Efficiency of AAL2 for Voice Transport: Simulation Comparison with AAL1 and AAL5

*David W. Petr, Raghushankar R. Vatte, Prema Sampath, Yong-Qing Lu**

Information & Telecommunication Technology Center
Department of Electrical Engineering and Computer Science
The University of Kansas
Lawrence, KS 66045-2228
Phone: 785-864-7762 FAX: 785-864-7789

*Sprint Corporation, Overland Park, KS 66212

Abstract

We evaluate the efficiency of AAL2 for voice transport in ATM networks using a simulation approach. After highlighting some of the general characteristics of AAL2, we adopt as the QoS requirement a bound on the K^{th} percentile delay in the AAL transmitter. We then determine the maximum number of homogeneous voice sources of a given bit rate that can be supported on a given bandwidth without violating this QoS requirement for AAL1, AAL2 and AAL5. Comparisons with simple analytical approximations are also provided. We conclude that AAL2 can support approximately 5 times as many 8 kb/s voice users as either AAL1 or AAL5, but all AALs can support approximately the same number of 64 kb/s users.

1: Introduction

1.1: Background

There has been a recent resurgence of interest in transporting voice over Asynchronous Transfer Mode (ATM) networks, evidenced by the introduction and continuing standardization of the redefined ATM Adaptation Layer type 2 (AAL2) that fills a need for multiplexing several small data units associated with different voice or data connections (known as logical link connections or LLCs) into a single ATM cell stream for transport over a single ATM virtual channel connection (VCC). This in-cell multiplexing limits packetization delay for compressed voice without wasting transmission bandwidth due to partial filling of ATM cells. The original motivation for AAL2 was to support compressed voice as found in digital cellular systems, but voice trunking and ATM to the desktop have since been added as potential AAL2 applications. For more details on potential AAL2 applications, see [1, 2]. AAL2 was originally designated AAL-CU (AAL for Composite Users).

Other alternatives for voice transport include AAL1 and AAL5 [1]. AAL1 is primarily intended for real-time constant bit rate (CBR) traffic. AAL5 is primarily intended for non-real-time data communications with relatively large data units. Neither supports the type of small data unit multiplexing that AAL2 provides. AAL1 can be very efficient if all cells are completely filled, but this would result in unacceptably large cell-formation delays (packetization delays) in most applications involving compressed voice. The large data unit overhead of AAL5 makes it quite inefficient when used with small data units, such as could fit in a single ATM cell.

AAL2 consists of a Common Part Sublayer (CPS) shared by all LLCs using a given ATM VCC combined with a Service Specific Convergence Sublayer (SSCS) for each LLC. The AAL2 Common Part Sublayer defines procedures for multiplexing small data units (called CPS-Packets) into ATM cells and for detection of such transmission impairments as bit errors and ATM cell losses. Each Service Specific Convergence Sublayer (several of which may eventually be defined) provides end-to-end services required by a specific application, such as compressed voice transport or data transport. Several different SSCS layers can be used within a single ATM VCC. In this paper, we are concerned with the performance of the AAL2 CPS for carrying voice LLCs in which the SSCS is assumed to provide only segmentation and reassembly of the coded bits in each voice talkspurt.

1.2: Previous Work on AAL2

Most of the technical community's efforts relative to AAL2 have been within the context of standards organizations such as ITU-T and the ATM Forum. This work has resulted in a very mature specification for the AAL2 CPS in ITU-T Recommendation I.363.2 [3]. This specification includes a timer (called Timer_CU) that allows partially filled cells to be transmitted, a permit arrival process governing the transmission of all cells, and a 3-octet header for each CPS-Packet. In addition, AAL2 includes a 1-octet overhead in each ATM cell which, combined with the 5-octet header already defined for each 53-octet ATM cell, reduces to 47 the number of octets in each ATM cell that can be used for CPS-Packets.

For the most part, efficiency analyses for AAL2 have been based on overhead percentages, bounding techniques or simple statistical calculations. For example, [4] compared AAL2 proposals in terms of overhead penalty (or percentage), defined as the ratio of overhead bytes to total AAL2 bytes (48 per ATM cell), $assuming$ that every ATM cell is completely filled (no AAL2 padding). Similarly, one can easily calculate an upper bound on the number of voice LLCs that can share an ATM VCC based purely on protocol overhead values. First define the AAL2 expansion factor E_{AAL2} as the ratio of the total AAL2 bytes to voice sample bytes ("real" data), given that every ATM cell is completely filled (no AAL2 padding used). If all CPS-Packets are of maximum size (CPS_{max}), then we have:

$$E_{AAL2} = \frac{CPS_{max} + 3}{CPS_{max}} \times \frac{48}{47}. \qquad (1)$$

Note that the AAL2 overhead percentage previously defined is identical to $\frac{E_{AAL2}-1}{E_{AAL2}}$.

Now if we let R be the bit rate of each voice LLC when active, SAF be the speech activity factor of each voice LLC (no packets are generated during silence), and P be the peak allowed rate of

The authors can be contacted via e-mail at dwp@ittc.ukans.edu. This research is sponsored by Sprint Corporation.

the VCC, then the maximum number of LLCs N must certainly obey the following inequality.

$$N \times R \times SAF \times E_{AAL2} \times \frac{53}{48} < P \qquad (2)$$

where the factor $\frac{53}{48}$ could be considered the ATM expansion factor.

Such a simple calculation only provides an upper bound because there is no specification for Quality of Service (QoS), in particular, for delay caused by queuing. The authors in [2] go one step further, using some simple statistical modeling based on a Gaussian approximation to calculate the maximum number of LLCs that can be supported subject to a QoS requirement that there is no more than a 0.1% probability that packets will be "lost" because (in effect) the queuing delay exceeds a certain limit. For $R = 4$ kb/s, $SAF = 0.50$, $CPS_{max} = 10$ octets, and $P = 1.536$ Mb/s (T1 rate), their calculations would show that 459 LLCs can be supported if the queuing limit (excluding packetization delay) is 20 ms. For the same parameters, the upper bound from equation 2 is 524 LLCs. While this analysis is certainly an improvement, the simplicity of the techniques and the assumptions made raise questions about the accuracy of the estimate. This will be revisited in a later section.

1.3: Current Research

Our specific goal in the current research is to provide a useful and realistic evaluation of the efficiency of AAL2 for voice transport compared to AAL1 and AAL5. Previous efforts either have adopted performance metrics (e.g. percentage overhead) that cannot be readily used for service provisioning or have been too simplistic (as indicated in the previous section). In contrast, we establish a QoS metric (a delay bound) for the AAL transmitter(s) and then determine (via simulation) the maximum number of voice sources with given characteristics that can be carried over a link with a given bit rate while still meeting the QoS (delay) objective. With this simulation-based approach, we expect to obtain more accurate efficiency estimates; we are also able to characterize and illustrate the performance of AAL2 in a number of other ways.

2: Models

2.1: Overview

All modeling and simulation was done with the BONeS Designer simulation package. At the highest level, the system under consideration consists of multiple on-off voice sources, an AAL2 transmitter or multiple AAL1/AAL5 transmitters, and a simplified, emulated receiver.

2.2: AAL2 Model

The AAL2 transmitter block is modeled as a finite state machine (FSM) based on the model specified in the ITU-T Recommendation I.363.2 [3]. Incoming data units (talkspurt bit streams, in our case) are segmented into CPS-Packet payloads with a specified size (in octets) denoted as CPS_Packet_Size. [1] The packetization delay (time to accumulate one CPS-Packet payload) is given by:

$$Packetization\_Delay = \frac{CPS\_Packet\_Size \times 8}{Voice\_Coding\_Bit\_Rate} \qquad (3)$$

The CPS-Packets are then packed into ATM cell payloads. We focus in this paper on Non-Deterministic Bit Rate (Non-DBR) operation, in which partially filled cells are sent only if the Timer_CU has expired, and null cells are never sent. The value of Timer_CU effectively gives an upper bound on the time any CPS-Packet will wait before the corresponding ATM cell will be declared "ready" to be sent. In the model developed it is assumed that AAL2 permit arrival (each "ready" ATM cell must obtain a permit in order to be transmitted) is determined by the peak VCC rate. Under these conditions, the CPS-Packet delay incurred in the transmitter is the sum of the packetization delay, the waiting delay for declaring an ATM cell ready to send (influenced by traffic load, Timer_CU and CPS_Packet_Size), the ATM cell transmission (or clocking) delay (influenced by rate of permit arrival and hence by peak VCC rate), and the queueing delay that results when CPS-Packets momentarily arrive faster than ATM cells can be sent.

[1] The last CPS-Packet in a talkspurt may be shorter than all the other CPS-Packets.

2.3: AAL1 and AAL5 Models

For voice transport applications, AAL1 and AAL5 would require an ATM VCC for each LLC (voice source). Hence the simulation model for each of these two AALs [9]associates an independent AAL (and associated ATM VCC) with each voice source, with the resulting ATM cells from all sources being multiplexed with a FIFO buffer onto a Virtual Path Connection (VPC). Fair comparisons can be made when the peak rate of the VPC for the AAL1/AAL5 case is equal to the peak rate of the VCC for the AAL2 case. In addition, the parameter Assembly_Size for AAL1/AAL5 plays the same role as CPS_Packet_Size for AAL2, including its role in calculating packetization (assembly) delay. The primary difference, of course, is that if the Assembly_Size plus the AAL overhead is less than the ATM payload, any remaining octets always will be "wasted" with padding and the ATM cell will be sent immediately.

Although larger assembly sizes tend to trade increased packetization delay for increased transmission efficiency, there is a limit to this tradeoff. Due to the AAL1 overhead (1 octet per cell), an AAL1 Assembly_Size greater than 47 octets would actually reduce transmission efficiency since 2 ATM cells would be required. The corresponding Assembly_Size limit for AAL5 is 40 octets. In fact, for Assembly_Size of 40 octets or less, the performance of AAL1 is identical to the performance of AAL5, since both AALs produce exactly one ATM cell for each assembly unit.

2.4: Voice Sources Model

We concentrate here on voice trunking applications with relatively high voice coding rates. Each voice source is considered to generate a constant rate bit stream interrupted by silence intervals. Hence each voice source is modeled as an on-off source with independent exponentially distributed on and off times. Speech activity factor is given by the ratio of the mean on time to the sum of the mean on and off times. In most of the simulations reported here, mean on and mean off times were 420 ms and 580 ms, respectively (42% speech activity). In some cases, we report results without silence elimination, in which case the speech activity is 100%.

The model allows any number (up to the AAL2 limit of 256) of independent voice sources to be simulated as the user population. After segmentation, the voice segments are either sent directly to an AAL1 or AAL5 transmitter, or they are multiplexed together and presented to a single AAL2 transmitter. The initial generation of talkspurts by the sources is randomized, and the start of data collection is delayed to allow for the system to reach steady state. This paper reports results for homogeneous sources.

2.5: Performance Metrics

The primary performance metric used here is delay for the CPS-Packets (or assembly units for AAL1/AAL5). Delay is measured from the generation of the first bit of a CPS-Packet (or assembly unit) until the last bit is received at the AAL receiver. This delay consists of packetization delay, Timer_CU delay (if any), permit delay (cell clocking delay), and queuing delay. Delay histograms can be collected in addition to statistical measures such as the mean.

In addition to delay, the simulation model also allows for the measurement of two other performance metrics. We define AAL (or AAL+ATM) efficiency as the mean of the AAL payload octets in an ATM cell divided by 48 (or 53), so that AAL efficiency and AAL overhead penalty (as defined in section 1) sum to 100%. We define bandwidth efficiency relative to 64 kb/s TDM as 64 kb/s divided by the measured, per-user bit rate of the ATM cell stream. Bandwidth efficiency reflects voice coding gain, speech activity factor, and AAL+ATM efficiency.

3: Preliminary AAL2 Performance Characterization

The simulation models outlined in the previous section have allowed us to perform a thorough performance characterization of AAL2 as various parameter values are changed. From the baseline parameter values listed in Table 1, each parameter value was varied individually as in Table 2.

This general characterization has been useful for illuminating some important properties of AAL2. For example, in most multiplexing systems, the delay increases as the number of users increases. For AAL2, however, packets may be delayed under low loads while waiting for other packets to arrive and help fill an ATM cell. This waiting delay decreases as the number of users increases, as illustrated in Figure 1. Note that the mean delay at low loads approaches the sum of the packetization delay (12 ms) and the Timer_CU value (3 ms), since nearly every ATM cell will wait for Timer_CU to expire. As the load increases, the mean delay approaches the packetization delay because the abundance of CPS-Packet arrivals ensures a short waiting time until enough have arrived to fill an ATM cell. Although not shown in Figure 1 due to the limit on number of users, this trend would eventually be reversed as the effects of queuing delay begin to dominate at higher loads. We will return to this point in the next section.

An appropriate choice of the Timer_CU value also depends on the traffic load, as illustrated in Figure 2. Mean delay is insensitive to Timer_CU value beyond approximately 1 ms for the following reason. For the given parameter values at 100% speech activity, a CPS-Packet is generated for each of the 64 sources every 12 ms, so that the mean time between CPS-Packet arrival at the AAL2 transmitter is 12/64 ms. With a 12-octet CPS_Packet_Size (15-octet total size of CPS-Packets), just over 3 CPS-Packets can fit in each ATM cell. In order for the ATM cell to be filled, one of these CPS-Packets must wait for one other CPS-Packet arrival, one must wait for two arrivals and the third must wait for three arrivals, for a mean waiting time of $\frac{1}{3} \times \frac{12}{64} \times (1+2+3) = 0.38$ ms. If the Timer_CU value is significantly greater than this, it has virtually no effect on the AAL2 operation, and the total delay saturates at 12 ms (packetization delay) + 0.38 ms (waiting delay), as shown in Figure 2. For 42% speech activity, the mean waiting time is multiplied by a factor of $\frac{1}{0.42}$, yielding a mean waiting time of approximately 0.9 ms.

Choice of CPS_Packet_Size is also important to performance, as illustrated for efficiency measures in Figure 3. As one would expect, this figure shows that efficiency increases with CPS_Packet_Size through a CPS_Packet_Size (payload) of 44 octets. At this value, efficiency peaks because the total size of each CPS-Packet is then 47 octets, and each CPS-Packet would completely fill an ATM cell. However, operating with this CPS_Packet_Size would essentially defeat the purposes of AAL2. Increasing CPS_Packet_Size to 45 octets results in a drastic drop in efficiency since each CPS-Packet would require slightly more than one ATM cell.

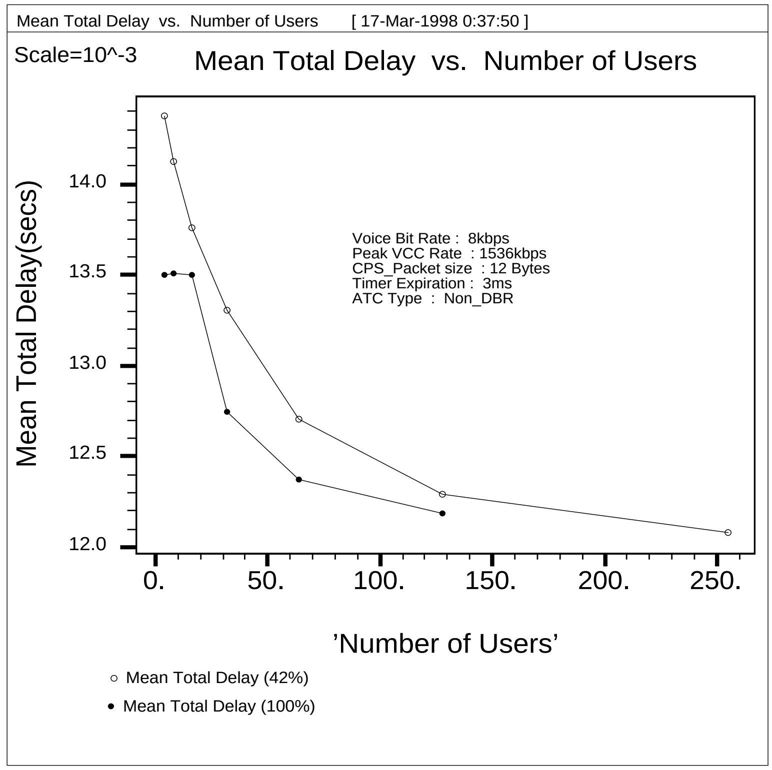

Figure 1. Mean Total Delay vs. Number of Users for AAL2

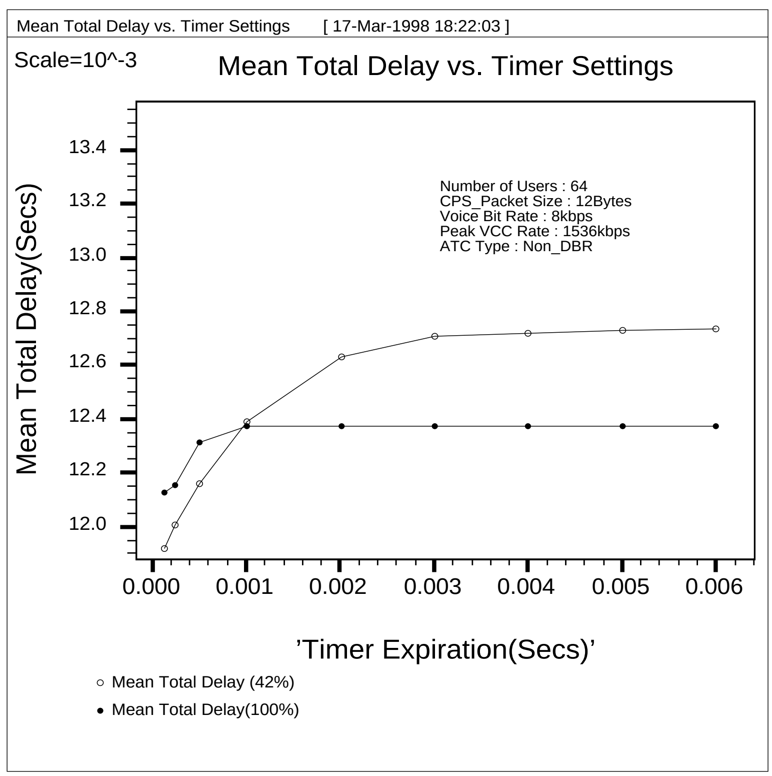

Figure 2. Mean Total Delay vs. Timer_CU Values for AAL2

Table 1. Baseline AAL2 Parameter Set

Parameter	Value
CPS_Packet_Size (excludes 3-octet CPS-Packet header)	12 octets
Voice Bit Rate	8 kb/s
VCC Peak Rate	1.536 Mb/s
Number of Users	64
Timer_CU	3 ms

Table 2. AAL2 Parameter Variations

Parameter	Values
CPS_Packet_Size [Octets]	8, 10, 12, 24, 28, 32, 40, 44, 45, 64
Voice Bit Rate [kb/s]	4, 8, 16, 24, 32
VCC Peak Rate [Mb/s]	0.384 (H0), 0.768, 1.536 (T1/H11), 1.92 (E1/H12), 40.7 (DS3)
Number of Users	4, 8, 16, 32, 64, 128, 256
Timer_CU [ms]	0.125, 0.25, 0.5, 1, 2, 3, 4, 5, 6

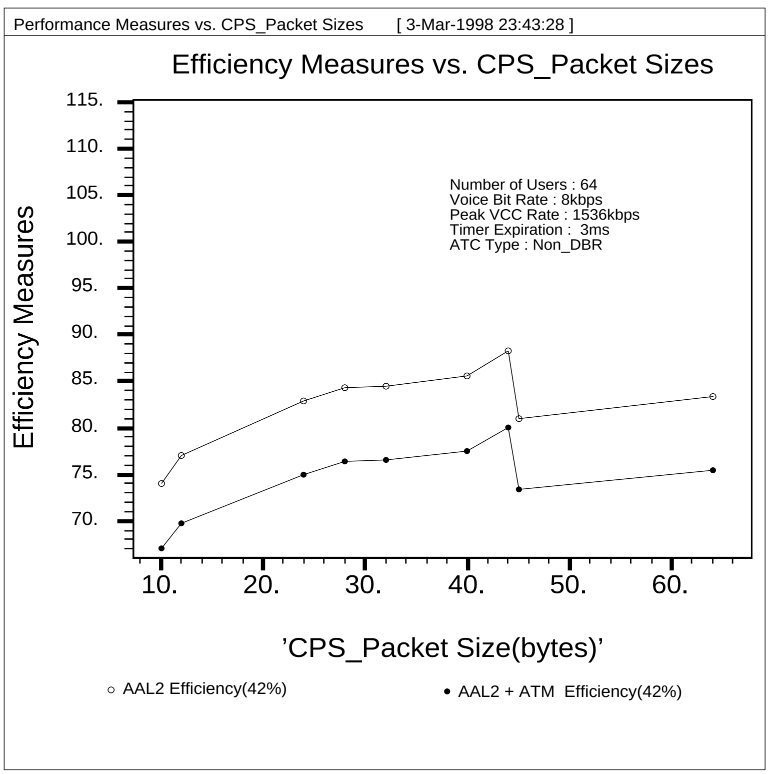

Figure 3. Efficiency Measures vs. CPS_Packet_Size for AAL2

4: AAL Efficiency Comparison Under QoS Constraint

4.1: Rationale for QoS Constraint

Although the performance characterization in the previous section has value, e.g., for illustrating the general behavior of AAL2, it is not very useful in an operational sense. One important operational question is:

Given a certain peak VCC rate and a set of homogeneous sources with given voice bit rate and speech activity factor, how many LLCs can be supported on the VCC while maintaining QoS for all the LLCs?

Notice that this question is not expressed in terms of efficiency or bandwidth gain, metrics which are often used to describe AAL2 performance. Nor is any direct mention made of CPS_Packet_Size or Timer_CU value. Furthermore, to address this question, one must define the QoS requirements in a meaningful way. We proceed to argue that mean delay is not a meaningful QoS measure for voice transport and propose another delay QoS.

In packet-mode voice transport, the timing relationships between successive packets are distorted as the packets experience variable delay across the network, and it is desirable to restore these timing relationships at the voice receiver. The voice receiver typically attempts to accomplish this by "building out" the overall delay for every voice packet to some fixed value by delaying packets at the receiver such that their overall delay is the "build-out" delay. [2] In this context, what is important is not mean delay, but the tail of the delay distribution.

This overall fixed delay can be taken to be the K^{th} percentile delay (if known). The K^{th} 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. Thus it can be expected that only (100-K)% of the packets would arrive "late" at the receiver (that is, with delay larger than the build-out delay). 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. This K^{th} percentile delay in fact corresponds exactly to the maxCTD (maximum cell transfer delay) QoS parameter defined by the ATM Forum [6], in which the value of K is left as a parameter. We have chosen $K = 95$ in this study because voice can be made relatively tolerant of losses of a few percent (e.g., see [7, 8]) and because larger values of K would have required

[2] Specific techniques for accomplishing this are beyond the scope of this paper, but see [5] for a classical discussion of these.

considerably longer simulation times to yield high-quality estimates. Furthermore, we are only dealing with AAL transmitter delay, and the end-to-end delay QoS could have a larger value of K.

4.2: Simulation Parameters

We will use simulation to answer the question posed above for AAL1, AAL2, and AAL5 with the following set of parameters. The peak VCC rate is fixed at 1.536 Mb/s, the mean on and off times of the voice sources are 420 and 580 ms, respectively, and the QoS metric is that the 95^{th} percentile delay in the AAL transmitter (packetization delay plus queueing delay) must be 10 ms or less, based on the following consideration. Echo cancellers are usually required when the total 1-way delay in a connection exceeds approximately 25 ms, so the 10 ms delay bound in the transmitter leaves 15 ms for other network queuing delays, propagation delay, and coder/decoder delay.

Four voice bit rates are considered: 8, 16, 32, and 64 kb/s [3]. For each voice bit rate, we run multiple simulations with different CPS_Packet_Sizes and different numbers of LLCs to find the optimal value of CPS_Packet_Size. For each value of CPS_Packet_Size, the AAL2 Timer_CU value is set to its maximum reasonable value, which is the difference between the delay bound and the packetization delay.

4.3: Results

The results of the simulations can be displayed by plotting 95^{th} percentile delay vs. number of LLCs (users) for each CPS_Packet_Size considered, as in Figure 4. Here, as opposed to Figure 1, we are operating in the high load region where queuing delay begins to dominate. [4] Note the following trends as CPS_Packet_Size increases. First, in the region before queueing delay dominates (e.g., 200 users in Figure 4), 95^{th} percentile delay is dominated by packetization delay, so larger CPS_Packet_Sizes produce larger 95^{th} percentile delays. However, as the number of users increases, larger CPS_Packet_Sizes postpone the onset of queueing delay dominance because larger CPS_Packet_Sizes are more efficient and thus generate less effective load. The general result is that the optimal CPS_Packet_Size is usually very close to the largest possible value that keeps packetization delay below the delay bound.

From graphs such as Figure 4, one can easily determine the maximum number of users for each CPS_Packet_Size that meets the QoS objective by noting where the curve crosses the 10 ms limit. For example, in Figure 4 AAL2 with a 4-octet CPS_Packet_Size can support approximately 207 8 kb/s users while a 6-octet CPS_Packet_Size allows approximately 243 8 kb/s users. In the particular case of 8 kb/s voice coding (Figure 4), the maximum number of users with 7-octet, 8-octet and 9-octet CPS_Packet_Sizes exceeds the AAL2 limit of 256 users.

The maximum number of users can be collected and plotted as a function of CPS_Packet_Size for each AAL, as shown by the solid curves in Figure 5 for AAL2. The points where the curves of Figure 5 return to zero users supported (CPS_Packet_Sizes of 10, 20 and 40 octets) are not generated from simulations since for these values, the packetization delay itself is equal to the delay bound, so even very low loads will exceed the delay bound.

Figure 5 illustrates that the maximum number of users supported is very sensitive to CPS_Packet_Size at low voice coding rates but quite insensitive at higher coding rates. Note also that if we were to choose an 8-octet CPS_Packet_Size for 8 kb/s

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

[4] The saturation of the curves at 15 ms is a simulation artifact: the delay histogram probe in the simulation was set to have a maximum value of 15 ms.

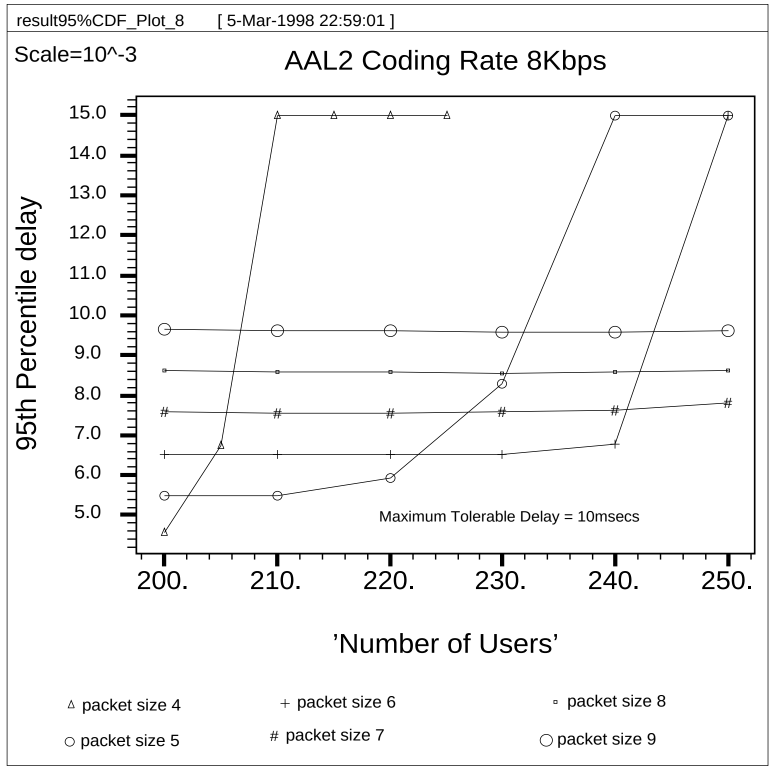

Figure 4. 95th Percentile Delay Plot for 8 kb/s Voice Coding Rate with AAL2

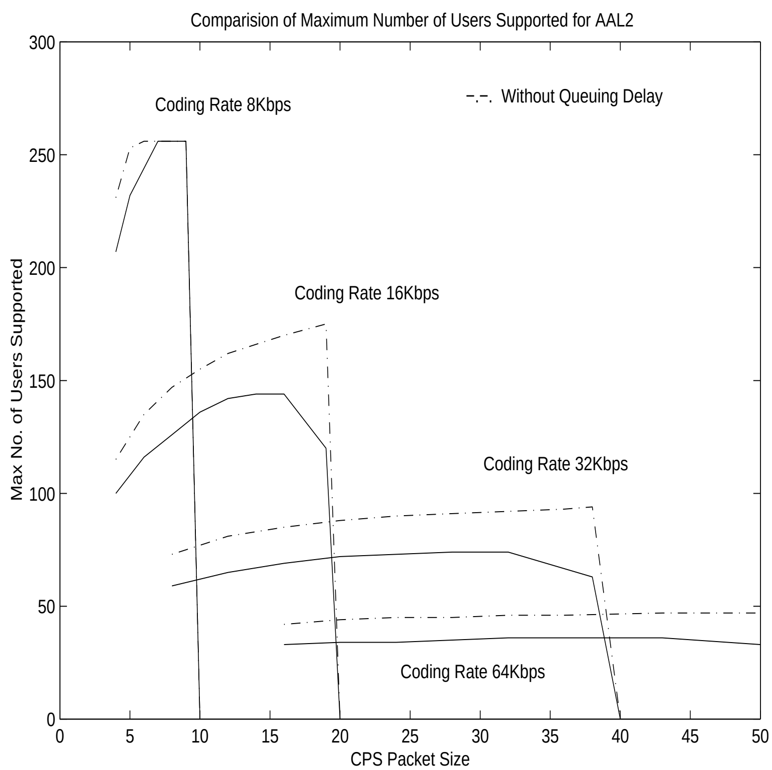

Figure 5. Plot for Maximum Number of Users as a Function of CPS_Packet_Size and Voice Coding Rate for AAL2

coding, 16 octets for 16 kb/s, and 32 octets for 32 kb/s, the CPS_Packet_Size would be near the optimal 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 in a heterogeneous voice coding scenario.

The dashed curves in Figure 5 are calculated without taking the queueing delay into account, i.e., using equation 2. We see that using a more realistic, simulation-based evaluation produces significantly smaller maximum number of users (approximately 80% of the values produces by equation 2).

It is also interesting to compare the values obtained by simulation with the values obtained using the simple analysis in [2]. For example, with 16 kb/s coding and a 16-octet CPS_Packet_Size, CPS_Packets can be generated every 8 ms by each active source, and the queueing delay must not exceed 2 ms for 95% of the packets in order to meet the 95^{th} percentile transmitter delay bound of 10 ms. Adjusting the analysis in [2] for 5% "loss" and applying a 2 ms queueing delay limit and 42% speech activity with 16-octet CPS_Packets yields an allowed number of users of only 32, compared to 144 indicated by the simulations! A primary reason for the discrepancy is the worst-case assumption in the analysis that all active sources generate CPS_Packets simultaneously, as opposed to the random generation times used in the simulation.

Using the optimal CPS_Packet_Size for each voice coding rate yields Figure 6, which plots maximum number of users supported as a function of voice coding rate for the three considered AALs. Here we see that AAL1 and AAL5 are nearly identical in their performance, in fact, exactly identical for Assembly_Sizes up to 40 octets, as observed earlier. Furthermore, unlike AAL2, the maximum number of users for AAL1 and AAL5 is not strongly influenced by voice coding rate. For 64 kb/s coding, all three AALs can support just under 50 users, but the advantage of AAL2 emerges dramatically as the voice coding rate decreases. The advantage of AAL2 over the other AALs is approximately 1.5:1 for 32 kb/s coding, 3:1 for 16 kb/s coding, and approximately 5:1 for 8 kb/s coding. The basic reason is simple: AAL2 multiplexing remains relatively efficient as voice coding rate (and CPS_Packet_Size) decreases, whereas AAL1 and AAL5 become less and less efficient.

5: Conclusions

This simulation study allows us to draw several important conclusions about the efficiency of AAL2 for voice transport, especially relative to AAL1 and AAL5.

- Efficiency metrics such as overhead penalty or bandwidth efficiency are not as useful in an operational context as the maximum number of users supportable for a given QoS.
- A bound on K^{th} percentile delay is a more appropriate QoS measure for voice than mean delay or absolute maximum delay.
- For AAL2 the total packet delay decreases toward the packetization delay as the load increases (due to the effect of Timer_CU) until it reaches a queuing saturation point after which even a slight increase in load causes the total packet delay to increase sharply.
- AAL1 and AAL5 perform identically for voice transport up to an Assembly_Size of 40 bytes.
- 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.
- The maximum number of users supported on a VCC with AAL2 is almost 5 times the number supported by 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 at 64 kb/s is nearly the same for all three AALs.

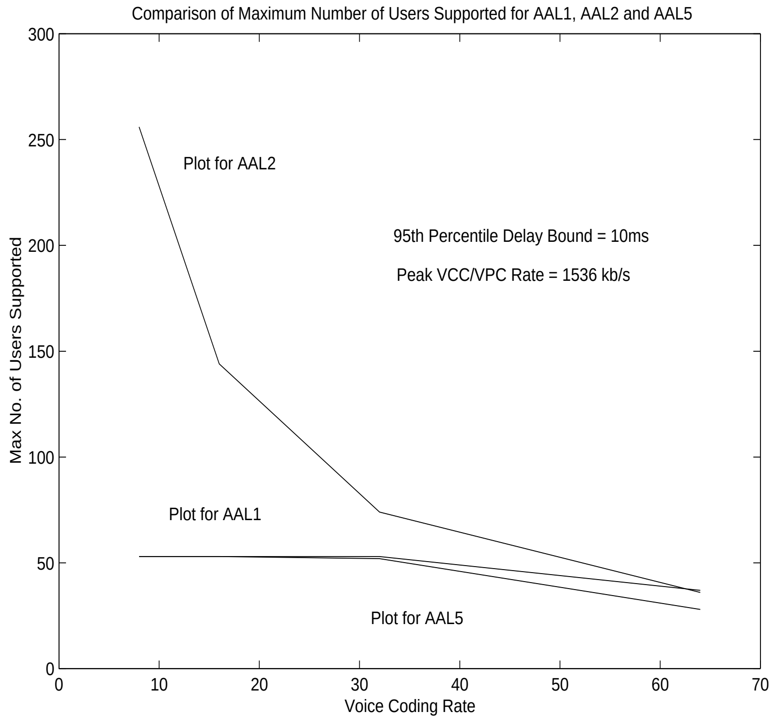

Figure 6. Maximum Number of Users Supported on a VCC for AAL1, AAL2 and AAL5 as a function of Voice Coding Rate

References

[1] D. J. Wright, "Voice over ATM: An Evaluation of Implementation Alternatives," *IEEE Communications Magazine*, Vol. 34, No. 5, May 1996, pp. 72-80.

[2] J. H. Baldwin, G. H. Bharucha, B. T. Doshi, S. Dravida, S. Nanda, "AAL-2 – A New ATM Adaptation Layer for Small Packet Encapsulation and Multiplexing," *Bell Labs Technical Journal*, Spring 1997, pp.111-131.

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

[4] G. Eneroth and T. Westholm, "Comparison of AAL-CU Protocol Proposals," ATM Forum Contribution 96-1330, October 1996.

[5] W. A. Montgomery, "Techniques for Packet Voice Synchronization," *IEEE Journal on Selected Areas in Communications*, Vol. SAC-1, No. 6, December 1983.

[6] ATM Forum Technical Committee and Traffic Management Working Group, " ATM Forum Traffic Management Specification Version 4.0"

[7] D. W. Petr, L. A. DaSilva and V. S. Frost, "Priority Discarding of Speech in Integrated Packet Networks," *IEEE Journal on Selected Areas in Communications*, Vol. 7, No. 5, pp. 644-656, June 1989.

[8] D. W. Petr, W. G. Huang, M. P. Mischler, "Subjective Effects of Cell Losses in Voice Over ATM," *IEEE Communications Letters*, Vol. 2, No. 7, pp. 186-188, July 1998.

[9] R. Yellisetti, D.W. Petr, " Development of simulation models for AAL1 and AAL5" *ITTC-FY98-TR-13110-02*, Feb 1998.