for the end-system (source/destination) consists of AAL clients sending/receiving traffic to/from the AAL/ATM/PHY protocol stack. The AAL layer is responsible for segmentation of data traffic into AAL PDUs. The ATM layer (represented as four modules: management, layer, translation and switching) segments the AAL PDU into ATM cells and transmits the cells to the network. The management module is responsible for signaling. The translation module receives incoming traffic and directs it to the higher layer or back to the network based on the destination address.

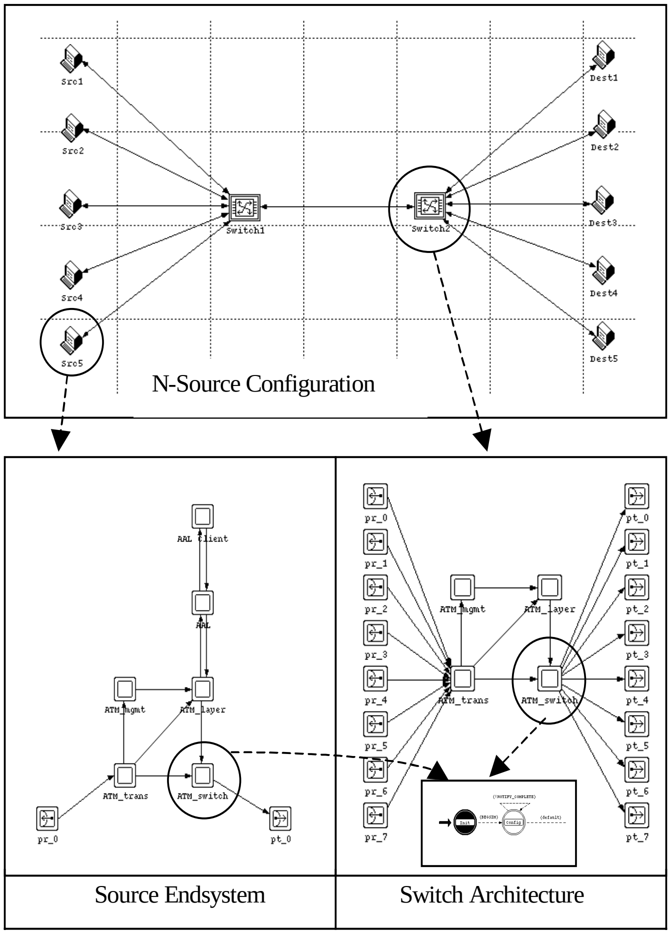

Figure 1 The OPNET ATM Model

The node architecture for the switch consists of the ATM layer functions modeled as four modules as described above. The switch can have several input and output ports. ABR traffic management and feedback functions are implemented within the ATM switch module in the form of a finite state machine.

Specification of QoS

AAL clients representing traffic sources specify their QoS requirements using the application traffic contract attribute. This requirement is a combination of service category, traffic parameters and QoS parameters that the source would like the network to provide for both incoming and outgoing directions. Traffic parameters include the PCR, MCR, SCR and MBS. QoS parameters include the CTD, CDV and CLR for both directions.

Service Category	Requested Traffic Parameters	Requested QoS Parameters

Application Traffic Contract

In order to be able to provide the requested QoS for a connection, intermediate devices may be configured to support various QoS levels. The switch buffer configuration attribute allows specification of QoS levels for each buffer. Cell streams belonging to different QoS levels may be buffered and serviced according to their QoS. The buffer configuration defines the buffer size, the maximum allocated bandwidth and minimum guaranteed bandwidth. The supported traffic parameters include PCR, MCR, SCR and MBS. The supported QoS parameters include CTD, CDV and CLR.

Buffer Configuration	Supported Traffic Parameters	Supported QoS Parameters

Switch Buffer Configuration

ABR Traffic Management in OPNET

ABR mechanisms allow the network to divide the available bandwidth fairly and efficiently among the active traffic sources. In the ABR traffic management framework, the *source end systems* limit their data transmission to rates allowed by the network. The network consists of switches that use their current load information to calculate the allowable rates for the sources. These rates are sent to the sources as feedback via *resource management (RM)* cells. The ABR traffic management model is a *rate-based end-to-end closed-loop* model.

There are three ways for switches to give feedback to the sources. First, each cell header contains a bit called Explicit Forward Congestion Indication (EFCI), which can be set by a congested switch. Such switches are called *binary* or *EFCI* switches. Second, RM cells have two bits in their payload, called the Congestion Indication (CI) bit and the No Increase (NI) bit, that can be set by congested switches. Switches that use only this mechanism are called relative rate marking switches. Third, the RM cells also have another field in their payload called explicit rate (ER) that can be reduced by congested switches to any desired value. Such switches are called Explicit Rate switches. RM cells are generated by the sources and travel along the data path to the *destination end systems*. The destinations simply return the RM cells to the sources.

Switches can use the virtual source/virtual destination (VS/VD) feature to segment the ABR control loop into smaller loops. In a VS/VD network, a switch can additionally behave both as a (virtual) destination end system and as a (virtual) source end system. As a destination end system, it turns around the RM cells to the sources from one segment. As a source end system, it generates RM cells for the next segment. This feature can allow feedback from nearby switches to reach sources faster, and allow hop-by-hop control.

At the time of connection setup, ABR sources negotiate several operating parameters with the network. The first among these is the peak cell rate (PCR). This is the maximum rate at which the source will be allowed to transmit on this virtual circuit (VC). The source also requests a minimum cell rate (MCR) which is the guaranteed minimum rate. The network has to reserve this bandwidth for the VC. During the data transmission stage, the rate at which a source is allowed to send at any particular instant is called the allowed cell rate (ACR). The ACR is dynamically changed between MCR and PCR. At the beginning of the connection, and after long idle intervals, ACR is set to initial cell rate (ICR).

Most resource management cells generated by the sources are counted as part of their network load in the sense that the total rate of data and RM cells should not exceed the ACR of the source. Such RM cells are called ‘‘in-rate” RM cells. Under exceptional circumstances, switches, destinations, or even sources can generate extra RM cells. These ‘‘out-of-rate” RM cells are not counted in the ACR of the source and are distinguished by having their cell loss priority (CLP) bit set, which means that the network will carry them onl tif there is plent tof bandwidth and can discard them if congested. The out-of-rate RM cells generated by the source and switch are limited to 10 RM cells per second per VC. One use of out-of-rate RM cells is for BECN from the switches. Another use is for a source, whose ACR has been set to zero by the network, to periodically sense the state of the network. Out-of-rate RM cells are also used by destinations of VCs whose reverse direction ACR is either zero or applies only to RM cells. All data cells in ABR should have CLP set to 0 and must always be within the rate allowed by the network.

Resource Management cells traveling from the source to the destination are called Forward RM (FRM) cells. The destination turns around these RM cells and sends them back to the source on the same VC. Such RM cells traveling from the destination to the source are called Backward RM (BRM) cells. Note that when there is bi-directional traffic, there are FRMs and BRMs in both directions on the VC. A direction bit (DIR) in the RM cell payload indicates whether it is an FRM or BRM.

The ERICA Switch Scheme Implementation in OPNET

The ERICA algorithm [SHIV] operates at each output port (or link) of a switch. The switch periodically monitors the load on each link and determines a load factor (z), the available ABR capacity, and the number of currently active virtual connections or VCs (N). A measurement or ‘‘averaging” interval is used for this purpose. These quantities are used to calculate the feedback which is indicated in RM cells. The feedback is given to the RM cells travelling in the reverse direction. Further, the switch gives at most one new feedback per source in any averaging interval. The key steps in ERICA are as follows:

At the End of at Averaging Interval, total ABR Capacity is computed as the difference between the link capacity and the bandwidth used by higher priority traffic. The Target ABR Capacity is then computed as a fraction (typically a function of the queuing delay) of the total ABR capacity. The overload (z) and the fair share (FS) are calculated as:

```
z  ← ABR Input Rate / Target ABR Capacity
FS ← Target ABR Capacity / N
```

Where N is the number of active VCs. The maximum allocations given in the previous and current intervals are maintained as:

```
MaxAllocPrevious ← MaxAllocCurrent
MaxAllocCurrent ← FS
```

When an FRM is received, the Current Cell Rate (CCR) in the RM cell is noted for the VC:

```
CCR[VC] ← CCR_in_RM_Cell
```

When a BRM is received
Feedback is calculated as follows and inserted in the ER field of the cell:

```
VCShare ← CCR[VC] / z
IF z > 1+ Δ
THEN ER ← Max (FairShare, VCShare)
ELSE ER ← Max (MaxAllocPrevious,
                FairShare, VCShare)
MaxAllocCurrent ← Max(MaxAllocCurrent,ER)
IF (ER > FairShare AND CCR[VC]<FairShare)
      THEN ER ← FairShare
ER in RM Cell ← Min (ER in RM Cell, ER,
                      Target ABR Capacity)
```

Details of the ERICA algorithm are available from [SHIV].

The Switch Process Model

The OPNET process modeling methodology was used in the development of the switch process model that delivered basic capabilities of the core ATM switching fabric, ABR feedback control, buffer management and scheduling. The key steps of this modeling methodology include: definition of the system context, identifying interdependent modules, enumeration of events, selection of states of a process, construction of an event response table and construction of the finite state machine. The development of the OPNET switch process model is described in the paragraphs below. A simple switch process receives cells on its input port. Cells are switched via the switching fabric to an output port based on its destination address. Cells may be enqueued at the output port and transmitted based on a scheduling algorithm. The functionality of a simple switch is illustrated in Figure 2.

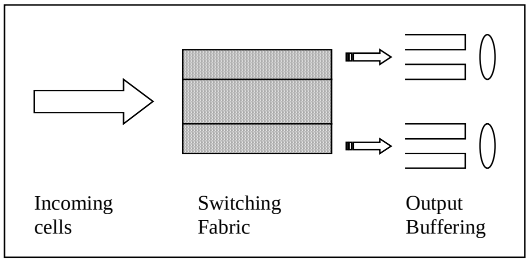

Figure 2. A Simple ATM Switch

Logical events that can occur at the switch include a cell arrival, time to transmit as indicated by the scheduler, and the end of fabric delay. Table 1 enumerates the events that can occur at the switch and the associated interrupt types.

Table 1. ABR switch events

Logical Event	*Event Description*	*Interrupt Type*
Cell_Arrival	Arrival of an ATM cell at the switch	Stream
Time_to_send	Indication from the scheduler that it is time to transmit a cell	Self
End_of_fabric _delay	Indication that a cell has completed processing via the switch fabric	Self

Table 2 outlines the actions taken when an event occurs within the switch. Each row of this table represents a combination of a state and an event and their associated conditions. Different actions performed for each combination and the resulting next state are listed.

Table 2. ATM ABR event response table

Current State	Logical Event	Condition	Action	Next State
None	Begsim	None	None	Init
Init		None	Initialize	Config
Config	Cell_Arrival	Neighbor notification not complete	Queue cell	Config
	Notify Complete	None	Process enqueued cells	Wait
Wait	Cell_arrival	Application traffic	Apply source rules before enqueue, schedule fabric delay	Wait
		Link Traffic and VSVD_ON	Apply destination rules, apply source rules before enqueue, schedule fabric delay	Wait
		Link_Traffic and VSVD_OFF	Schedule fabric delay	Wait
	End_of_fabric_delay	Cell can be buffered	Enqueue cell, activate scheduler	Wait
		Cell cannot be buffered	Destroy cell	Wait
	Time_to_send	More cells waiting to be sent	Dequeue and send cell, re-activate scheduler	Wait
		No more cells waiting	None	Wait

Figure 3 illustrates the state machine obtained as a result of the event response table. Multiple state machines are used for modularity. The *Init* state is entered when the process receives a *begin simulation* interrupt. The switch buffer configuration specified by the user and the ABR attributes (VSVD mode, feedback scheme) are obtained and initialization functions are executed. The ATM models go through a configuration phase where topologies and interconnections are verified. The process then goes into the *wait* state where all subsequent processing of cells takes place. When a cell arrives, the process checks if this is application traffic arriving from the higher layer or if this is link traffic. Application traffic for an ABR connection goes through the source rules described in [JAIN]. Link traffic for ABR connections goes through destination rules [JAIN] if the VSVD mode is ON. Otherwise, it goes through the switching fabric. Once all cells have been through the switching fabric, they are processed by the output buffer management function where they may be enqueued or dropped. A scheduler sends out cells from the buffers onto the link based on the cell rate for the connection and the scheduling scheme.

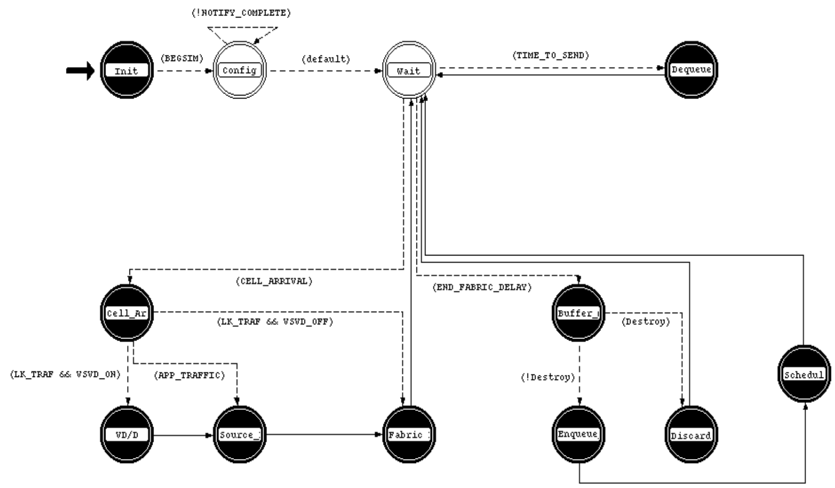

Figure 3. OPNET ATM ABR state machine implementation

4 The Simulation Experiment

We use the OPNET model previously discussed to analyze the effect of the RM cell frequency on ABR feedback. We vary Nrm and examine the allowed cell rates at the sources, as well as the queue lengths at the switches, the link utilizations and the throughput at the destinations. Since the Nrm value must be a power of two that is allowed to range between 2 and 256 (according to the current specifications), we have conducted experiments with all the allowed Nrm values (2, 4, 8, 16, 32, 64, 128 and 256). However, we only show the simulation results for Nrm = 8, 32 and 256 here. The reason why we have selected these values it that values smaller than 8 incur a very high control cell overhead and are not very realistic. 32 is the default value, and 256 is the maximum allowed value.

In our simulations, all links are 155.52 Mbps links. The initial cell rate (ICR) of all sources is set to 150 Mbps, while the remaining the ABR parameters are set to their default values as specified in the specifications. In particular, note that the value of the rate increase factor (RIF) parameter is set to 1/16. The ERICA switch averaging interval is set to a fixed time of 5 ms. ERICA target utilization is set to 90% of the link capacity.

Figure 4 and Figure 5 illustrate the test configurations used for the study. Figure 4 shows a configuration that consists of two ABR sources. Source 1 sends data from t=0.5 sec to t=1.5 sec, while source 2 is a transient source that comes on at t=0.7 sec and sends data for about 200 ms. The entire simulation time is 1.5 secs, where the first 0.5 secs are used to exchange OPNET signaling messages for connection setup. Both sources send persistent traffic at 150 Mbps. The grid shown in the figure is spaced at 1000 km, i.e., link lengths are close to 900 km or so, and hence the round trip time is around 23 ms. The main aim of this configuration is to test how the responsiveness of the system is affected by the Nrm value, both during rate increases and rate decreases.

Figure 5 shows the fairness configuration used to test the effect of Nrm in the presence of an upstream bottleneck. The grid shown in the figure is spaced at 100 km. As seen in the figure, the first link is shared by 15 connections, while the second link is shared by 3 connections. This configuration illustrates how the capacity left over by the connection bottlenecked upstream (Source 1 to Destination 1) is shared by the two non-bottlenecked connections between Switch 2 and Switch 3 (16 and 17). In this configuration, sources 1 through 15 are not sending at full load, but are bottlenecked at 10 Mbps. Hence, although the initial ACR values are high, the initial load at switch 1 is close to the ideal. Sources 16 and 17 start sending at 100 Mbps load, so Switch 2 is initially overloaded. All sources start transmission after 0.5 seconds, and the simulation time is one second.

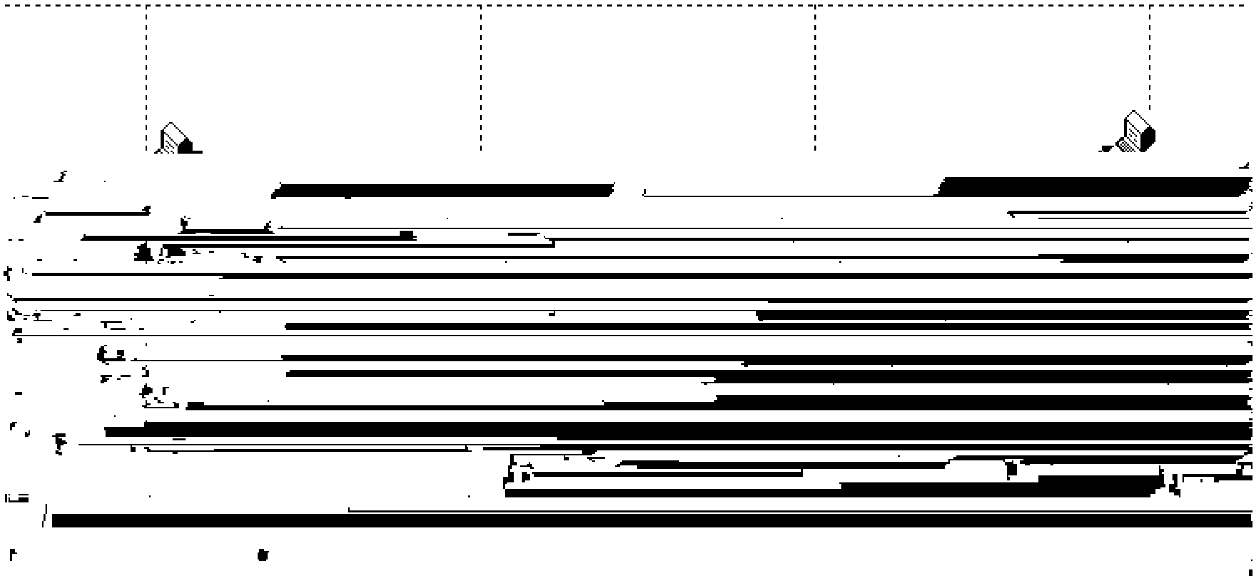

Figure 4 Two Source Transient Configuration

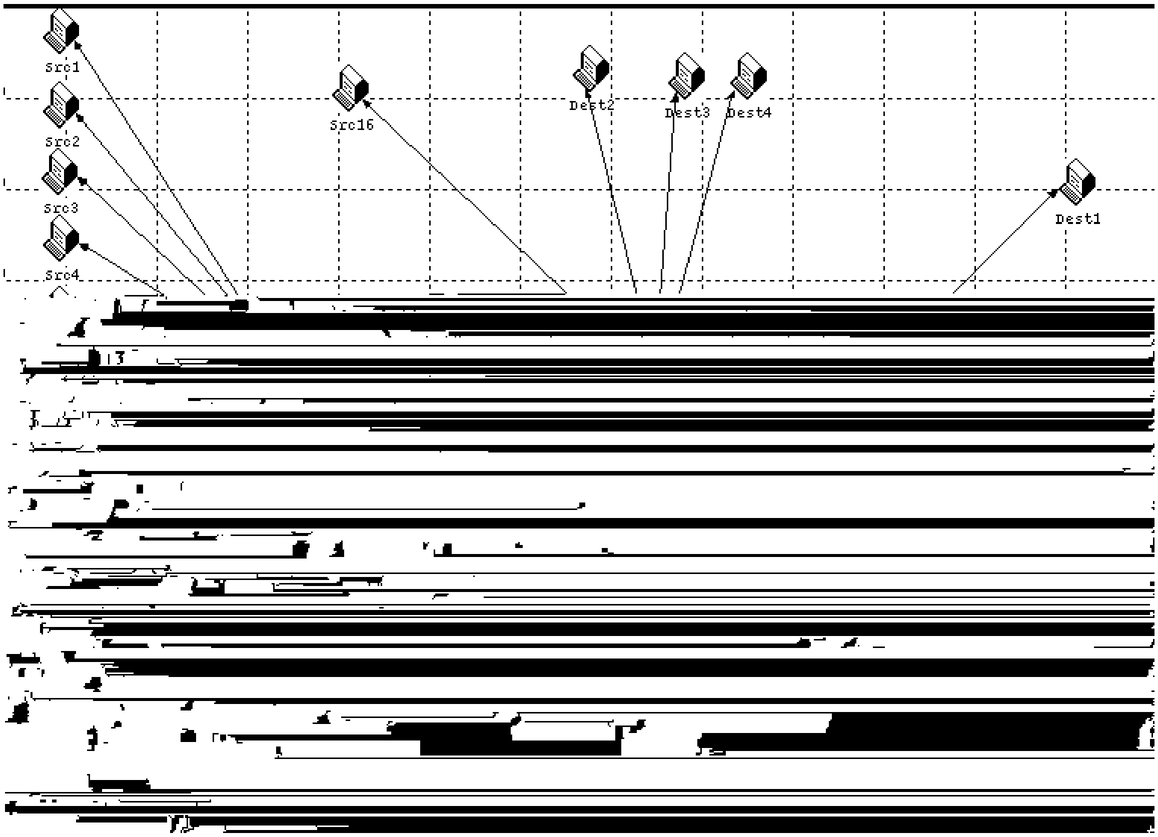

Figure 5 The Fairness Configuration

5 Simulation Results

ABR performance for the two source transient configuration is shown in the figures below (see Figure 6, Figure 7, and Figure 8 for Nrm=8, 32, and 256 respectively). The figures show the ACRs of the two sources, the link utilization at the bottleneck link and the queue length for switch 1. Both sources start at an ICR of 155.52 Mbps. In all cases, source 1 ACR quickly comes down to its target value of about 140 Mbps. When source 2 starts to send data, the ACRs of both sources are brought down to 70 Mbps. When source 2 stops sending data, the ACR for source 1 comes back up to 140 Mbps. There is a difference in the rate of increase of ACR for the three Nrm values. Since RIF is set to 1/16, the ACR comes up in steps on the receipt of every BRM cell. With Nrm=8, the source receives BRMs more frequently than with Nrm=256. As a result, the ACR for source 1 reaches 140 Mbps faster for Nrm=8.

The overhead with small Nrm values is quite high, however. This can be clearly observed by measuring the throughput at the application layer at the destinations (these plots are not shown here). Another interesting observation is that for smaller Nrm values, Source 1 does not *start* rising as fast as with larger Nrm values because the high RM cell overhead causes the data of the second source to take a longer time to be transmitted, and hence the two sources must share the bottleneck link for a longer time.

Figures 9, 10 and 11 show the results for the upstream bottleneck configuration for Nrm=8, 32, and 256 respectively. The figures show the ACRs of the three sources that share the link between Switch 2 and Switch 3 (sources 1, 16 and 17), the link utilization at the link between Switch 2 and Switch 3, and the queue length for Switch 2. Again, all sources start at an ICR of 155.52 Mbps, but the actual load for sources 1 through 15 is only 10 Mbps. Clearly, the ACRs for all three sources shown stabilize at the correct rates faster for smaller Nrm values, since the RM cells are more frequent. For source 1, since all the sources sharing the link between Switch 1 and 2 are only sending at 10 Mbps, the ACR values decrease slowly, since the load factor is small (starts at around 1.3 and eventually decreases till it is very close to 1). It is clear that the ACR for source 1 reaches its correct value much more rapidly with smaller Nrm values due to the high frequency of RM cells that convey to the switch the ACRs of the sources (in the CCR field of RM cells), and convey to the sources the explicit rate computed by the switches. The ACRs of sources 16 and 17 also converge much faster with smaller Nrm values.

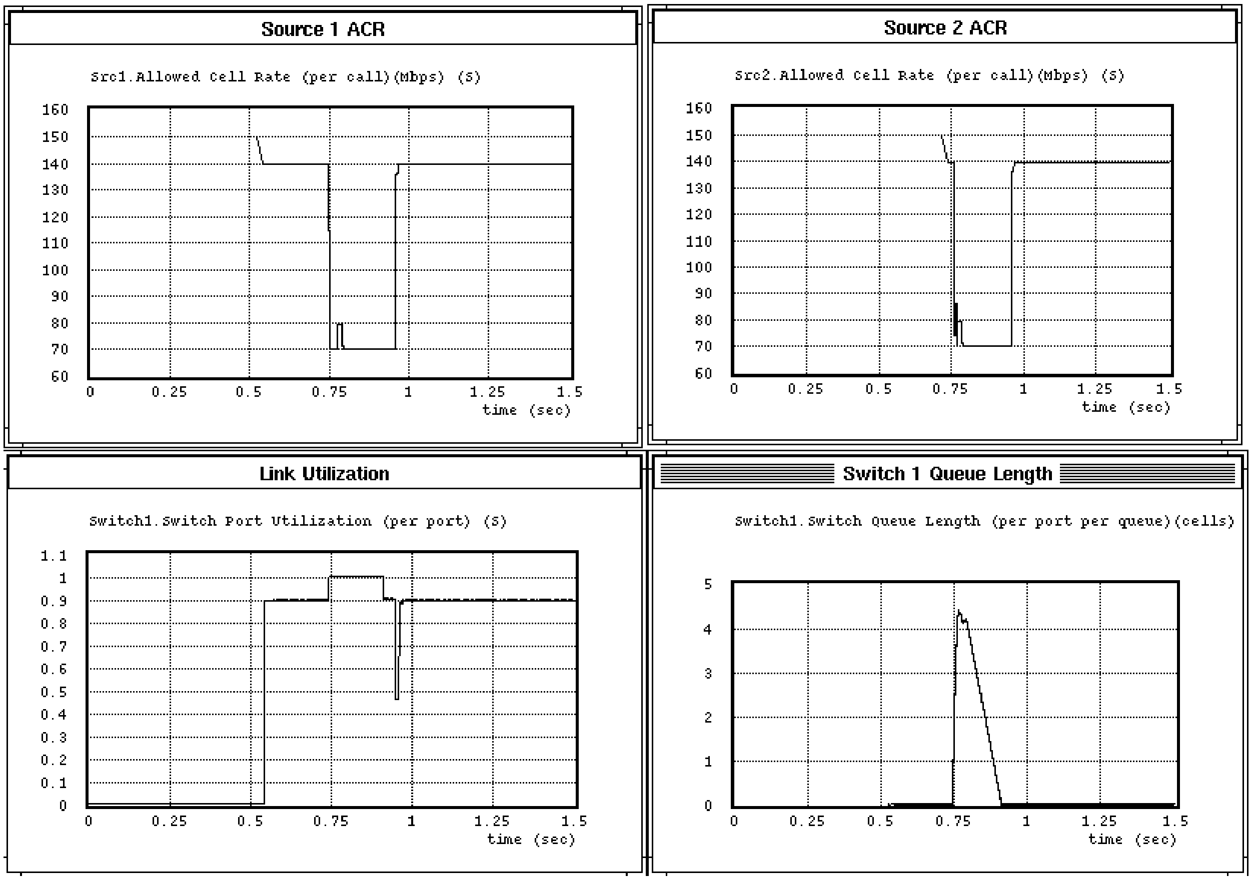

Figure 6 Transient configuration, Nrm=8

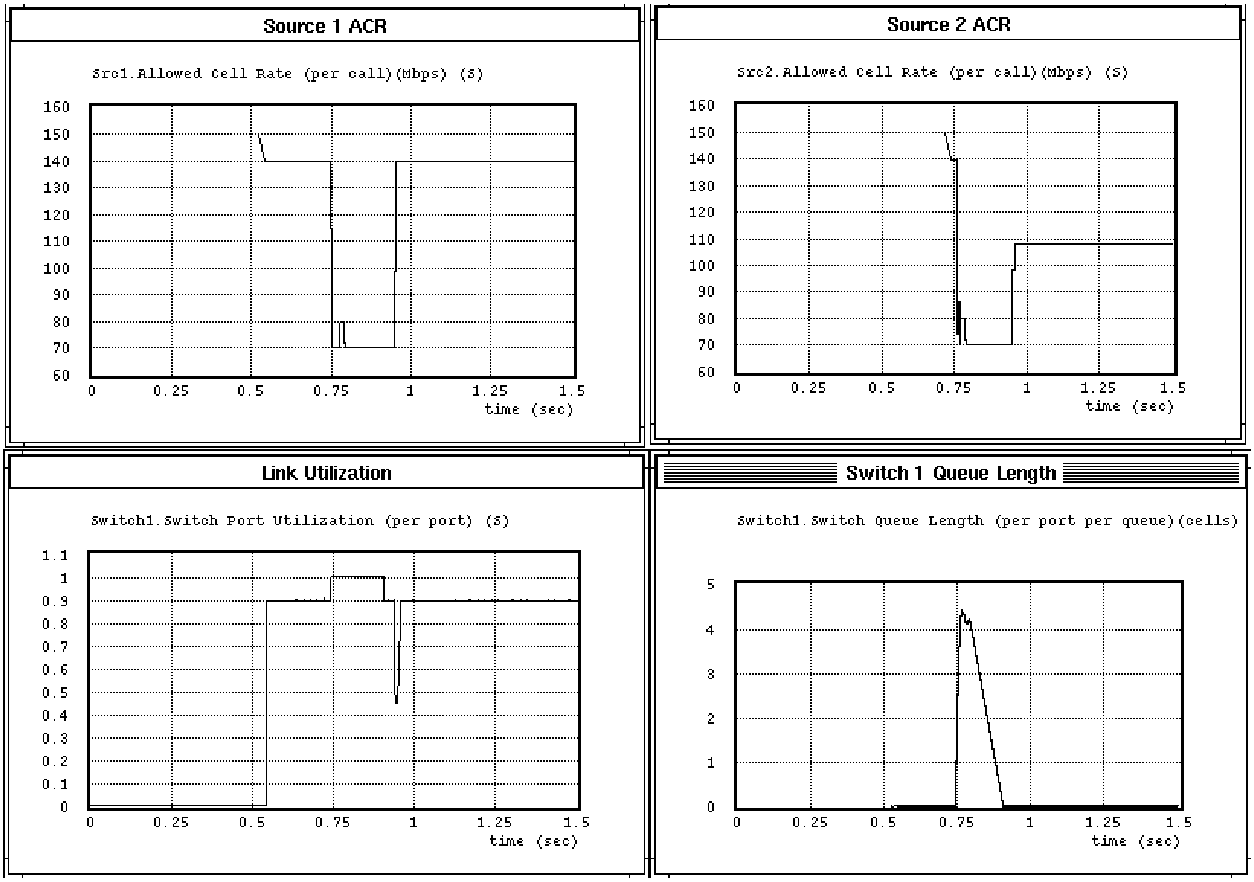

Figure 7 Transient configuration, Nrm=32

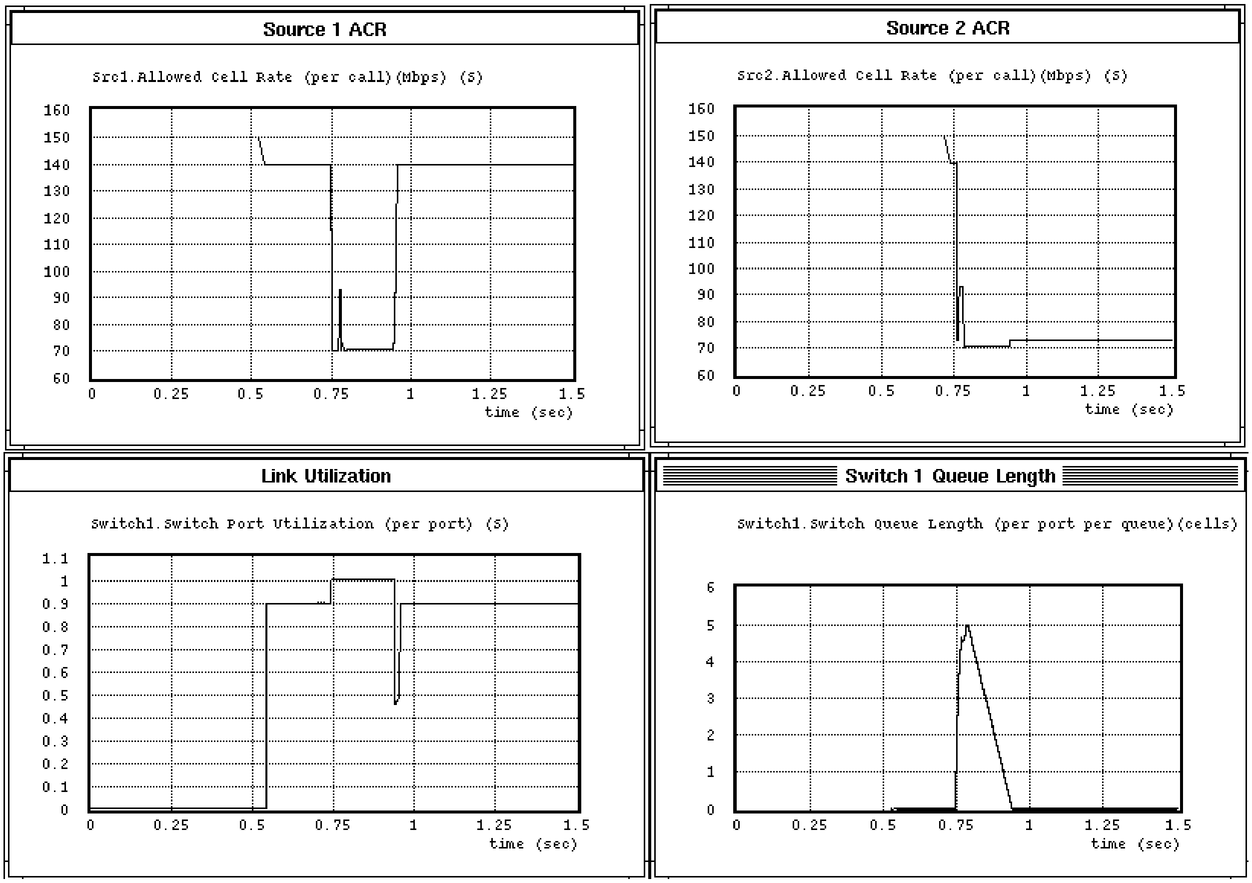

Figure 8 Transient configuration, Nrm=256

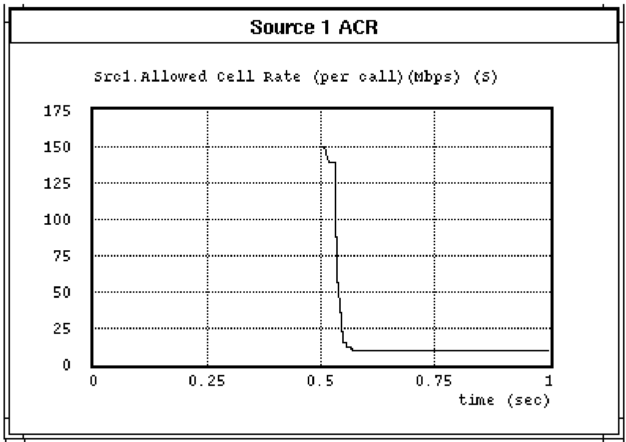
Source 1 ACR

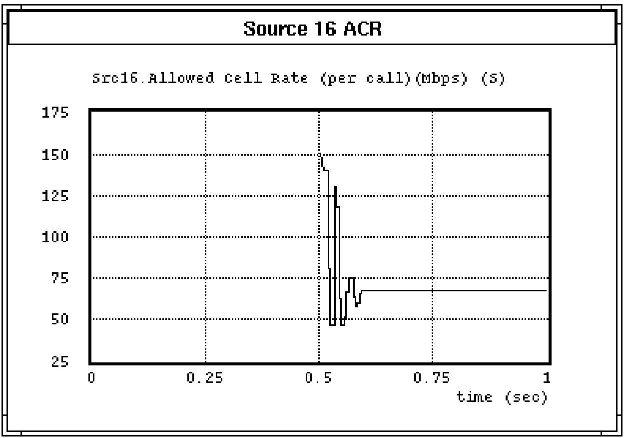
Source 16 ACR

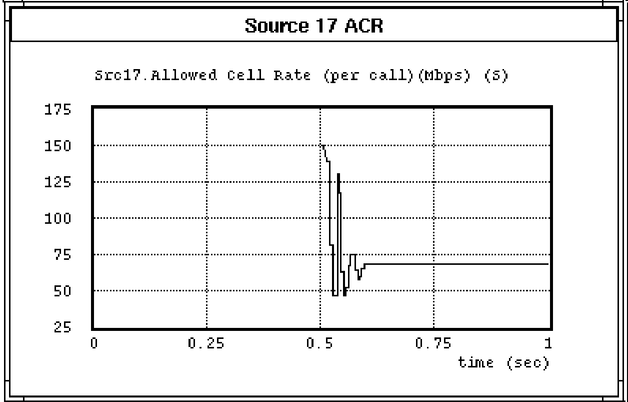
Source 17 ACR

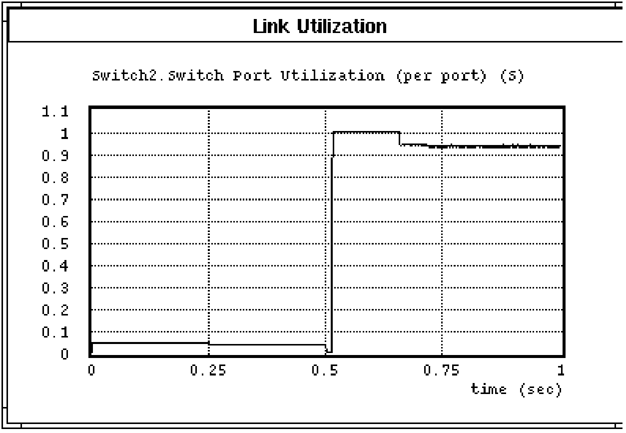
Link Utilization

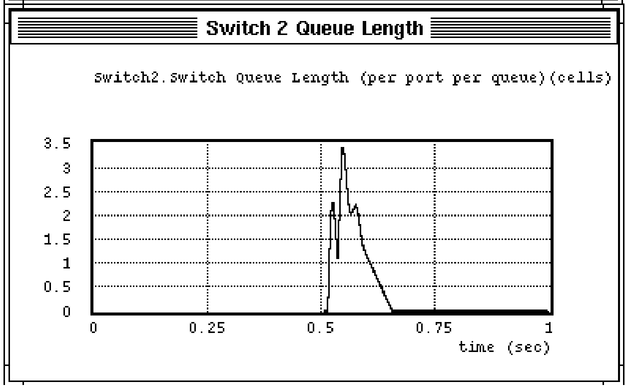
Switch 2 Queue Length

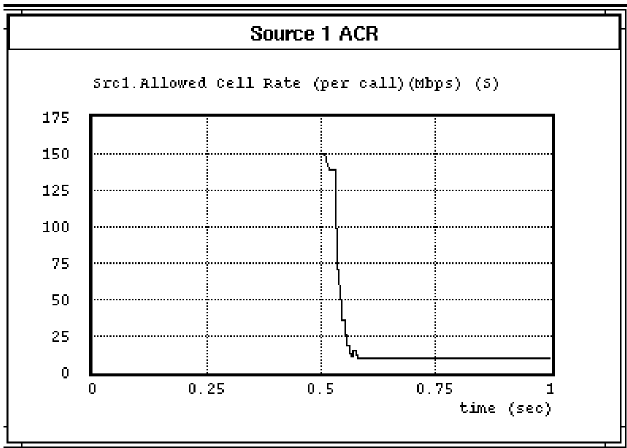
Source 1 ACR
Src1.Allowed Cell Rate (per call)(Mbps) (S)
175
150
125
100
75
50
25
0
0
0.25
0.5
0.75
1
time (sec)

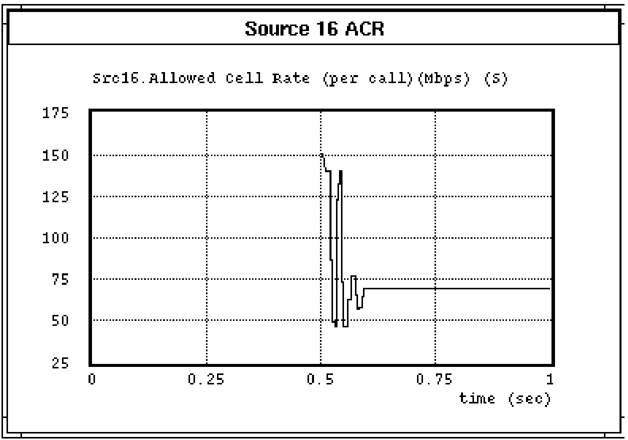
Source 16 ACR
Src16.Allowed Cell Rate (per call)(Mbps) (S)
175
150
125
100
75
50
25
0
0.25
0.5
0.75
1
time (sec)

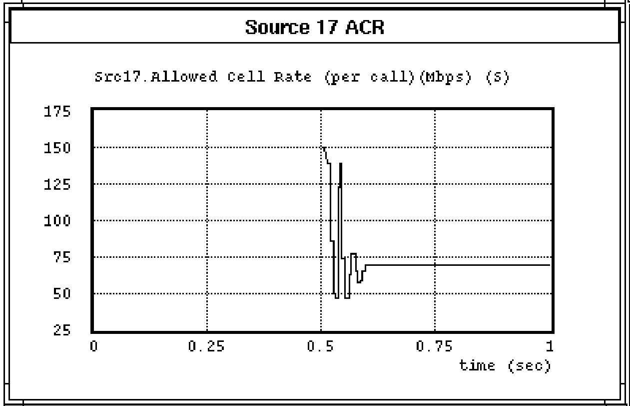
Source 17 ACR
Src17.Allowed Cell Rate (per call)(Mbps) (S)
175
150
125
100
75
50
25
0
0.25
0.5
0.75
1
time (sec)

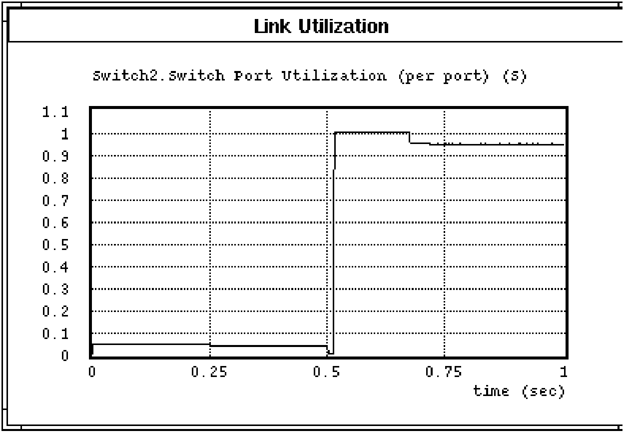
Link Utilization
Switch2.Switch Port Utilization (per port) (S)
1.1
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
0.25
0.5
0.75
1
time (sec)

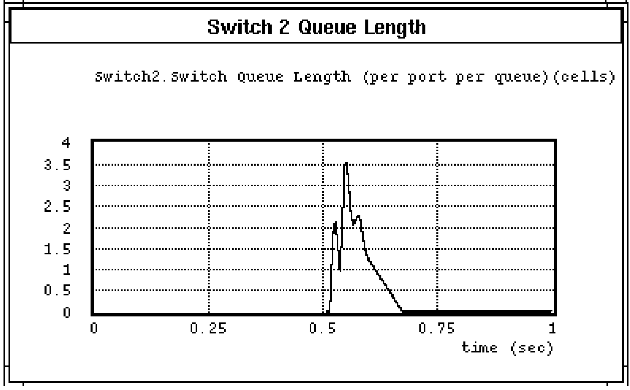
Switch 2 Queue Length
Switch2.Switch Queue Length (per port per queue)(cells)
4
3.5
3
2.5
2
1.5
1
0.5
0
0
0.25
0.5
0.75
1
time (sec)

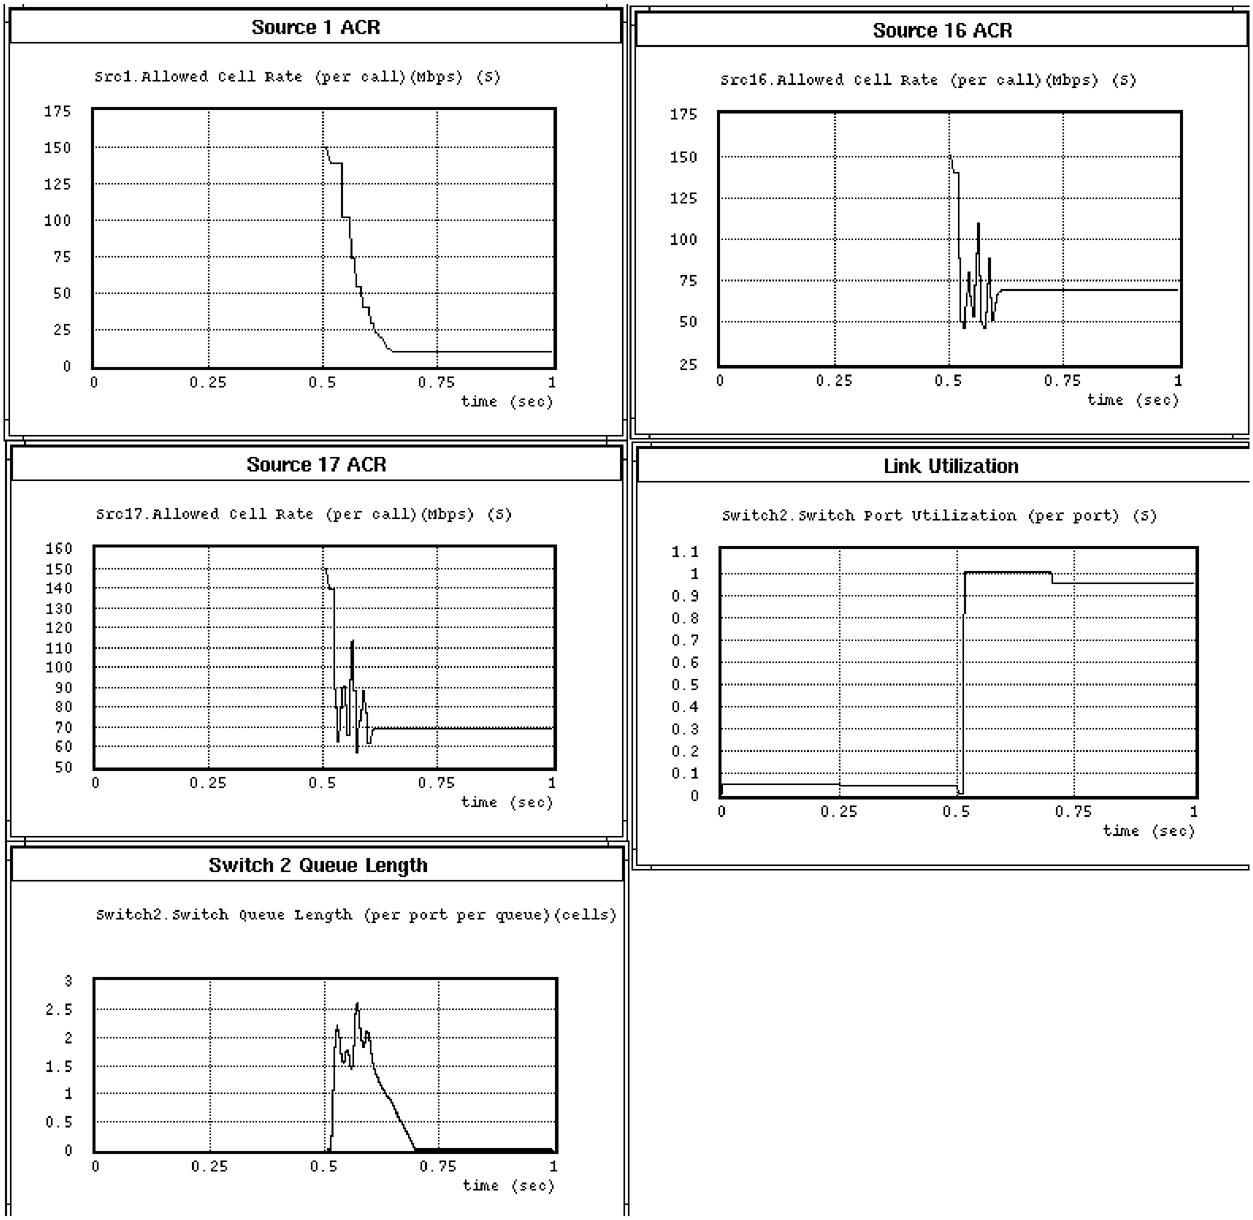

6 Summary

We have discussed the ATM ABR traffic management model, and its implementation in OPNET. This model will replace the existing ATM model in OPNET. Simulation results on ATM ABR performance show that the rate changes and throughput values are affected by the value of the ABR parameter Nrm. The Nrm value affects the control cell overhead, as well as the responsiveness of the ABR congestion avoidance mechanism to changing conditions.

7 References

[DUFF] N.G. Duffield, K. K. Ramakrishnan, A. R. Reibman, “SAVE: An algorithm for smoothed adaptive video over explicit rate networks”, In Proc. of INFOCOMM, April 1998.

[JAIN] Raj Jain et. al, "Source Behavior for ATM ABR Traffic Management: An Explanation," IEEE Communications Magazine, November 1996.

[KANAK1] H. Kanakia, P. Mishra, A. Reibman, “Packet Video transport in ATM networks using single-bit binary feedback”, In Proc. of Sixth International Workshop on Packet Video, Portland, Oregon, Sept. 1994.

[KANAK2] H. Kanakia, P. Mishra, A. Reibman, “An adaptive congestion control scheme for real-time packet video transport”, In Proc. of ACM SIGCOMM'93 conference, Sept. 1993.

[LAKSH] T. V. Lakshman, P. Mishra, K. K. Ramakrishnan, “Transporting compressed video over ATM networks with explicit rate feedback control”, In Proc. of IEEE INFOCOMM, 1997.

[MIL31] OPNET Modeling Manual Vol 1, OPNET Version 3.5, MIL 3 Inc., 1997.

[MIL32] OPNET Models/Protocols Manual, OPNET Version 3.5, MIL 3 Inc., 1997.

[MPEG2] “MPEG-2 Systems”, ISO/IEC 13818-1, 13818-2, 13818-3, November, 1994.

[SHIV] Shivkumar Kalyanaraman, et. al, "The ERICA switch algorithm for ABR traffic management in ATM networks," submitted to IEEE/ACM Transactions on Networking.

[TM40] ATM Forum Traffic Management Specification, Version 4.0, ATM Forum, April 1996.

[VICKERS] B. J. Vickers, M. Lee, T. Suda, “Feedback control mechanisms for real-time multipoint video services”, IEEE Journal Selected Areas Comm, Vol. 15, No. 2, April 1997.

[VOD] The ATM Forum, “Audiovisual Multimedia Services: Video on Demand Specification 1.1,” http://www.atmforum.com/approved-specs/af-saa-0049.001.ps, March 1997.