The University of Kansas

Information and Telecommunication Technology Center

Technical Report

Performance Comparison Between AAL1, AAL2 and AAL5

Raghushankar R. Vatte and David W. Petr

ITTC-FY1998-TR-13110-03

March 1998

Project Sponsor:
Sprint Corporation

Contents

Abstract

Presently efforts are being made to standardize the AAL type 2 layer specifications by the ATM forum and the ITU-T. This simulation is used to compare the performance of AAL1, AAL2 and AAL5 for Voice Transport over ATM (VTOA). We compare the maximum number of users with a given voice coding rate that can be served by a DS1 link (1.536 Mb/s), subject to an upper bound on the 95th percentile of voice packet delay, which includes transmitter queuing delay. The optimal voice packet size (CPS_packet size) is also found for each case. The simulations are performed for homogeneous sources only.

1 Introduction

The objective of this study is to compare the performance of AAL1, AAL2 and AAL5 for Voice Transport over ATM (VTOA). Specifically, our goal is to compare the maximum number of users with a given voice coding rate that can be served by 1.536 Mb/s of available bandwidth (corresponding to a DS1 link) subject to an upper bound on the 95th percentile of voice packet delay (explained below). For AAL2, this bandwidth would be in the form of a Virtual Circuit Connection (VCC), while for AAL1 and AAL5 it would be a Virtual Path Connection (VPC). Previous studies (such as [6]) have made similar comparisons, but with simplified analyses. In particular, they have not considered the effects of queuing delay in the AAL transmitters. The primary contribution of this report is to include these effects and to show that they can significantly reduce the number of connections that can be supported. As part of this process, we also determine the optimal voice packet size for each voice coding rate.

The Kth percentile of voice packet delay is the delay D_K such that K% of voice packets have a delay less than or equal to D_k. This is an important performance measure for voice transport because the voice receiver typically attempts to "build out" the overall delay for every voice packet to some fixed value, so that the time relationships between the packets are preserved across the network. This overall fixed delay is often taken to be the Kth percentile delay (if known). In this case, it can be expected that only (100-K)% of the packets would arrive "late" at the receiver. There are several options for dealing with the late packets, but the simplest one is to simply discard them. Larger values of K reduce the number of late (discarded) packets , but produce larger values of D_k. We have chosen K equal to 95 as a reasonable value for this study, especially considering that we are dealing only with transmitter delay and not end-to-end network delay.

The simulation models for the AAL2 system are the same as in [1] and the simulation models for AAL1 and AAL5 are the same as in [2]. The simulation strategy, including a discussion of simulation parameters, is explained in section 2. Results are presented and discussed in section 3, followed by conclusions in section 4.

2 Simulation Strategy

We use BONeS [9,10] simulation models that were previously developed for voice sources and AAL2 as described in [1] and for AAL1 and AAL5 as described in [2]. Each individual source is modeled as a fixed bit rate, on-off voice source with exponentially distributed on and off periods. For this study, all sources for a given simulation are statistically independent and have the same voice bit rate (homogeneous sources), with a mean on time of 420 ms and a mean off time of 580 ms. Each voice bit stream is packaged into packets of fixed maximum size (referred to below as CPS_Packet Size, where CPS stands for Common Part Sublayer). The AAL2 model is a very detailed and accurate representation of the transmitter processing specified in [3]. For AAL2, all sources are multiplexed into a single VCC using a single AAL2 transmitter, allowing packets from different sources to be included in a single ATM cell. For AAL1 and AAL5, each source has its own AAL transmitter (each cell contains packet data from a single source), and the transmitter outputs (cells) are multiplexed onto a single VPC. Packet delay is measured from the time the first bit of a packet arrives at the packetization function until the last bit of the packet reaches the emulated receiver. The network delay between the transmitter and receiver was set to zero, so the measured delay is primarily composed of packetization delay (time to accumulate a packet at the voice bit rate) and transmitter queuing (multiplexing) delay. More details on the simulation models and parameters can be found in [1] and [2].

The main aim of these simulations is to find the maximum number of users which can be supported on a VCC or VPC with a given set of parameters. This is done in 2 phases:

- The first phase consists of narrowing the range of number of users where the maximum occurs. Here the CDF plots and the 95th percentile delay plots are used (explained in section 2.2).
- The second phase consists of zeroing in on the value of the maximum number of users supported.

For this to happen we first need to specify the fixed and variable parameters for the simulations.

2.1 Parameters

The parameters which are fixed across all simulations are :

- Peak VCC Rate (1536 kb/s).
- 95th Percentile Delay Bound (10 ms).
- Speech Activity Factor (silence detection is enabled, with 42% speech activity).
- Trunking (as opposed to wireless).

The parameters which are varied are :

- Voice Bit Rate (Coding Rate).
- CPS_Packet Size.
- Number of Users.
- Timer Setting.

Each of the parameters is discussed briefly below.

- PEAK VCC RATE: The Peak VCC Rate [1] determines the permit arrival rate for AAL2 (each cell requires a permit to leave the transmitter) given by

$$\frac{8 * 53}{p} \, ms \tag{1}$$

where the peak VCC rate is p kbps. The Peak VCC Rate in our case is 1536 kb/s which corresponds to the DS1 (T1) Rate. This gives us a value of 276 μs for the permit arrival rate (minimum time between ATM cells on the VCC or VPC).

- 95th PERCENTILE DELAY BOUND: One concern associated with packetized voice is the extra packetization and queuing delay incurred, which must be controlled to avoid excessive requirements for echo cancellation. Echo cancellers are usually required when the total 1-way delay in a connection exceeds approximately 25 ms. In this study, we allow 40% of that delay (10 ms) to be incurred in the AAL transmitter (packetization and queuing). This leaves 15 ms for other network queuing delays , propagation delay, coder/decoder delay[2], etc. Values of 95th Percentile Delay Bound larger or smaller than 10 ms would affect the absolute number of connections that could be supported, but are not likely to significantly affect the comparisons between AAL1, AAL2, and AAL5.

- SPEECH ACTIVITY FACTOR: Since voice contains considerable amounts of silence periods, data need not be transmitted during these periods and instead a silence identifier is transmitted to identify the beginning of a silence period. This process results in the efficient use of the bandwidth. The percentage of actual speech present in voice is known as Speech Activity Factor. This has been experimentally determined to be about 42%.

- TRUNKING: For the trunking scenario, we assume that the speech has not incurred any packet-formation delay prior to the AAL transmitter.

- VOICE BIT RATE: 4 Voice Bit Rates (Coding Rates) are considered: 8 kb/s, 16 kb/s, 32 kb/s, 64 kb/s. These bit rates are chosen so that they cover most of the general purpose coders available now.

- CPS_PACKET SIZE: The CPS_Packet Size is varied in accordance with the Coding Rate chosen. Both Coding Rate and CPS_Packet Size influence the packetization (segmentation) delay in the transmitter. The total packet delay in the transmitter is the sum of the packetization delay and the queuing delay in the transmitter . The packetization delay is given by

$$\frac{(CPS\_PacketSize * 8)}{(CodingRate)} = PacketizationDelay \tag{2}$$

[1]For AAL1 and AAL5 this would be the peak VPC (Virtual Path Connection) rate since AAL1 and AAL5 do not allow multiplexing within a VCC.

[2]Note that the coder/decoder delay of the 8kb/s coder is large enough by itself to require echo cancellation, while the 16, 32, and 64 kb/s coders all have delays less than 1 ms.

which is the time taken by the coder to accumulate data of size CPS_Packet Size (in bytes) at Coding Rate. The maximum CPS_Packet Size is then obtained when the packetization delay equals the 95th percentile delay bound in the transmitter (in our case 10 ms). This is given by

$$\frac{(95\% DelayBound(ms) * CodingRate(kb/s))}{8} = max.CPS\_PacketSize(bytes) \quad (3)$$

- NUMBER OF USERS: From the remaining parameters the maximum number of users supported is calculated by using the equation (for AAL2)

$$n * R * 0.42 * ((CPS\_Packetsize + 3)/CPS\_Packetsize) * (53/47) \ < p \quad (4)$$

  where we are assuming 42% speech activity, R kb/s Coding Rate, Peak VCC Rate p, and the number of users being n. The left side of the inequality is the approximate aggregate mean arrival rate at the transmitter, including ATM and AAL2 overhead. Similar inequalities hold for AAL1 and AAL5. Then the parameter Number of Users is varied from a low value of 5 to its maximum value or a value where its 95th percentile delay exceeds the bound.

- TIMER SETTING: The Timer_CU Setting (applicable only to AAL2) gives the maximum time the AAL2 transmitter waits before sending a partially full cell. Timer_CU has a profound effect on the total packet delay at low loads, as the packets do not arrive quickly enough to fill cells before the timer expires (see results in [1]). In our case the timer is set to the maximum possible value, which is the difference between the 95th percentile delay bound in the transmitter and the packetization delay. Hence its value is adjusted whenever the delay bound or packet size changes. This strategy could present difficulties with heterogeneous sources (different sources have different coding rates), unless the packet sizes of the different sources could be adjusted so that the packetization delay (packet size divided by coding rate) is held constant for all sources (as illustrated in the discussion of results). Whether or not this is possible would probably depend on implementation flexibility at layers above the AAL2 Common Part Sublayer (CPS).

2.2 Simulation Strategy

The whole simulation process can be summarized in 10 steps :

1. Fixing all the parameters except number of users, then simulate the complete range with large steps (for numbers of users).[3] This is phase 1.
2. Look at the graph of 95th percentile delays vs. number of users produced in the simulations and decide on the range of values where the delay is going from an acceptable value (less than 10 ms in our case) to an unacceptable value (more than 10 ms in our case).

[3] Unless explicitly informed all the simulations in this section are run for 300,000 Packets.

3. Simulate in the range decided in step 2 in small steps (1, 2 or 3 users depending on the range) for number of users. This is phase 2.

4. From the 95th percentile graph obtained from simulations in step 3 find the value of the maximum number of users which satisfies the condition of having a delay of less than 10 ms.

5. Repeating steps 1 through 4, find these values for different CPS_packet sizes for a given voice bit rate.

6. Now pick the maximum of all the values obtained in step 5 to get the maximum number of users supported for a given voice bit rate using AAL2.

7. Repeat steps 1 through 6 for different voice bit rates using AAL2.

8. Now plot the graph of maximum number of users supported vs voice bit rate using AAL2.

9. Repeat steps 1 through 8 for AAL1 and AAL5.

10. Now put the 3 plots on a single graph for comparison of AAL1, AAL2 and AAL5 with respect to the maximum number of users supported vs. voice bit rate on a single VCC or VPC with a peak rate of 1536 kb/s.

2.3 95th Percentile Delay Graphs

All through the former section we were talking about the 95th percentile delay graphs. We will now briefly look into the process in which the 95th percentile delay graphs are generated.

1. First simulate for a given set of parameters with a CDF probe.

2. A CDF probe gives us the cumulative distribution function of the delay of all the packets leaving the transmitter. A CDF probe divides the delay range (in our case 0 to 15 ms) in to a certain number of bins (in our case 10000). Whenever a packet experiences a delay which lies in one of the bins, the count for that bin is increased by one. The percent of packets in a bin gives the probability mass function (PMF) of that bin. From the PMF the CDF plots are obtained by summing the PDF values for all bins with delay less than or equal to a given delay. Fig. 1 gives an example of a CDF plot.

3. From the CDF plot obtained we choose the minimum of all the bins which have a cumulative value of 0.95 or more. This gives us a single value corresponding to the 95th percentile delay. In Fig. 1, this delay is approximately 2.75 ms.

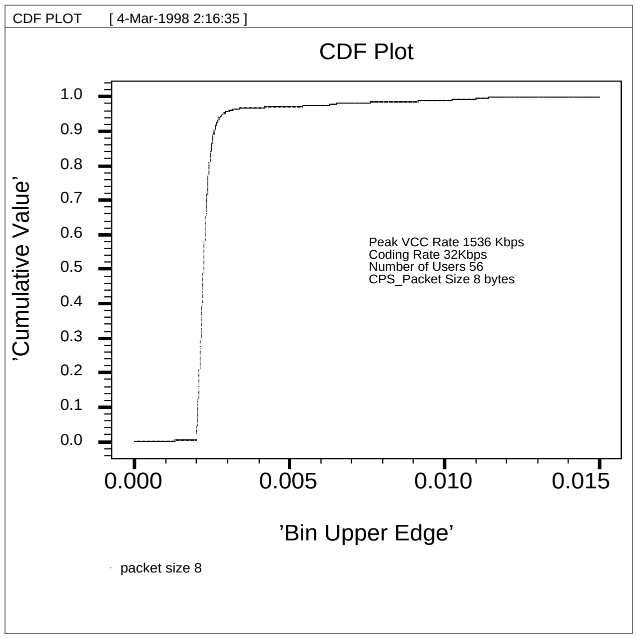

Figure 1: An Example of a Cumulative Distribution Function

4. Now that we have the 95th percentile delay for given parameters (CPS_Packet Size, Number of Users), we change the parameter Number of Users to get different values for different loads. We now have a plot of 95th percentile delay vs. number of users for a given CPS_packet size.

5. Steps 1 and 4 are repeated for different CPS_packet size to get a plot as shown in Fig. 2.

3 Results and Explanations

- Figures 2 through 5 are the 95th percentile delay plots for the first phase simulations for AAL2. Here each curve corresponds to a CPS_packet size, and these are used to narrow the region of search as mentioned earlier. For example in Fig. 2 for the CPS_Packet Size of 4 bytes the region of interest is from 200 to 220 users. Similarly for CPS_packet size of 6 bytes the region of interest is 200 to 250 users. Figures 3, 4 and 5 have similar plots for voice coding rates of 16 kb/s, 32 kb/s, 64 kb/s respectively for AAL2.

- Figures 6 through 9 are phase 2 plots for the voice regions of interest of Figures 2 through 5. We can see that in Fig. 6 the curves for CPS_packet sizes of 7, 8 and 9 bytes are well below the 10 ms mark even for the maximum number of users of 256.

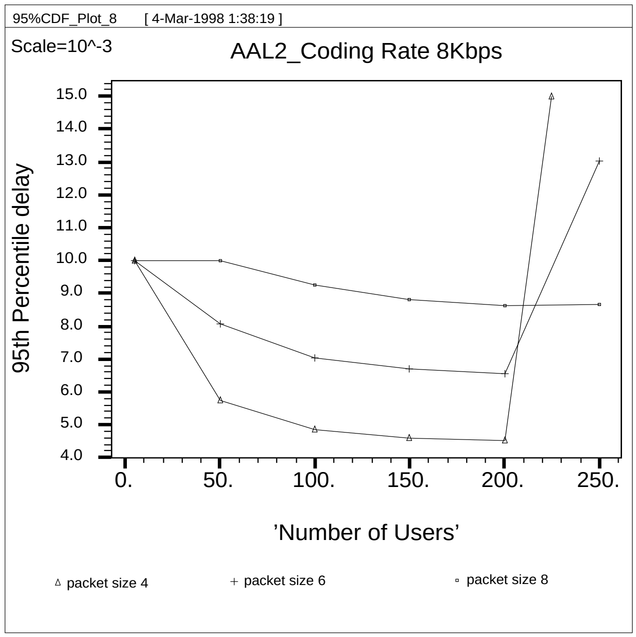

Figure 2: 95th Percentile Delay Plot for 8 kbp/s Voice Coding Rate with AAL2, First phase

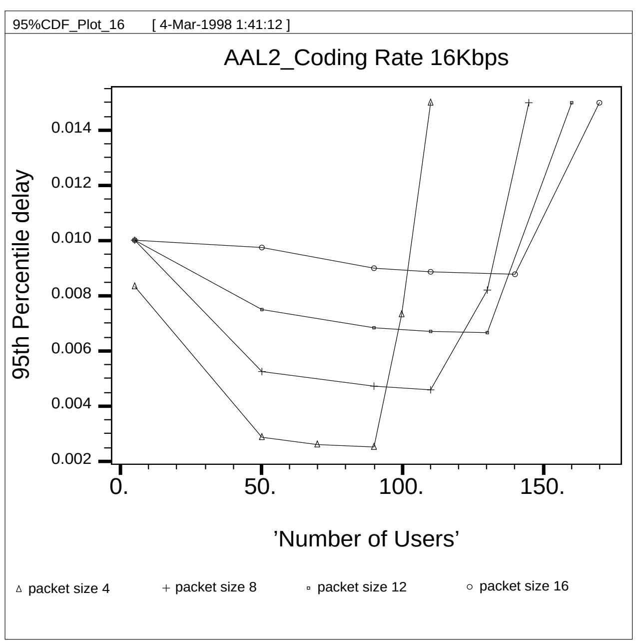

Figure 3: 95th Percentile Delay Plot for 16 kb/s Voice Coding Rate with AAL2, First phase

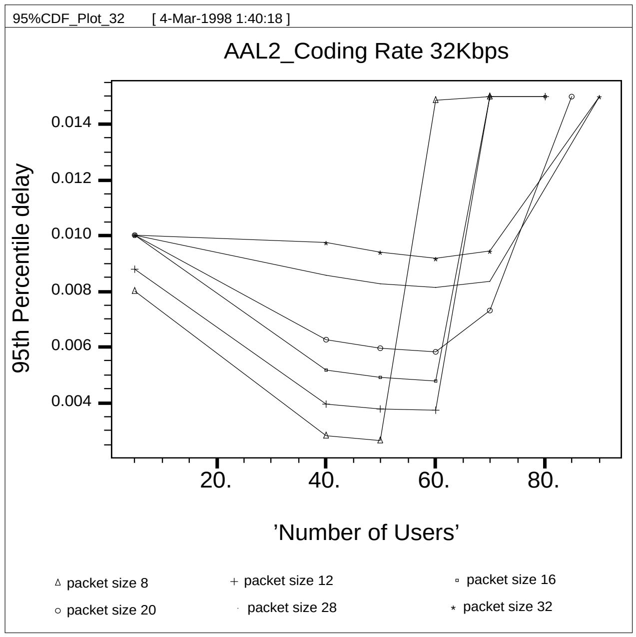

Figure 4: 95th Percentile Delay Plot for 32 kb/s Voice Coding Rate with AAL2, First phase

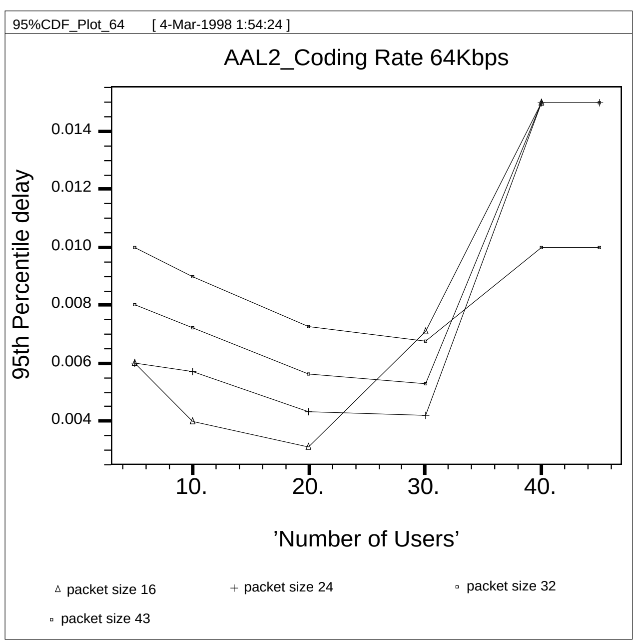

Figure 5: 95th Percentile Delay Plot for 64 kb/s Voice Coding Rate with AAL2, First phase

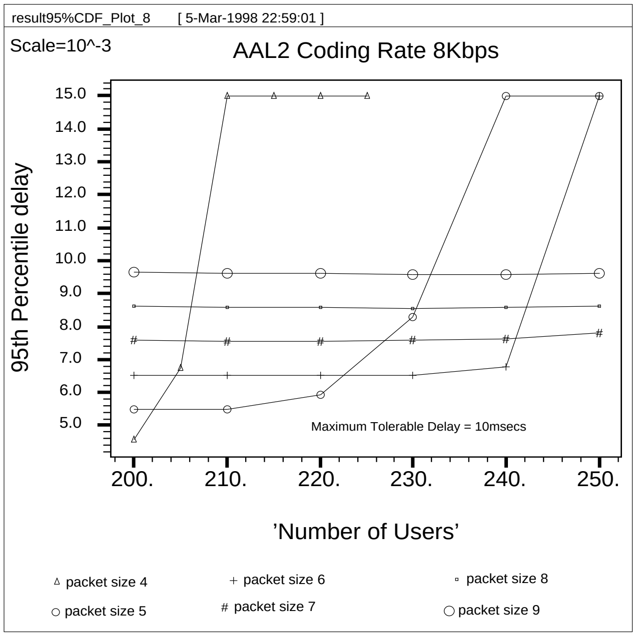

Figure 6: 95th Percentile Delay Plot for 8 kb/s Voice Coding Rate with AAL2, Phase 2

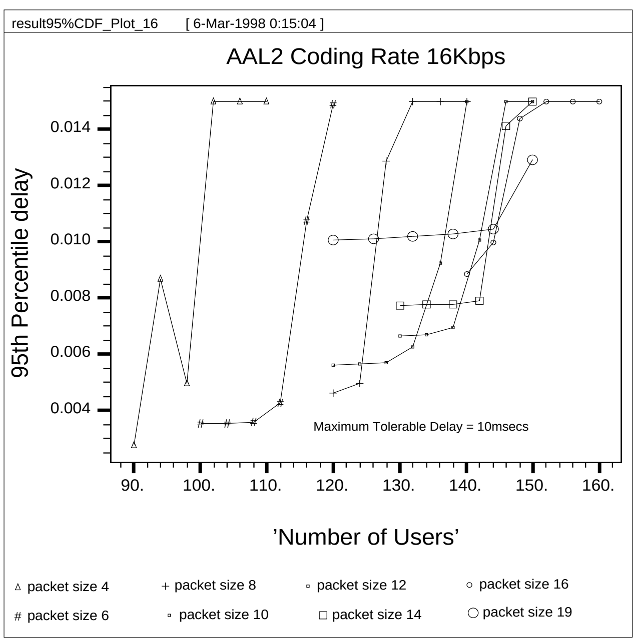

Figure 7: 95th Percentile Delay Plot for 16 kb/s Voice Coding Rate with AAL2, Phase 2

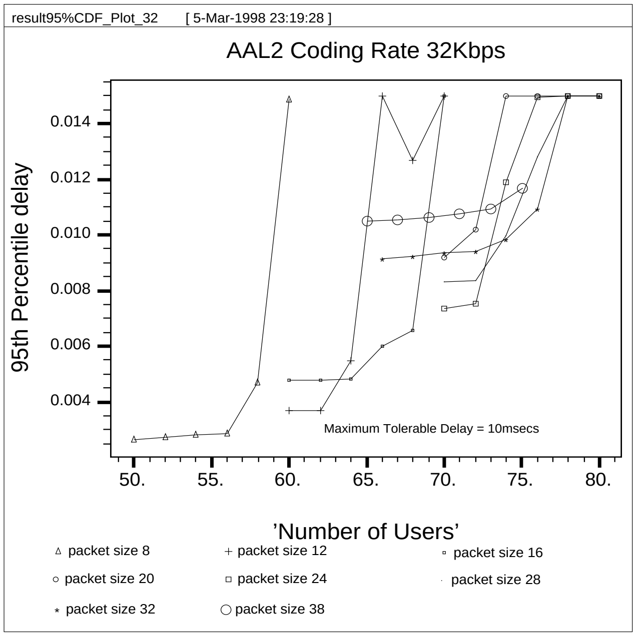

Figure 8: 95th Percentile Delay Plot for 32 kb/s Voice Coding Rate with AAL2, Phase 2

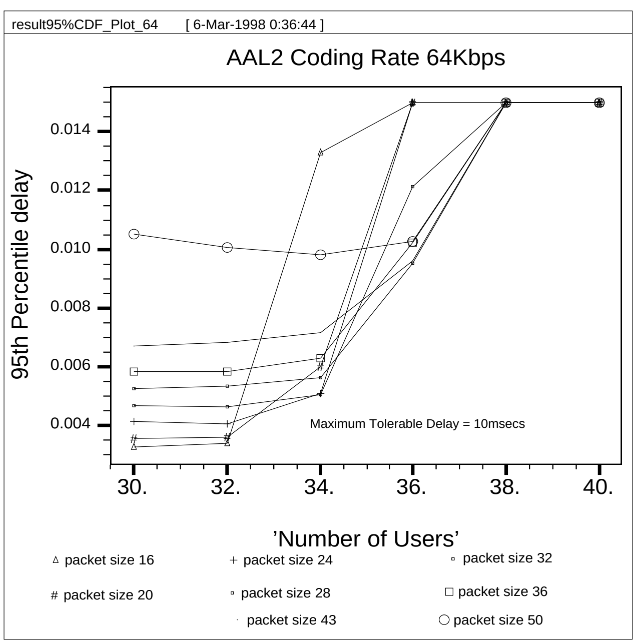

Figure 9: 95th Percentile Delay Plot for 64 kb/s Voice Coding Rate with AAL2, Phase 2

This is because the simulations were stopped at that point, as 256 is the maximum allowed limit for multiplexing in a single VCC in AAL2. Figures 7, 8 and 9 give us similar plots for voice coding rates of 16 kb/s, 32 kb/s and 64 kb/s respectively.

- Figures 10 through 13 give us the phase 2 plots for different voice coding rates for AAL1. These plots were generated following the same procedure as in the case of AAL2. We can clearly see that AAL1 supports fewer users on a single VCC when compared to AAL2.

- Figures 14 through 17 refer to similar phase 2 plots for AAL5 with different coding rates. These plots were generated following the same procedure adopted for generating AAL1 and AAL2 plots.

- Figure 18 is the summary of the plots 6 through 9 with respect to the maximum number of users supported for AAL2. It plots maximum number of users supported vs. CPS_packet size for different voice coding rates for AAL2. From Figures 6 through 9 the maximum number of users supported for 95th percentile delay bound of 10 ms is obtained for each CPS_packet size.

  The points where the curves of Fig. 18 go back to zero on the X axis (CPS_packet size equal to 10, 20, and 40 bytes) are not generated from simulations but are theoretically obtained. If we consider the point 20 bytes the Timer_CU will be ,

$$10 - \frac{20 * 8}{16} = 0 \tag{5}$$

  i.e. the packets have to be sent as soon as they arrive and there is no multiplexing in the cells. Moreover the packetization delay would be

$$\frac{20 * 8}{16} = 10ms \tag{6}$$

  which equals the tolerable limit set. Therefore in these situations even very low loads will have a delay of at least 10 ms[4]. The remaining curves have the same behavior. The curve for voice coding rate of 64 kb/s ends at CPS_Packet Size of 50 bytes as the simulations are conducted only until that point[5].

  Here we see that a maximum of 256 users can be supported for the 8 kb/s voice coding rate at CPS_packet size of 7, 8 and 9 bytes. For 16 kb/s the maximum is 144 users at a CPS_packet size of 16 bytes, for 32 kb/s it is 53 users and for 64 kb/s it is 36. An important point to be observed here is that at low voice coding rates the maximum number of users supported is very sensitive to the CPS_Packet Size and at higher voice coding rates the the maximum number of users supported does not vary much with

[4] In the case when Timer_CU value is non zero , the delay would only increase because of the timer. Therefore the Maximum number of Users supported is taken to be zero.

[5] Any value of CPS_packet size greater than 44 bytes results in spilling of packets into the next cell and thus causes the delay for each packet to increase.

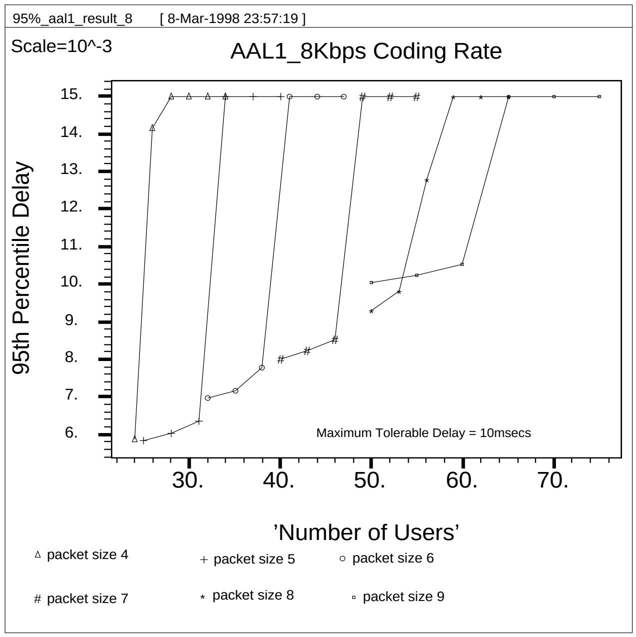

Figure 10: 95th Percentile Delay Plot for 8 kb/s Voice Coding Rate AAL1, Phase 2

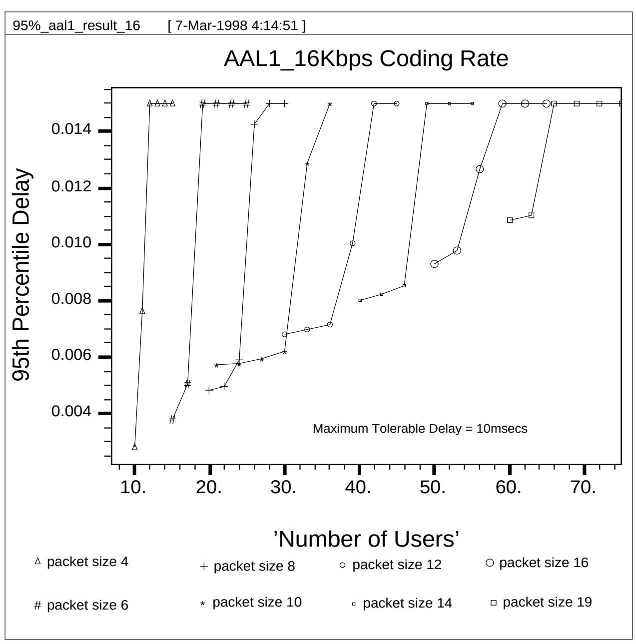

Figure 11: 95th Percentile Delay Plot for 16 kb/s Voice Coding Rate AAL1, Phase 2

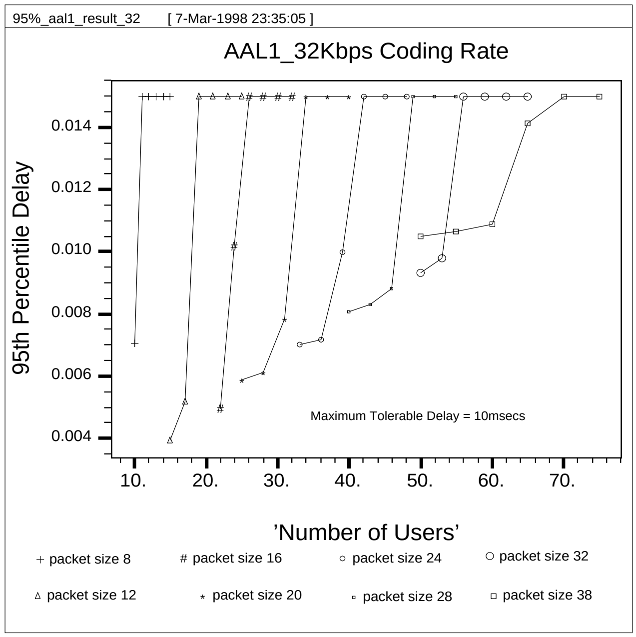

Figure 12: 95th Percentile Delay Plot for 32 kb/s Voice Coding Rate AAL1, Phase 2

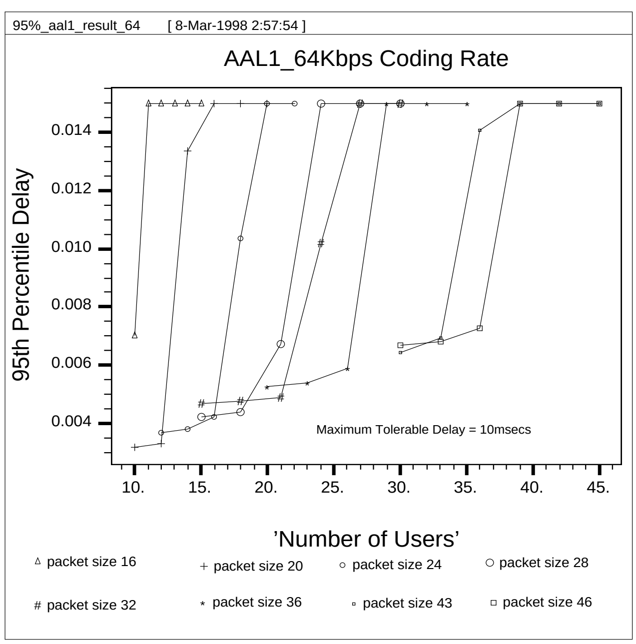

Figure 13: 95th Percentile Delay Plot for 64 kb/s Voice Coding Rate AAL1, Phase 2

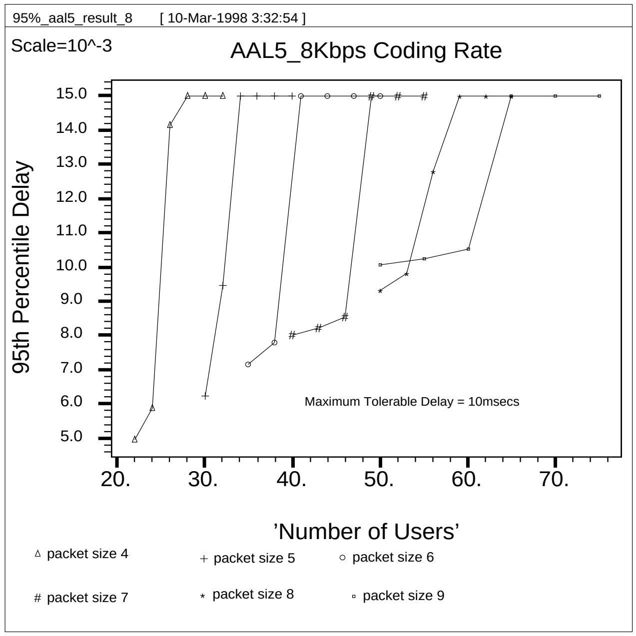

Figure 14: 95th Percentile Delay Plot for 8 kb/s Voice Coding Rate AAL5, Phase 2

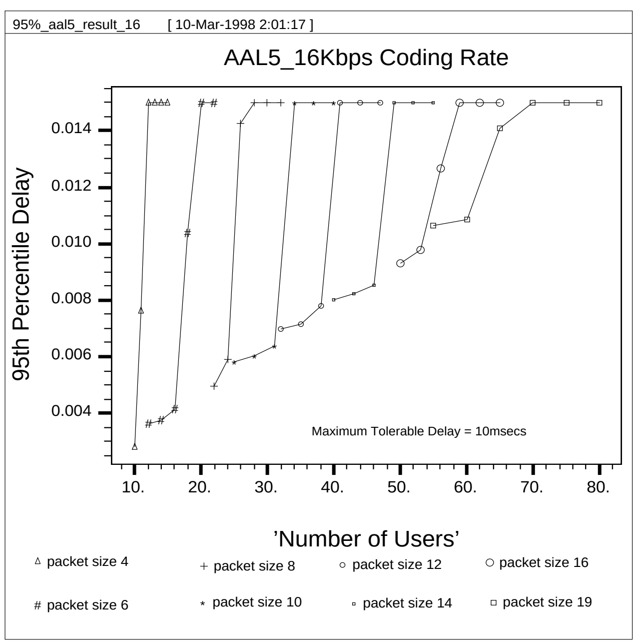

Figure 15: 95th Percentile Delay Plot for 16 kb/s Voice Coding Rate AAL5, Phase 2

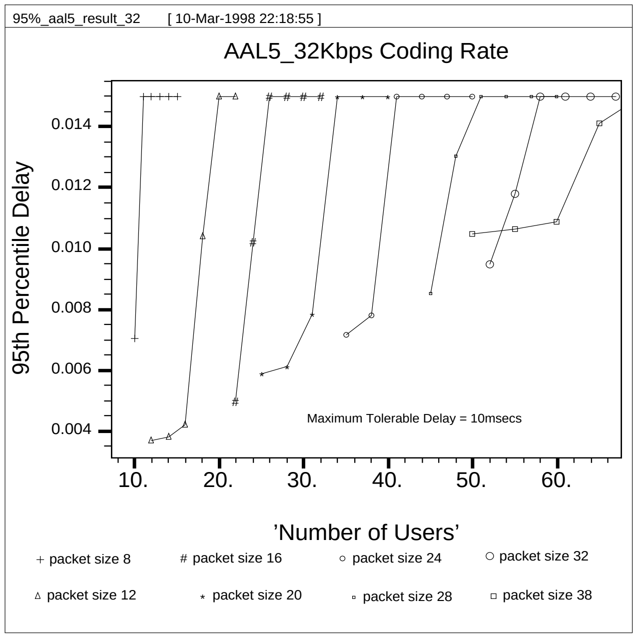

Figure 16: 95th Percentile Delay Plot for 32 kb/s Voice Coding Rate AAL5, Phase 2

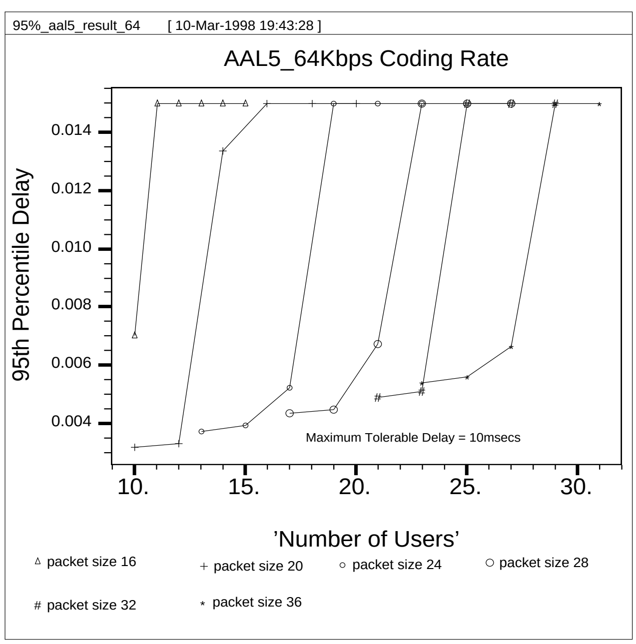

Figure 17: 95th Percentile Delay Plot for 64 kb/s Voice Coding Rate AAL5, Phase 2

the CPS_packet size chosen. Note that if we were to choose an 8-byte CPS_packet size for 8 kb/s, 16 bytes for 16 kb/s, and 32 bytes for 32 kb/s, the CPS_packet size would be near the optimum value and the packetization delay would be 8 ms for all of these voice coding rates, allowing a single Timer_CU value of 2 ms regardless of coding rate.

The other curves (in a dashed line) give the maximum number of users supported vs. CPS_packet size without taking the queuing delay into account, i.e., using equation 4, corresponding to 100% queuing load. This simplistic analysis has been often used in the literature, but we see that using a more realistic, simulation based evaluation produces significantly smaller maximum number of users. (approximately 80% of the numbers produced by equation 4)

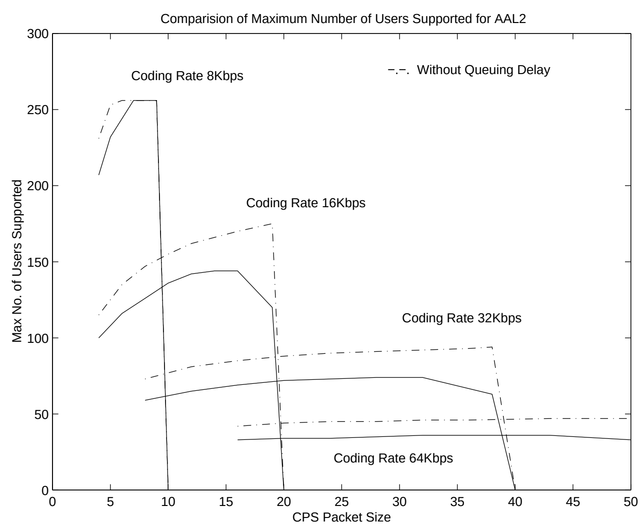

Figure 18: Plot for Maximum Number of Users as a Function of CPS_Packet Size and Voice Coding Rate for AAL2

- Figure 19 gives the maximum number of users supported vs. CPS_packet size for different voice coding rates for AAL1. This figure is obtained from Figures 10 through 13 with a maximum tolerable delay of 10 ms in the transmitter. We can see that unlike Fig. 18 (AAL2), here the maximum number of users supported for a voice coding rate does not change significantly with the coding rate. The curve for 64 kb/s ends at CPS_packet size of 46 bytes as any packet more than 46 bytes would be accommodated in 2 cells and thus would only increase the delay and lower the maximum number of users supported. The points where the curves go back to zero have the same explanations as in the AAL2 case. The other curves (in a dashed line) give the maximum number of users supported vs. CPS_packet size without considering the queuing delay. The maximum number of users supported for coding rates of 8, 16, 32 kb/s is 53 users; for 64 kb/s the value is 32. Notice that the realistic numbers (from simulation) are about 75% of the numbers produced from theoretical calculations in the AAL1 case [2].

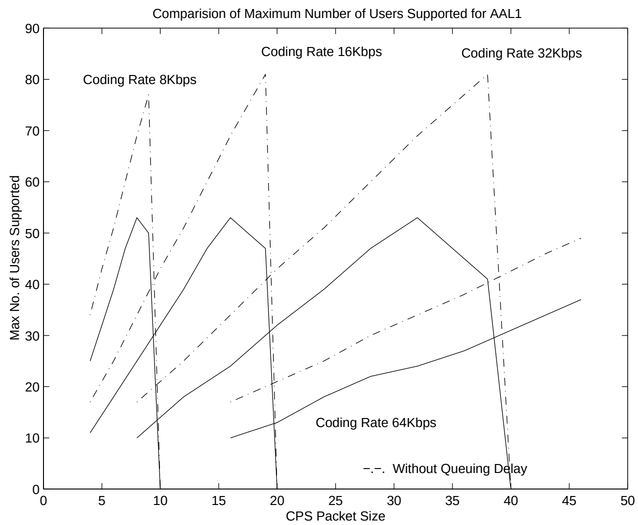

Figure 19: Plot for Maximum Number of Users as a Function of CPS_Packet Size and Voice Coding Rate for AAL1

- Figure 20 gives us the maximum number of users supported vs. CPS_packet size for different voice coding rates for AAL5. This plot is generated from Figures 14 through 17, with the 95th percentile delay bound of 10ms in the transmitter. This plot is almost the same as Figure 19 until the CPS_packet size of 40 bytes, as both AAL1 and AAL5 behave in a similar manner with respect to the overhead and performance until that point [2]. The maximum number of users supported for coding rates of 8, 16, 32 kb/s is 53 users; for 64 kb/s the value is 29. The other curves (in a dashed line) give the maximum number of users supported vs. CPS_packet size without considering the queuing delay. Notice that the realistic numbers (from simulation) are about 75% of the numbers produced from theoretical calculations in AAL5 case[2].
- Figure 21 summarizes the Figures 18 through 20 and has the plots for the maximum number of users supported for different voice coding rates for AAL1, AAL2 and AAL5. We can see that AAL2 supports substantially more users on a single VCC than AAL1 and AAL5 support on a VPC at low voice coding rates and is almost equal to AAL1 and AAL5 at high voice coding rates. We can also conclude that AAL2 is more sensitive to coding rate than AAL1 and AAL5 with respect to the maximum number of users supported.

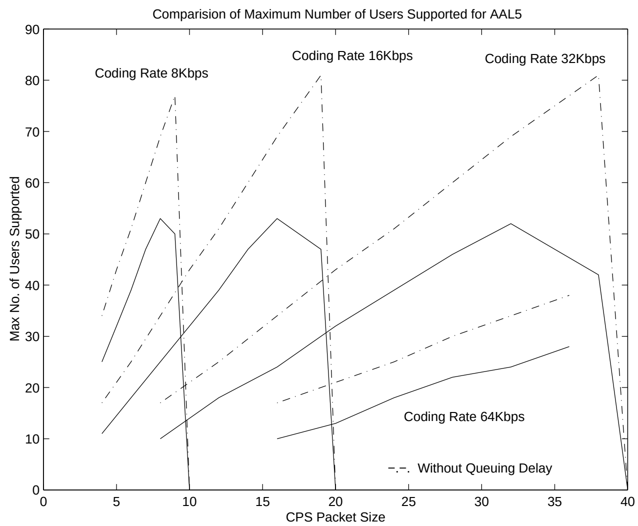

Figure 20: Plot for Maximum Number of Users as a Function of CPS_Packet Size and Voice Coding Rate for AAL5

4 Conclusions

This simulation study allows us to draw several important conclusions about the performance of AAL2, especially relative to AAL1 and AAL5. Before enumerating these conclusions, however, we caution the reader to keep in mind some of the limitations of this study, particularly regarding the assumptions which include homogeneous sources, specification of service quality in terms of 95th percentile delay, 10 ms value for 95th percentile delay bound in the AAL transmitter, available bandwidth of 1.536 Mb/s, and a trunking (as opposed to wireless) application scenario. We also remind the reader that only transmitter-induced delays have been considered.

- For a realistic comparison, 95th percentile delay (or 96th to 99th percentile delay) is more appropriate than absolute maximum delay or average delay.

- The maximum number of users supported on a VCC for a given coding rate is almost 5 times more for AAL2 than AAL1 and AAL5 at 8 kb/s voice coding rate and this advantage decreases with the increase in the coding rate until eventually the maximum number of users supported on a VCC for AAL2 almost equals that of AAL1 and AAL5 at 64 kb/s coding rate.

- For AAL2 the total packet delay is influenced by Timer_CU to a great extent at low loads. In fact the delay is very near to the Packetization Delay + Timer_CU.

- For AAL2 the total packet delay decreases as the load increases until it reaches a queuing saturation point after which even a slight increase in load causes the total packet delay to increase sharply.

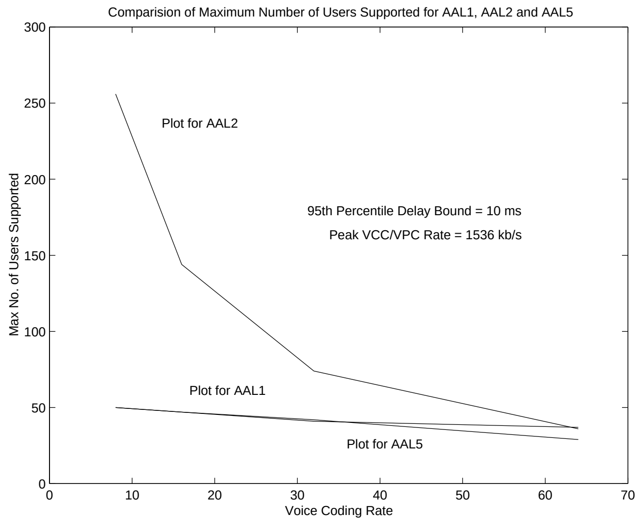

Figure 21: Maximum Number of Users Supported on a VCC for AAL1, AAL2 and AAL5 as a function of Voice Coding Rate

- For AAL2 as we increase the CPS_packet size (in turn decreasing the Timer_CU) the maximum number of users supported on a VCC increases until the Timer_CU approaches twice the Permit Arrival Rate after which it starts to decrease rapidly eventually falling back to zero at Timer_CU equal to 0.
- For AAL2 the maximum number of users supported on a VCC is very sensitive to the CPS_packet size at low voice coding rates and is quite insensitive at high voice coding rates. In other words, at high coding rates it does not really matter what CPS_packet size we are using but it can significantly affect the efficiency at low coding rates.
- For all the AALs the maximum number of users supported on a VCC decreases with the increase in voice coding rate (with CPS_Packet Size optimized), but the decrease is much more rapid for AAL2.
- For AAL1 and AAL5 the maximum number of users supported on a VPC increases with the CPS_packet size until a point where it saturates and then decreases rapidly. This sensitivity to CPS_packet size is similar for all the voice coding rates (in the range of CPS_Packet Size 1 to 40 bytes).
- AAL5 behaves in the same way as AAL1 till the CPS_packet size of 40 bytes and has the same characteristics as the above result in that range.
- The maximum number of users supported on a VPC for a given coding rate is almost equal for AAL1 and AAL5 until 32 kb/s and then they differ by a small margin afterwards.

5 References

1. AAL2 Transmitter Simulation Study : Revised. Prema Sampath, Raghushankar R. Vatte, D.W. Petr. Technical Report ITTC-FY98-TR-13110-01, March 1998.

2. Development of Simulation Models for AAL1 and AAL5. R. Yellisetti, D.W. Petr. Technical Report ITTC-FY98-TR-13110-02, February 1998.

3. ITU-T Draft Specification I.363.2, February 1997.

4. ATM Forum Contribution ATM-f/96-0732.

5. ATM forum Contribution ATM-f/96-1134, "Additional simulation tests for AAL-CU," Hisakazu Tsuboya, NTT mobile communication Network Inc., August 1996.

6. ATM Forum Contribution ATM-f/96-1330, "Comparison of AAL2-CU Protocol Proposals," Ericsson Radio Systems AB, October 1996.

7. ATM Forum Contribution ATM-f/97-0575, "Robustness in an AAL2 Voice SSCS," Ericsson Telecom AB, July 1996.

8. ATM Forum Contribution ATM-f/97-0513, "Future Support of AAL2 to AAL5 Interworking," Nortel Telecom, July 1996.

9. Bones Designer Core Library Reference.

10. Bones Designer Users Guides: FSM Editor, Project Editor, & ISM.