VBR Voice over ATM: Analysis of Multiplexing Gain and Effect of Scheduling & Drop Policies

Jayaraman Iyer, Raj Jain, Sohail Munir

The Ohio State University

Sudhir Dixit, Nokia Research Center

Conta[illegible]e.Edu

http://w[illegible]Jain/

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

- Performance for Multiplexed VBR Voice
  - Separate Queues: Scheduling Policies
  - One Common Queue: Drop Policies
- Multiplexing gain due to silence suppression

Performance Requirements

- End-to-end delay of 0 to 150 ms most acceptable. [G.114]
- 100 ms end-to-end delay for highly interactive tasks.
- Cell Loss in the order of 10^{-3}. [Onvural]
- Buffering at receiving end can take care of the delay variation.

N-Source Configuration

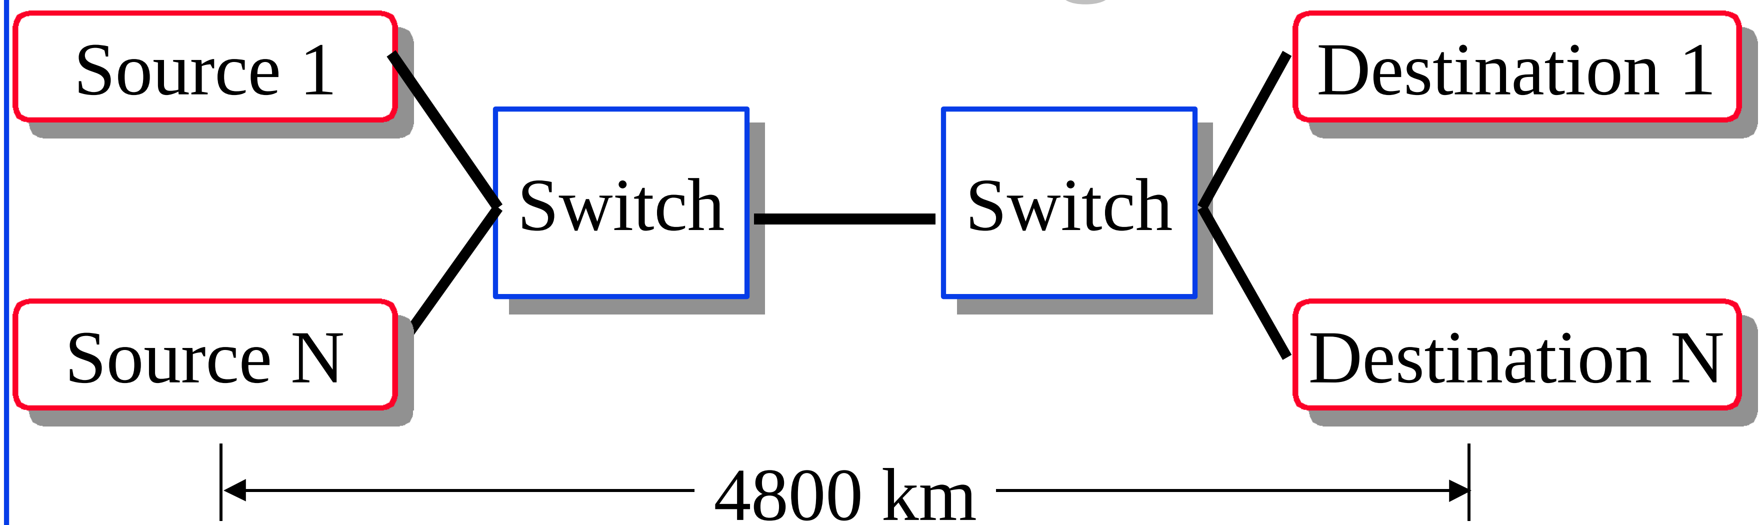

- Links between Switches = 1.544 Mbps (T1).
- N multiplexed 64-kbps VBR voice sources
  Silence suppression ⇒ VBR
- Per-VC Queuing at the Switch
  Multiple queues ⇒ need proper scheduling

Simulation configuration

- Propagation delay : 24 ms
- Avg packetization delays: 6 ms (PCM)
- Reception at the destination: 6 ms
- Assuming 5 switches on a typical path,
  delay variation allowed at each switch
  = (100 - 24 - 6 - 6)/5 = 12.8 ms
- For single switch bottleneck case,
  End-to-end delay = 12.8 + 24 = 36.8 ms ≈ 40 ms
- We tried end-to-end delay bounds of 40 ms and 30 ms.

CDV

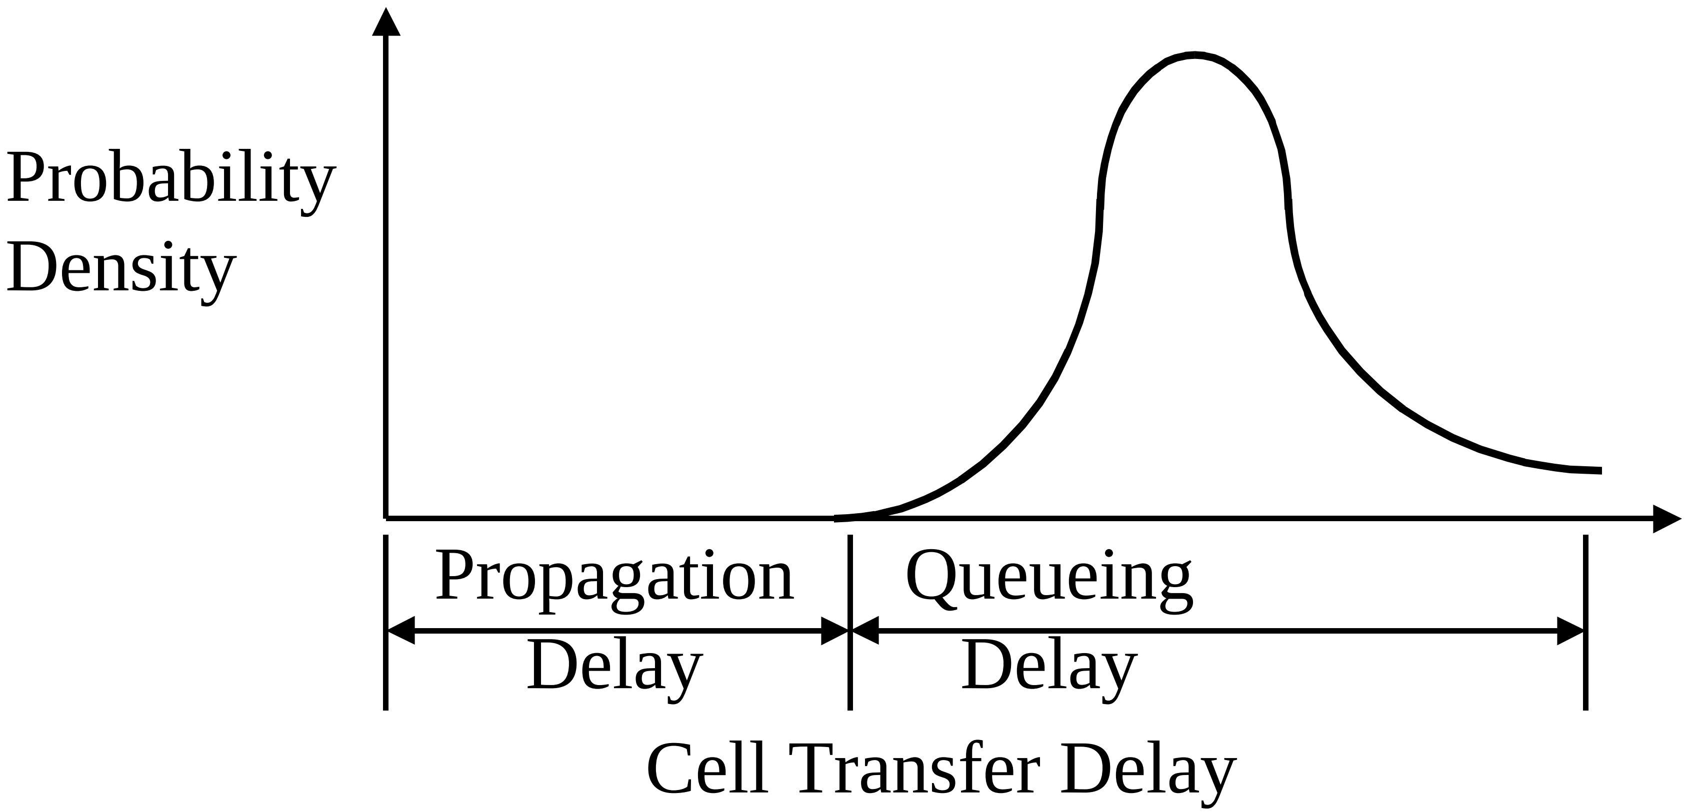

- For VBR voice, we need to specify Max CTD

Source Model

- 2-State Markov Model [Brady69]
- On-off times for silence and speech
- Exponential distribution for speech and silence state.
- Speech activity = 35.1%

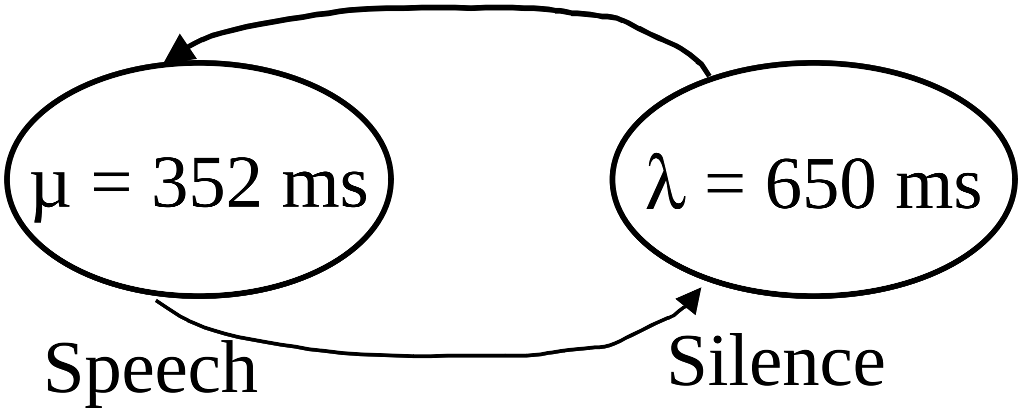

Performance Metrics

- Degradation in Voice Quality (DVQ) = Ratio of cells lost or delayed to total number of cells sent across.
- Cells lost or delayed = Cells dropped by switches + Cells arriving late.
- Fairness = $$\frac{(\Sigma x_i)^2}{n \Sigma x_i^2}$$

x_i is the DVQ for the ith source

Multiplexing Gain

NS	Load (%)	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

Scheduling Policies

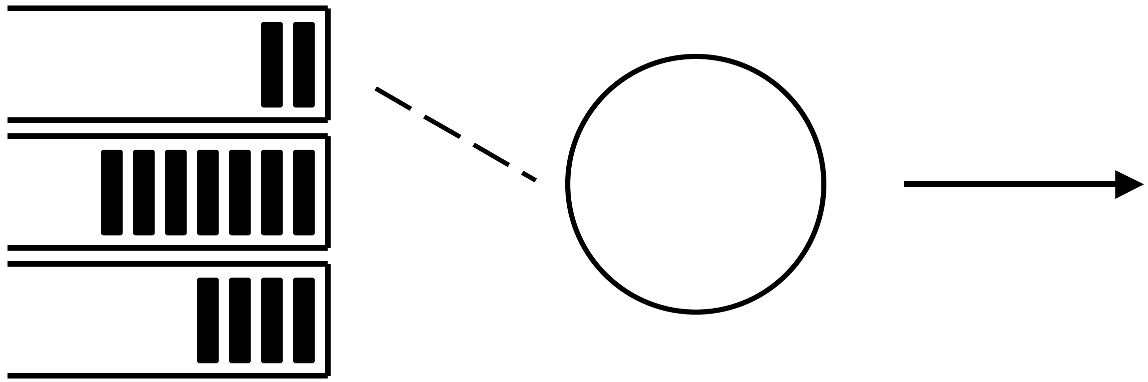

- Round Robin (RR)
- Earliest Deadline First (EDF)
- Longest Queue First (LQF)

Scheduling Results: 1 Buf/VC

NS	Q	Sched	CLR	DVQ	Fairness
24	1	rr	0.000000	0.000000	1.000000
24	1	lqf	0.000000	0.000000	1.000000
24	1	edf	0.000000	0.000000	1.000000
25	1	rr	0.000050	0.000050	1.000000
25	1	lqf	0.000050	0.000075	1.000000
25	1	edf	0.000050	0.000050	1.000000
26	1	rr	0.000218	0.000218	1.000000
26	1	lqf	0.000218	0.000291	0.999999
26	1	edf	0.000218	0.000218	1.000000
27	1	rr	0.000397	0.000397	1.000000
27	1	lqf	0.000397	0.000444	0.999998
27	1	edf	0.000397	0.000397	0.999999

Scheduling Results: 1 Buf/VC (Cont)

NS	Q	Sched	CLR	DVQ	Fairness
28	1	rr	0.000585	0.000585	1.000000
28	1	lqf	0.000585	0.000675	0.999996
28	1	edf	0.000585	0.000585	0.999999
29	1	rr	0.000830	0.000830	1.000000
29	1	lqf	0.000830	0.000939	0.999995
29	1	edf	0.000830	0.000830	0.999998
30	1	rr	0.001126	0.001126	1.000000
30	1	lqf	0.001126	0.001274	0.999991
30	1	edf	0.001126	0.001126	0.999997
35	1	rr	0.002400	0.002400	0.999999
35	1	lqf	0.002418	0.002655	0.999978
35	1	edf	0.002400	0.002400	0.999994

Scheduling Policies: Results I

- With more than 24 users, the cell loss rate is more than 10^{-3}
  - VBR does not allow overbooking based on average rate
  - It does save bandwidth that can be used by lower priority traffic
- At lower loads and low buffers, scheduling does not affect performance.

Scheduling Results: 2 Bufs/VC

NS	Q	Sched	CLR	DVQ	Fairness
24	2	rr	0.000000	0.000000	1.000000
24	2	lqf	0.000000	0.000000	1.000000
24	2	edf	0.000000	0.000000	1.000000
25	2	rr	0.000000	0.000000	1.000000
25	2	lqf	0.000000	0.000000	1.000000
25	2	edf	0.000000	0.000000	1.000000
26	2	rr	0.000000	0.000000	1.000000
26	2	lqf	0.000000	0.000000	1.000000
26	2	edf	0.000000	0.000000	1.000000
27	2	rr	0.000000	0.000000	1.000000
27	2	lqf	0.000000	0.000023	1.000000
27	2	edf	0.000000	0.000000	1.000000

Scheduling Results: 2 Bufs/VC (Cont)

NS	Q	Sched	CLR	DVQ	Fairness
28	2	rr	0.000045	0.000045	1.000000
28	2	lqf	0.000000	0.000135	1.000000
28	2	edf	0.000045	0.000045	1.000000
29	2	rr	0.000306	0.000306	1.000000
29	2	lqf	0.000197	0.000568	0.999998
29	2	edf	0.000306	0.000306	1.000000
30	2	rr	0.000616	0.000637	1.000000
30	2	lqf	0.000488	0.000998	0.999996
30	2	edf	0.000616	0.000637	0.999999
35	2	rr	0.001964	0.003127	0.999998
35	2	lqf	0.001764	0.002491	0.999983
35	2	edf	0.001964	0.003091	0.999993

Scheduling Policies: Results II

- With more buffers, scheduling does matter
- At low loads, scheduling affects efficiency but not fairness
- The number of users supportable is still close to 24 $\Rightarrow$ Buffering of time critical traffic does not help.
- With larger buffers, less cells are dropped in the switch but more cells arrive late and are dropped at the destination.

Scheduling Results: Medium Load

NS	Q	Sched	CLR	DVQ	Fairness
40	2	rr	0.003865	0.007365	0.999993
40	2	lqf	0.003579	0.004692	0.999956
40	2	edf	0.003865	0.007253	0.999981
48	2	rr	0.006423	0.013188	0.999967
48	2	lqf	0.006161	0.007839	0.999887
48	2	edf	0.006371	0.013004	0.999951
60	2	rr	0.025959	0.038383	0.999870
60	2	lqf	0.024932	0.035385	0.997050
60	2	edf	0.025353	0.035714	0.999874

Scheduling Results: Heavy Load

NS	Q	Sched	CLR	DVQ	Fairness
65	2	rr	0.049184	0.069259	0.999683
65	2	lqf	0.046462	0.063567	0.989938
65	2	edf	0.048210	0.064780	0.999776
70	2	rr	0.082518	0.123509	0.999439
70	2	lqf	0.079017	0.102732	0.973244
70	2	edf	0.081647	0.107465	0.999579
75	2	rr	0.127650	0.207901	0.998742
75	2	lqf	0.124222	0.154610	0.936282
75	2	edf	0.127535	0.188157	0.998999

Scheduling Policies: Results III

- At heavy loads, scheduling affects efficiency as well as fairness
- However, at such high loads, voice quality is not acceptable. The load may consist of lower priority data traffic.
- We expect scheduling to have even more impact for asymmetric loads (low bit rate and high bit rate voice sources together)

Drop Policies

- FIFO Discard: Any cell arriving to a full queue is dropped
- Selective Discard: If the queue is over a threshold,
  - Cells for VCs using more than the fair share are dropped.
  - Cell for VCs using less than the fair share are admitted.
- One queue for all VCs: Buffer size = 60
  No per VC queueing $\Rightarrow$ No scheduling required
- Buffer threshold: 80% (for selective drop)

Drop Policies Results

NS	Drop	CLR	DVQ	Fairn
20	tail	0.000000	0.0000	1.0000
20	sel	0.000000	0.0000	1.0000
24	tail	0.000000	0.0000	1.0000
24	sel	0.000000	0.0000	1.0000
30	tail	0.000361	0.0011	1.0000
30	sel	0.000361	0.0011	1.0000
35	tail	0.001746	0.0027	1.0000
35	sel	0.001746	0.0027	1.0000
40	tail	0.003611	0.0049	1.0000
40	sel	0.003611	0.0049	1.0000

Drop Polices Results: Heavy Load

NS	Drop	CLR	DVQ	Fairn
48	tail	0.005938	0.0075	1.0000
48	sel	0.005938	0.0075	1.0000
60	tail	0.023042	0.0772	0.9990
60	sel	0.023042	0.0772	0.9990
65	tail	0.044562	0.1901	0.9971
65	sel	0.046682	0.0484	0.9998
70	tail	0.078797	0.3257	0.9861
70	sel	0.080486	0.0826	0.9994
75	tail	0.124850	0.4631	0.9636
75	sel	0.126091	0.1315	0.9991

Drop Policies: Results

- The multiplexing gain conclusions apply to single queue also.
- At low loads (up to 60%) both drop policies behave identically.
- At higher loads, selective drop is better over plain FIFO drop.
- Fairness of selective discard is very close to 1.

Summary

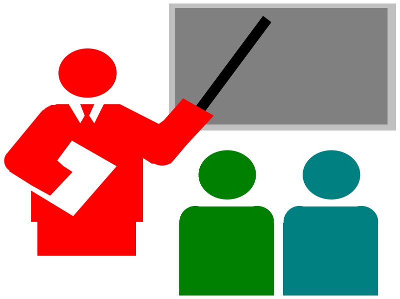

- Overbooking VBR voice causes queueing and performance becomes unacceptable.
- Instead of overbooking, it is better to fill the left-over bandwidth by ABR or UBR.
- Small buffering (1 or 2 cells ok). Larger buffering makes delay unacceptable.
- Scheduling or drop policies are important at higher loads or for asymmetric loads.

Future Work

- Higher link speeds: Relative multiplexing gain is expected to be higher for T3 and OC-3
- Sensitivity Analysis: Talk and silence times, allowed maximum delays
  Lower allowed delay $\Rightarrow$ Lower queueing
  $\Rightarrow$ Lower multiplexing
  $\Rightarrow$ Less overbooking
- AAL2: Low bit rate voice
- Downspeed

References

- Jayaraman Iyer, Raj Jain, Sohail Munir, Sudhir Dixit, "Performance of Compressed Voice Sources over VBR," ATM Forum/97-0608, July 1997, http://www.cis.ohio-state.edu/~jain/atmf/a97-0608.htm
- All our papers and ATM Forum contributions are available on-line at http://www.cis.ohio-state.edu/~jain