*******************************************************************

ATM Forum Document Number: ATM_Forum/96-0179

*******************************************************************

Title: TCP/IP over UBR

*******************************************************************

Abstract:
Performance of TCP with UBR service is examined. We study the effect TCP timer granularity, switch buffer size, and cell drop policy on TCP performance.

*******************************************************************

Source:

Raj Jain, Rohit Goyal, Shiv Kalyanaraman, Sonia Fahmy, and Fang Lu
The Ohio State University
Department of CIS

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

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

*******************************************************************

Date: February, 1996, Los Angeles

*******************************************************************

Distribution: ATM Forum Technical Working Group Members (Traffic Management)

*******************************************************************

Notice: This contribution has been prepared to assist the ATM Forum. It is offered to the Forum as a basis for discussion and is not a binding proposal on the part of any of the contributing organizations. The statements are subject to change in form and content after further study. Specifically, the contributors reserve the right to add to, amend or modify the statements contained herein.

************************************************************

INTRODUCTION:
------------

TCP is one of the few transport protocols that have their own congestion control mechanism. TCP's slow start mechanism allows it to reduce its load whenever the network is congested and drops its packets. This has lead to the possibility that TCP will work well on ATM networks without any further help from the ATM layer. In other words, TCP will work well on unspecified bit rate (UBR) or some slight improvement of it [1,2,3,4,5,6,7].

The purpose of this and our subsequent contributions on this topic is to characterize the performance of TCP over UBR and its various enhancements. This particular contribution concentrates on the effect TCP timer granularity, switch buffer size, and cell drop policy on TCP performance.

In an associated contribution [8], we have studied the performance of TCP/IP over ABR. There we characterize the effect of transient buffer exposure (TBE) parameter of ACR traffic on TCP's performance. In that contribution, we briefly explain TCP's slow start and timeout mechanisms. The terms "exponential increase phase," "linear increase phase," and "steady state" introduced there will be used here.

PARAMETERS:
----------

TCP timer granularity: Most current TCP implementations measure round-trip delays using a granularity of 100 or 500 ms. Larger granularity means that the sources have to wait longer before timing out and recovering from a lost packet. Lower granularity is, therefore, expected to improve performance. We study the effect of changing clock granularity on various performance metrics.

SWITCH BUFFER SIZE AND CELL DROP POLICIES:
----------------------------------------

We study TCP performance with finite and infinite switch buffers. The effect of various cell drop policies on TCP is examined. In its simplest form, a switch could implement a tail drop policy in which all cells arriving after the buffers are full are dropped. This is expected to result in excessive wasted bandwidth. If cells are dropped from multiple TCP packets, then all these packets need to be retransmitted by the source TCP.

Early Packet Discard (EPD) has been suggested to remedy the problems caused by tail drop. EPD tries to discard cells from the same TCP packet during congestion. A threshold is set at the switches, and when the switch queue length exceeds this threshold, cells from any new packets are dropped. The EPD algorithm is the one suggested by [3,7]. We present both the effect of simple tail drop as well as Early Packet Discard. However, EPD makes no attempt at achieving fairness among different VC's. It is also not known how to choose the threshold for the switch buffer. The choice of EPD threshold value may effect performance. We experiment with different buffer sizes and EPD thresholds.

CONFIGURATION:
-------------

We use a variety of simulation configurations. For our first simulation, we use the following parking lot configuration with six sources. The sources are indicated as S1 through S6 and the switches are indicated as Sw1 through Sw6. All sources go to the destination "Dest." This configuration was chosen to compare our results with those published earlier in [1,2] and, thus, validate our model.

```
 S1
  \
   \
    \
     \Sw1---->Sw2---->Sw3---->Sw4---->Sw5---->Sw6---->Dest[1..6]
     /         ^       ^       ^       ^
    /          |       |       |       |
   /           |       |       |       |
  /            |       |       |       |
 S2            S3      S4      S5      S6
```

We use both LAN and WAN configurations.

LAN:

ES -> Switch link propagation delay = 2.5 microseconds

Switch -> Switch link propagation delay = 10 microseconds

WAN:

ES -> Switch link propagation delay = 2.5 microseconds

Switch -> Switch link propagation delay = 2 milliseconds

* All link bandwidths are 155.52 Mbps

* PCR = 155.52 Mbps

* All sources are infinite sources sending data to the TCP layer at PCR.

* Only the sources send data. The destination sends only acks.

* TCP Reno model is used (Fast Retransmit and Recovery are implemented).

* TCP segment size = 512 bytes

* TCP maximum window size = 64 kB

METRICS:
-------

We measure performance based on the following metrics,

1) TCP sequence numbers: measured at the source TCP layer.

2) TCP throughput: This is measured at the destination TCP layer. The destination measures throughput of all the segments received in sequence. If a segment is received out of sequence, (due to error/loss of previous segments) it is not included in the throughput measurement until all the missing segments are received.

3) Switch queue length.

RESULTS:
-------

Simulation Results will be presented at the meeting.

REFERENCES:
----------

[1] Chien Fang, Arthur Lin: "On TCP Performance of UBR with EPD and UBR-EPD with a Fair Buffer Allocation Scheme,"

AF-TM 95-1645, December 1995.

[2] Chien Fang, Arthur Lin, "A Simulation Study of ABR Robustness with Binary-Mode Switches: Part II," AF-TM 95-1328R1, October 1995.

[3] Hongqing Li, Kai-Yeung Siu, and Hong-Ti Tzeng, "TCP over ATM with ABR service versus UBR+EPD service," AF-TM 95-0718, June 1995.

[4] AF-TM 95-0503, Hongqing Li, Kai-Yeung Siu, Hong-Yi Tzeng, Brian Hang, Wai Yang, "Issues in TCP over ATM," April 1995.

[5] Tim Dwight, "Guidelines for the Simulation of TCP/IP over ATM," AF-TM 95-0077r1, March 1995.

[6] Stephen Keung, Kai-Yeung Siu, "Degradation in TCP Performance under Cell Loss," AF-TM 94-0490, April 1994.

[7] Allyn Romanov, Sally Floyd, "Dynamics of TCP Traffic over ATM Networks."

[8] Raj Jain, S. Kalyanaraman, R. Goyal, S. Fahmy, F. Lu, and S. Srinidhi, "TBE and TCP/IP traffic," AF-TM 96-0177, February 1996.

All our past ATM forum contributions and presentations can be obtained on-line: http://www.cse.wustl.edu/~jain/