ATM_FORUM/95-1139

| | |
|---|---|
| ATM FORUM: | Traffic Management Subworking Group |
| TITLE: | Performance Evaluation of Response Time in ATM LANs |
| SOURCE: | Helen Chen, Jim Brandt<br>Sandia National Laboratories, California |
| DATE: | December 11, 1995 |
| DISTRIBUTION TO: | The ATM Forum Members |

**Abstract**

This contribution compares the response-time performance of ATM LANs using ABR EFCI, UBR FIFO, and UBR with per VC queuing switches. Our study is based on experimental as well as simulation results. We found that, with or without congestion, UBR switches with per VC queuing provide the best response times.

**Notice**

This contribution has been prepared to assist the ATM Forum. This document is offered to the Forum as a basis for discussion and is not binding on Sandia National Laboratories or any other company. The requirements are subject to change in form and content after study. Sandia National Laboratories reserves the right to add to, amend, or withdraw any statements contained herein.

**Acknowledgment**

We would like to thank Chien Fang for the use of his ATM switch and host simulators. In addition, we express our appreciation to Rich Palmer for reviewing and providing helpful suggestions to this contribution.

## Introduction

The purpose of our study is to evaluate the ATM technology with respect to LAN performance requirements. Typical LAN requirements include high throughput for bulk-transfer applications and fast response-time for interactive or distributed-computing applications. Since the throughput issue has already been addressed by numerous contributions, this study will concentrate on the response-time requirements of delay-sensitive applications. For packet sizes less than ATM's maximum transfer unit (MTU), of 9188 bytes, the dominant factor in performance is time-out clock granularity, therefore we limit our discussion to loss-less scenarios. Furthermore, we will not address the robustness of the enhanced proportional rate control algorithm (EPRCA) regarding its rate-control parameters. The ATM technologies being evaluated include: 1) ABR EFCI switch, 2)UBR FIFO switch, and 3) UBR switch with per VC queuing.

## Evaluation Topology and Methodology

For this study, we used a two-node configuration to inter-connect one UDP-based echo session and from one to three TCP-based TTCP sessions (Figure 1). We measured, experimentaly and via simulation, the round-trip-times (RTT) of UDP packets of 64, 128, 256, 512, 1024, and 2048-bytes. Measurements were made with from one to three background TTCP streams.

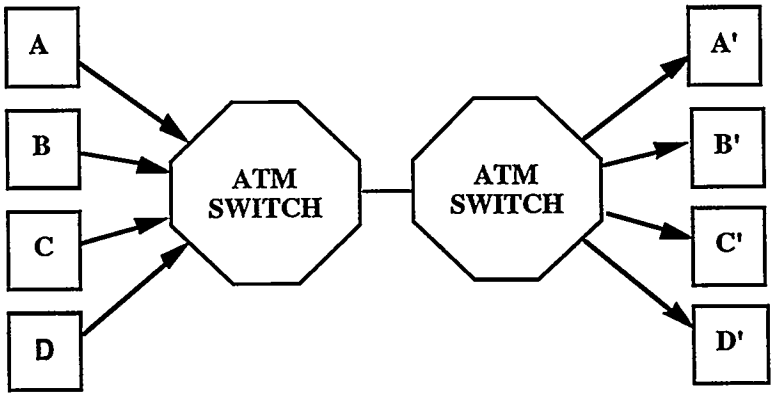


***Figure 1. Two-node configuration.***

We limited our experimental evaluation to equipment which we own, the DEC AN3 switch and the Fore ASX200 switch. Our evaluation is based on the RTTs of a UDP echo application running on a pair of DECstation 3000/600 with DEC OTTO ATM adapters. All background TTCP streams are also generated using identical DECstation pairs. Both the AN3 and the ASX200+ switches implement per VC buffer management using large buffer

resources. The AN3 has 2000 cells of *input* buffer per OC3c port and the ASX200+ has 13000 cells of *output* buffer which is dynamically shared among four OC3c ports.

In the absence of UBR FIFO and ABR EFCI hardware, we conducted our tests using simulated models developed at Sandia National Laboratories. The simulated UBR FIFO switch had 13000 cells of output buffer shared among four OC3c ports, which is identical to the ASX200+ buffer resources. We simulated the ABR EFCI switch with unlimited output buffer but with an EFCI threshold of 500 cells. This arrangement allows observation of EPRCA's rate control behavior as well as its switch buffer requirements under loss-less conditions.

To ensure a fair comparison between our experimental and simulation measurements, we calibrated pertinent simulation parameters against experimentally obtained values. We set the latency of our simulated switch to 12 µsec which is the measured latency of the AN3. We adjusted the parameters for protocol processing overhead as well as data-copying overhead for the simulated workstation until its RTT vs. packet size curve matched the experimental curve. Results are depicted in Figure 2.

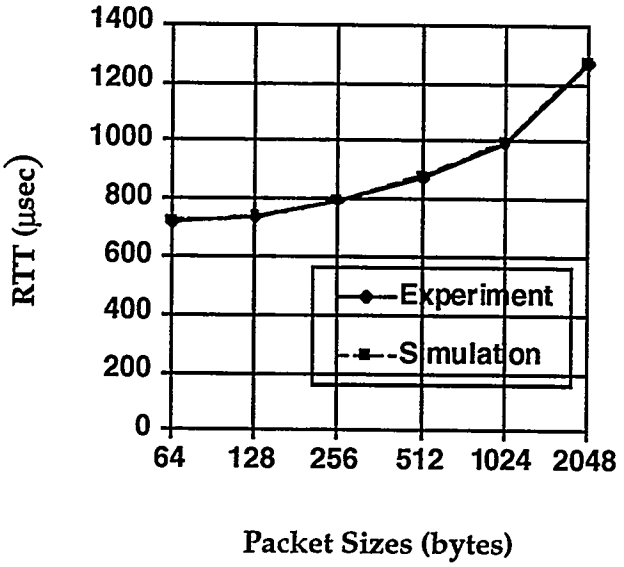


***Figure 2. Calibration of Workstation Simulation Parameters***

We adopted the "optimal" EPRCA parameters for LAN scenarios from the study by C. Fang and A. Lin, AF/95-1328R, they are listed in Table 1.

*Table 1. EPRCA Simulation Parameters*

| PCR | AIR | ICR | MCR | RDF | Nrm | TOF | Trm | TCR | Mrm | Xrm | Xdf | TDFF |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 155 | 0.01892 | 10 | 0.5 | 512 | 32 | 2 | 100 | 10 | 2 | 32 | 0.5 | 0.5 |

**Results and analysis**

The results of our study are summarized in Figure 3. Plots of the echo RTTs against UDP packet sizes, with one background TTCP session are shown in 3a. They are plotted for the scenarios with the DEC AN3, Fore ASX200+, UBR FIFO, and ABR EFCI switches. Similarly, 3b and 3c plot RTT vs. packet size, but with two and three background TTCP sessions respectively. These plots represent RTT values without congestion (3a), and with incremental increases in congestion (3b and 3c).

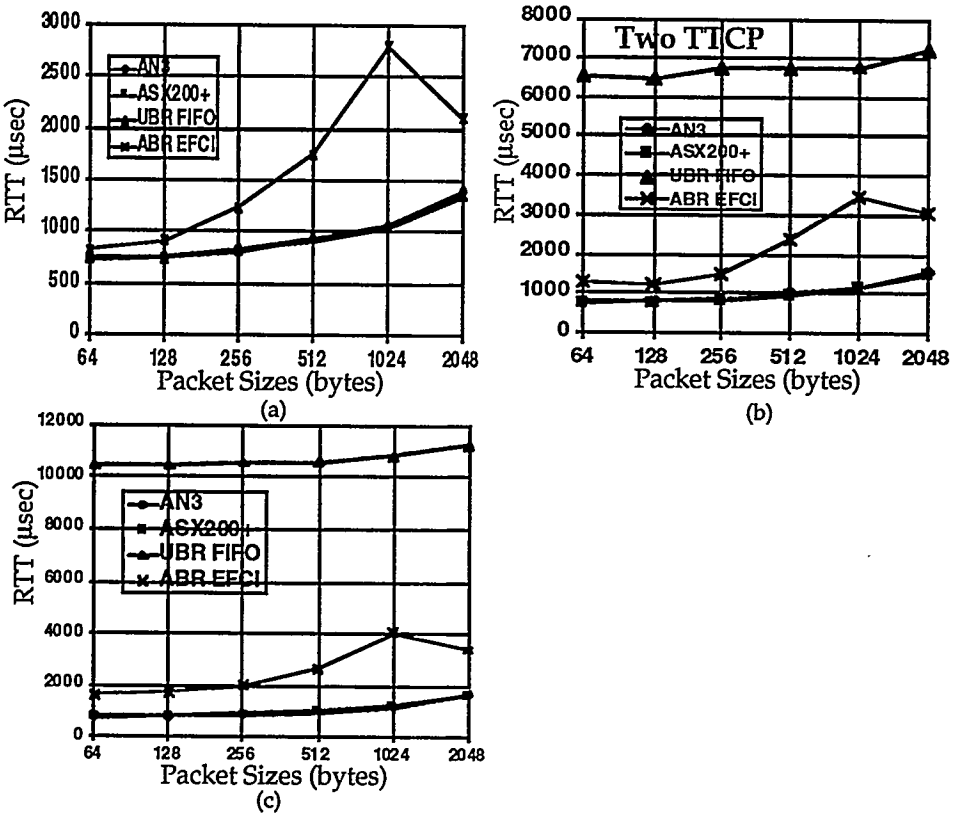


*Figure 3. UDP echo RTTs vs. Packet sizes with : a) one TTCP, b) two TTCP, and c) three TTCP*

As shown, in all three cases, the ATM LANs with UBR per VC queuing (AN3 and ASX200+ switches) provided the best RTT values. We believe that the switch's per VC queuing mechanisms provide the necessary VC isolation to ensure fairness and optimal response-time.

The performance of the UBR FIFO switch (Figure 3a), though optimal in the absence of congestion, degrades rapidly as congestion increases. Figure 4 depicts the switch queue states and reveals that large RTTs are introduced by large FIFO queues.

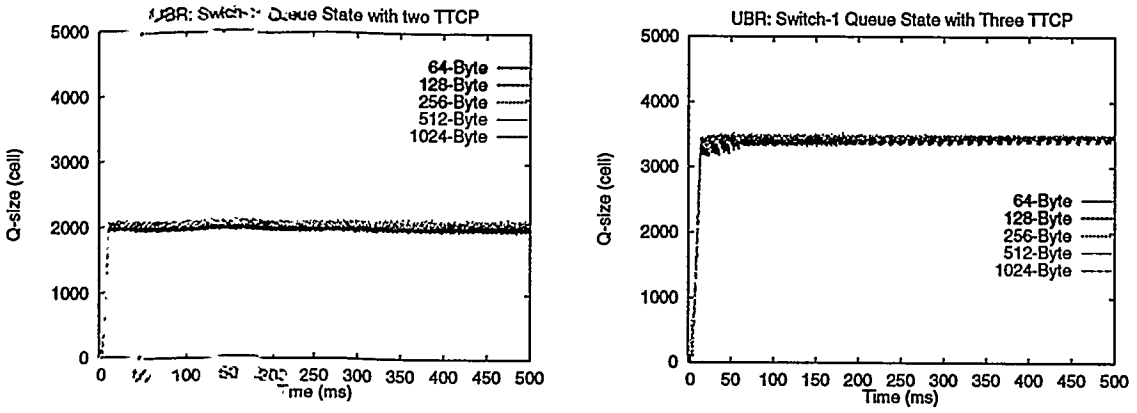


***Figure 4. UBR FIFO switch queue state***

Figure 3a also shows that without congestion the ABR EFCI switch caused larger RTTs than either the UBR FIFO or UBR per VC queuing switches. An examination of the plots in Figure 5 reveals that when the UDP packet-size is 1024-bytes or less, the EPRCA prevents the echo sessions allowable cell rates (ACR) from being adjusted beyond their initial cell rates (ICR) of 10 Mbps. With congestion (Figure 3b and 3c), the ABR EFCI RTTs are larger than those of the UBR per VC queuing configuration. Again, this is a result of FIFO queue buildup. Figure 6 plots the queue states of the ABR EFCI switch. As shown, the ABR EFCI queue size averaged around 500 cells thereby causing less than optimal RTTs. Nevertheless, the queue size is much smaller than that seen in the UBR FIFO switch (Figure 4) and thus yields much better RTT values.

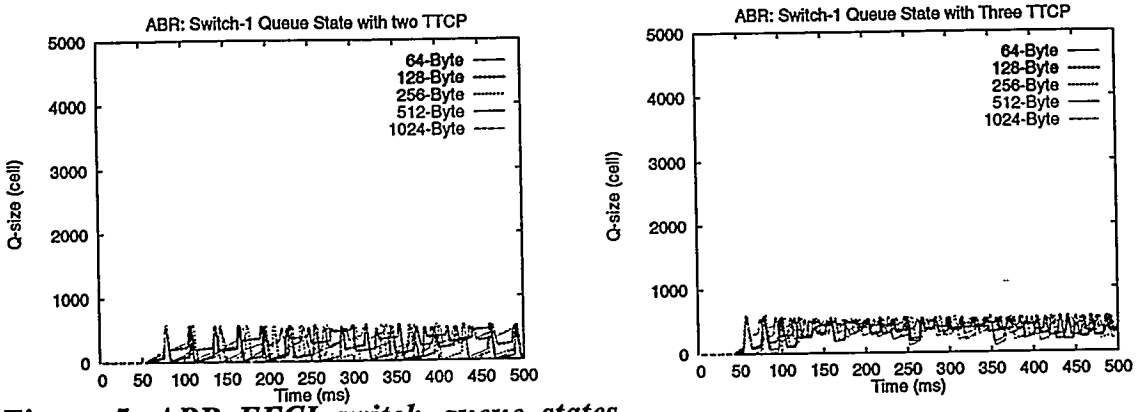


*Figure 5. ABR EFCI switch queue states*

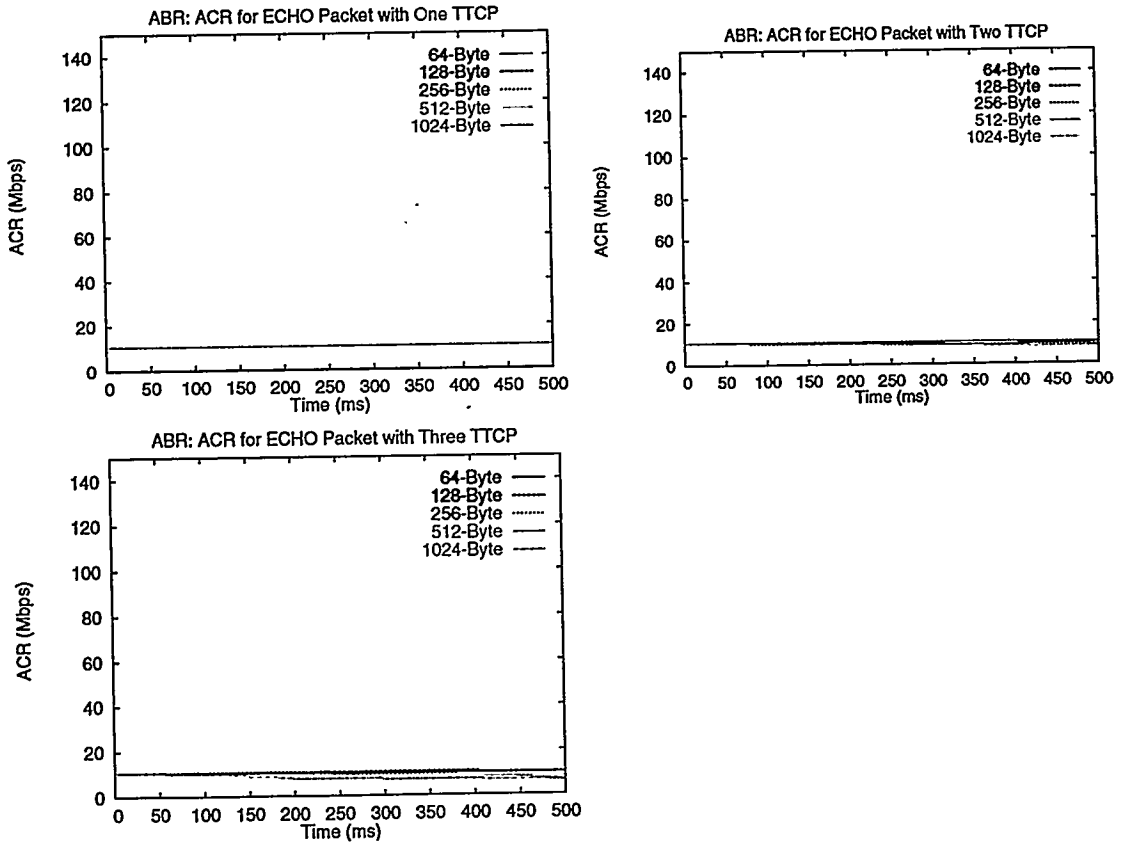


*Figure 6. UDP echo ACR vs. time*

Without congestion, it would be desirable to allow applications to have immediate full-link access, especially to those that are delay sensitive. Therefore, we increased EPRCA's ICR value to 155 Mbps and plotted the results in Figure 7. As shown, for our simple configuration, peak queue sizes reached two to three times the EFCI threshold value before reaching steady state. Moreover, in the presence of bursty traffic, switch queues had large

oscillations which resulted in large response-time jitters and occasionally inefficient link utilization. Due to the dynamic complexity of real life traffic it would be extremely difficult to provision a switch with adequate buffer resources to accommodate bursty traffic while allowing immediate full link access.

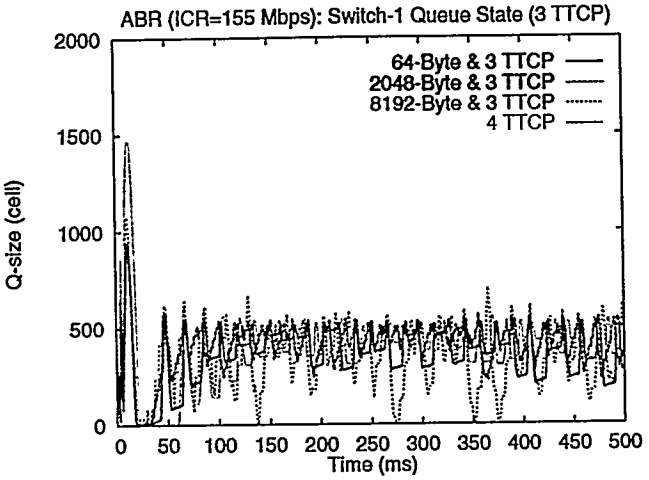


***Figure 7. ABR EFCI switch queue states with ICR equal PCRß***

**Summary and Conclusions**

We evaluated the response-time performance in ATM LANs using UBR FIFO queuing, ABR EFCI, and UBR per VC queuing switches. We found that during congestion large buffers in UBR FIFO switches incur undesirable queuing delays on all VCs, including VCs of delay sensitive applications. Similarly, ABR EFCI switches, which also maintain FIFO queues, experience delay during congestion. However, EPRCA reduces queuing delay by limiting queue growth. Even so, their RTTs are still large when compared with the UBR per VC queuing switches. Another disadvantage of the EPRCA-based EFCI switches is their inability to quickly adjust rates for applications with small traffic bursts. This results in their ACRs remaining at their ICRs. Larger response times result due to larger transmission delays (ICR much less than PCR) even when there is no congestion. With or without congestion, the UBR switches with per VC queuing (AN3 and ASX200+) have the best response-time performance.

In conclusion, we believe that, as an interim solution, UBR switches with per VC queuing are the most suitable technology for the LAN. Per VC queuing provides VC isolation which supplies the fairness necessary for interactive applications to obtain optimal response-times as well as a fair share of bandwidth. It is intuitive that, for interactice

sessions, cell-loss could be prevented using per VC queuing. Furthermore, cell-loss in bulk transfer applications could be reduced without affecting response-times of interactive sessions by simply increasing buffer size. Though we have not addressed the robustness of the EPRCA, it is an issue that has dominated the TM sub-working group's effort for the past year or so and thus should not be taken lightly. It is our opinion that the TM sub-working group should, in the interest of succesfully ATM deployment, take a hard look at what they are actually trying to accomplish and perhaps try to put politics aside and come up with a truly workable solution such as maybe credit based flow control schemes.