```
*******************************************************************
```

ATM Forum Document Number: ATM_Forum/97-0608

```
*******************************************************************
```

Title: Performance of VBR Voice over ATM: Effect of Scheduling
and Drop Policies

```
*******************************************************************
```

Abstract:

Voice in its inherent nature is variable. It is natural to use the
VBR class for voice, and transmit voice in its bursty form. On an
average, human voice has an speech activity factor of about 42%. This
allows us to multiplex a lot of voice sources together. We study the
degradation in voice quality under various scheduling policies for
VBR voice. Our simulations showed that it is desirable to keep the
total VBR-voice load to about 50% of the total link capacity, in
order to maintain good quality voice. The remaining link capacity
could then be filled up by data.

```
*******************************************************************
```

Source:
Jayaraman Iyer, Raj Jain, Sohail Munir
Department of CIS, The Ohio State University (and NASA)
395 Dreese lab, 2015 Neil Ave, Columbus, OH 43210-1277

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

Sudhir Dixit
Nokia Research Center
3 Burlington Woods Dr., Suite 250
Burlington, MA 01803
Phone: 617-238-4915, Email: sudhir.dixit@research.nokia.com

```
*******************************************************************
```

      Date: July 1997

```
*******************************************************************
```

Distribution: ATM Forum Technical Working Group Members (AF-TM, AF-VTOA)

```
*******************************************************************
```

```
*******************************************************************
```

A postscript version of this contribution including all figures and
tables has been uploaded to the ATM forum ftp server in the incoming
directory. It may be moved from there to the atm97 directory. The
postscript version is also available on our web page as:

http://www.cse.wustl.edu/~jain/atmf/a97-0608.htm

```
*******************************************************************
```

1 Introduction:

ATM has the capability to transport both voice and data. However, the characteristics of both these are very different. Voice is sensitive to end-to-end delay as well as delay variation, and less sensitive to cell loss. Data, on the other hand is not sensitive to delay and delay variation, but cannot tolerate cell loss. Voice places real-time constraints on the network to guarantee tight delay bounds. It is of course easy to find a solution to support voice using the CBR class. However, this solution does not take into account the characteristics of voice such as low speech activity and hence does not utilize the bandwidth effectively. Voice has an average activity factor of about 42%. This means that if we do not use silence suppression, we would be wasting 58% of the link. We could statistically multiplex multiple voice sources and thereby utilize the available bandwidth effectively. At the same time, we need to fulfill certain guarantees to maintain good voice quality. In this contribution, we explore the possibility of transmitting voice with strict delay bounds as multiplexed VBR traffic.

2 Delay characteristics of voice

Voice is highly sensitive to delay, and mildly sensitive to delay variations. In this section we look at required quality of service parameters to support voice.

2.1 Effect of end-to-end delay.

Studies have been conducted to observe the effect of large end to end delay as well as the variation in delay. For applications including support of interactive voice communication, delay can have two effects on connection performance. In the absence of noticeable echo, delay can interfere with the dynamics of voice communication. In the presence of noticeable echo, increasing delay makes echo effects worse. At a certain amount of delay, echo cancelers are required to control the echo. Quality ratings for mean delays of about 165 ms and a variation of 35 ms has been poor.

ITU-T Recommendation G.114 recommends the following general limits for one-way transmission time for connections with adequate echo control.

0 to 150 ms,   acceptable to most user applications.
150 to 400 ms,   acceptable under awareness of impact on quality
Above 400 ms,   unacceptable

[ITU-G.114] also states that some highly interactive tasks may suffer degradation even at delays of the order of 100 ms.

The end-to-end delay in an ATM network consists of two components: queuing delay and re-assembly delay. Queuing delay is caused by the presence of other traffic in the switches. Proper scheduling policies and drop policies are required in the switches to minimize this delay.

The re-assembly delay is caused by the need to fill the ATM cell. At 64 kbps, it takes about 6 ms to fill up the entire payload of an ATM cell. The re-assembly delay increases with higher

compression rates. With multiple voice channels multiplexed into
a single ATM cell, the time is smaller.

One can divide the end-to-end path into three: access loop between
the source and the (virtual)PBX, the loop between the source
PBX and the destination PBXs, and the loop between the destination
PBX and the destination. The end loops carry one conversation per VC
and use partially filled cells. The middle loop carries multiple
conversations and has fully filled cells. Since the bandwidth
is easily and cheaply available in the access loops, loosing
bandwidth in partially filled cells might not be a serious concern.

This contribution focuses on the queuing delay segment arising out of
interaction between the traffic in the switches.

2.2 Cell Loss

Voice can tolerate a small amount of isolated cell discarding. If
there is correlation (bursty loss) either due to compression or
otherwise between the dropped voice cells, then reconstruction of
voice at the destination suffers a degradation in quality. Using a
two-priority coding scheme can reduce the correlation between
successive voice cells as well as reduce the degradation suffered
[Sriram91]. The two-priority coding scheme, consists of sending voice
as two-cell pairs, the first one a high priority cell and the second
a low-priority one. The low priority cell has the CLP bit set. During
congestion, the network can choose to drop the low-priority cell.

The effect of cell loss depends on the number of voice calls
multiplexed on a single ATM virtual circuit. A single cell on a VC
carrying 64 kbps without any compression will lose about 6 ms of
voice.

The Cell Loss Ratio's for CLP=1 and CLP=0 are different. In our
simulations, we assume that a scheme exists that can set the CLP bit
on cells, thereby reducing the correlation between subsequent cells.

Telephony voice can handle a higher cell loss ratio of the order of
10-3 [Onvural95]. If a two-priority scheme is used as outlined in
[Sriram91], losses of up to 10% could be acceptable.

2.3 Delay Variation and Delay Equalization

Voice is an isochronous application [RFC1257]. The end-user is
sensitive to variations in inter-arrival times. The ATM cells incur
variable queuing delays as they traverse through the network. This
variation in the arrival times is caused by the queuing delay, the
delay in the switching fabric, and the propagation delay. From an
application viewpoint, there is the need to ensure a synchronous
playout. The is typically achieved by buffering the cells for enough
time so that the variation in inter-arrival times is kept to a
minimum. This additional artificial delay until playout can either be
fixed or can vary adaptively during a call's lifetime. In order that
there is no buffer underflow, each cell has a scheduled arrival time.
The cells that arrived later than the scheduled arrival time cause
buffer underflow and are considered lost. The buffer functions as a
mechanism to smooth out the variation in the network delays.

For very high quality voice (256 kbps MPEG), an upper bound on the
cell delay variation of 9 ms, while for low bit rate (16 kbps) voice
a variation of 130 ms may be acceptable [Onvural95]. Given an upper

bound on the delay variation, and the propagation delay, the buffer
delay necessary for adaptive playout can be estimated. [Ramjee94,
ITU-G.764A].

The buffer size for the delay equalization increases in order to
accommodate higher delays in the network.

3. Models for Voice Traffic

 It has been determined that both talkspurt and silence periods of
digitized voice are exponentially distributed. [Brady69]. Therefore,
a commonly accepted model for a speaker is a continuous time,
discrete   state Markov chain with two states. The two states
essentially correspond to the speech and silence periods. The holding
time in each state is assumed to be exponentially distributed. The
commonly accepted values for holding time in the silent state is 650
ms and that in the speech state is 352 ms. These values depend on the
sensitivity of the silence detectors [Deng95, Gruber81].

 The two state model has some known limitations. In particular, the
simple two-state model does not model a two-way conversation, since
two-way conversations cannot be modeled by merely superimposing
multiple single source speech generators. Events like interruptions
and double talking are possible in two-way conversations. Such events
will affect any model that tries to approximate speech patterns
[Vickers94]. Some researchers have proposed a four-state Markov Model
to describe the behavior of such a system. The four states represent
who is doing the talking; no one, one person, the other person, or
both. This is only a crude approximation since such a Markov model
has exponential distribution for each state, which may not be
realistic. A better model is to add two more states making a total of
six states [Brady69]. Mutual silence and double-talk states are split
into two states, with the identity of the last lone speaker
differentiating them.

The exponential distribution holds good only for digitized voice and
not for packetized voice [Deng95]. For a large number of voice
sources, the generated traffic inter-arrival times approach that of a
Poisson process[Vickers94].   We chose the two-state Markovian model
with mean speech and silence times 352 ms and 650 ms [Brady69].

4. The Network Model

The "N VBR Source" configuration shown in Figure 1 has two ATM
switches connected by "LINK1". "LINK1" multiplexes all the voice
sources (N) onto the single link and then sent to the respective
destination. Each voice source is a 64 kbps source following the
two-state Markov model. The mean on-times and off-times are 352 ms
and 650 ms. The end-to-end distance between the source and the
destination is kept at 4800 km (coast-to-coast distance in the United
States). "LINK1" operates at 1.544 Mbps. This was chosen so that the
values match the current telephone network scenario.

```
        --------                                          --------
       │ Src 1  │\                                      /│ Dest 1 │
       │        │ \                                    / │        │
        --------   \    --------   LINK1  --------    /   --------
           .        ----│  Sw1   │-------│  Sw2   │-------       .
           .         ---│        │       │        │-------       .
           .        / /  --------         --------     \         .
```

```
                          / /                              \        \
      --------         / /                                  \        \--------
     | Src N-1|/ /                                           \       | Dest N-1|
     |        |/                                              \      |        |
      -------- /                                               \      --------
      --------/                                                 \      --------
     | Src N  |                                                -- |Dest N  |
     |        |                                                   |        |
      --------                                                     --------
```
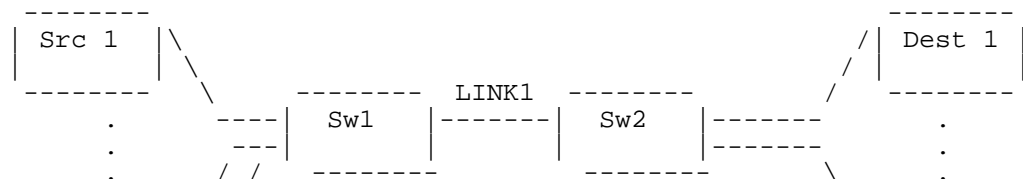
                Figure 1: The "N VBR Source" Configuration

At each switch, we assume per-VC queuing at the output port, and
supporting multiple classes of service.

End-to-end-propagation delay: 25 ms

Packetization/depacketization delays: 6 ms + 6 ms (for plain PCM)

We noted in section 2.1 that we require to guarantee an end-to-end
delay of about 100 ms (with echo-cancelers) for interactive tasks.
Allowing for about 5 switches on an average, the delay variations
that can be introduced by a switch will be (100 - 25 - 6 - 6)/5 or
about 12.6 ms. In order to support other applications with more
stringent delay requirements, we also consider the support of a delay
variation bound of 5 ms per switch.

In order that we support high-quality voice, we fixed the upper bound
on the end-to-end delay variations to be 5 ms and 15 ms. Hence we
defined two kinds of high-quality voice traffic based on the delay
bounds i) that which can tolerate an end-to-end delay of 40 ms, and
ii) that which can tolerate an end-to-end delay of 30 ms. Note that
these figures do not take into account the packetization delays at
the source or the destination, as well as the additional delay
incurred due to compression.

5. QoS Metrics for voice

In order to study the effect of voice, we need to define certain
metrics that give us a quantitative measure of the voice traffic
quality.

Voice is extremely sensitive to delay. Usually, medium delays would
require echo cancelers, and long delays would be unsuitable for most
purposes. The delay usually depends on the delay at the intermediate
nodes in addition to the propagation delay of the network.

o The mean cell transfer delay (CTD) gives us a measure of the end-
to-end delay incurred in an ATM network. CTD is a function of the
propagation delay, the queuing delay and the switching delay.

o The peak-to-peak cell delay variation (p-p CDV) is another
important QoS parameter. The CDV is a function of the number of
multiplexed connections, the type of the multiplexed connections
through the switches, and the switching variability. CDV varies with
the mixture of the traffic. However, as cited above, CDV is not a
major concern for voice. ATM networks typically have a low cell-delay
variation, and this can be taken care of by a playout buffer at the
receiver. It may not be always possible to delay cells to compensate
for the maximum network delay, in which case it is preferable to drop
cells delayed more than an acceptable value rather than attempting to

handle the CDV.

o Cell Loss Ratio (CLR) is defined as the number of cells lost divided by the total number of transmitted cells. To guarantee high quality voice, it is desirable to keep the cell loss ratio to a minimum.

5.1 Overall Quality of Voice

Voice quality is a function of the the cell-loss and the transfer delays. Voice can tolerate a low cell-loss-ratio. At 64 kbps, a cell loss amounts to a loss of about 6 ms of voice content. The cell loss varies with the amount of buffer available at the switches. However, adding more buffer and thereby reducing the cell loss ratio does not necessarily improve the overall voice quality. The end-to-end delay requirements and the end-to-end delay variation requirements also need to be satisfied. At the receiving end, not all cells that are received are useful. The useful cells are only the ones that arrive within a specified time. The cells arriving later will be discarded.

We define degradation in voice quality (DVQ) as a function of these two parameters.

DVQ = (number of cells lost + number of cells above the delay
      threshold) / total cells sent

We define useful cell ratio as (1 - DVQ).

The assumption here is that the cell loss (or the delayed cells) is going to be distributed uniformly during the duration of the call. The impact of having two consecutive cell losses every 2 t msec, is going to have a more significant impact than a cell loss every t msec. We studied DVQ, and CLR under various scheduling and cell-dropping schemes. We also compared the fairness of each scheduling technique. If n users contend for a shared resource and ith user obtains a performance of Xi, the fairness index [Jain91] for this resource allocation is defined as

              (sum total of Xi s)*(sum total of Xis)
               ----
               n * (sum of Xi*Xi)

In our simulation, we evaluated 64 kbps (voice under no compression). We can easily extend the results for compressed voice. The delay threshold bounds will change, and the Degradation in Voice Quality (DVQ) can be defined with a multiplicative factor based on the compression ratio.

6. Scheduling Policies

Scheduling Policies play a very important role towards guaranteeing fairness in user level voice quality. As defined earlier, the primary metric for evaluating the scheduling policies is by looking at the fraction of good cells sent across. The number of cells discarded is the sum of the cells making it to the destination after the deadline has expired or if dropped by the switches due to a lack of buffers. For the purpose of this simulation, we are not evaluating other important performance measures such as the bandwidth overhead or computational cost. These considerations should also enter into

evaluating the merits of the candidate policies. We studied the following scheduling policies.

i) Earliest Deadline First (EDF): Each cell has a deadline when it enters its VC-Queue. The cell with the earliest deadline from the head of all the VC Queues is chosen to be sent out on the output link.

ii) Longest Queue First (LQF): The longest queue is chosen, in order to prevent growing of the queues. The cell at the head of that queue is sent out on the output link.

iii) Round Robin (RR): The queues are chosen in a round-robin fashion.

We assumed a class-based priority scheduling policy. This is required for real-time traffic such as voice to be prioritized over data. Class-based scheduling is required also to isolate the interference from other lower priority traffic classes. Since we already have a priority class scheduler, we do not consider any priority-based scheduling scheme. Since we have symmetric voice sources over ATM, and each source transmits only cells of equal size at 64 kbps, we did not consider any weighted fair schemes.

7. Drop policies

In addition to the scheduling policies, we studied the following drop policies.

i) Tail-drop: This is the simple FIFO drop. The cells are dropped if there is no buffer to add them in.

ii) Selective discard: In this scheme, we use per-VC accounting to maintain buffer utilization of each active VC in the switch. When the buffer occupancy exceeds a preset threshold, further cells from that VC are dropped, thereby ensuring fairness in buffer usage. We do not consider the push-out selective discard scheme, since it is very costly for implementation, and buffer-threshold scheme gives the same results if the threshold values are appropriately chosen.

8. Simulation results.

In this set of simulations, we varied the number of voice sources from 20 to 75, and studied CLR, DVQ(with a threshold of 40 ms and 30 ms), as well as the Fairness corresponding to each threshold. The switches supported a per VC queue of size 2. We studied the effect of 3 scheduling policies using per VC queuing - Round Robin (RR), Longest Queue First (LQF) and Earliest Deadline First (EDF). We also studied the effect of drop policies using a common buffer for Tail Drop and Selective Discard.

NS	Offered Load (%)	Multiplexing gain
20	29.26	0.83
24	35.12	1.00
30	43.90	1.25
35	51.21	1.45
40	58.53	1.66
48	70.24	2.00
55	80.48	2.29

```
60        87.80               2.50
65        95.11               2.70
70       102.43               2.91
75       109.75               3.12
```

Table 1: Offered Load and the multiplexing gain got

  Table 1 shows the offered load with increasing number of sources,
and the utilization gain that we get over the simple CBR case with a
64 kbps VC. The offered load is in terms of the total link capacity.
In our simulations the link capacity between the switches is 1.544
Mbps. We calculate the multiplexing gain as the ratio of the number
of voice sources to the number of 64 kbps voice channels that could
be supported (equal to 24 in our case). Note that this gain is the
gain only for the voice usage.


```
NS: Number of Sources
Q:  Per-VC Queue Length
S:  Scheduling Policy
        rr: Round Robin
        lqf: Longest Queue First
        edf: Earliest Deadline First
CLR:   Cell Loss Ratio(%)
DVQ40:   Degradation in Voice Quality (with threshold 40 ms)
DVQ30:   Degradation in Voice Quality (with threshold 30 ms)

F40:    Fairness Index using DVQ40
F30:    Fairness Index using DVQ30
```

NS	Q	S	CLR	DVQ40	DVQ30	F40	F30
20	2	rr	0.0000	0.0000	0.0000	1.0000	1.0000
20	2	lqf	0.0000	0.0000	0.0000	1.0000	1.0000
20	2	edf	0.0000	0.0000	0.0000	1.0000	1.0000
24	2	rr	0.0000	0.0000	0.0005	1.0000	1.0000
24	2	lqf	0.0000	0.0000	0.0005	1.0000	1.0000
24	2	edf	0.0000	0.0000	0.0005	1.0000	1.0000
30	2	rr	0.0616	0.0006	0.0121	1.0000	1.0000
30	2	lqf	0.0488	0.0010	0.0123	1.0000	1.0000
30	2	edf	0.0616	0.0006	0.0121	1.0000	1.0000
35	2	rr	0.1964	0.0031	0.0165	1.0000	0.9999
35	2	lqf	0.1764	0.0025	0.0175	1.0000	0.9999
35	2	edf	0.1964	0.0031	0.0164	1.0000	0.9999
40	2	rr	0.3865	0.0074	0.0187	1.0000	0.9999
40	2	lqf	0.3579	0.0047	0.0201	1.0000	0.9999
40	2	edf	0.3865	0.0073	0.0190	1.0000	0.9999
48	2	rr	0.6423	0.0132	0.0429	1.0000	0.9996
48	2	lqf	0.6161	0.0078	0.0469	0.9999	0.9994
48	2	edf	0.6371	0.0130	0.0445	1.0000	0.9997
60	2	rr	2.5959	0.0384	0.2717	0.9999	0.9947
60	2	lqf	2.4932	0.0354	0.3147	0.9971	0.9855
60	2	edf	2.5353	0.0357	0.2945	0.9999	0.9951
65	2	rr	4.9184	0.0693	0.4232	0.9997	0.9880
65	2	lqf	4.6462	0.0636	0.4763	0.9899	0.9744
65	2	edf	4.8210	0.0648	0.4529	0.9998	0.9874
70	2	rr	8.2518	0.1235	0.6040	0.9994	0.9772
70	2	lqf	7.9017	0.1027	0.6781	0.9732	0.9366
70	2	edf	8.1647	0.1075	0.6420	0.9996	0.9731
75	2	rr	12.7650	0.2079	0.7651	0.9987	0.9552
75	2	lqf	12.4222	0.1546	0.8380	0.9363	0.8477
75	2	edf	12.7535	0.1882	0.8085	0.9990	0.9405

Table 2: Scheduling schemes with a per VC Queue of length 2.

Table 2 shows the comparison of various scheduling schemes for the
voice sources varying from 20 to 75. Referring to Table 2, Column 2
shows the Cell Loss Ratio. We find that the CLR increases with the
increase in the number of voice sources. The Cell Loss Ratio
increases sharply after 40 sources.
 40 sources gives a CLR of 0.35%. A higher CLR might be unacceptable
for good quality voice. At higher loads such as 60 sources and above,
the Cell Loss Ratio is very high.

Column 3 and 4 show the Degradation in Voice Quality (DVQ)
corresponding to thresholds of 40 ms and 30 ms. We find that at the
load (number of voice sources) equal to the link capacity, DVQ(40 ms)
and DVQ(30 ms) are very high. DVQ(40 ms) increases beyond 8% and
DVQ(30 ms) reaches very high values such as 50% and above. Both
values are practically not acceptable.

If we allow 35 sources, we can fill up the bandwidth with 50% voice
and still get a multiplexing gain of 1.45 (from Table 1).If we load
the link with VBR voice of upto 50%, then we can give upper bound
guarantees for supporting good quality voice, such as a low CLR and a
low DVQ value. The remaining bandwidth could then be filled with
data.

Looking at the scheduling policies, we find that for a 40 ms network
delay threshold, Longest Queue First (LQF) performs better than
Round-robin or Earliest Deadline First (EDF), giving a low DVQ value.
However, for a 30 ms network delay threshold, Round Robin is better.
In other words, LQF gives a higher average delay than RR, but given a
higher threshold a smaller percentage of cells exceed that. This
shows that the quality of voice depends not only on the average
delay, but also on the delay distribution.

Table 3 shows the scheduling schemes with a max. per VC Queue of size
1. With a per VC queue of size 1, there is less buffering at the
switches.

NS	Q	S	CLR	DVQ-40	DVQ-30	F40	F30
20	1	rr	0.0000	0.0000	0.0000	1.0000	1.0000
20	1	lqf	0.0000	0.0000	0.0000	1.0000	1.0000
20	1	edf	0.0000	0.0000	0.0000	1.0000	1.0000
24	1	rr	0.0000	0.0000	0.0005	1.0000	1.0000
24	1	lqf	0.0000	0.0000	0.0005	1.0000	1.0000
24	1	edf	0.0000	0.0000	0.0005	1.0000	1.0000
30	1	rr	0.1126	0.0011	0.0037	1.0000	1.0000
30	1	lqf	0.1126	0.0013	0.0026	1.0000	1.0000
30	1	edf	0.1126	0.0011	0.0037	1.0000	1.0000
35	1	rr	0.2400	0.0024	0.0063	1.0000	1.0000
35	1	lqf	0.2418	0.0027	0.0040	1.0000	1.0000
35	1	edf	0.2400	0.0024	0.0063	1.0000	1.0000
40	1	rr	0.4183	0.0042	0.0102	1.0000	1.0000
40	1	lqf	0.4215	0.0045	0.0060	1.0000	0.9999
40	1	edf	0.4183	0.0042	0.0101	1.0000	0.9999
48	1	rr	0.7446	0.0074	0.0180	1.0000	0.9999
48	1	lqf	0.7983	0.0086	0.0105	0.9998	0.9998
48	1	edf	0.7223	0.0072	0.0191	1.0000	0.9999
60	1	rr	3.1103	0.0312	0.0844	1.0000	0.9997
60	1	lqf	3.4605	0.0377	0.0437	0.9949	0.9939
60	1	edf	3.0035	0.0301	0.0917	0.9999	0.9993
65	1	rr	5.3873	0.0544	0.1441	0.9999	0.9991

```
65      1  lqf  5.7416  0.0620  0.0703  0.9861  0.9837
65      1  edf  5.2737  0.0533  0.1550  0.9999  0.9990
70      1  rr   8.7110  0.0880  0.2268  0.9998  0.9986
70      1  lqf  9.0092  0.0962  0.1059  0.9686  0.9639
70      1  edf  8.6151  0.0870  0.2471  0.9998  0.9980
75      1  rr   13.1494 0.1329  0.3266  0.9998  0.9983
75      1  lqf  13.4008 0.1419  0.1535  0.9336  0.9257
75      1  edf  13.0457 0.1319  0.3466  0.9997  0.9977
```

Table 3: Scheduling schemes with a per VC Queue of length 1

Comparing the simulation results for per-VC queues of length 1 and
length 2, we note that the CLR values are higher for a per-VC queue
of size 1 since we now have one cell-buffer less for each VC.
However, the percentage of delayed cells reduce a lot. DVQ(40 ms)
values are approximately the same as with a per-VC queue of length 2,
but the DVQ(30 ms) for a per VC Queue of length 1 are lower than
those of a VC queue of length 2.

A per VC Queue length of 1 gives a lower DVQ value for a threshold
of 30 ms, than with a per VC Queue length of 2. This means that it
may not be always possible to account for the delay variations
arising out of the network with buffers at the receiving end. It is
better to have smaller queues at the switches in order to guarantee
strict delay requirements. Since the delayed cells are going to be
dropped, it is better to drop them right at the switches. This could
also avoid some congestion occurring at the switches.


```
NS Drop  CLR(%) DVQ40   DVQ30    F40      F30
-----------------------------------------------
20 tail 0.0000  0.0000  0.0000  1.0000  1.0000
20 sel  0.0000  0.0000  0.0000  1.0000  1.0000
24 tail 0.0000  0.0000  0.0005  1.0000  1.0000
24 sel  0.0000  0.0000  0.0005  1.0000  1.0000
30 tail 0.0361  0.0011  0.0134  1.0000  1.0000
30 sel  0.0361  0.0011  0.0134  1.0000  1.0000
35 tail 0.1746  0.0027  0.0185  1.0000  0.9999
35 sel  0.1746  0.0027  0.0185  1.0000  0.9999
40 tail 0.3611  0.0049  0.0205  1.0000  0.9999
40 sel  0.3611  0.0049  0.0205  1.0000  0.9999
48 tail 0.5938  0.0075  0.0475  1.0000  0.9996
48 sel  0.5938  0.0075  0.0475  1.0000  0.9996
60 tail 2.3042  0.0772  0.3218  0.9990  0.9927
60 sel  2.3042  0.0772  0.3218  0.9990  0.9927
65 tail 4.4562  0.1901  0.4870  0.9971  0.9812
65 sel  4.6682  0.0484  0.4684  0.9998  0.9819
70 tail 7.8797  0.3257  0.6827  0.9861  0.9251
70 sel  8.0486  0.0826  0.6554  0.9994  0.9397
75 tail 12.4850 0.4631  0.8525  0.9636  0.0869
75 sel  12.6091 0.1315  0.8302  0.9991  0.3541
```

Table 4: Comparison of Drop Policies : Buffer Size: 60, Threshold=80%

A comparison of using a Selective Drop Scheme with a plain Tail-end
drop with a common buffer is shown in Table 4. For the Selective Drop
scheme, we used a buffer-threshold of 80%, and a total buffer size of
60. We find that the figures for both schemes are identical for low
loads, or in other words selective drop does not make any impact up
to 60% load with 80% threshold and a queue length of 60. But at
higher loads, selective drop performs better over the plain tail

drop. Since the Cell Loss Ratio increases to unacceptable limits with higher loads (higher than 50%) we are not concerned with operating the network at high loads with only voice. The fairness values for selective discard are identical for low loads, and at very high loads, selective discard performs better. However, as noted earlier, operating at that region is not desirable.

Summary

We studied characteristics of VBR voice by varying the number of voice sources. Scheduling does play an important role in determining the quality as well as fairness. However, in order to support high-quality voice as well as data on the same link, it is necessary to have prioritized service classes, and voice could occupy up to 50% of the link bandwidth. The remaining could be filled by data using ABR or UBR services. The voice quality also depends on the thresholds required by the specific applications as well as the delay distributions. At low loads, the scheduling schemes or the drop policies have comparable fairness and promise similar DVQ values. At high loads, round robin gives a better fairness. However, operating the network at high loads gives larger values of CLR and DVQ. It is better to keep the queues at the switches smaller in order to give strict delay guarantees. This helps also in avoiding congestion in the network.

References

[Brady68] Brady, P.T., "A Statistical Analysis of ON-OFF Patterns in 16 Conversations", The Bell System Technical Journal, Jan 1968,pp 73-91.

[Brady69] Brady, P.T., A Model for Generating ON_OFF Speech Patterns in Two-Way Conversations. Bell System Technology Journal, Vol 48, Sept. 1969, pp 2445-2472.

[Deng95] Deng,S., Traffic Characteristics of Packet Voice, International Conference on Communications, Number 3, 1995, pp 1369.

[Gruber81] Gruber,J.G, Delay Related Issues in Integrated Voice and Data Networks, IEEE Transactions on Communications, Vol. COM-29, No. 6, June 1981, pp 786.

[ITU-G.114] ITU Recommendation G.114, "One-Way Transmission Time", 6 Feb 1996.

[ITU-G.764A] ITU Recommendation G.764, Appendix 1, "Voice Packetization Guide", 13 Nov 1995.

[Jain90] R. Jain, "Congestion Control in Computer Networks: Trends and Issues", IEEE Network, May 1990, pp. 24-30.

[Jain91] R. Jain, "The Art of Computer Systems Performance Analysis", Wiley, 1991.

[Jain92] R. Jain, "Myths about Congestion Management in High Speed Networks", Internetworking: Research and Experience, Vol 3, 1992, pp. 101-113.

[Sriram91] Sriram,K, McKinney,R.S, Sherif, M.H., Voice Packetization and Compression in Broadband ATM Networks. IEEE Journal of Selected Areas in Communication, Vol 9, Num 3, April 1991, pp 294

[Vickers94] Vickers, Brett. Suda,Tatsuya. Some Measured Characteristics of Data and Voice Traffic: A brief Survey, Technical Report #94-9, Dept of Information and Computer Science, University of California, Irvine.

[Onvural95] Onvural, Raif., Asynchronous Transfer Mode Networks, Performance Issues, Artech House, 1995.

[Ramjee94] R. Ramjee, J. Kurose, D. Towsley, and H. Schulzrinne, "Adaptive playout mechanisms for packetized audio applications in wide-area networks," in Proceedings of the Conference on Computer Communications (IEEE Infocom), (Toronto, Canada), June 1994.

[VTOA] Voice and Telephony over ATM to the Desktop Specification, Versi/95-091, June 1996.

[Yates94] On per-session end-to-end delay and the call admission problem for real-time applications with QoS requirements, Yates, Kurose, Towsley, and Hluchyj, CMPSCI Technical Report, 93-20, May 31,1994.

[Gedeon94] Cell Voice Transport Consideration for ATM, Gedeon,I; Narayanan, R; Vreugdenhil, Nick; Candadian Conference on Electrical and Computer Engineering, IEEE 1994.

[ATMF96-0340] Voice Trunking v cell level switching

[ATMF96-0460] Composite User AAL (AAL-CU) for multiplexing voice.

[ATMF96-1563] Voice and Telephony over ATM for Landline Access at E1/T1 Rates.

[RFC1257] Isochronous Applications do not need Jitter Controlled Networks, Craig Partridge