Experiments for Analysis of Video Transmission via ATM*

Mauricio Pissioli [1], Edson dos Santos Moreira [2], Rudinei Goularte [2]

1-Avaya Communications, São Paulo, SP, Brazil
mpissioli@avaya.com.
2 - São Paulo University – Institute of Mathematics and Computer Science
P.O. box 668, 13560-970 - São Carlos, SP, Brazil
{edson, rudinei}@icmc.sc.usp.br

Abstract

A large quantity of data of different types and priorities pass through on high-speed networks. A relevant part of these data, like continuous media (audio, video, and so), need a control of its Quality of Service (QoS). This paper describes experiments on the characteristics of the ATM network technology that provide support to quality video streaming transmission. The experiments were realized considering QoS, connection admission control, network congestion, transmission priorities and bandwidth demand associated to applications and network configuration. Three basic kinds of videos were analyzed – movies (smooth movies), sports (action movies) and talking head movies. Data about video transmission were collected from an ATM switch, showing results that can be used in the future on ATM channel mapping to multimedia data distribution (in applications like Interactive TV and Video-on-Demand, for example).

Key words : computer networks, ATM, high-speed networks, video-on-demand.

1 – Introduction

One of the most important issues in telecommunications is the enhancement and the utilization of network infrastructures (the base for telecommunications services) that are capable to ensure Quality of service (QoS). The control of the network resources in a way to preserve the QoS is reached through the definition of a set of parameters that, at the moment of establishing connections, describe the relevant characteristics of the traffic. Another important aspect is the management of the bandwidth used in connections that carry out a variety of data with different priorities, using the same network infrastructure.

The ATM network concept has been designed with the purpose of making good use of the available bandwidth (using statistic multiplexing instead of time division multiplexing – inefficient to multimedia data transport) and at the same time, providing mechanisms to QoS maintenance 9 (Tanembaum, 1996).

The conditions under which the statistic multiplexing works efficiently in ATM networks has been an active area of research and experimentation in the academy and industry. It has also been a source of technical publications and speculations. The telecommunication companies in the U.S.A., Europe and Japan, research centers and standardization commissions are actively investigating what is the best way to do

statistic multiplexing in a way that the bandwidth in the ATM networks be used efficiently. Standardization commissions are also engaged in similar pursuits investigating what is the better way to do statistic multiplexing. These commissions aim to make the bandwidth in the ATM networks more efficiently used, and that the QoS needed be granted during congestion periods. The reason for the attention given to this problem is the following: if the peak bandwidth needed by all connections is allocated, the ATM degenerated to STM and there is not statistic advantage with the anticipation of the bandwidth use through the description of the traffic to be transported before the connection establishment (Walke et al., 1996; Umehira et al., 1996; Murphy, 1997; The ATM Forum, 1996).

The cell commuting technology uses units with a fixed size of 53 bytes: five of header and forty-eight of data. The idea is to quickly switch these cells while on route each cell carrying only the identifier of the connection and the data. The small and fixed size of the cells facilitates the recovery of losses because the amount of data in each cell is also small (Tanembaum, 1996).

This approach is called "band on demand" once the network provides exactly the effective bandwidth needed by the applications. The ATM switches cells in the same speed that the cells are submitted to transmission (all the switching is made by a specific hardware). In essence, this mechanism is very similar to packet switching and for this reason was denominated "fast switching of small and fixed size packets".

Distribute applications that involve multimedia data transmission must provide the necessary support to ensure that the data generated by a source can be reproduced at the end with the same quality and this requires QoS parameters (Nahrstedt et al., 1995).

When a host connects an ATM network it is essentially doing a contract with the network, based on QoS parameters. This contract describes the flow of the traffic to be transported through parameters such as bandwidth peak, the average bandwidth, the average burst duration, the maximum delay between cells, etc. It is the responsibility of the ATM device to accept this contract using queues in order to arrange the data burst, to limit the data rate peak and to soften the delay between cells in such a way to adequate the traffic to the contract specifications (Sirivara, 1997; Raychaudhuri, 1996; Tanembaum, 1996).

All of these concepts are also valid to video applications. In order to transport video with acceptable quality, the network must provide a certain level of QoS. Variations on the cells delay, bit errors and cell losses can have strong effects on the quality of the received video stream. For example, a link of transmission with a bit error rate of 5 to 10% can be acceptable in the case of data transmission that is not be transmitted in real time (using some error correction mechanism). However, for a video stream, this rate of error can cause a serious degradation of the video quality (Moreira et al., 1995; Moreira et al., 1997; Riley et al., 1997). In the same way, cell delay, cell losses and control of the transmission rate have significant impact on the video quality. ATM supports the control of latency and delay at the moment of contracting the service. It is also possible to control the cell loss and the error rate keeping both within acceptable levels (The ATM Forum, 1996; Tanembaum, 1996).

This paper presents a series of experiments on video transmission in a ATM network, trying to analyze the characteristics that this network offers to transport video streams, considering QoS, admission control of connections, network congestion, transmission priorities, demand of application bandwidth and network configuration.

2 – Evaluations of the ATM video transmission

The environment for the tests was composed by a ForeRunner ASX200 BX ATM switch (Fore System), two Fore Systems AVA300 and a Fore systems ATV300 (video signals coding in ATM cells and ATM cells decoding in video signals, respectively), a VCR, a 29" TV and a Pentium 133 (WinNT 4.0) PC with a ATM interface to control the switch. The UPC contracts were applied through a interface software provided by the Fore Systems.

2.1 – Evaluations of the VBR video streams transmission

Characterization of the VBR video streams used in the tests

Each type of video presents its own characteristics related to the nature of its content. These characteristics are responsible for evident effects on the rates of compressed data. These rates, when observed in constants time ranges, follow patterns of variation that are different for each type of video. This fact defines patterns of bandwidth utilization, as shown in figure 1.


Figure 1 – Variations in the bandwidth utilization for movie, sports and talking head video.

Figure 1 shows the amount of transmitted cells for the three types of video analyzed in this work (movie, sports and talking head), evidencing the bandwidth utilization. For comparison purposes, the conditions in which the data was collected was the same for the three videos: MJPEG compression with constant quality, resolution of 608x224, 30 frames per second, transmission in a ATM using VBR connection with peak cell rate (PCR) of 200000 cells per second. No UPC contracts were applied to the

connection and there was no other data source connected to the switch, granting that no one cell was rejected due to traffic policy or congestion. All of the videos were 20 minutes NTSC coded. The data were collected at a rate of one sample per second.

Tests with VBR streams and UPC contracts

The aim of these tests is to demonstrate the effects generated by the cells discard on the quality of a video transmitted in ATM, when using the service class VBR in the case of the three videos previously described. Another aim is to attempt to define policies for the UPC contracts applications.

Two distinct scenarios motivate these tests. The first one is a situation where an application requests a connection, gives the connection parameters to the CAC (Connection Admission Control, responsible for the application of the contracts), the CAC accepts and establishes the connection, but the application doesn't respect the established parameters and sends cells to the network at rates that is higher than the negotiated parameters. If the network is using an UPC contract to police the traffic, and starts to discard cells, what would be the impact on the quality of the received video?

The second scenario refers to the optimization of the network resources, which can be achieved reducing the amount of cells transmitted applying UPC contracts that make a "calculated" discard. What is the amount of cells that can be discarded in a VBR connection without compromising the perceptible video quality? And how to carry out this calculated discard?

Some considerations

The configuration of the transmitted streams are the same used in the analysis done in to the characterization of the three types of video previously examined: MJPEG video with constant compression (Q-factor 16), resolution of 608x224, 30 fps, ATM transmission using VBR connection, PCR of 200000 cells per second (to the SCR and MBS discard tests) and of 105000 to 110000 cells per second (to the PCR discard tests). All the three videos were coded in NTSC and are 20 minutes long. Each sample consists of a number of cells transmitted per second.

In the parameters definition for each type of UPC contract (PCR, SCR, MBS and CDTV) basic statistics principles were used and applied to the collected data. For each one of the three videos, we calculated what follows:
- the total number of transmitted cells;
- the average of all samples;
- the maximum and the minimum values founded among the samples;
- the observed frequency of samples, taking in consideration ranges of 1000 cells, starting with the minimum of 10000 cells up to a maximum of 32000 cells.

Discard Policies

The cell discard was realized in two different ways: by PCR and CDVT violation (this is verified at the fist stage of the GCRA), and by SCR and MBS violation (verified at the second stage of the GCRA). The aim is to verify what are the effects generated by the PCR and SCR violation in the received video quality. The UPC contracts utilized were VBR0+1, that is, the monitoring is applied to all the cells (CLP=0 or CLP=1) to stream VBR. To make sure that there was no discard by SCR

violation in the violation PCR test, the contracts were created with SCR and MBS values much superior to the ones required in the transmission. The same principle was used in the SCR tests, that is, PCR and CDVT were calculated with higher values.

The policies used to measure the quantity of cells that should be discarded to the SCR test are as follows:

- to cut the cells; objecting that the number of transmitted cells per second was slightly lower than the higher sample observed in the original stream (cut the peaks).
- to cut the cells; objecting that the number of transmitted cells per second was the same for the average of all the samples of the original stream (cut considering the average).

The objective were to observe the light effects in the quality of the received video and to observe more severe effects too. The data calculated using the collected samples were utilized in the definition of the parameters of the UPC contracts to the SCR discard tests.

Table 1 shows the maximum and the minimum values of the samples that express the number of per second transmitted cells for each one of the three videos, and the medium value of these samples (1164 samples for each kind of video). The maximum value of per second transmitted cells was 20566 for type movie, 29551 for sports and 18933 for talking head.

Table 1 – Number of cells per second

Type of video	Minimum (cps)	Medium (cps)	Maximum (cps)	Samples
Movie	10559	15058	20566	1164
Sports	11228	19083	29551	1164
Talking head	13069	15179	18933	1164

In order to implement the cuts by the pecks, the observed frequency of the samples was used. It was noted that in the case of the type movie, the majority of the samples (98.46%) have values in the range of 0 to 19000 cells per second. It was therefore decided to discard cells in a way that, a maximum of 19000 cells per seconds would be transmitted (the peak was 20566 cps). If the discard were limited to 18000 cps, 11.54% of the sample would be affected.

The same method was used to the sports video (cuts at 25000 cps, peak of 29551, affecting 5.24% of the samples) and to the talking head video (cuts at 16000 cps, peak of 18933, affecting 3.61% of the samples).

The tests "cuts by the average" had the following values: 15058 cps to movie, 19083 cps to sports and 15179 to talking head video. From these values, we calculated the UPC contract parameters used in the tests. It should be noted that the rate of 30 fps implies on intervals between cells of approximately 33.3 milliseconds.

This fact must be considered in the MBS calculation of the UPC contracts. For example, in the case of the movie video type, it is desirable to limit the transmission rate to a maximum of 19000 cps. This value represents the SCR, which is defined like the medium rate of cells during all transmission (in cps). To calculate the MBS (defined as the maximum burst size of data in cells) it is sufficient to divide the SCR by 30 (the number of bursts in one second) which results in approximately, 633 cells. It must be borne in mind that this is the medium number for the sizes of the bursts. The size of bursts can, in fact vary. Still, we can consider that this calculation is valid given that in

one second of a video sequence, the frames probably will be very similar thus the compression of these frames would generate the same amount of data (and consequently of cells) at every 33.3 milliseconds. As such, the bursts that have a number of cells higher than 633 will suffer discards, keeping the average of 19000 cps to the stream.

Limiting the peaks through the SCR

Figure 2 – Cell discard caused by SCR violation (movie, limiting the peaks)

The limitation by peaks tests results for the three types of video was very similar. In view of that, the movie vide type results will be presented as a demonstration. Similar results were also observed with the cut by average tests to the movie and sports video types. Only the cut by average test for the talking head video type shows important differences and will be fully presented.

Test data:

- Video type: **movie**;
- Stream type: **MJPEG,** *minimal compression***, 604x228, 30 fps**
- Service class used in the transmission: **VBR**;
- Discard type: **SRC violation** through UPC contract applied to the VCC used in the transmission (VPI=0, VCI=200);
- UPC contract: **VBR0+1**, **PCR=110000 cps**, **SCR=19000 cps**, **MBS= 633 cells**, **CDVT = 10 microseconds;**
- Duration: 1164 seconds.

Figure 2 shows the cells discard in the video type movie with peak limitation at 19000 cps. As was observed in practice, the discard effectively occurs around 19600 cps. The observed consequences pointed out that: the video transmission was paralyzed (not continuous) during a couple of seconds, starting with a freezing image and being replaced by a dark screen, during approximately six seconds. These periods were the same when the cell discard by SCR violation occurred. The rest of the video remained

without alterations. From 17469671 cells, 2606 were discarded, representing 0.014% of the total.

Limiting the cell rate by the average through the SCR – "movie" video type

Figure 3 - Cells discard by SCR violation (movie, limiting by average)

Test data:
- Video type: **movie**;
- Stream type: **MJPEG, *Q-Factor* 16, 604x228, 30 fps**
- Service class used on the transmission: **VBR**;
- Discard type: **SRC violation** through UPC contract applied to the VCC used on the transmission (VPI=0, VCI=200);
- UPC contract: **VBR0+1**, **PCR=110000 cps**, **SCR=15058 cps**, **MBS= 502 cells**, **CDVT = 10 microseconds**
- Duration: 1164 seconds.

Figure 3 outlines the cell discard for the movie video type with rate limitation of cells at 15058 cps (sample average of the original video). In practice it was observed that the discard occurs effectively, around 15700 cps.

Observed consequences: the video transmission presented on the whole, about 8 minutes (non-continuous) of paralysis, starting with freezing images and then being replaced by dark screen. These periods were the same when the cell discard by SCR violation occurred. The normal periods (without discards) interchanged with the periods where discards occur making the video lose the logic sequence (between minute 0 and minute 12). The rest of the video remained like the original one. From a total of 17454087 transmitted cells were discarded 726128, representing 4.14% of the total.

Limiting the cell rate by the average through the SCR – "talking head" video type

Test data:
- Video type: **talking head**;
- Stream type: **MJPEG, *Q-Factor* 16, 604x228, 30 fps**

- Service class used on the transmission: **VBR**;
- Discard type: **SRC violation** through UPC contract applied to the VCC used on the transmission (VPI=0, VCI=200);
- UPC contract: **VBR0+1**, **PCR=110000 cps**, **SCR=15179 cps**, **MBS= 506 cells**, **CDVT = 10 microseconds;**
- Duration: 1164 seconds.

The cells discard in the case of the talking head video type was done with cell rate limitation at 15179 cps (average of the original video samples). In practice, it was observed that the discard occurs effectively, around 15600 cps. Observed consequences: the video presented 30 seconds of paralysis (non-continuous) with freezing images but without dark screens. The periods of paralysis were the same when the cells discard by SCR violation occurs. There were no observed consequences in the logic sequence: from 17750203 transmitted cells 29527 were discarded, representing 0.167% of the total.

Limiting the cell rate through PCR – all videos
All original streams were transmitted with PCR=110000 cps. This test had the aim to observe the cells discard effects due to PCR violation through the use of a UPC contract on the VCC used in the stream transmission.

As mentioned above, UPC contracts (in a total of 20) were created with PCR varying between 109000 cps and 90000 cps. The SCR and MBS values were set with values higher than the necessary, ensuring that cells are not discarded neither by SCR nor by MBS. The CDVT value was 10 microseconds on all contracts, to guarantee a low difference tolerance between the actual and the theoretical arrival times of the cells.

Observed consequences: for contracts with PCR between 100000 cps and 109000 cps, little random regions on the image were outdated by periods of time of a second fraction, being refreshed immediately. Another effect, which occurred jointly with the previous one is that the contract with the PCR was observed to be between 92000 and 99000 cps. It was clearly noted that a type of reduction in the fps rate results in that all image is outdated for periods of time perceptible by the human eye. Obviously, the intensity of this effect was proportionally inverse of the PCR of the contracts or in other words, the bigger the PCR of the contract used the smaller the intensity of the effects. For the contracts with the PCR below 92000, no single image appeared on the screen.

It must be observed that, for contracts with PCR higher than 92000 cps, there were no long periods of paralysis in the image sequence nor dark screens. The images keep a good level of dynamism, compatible with the PCR contracts used. The logic sequence of the video was kept making it possible to understand the video. Only some image details were lost but not the entire image.

Figure 4 shows the cell discard occurred during the consecutive transmission of all videos, applying an UPC contract with PCR of 95000 cps. It was observed that the cells discard occurred in a uniform manner during the whole transmission, once all videos had necessity of 110000 cps to be transmitted. There were discards during the whole test session and it is important to consider this: the differences in the quantity of cells discarded, observed between samples resulted from the variation of the cells burst size, and were not the result of large variation in the PCR of the stream. These variations exist but are small and in most of the cases don't cause discards due to the

tolerance achieved to the CDVT parameter. In other words, the bigger the burst size the bigger is the number of cells that are coming with PCR higher than the UPC contract and consequently, a higher number of cells are discarded.


Figure 4 – Discards by PCR violation - (movie, sports and talking head)

Analysis of the results

In the cases of the movie and the sports videos, the cells discard by SCR violation limiting only the peaks, presented minimal effects. These effects were discontinuities in the image (freezing images and dark screens) during short periods of time. These effects didn't disturb the logic of the video sequence. The most drastic effects were observed when the discard was realized by SCR violation limiting the cells rate to the samples average observed in the original videos. These discontinuities had long periods of time, compromising the logic of the video sequence and decreasing the video quality to unacceptable levels.

As to the talking head video, the discard by SCR violation presented minimal consequences in the peaks and average limitation in view of the fact that the difference between both is short. There was no compromising of the login sequence.

It was observed that the effects caused by the PCR violation for all videos, were delays of small regions of the image representing degradation on the image quality. This effect was observed during all the video presentation, but the sequence login was not affected. It was also was observed that for the talking head video, the discard by PCR caused more serious degradation than the discard by SCR, in both peak and average cases. The appliance of a UPC contract, with PCR of 16% (approximately) lower than the one used in the transmission (92000 cps against 110000 cps), makes the transmission impossible.

From these observation some conclusions can be drawn:

- Discard by SCR violation cause a non-constant effect due to the burst variation. In other words, the discard and its effects happen only at the moments when the SCR is violated (the burst size becomes bigger than the MBS). Depending on the duration of the violation, several consecutive

frames can be entirely discarded, affecting the logic of the video sequence. The quality of the individual images remains the same during the violations periods.

- Discard by PCR violation causes a constant effect (it means that the effect occurs every time) if it is considered that the transmission source always sends cells with the same PCR. The intensity of this effect depends on the quantity of discarded cells. This effect doesn't affect the logic sequence of the videos. There does not occur losses of entire video frames, but the quality of each frame is affected individually.

It was not possible to obtain sureness on why the SCR violation causes consecutive losses of entire frames and why the PCR violation cause losses of little portions of the image in each frame. We speculate that the observed effect on the SCR discard is possibly caused by the consecutive loss of AAL5 PDUs containing cells that violate the UPC contract (cells that are at the end of each burst). This can be the cause of the freezing images and the dark screens, a kind of synchronism loss that interrupts the video. In the PCR case, the violation causes the PDU loss but not in a consecutive manner, losing only pieces of the image.

Tests with VBR streams and Statistic Multiplexing (VBR Overbooking)

The switches used in the tests are by default, not configured to use statistic multiplexing in the VBR connections. The connection admission control algorithm (CAC) reserves to the VBR connections, bandwidth related to the specified PCR value. If the sum of these bandwidths surpasses the capacity of one of the links involved in the connection, the connection is rejected. VBR Overbooking makes the admission test be realized considering a capacity higher than the real capacity of each link. This is possible because in the course of the connections, in the most part of the time, the sum of the instantaneous bandwidth go beyond the capacity of any link.

Nevertheless, if in some instants the sum of the bandwidths actually used surpass the capacity of any one involved link, the link will transmit cells at the maximum capacity, using buffers to store the exceeding cells and introducing delays in the transmission. These delays can cause serious consequences if the cells transport audio or video.

To examine these effects, the following test was realized: two transmission systems (AVA300) were configured to transmit VBR streams, one of them with PCR of 150000 cps at the 1c4 switch port and the other with PCR of 250000 cps at the 1a3 switch port. Both of these systems were directed to the 1c3 switch port, which has 155.22 Mbps (353207 cps). To this port, an ATV300 was connected through which one of the videos was watched. It should be noted that the sum of the PCRs of the two streams (400000 cps) overcome the 1c3 port capacity. Was necessary to configure an overbooking around 15% at the 1c3 port for that the VCCs establishment works.

We observed serious effects in the watched video quality, clearly evidencing the delay imposed on the cells during the transmission. Many periods of interruption were observed and the scenes that could be watched were being played at a rate lower than 30 fps.

2.2 – Evaluations on the CBR stream transmission

The compressed video streams transmission using connection with VBR service class, present cell rates that vary due to the compression. In the non-compressed video streams transmission, using CBR connections, the bandwidth utilization (observed through the cells rate) depends only on the transmitted video format (resolution, frame rate, coding, etc.). It must also be taken into consideration that the overheads introduced by the AAL5 utilization. It is possible to estimate what cells rate will be used calculating the size (in bytes and in cells) for each video sequence frame, multiplying for the frame rate (per second) and summing the overhead.

For example, for a 15 bit RGB video (2 bytes per pixel), 304x224, 30 fps, non-interlaced, the rate will be (not considering the overhead and little variations on the cells rate):

N= 304 * 224 * 2 * 30 * 1 = 4085760 bytes per second.

In cells, considering the 48 bytes of payload for each cell:

Nc = N / 48 = 4085760 / 48 = 85120 cells per second.

Figure 5 shows the cell rate in the video transmission with the same format (from 00:00:00 to 00:03:24) which has its resolution reduced to 304x112 from 00:03:24 to the end. It was observed that the cps rate was approximately 95000 cells, 12.6% higher than the calculated 85120 cells. In the second half of the graphic below, it will be noted that a cell rate of approximately 54200 cps, that is, 27,3% higher than the 42560 cells previously calculated.


Figure 5 – Bandwidth utilization for CBR vide streams

Other statistical operations can be made: if 9880 (from 95000 – 85120) cells transport only the overhead in its payload, then, the overhead is 474240 bytes per second. This indicates that 59280 PDUs AAL5 were transmitted. We used a UPC CBR0+1 contract applied to the VCC of the connection with the aim of identifying which kind of degradation in the visual quality would occur. An extremely low tolerance related to the loss of cells was observed. Some few dozens of lost cells caused the interruption of the images. The CBR connection was established with PCR of

198100 cps, and the effect was observed when a UPC contract with PCR of 198030 cps was applied. Intermediate values also caused discards, but there were no image interruptions and no observed degradation in the visual quality.

One advantage of the CBR service class utilization, in terms of quality service, is that it doesn't have the delays observed in the VBR connection overloaded by the VBR overbooking utilization.

3 – Conclusions

The conclusions drawn from this work, related to the particularities involved in the video transmission with service guaranty on ATM networks, was presented throughout section -2-. With the available resources, it was possible to achieve a good level of detailing in the results and analysis. Clearly, more in-depth studies will demand the use of more powerful tools. Finally, the utilization of additional equipment, such as switches and transmission devices connected to the network, would increase the number of tests possible to be carried out.

Bibliographic references

ATM Fórum, The; "ATM Tutorial", 1996.

Moreira, E. D. S.; Reami, E. R.; Pissioli, M; "Live Video for Distributed Multimedia/Hypermedia Applications: Directly Connecting Video Sources to TCP/IP Networks", proceedings do MIPRO95 – Multimedia and Hypermedia Systems, Opatija, Croatia, Dec. 1995 - ISBN 953-6042-10-X

Moreira, E. D. S.; Encinas, W.S.; Ramos, A.C.B.; "Bringing wireless video to networked multimédia systems", Displays International Journal, Elsevier Science B. V., n. 17, p. 207-215, 1997.

Murphy, J., "Statistical Multiplexing", http://thorung.eeng.dcu.ie/~murphyj/the/the/node19.html

Nahrstedt, K., Steinmetez, R. "Resource Management in Networked Multimedia Systems". *IEEE Computer*, May 1995.

Raychaudhuri D., "Wireless ATM Networks: Architecture, System Design and Prototyping," IEEE Personal Communications, pp. 42-49, Oct 1996.

Riley, M. J., Richardson, I. E. G., "Digital Video Communications", Artech House, Boston, MA, 1997.

Sirivara, S.; "ATM Technology & Applications", http://www.eece.unm.edu/students/sudheer/atm/atm3.html

Tanembaum, A.; "Computer Networks", 3rd edition, Prentice-Hall, 1996

Umehira M., et al, "ATM Wireless Access for Mobile Multimedia: Concept and Architecture," IEEE Personal Communications, pp. 39-48, Oct 1996.

Walke, B., Petras D., Plassmann D., "Wireless ATM: Air Interface and Network Protocols of the Mobile Broadband System," IEEE Personal Communications, pp. 50-56, Oct 1996.