97-0608: Performance of VBR Voice over ATM: Effect of Scheduling and Drop Policies

Jayaraman Iyer, Raj Jain, Sohail Munir

The Ohio State University

Sudhir Dixit, Nokia Research Center

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
  - Scheduling Policies
  - 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

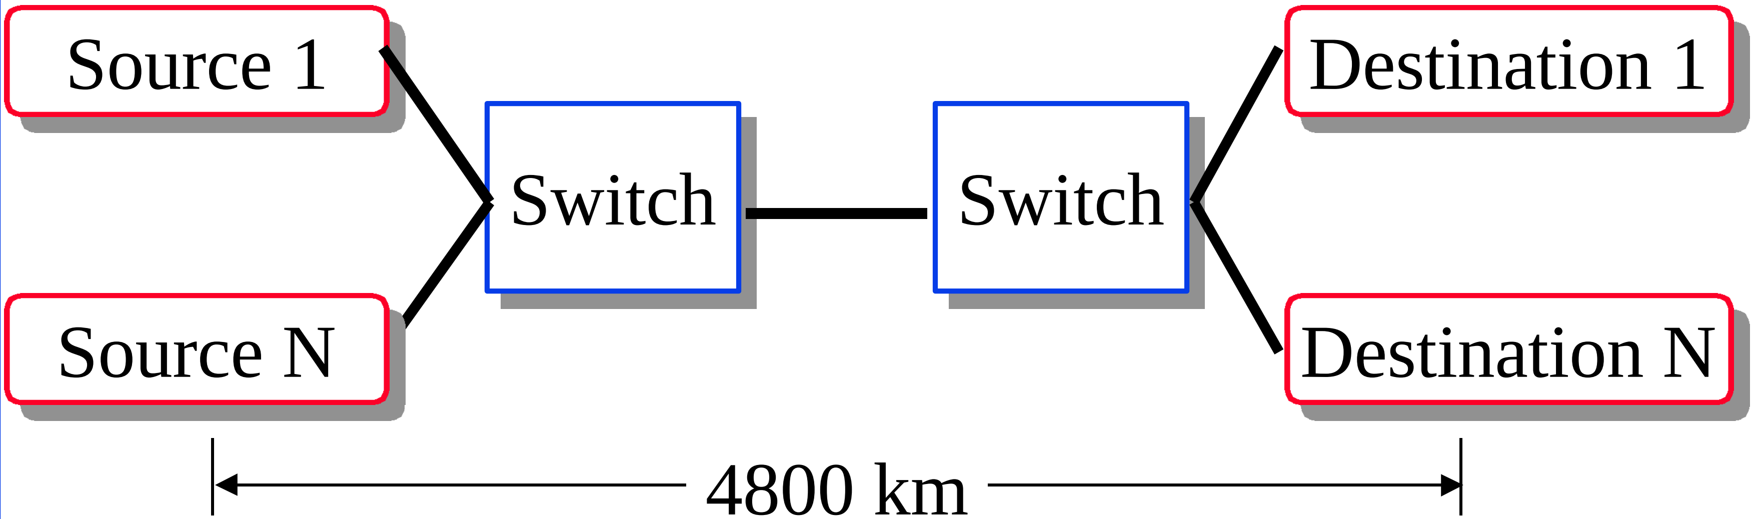

- Links between Switches = 1.544 Mbps (T1).
- N multiplexed 64-kbps VBR voice sources
  Silence suppression ⇒ VBR
- Per-VC Queuing at the Switch
  Multiple queues ⇒ need proper scheduling

Simulation configuration

- Propagation delay : 24 ms
- Avg packetization delays: 6 ms + 6 ms (PCM)
- 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

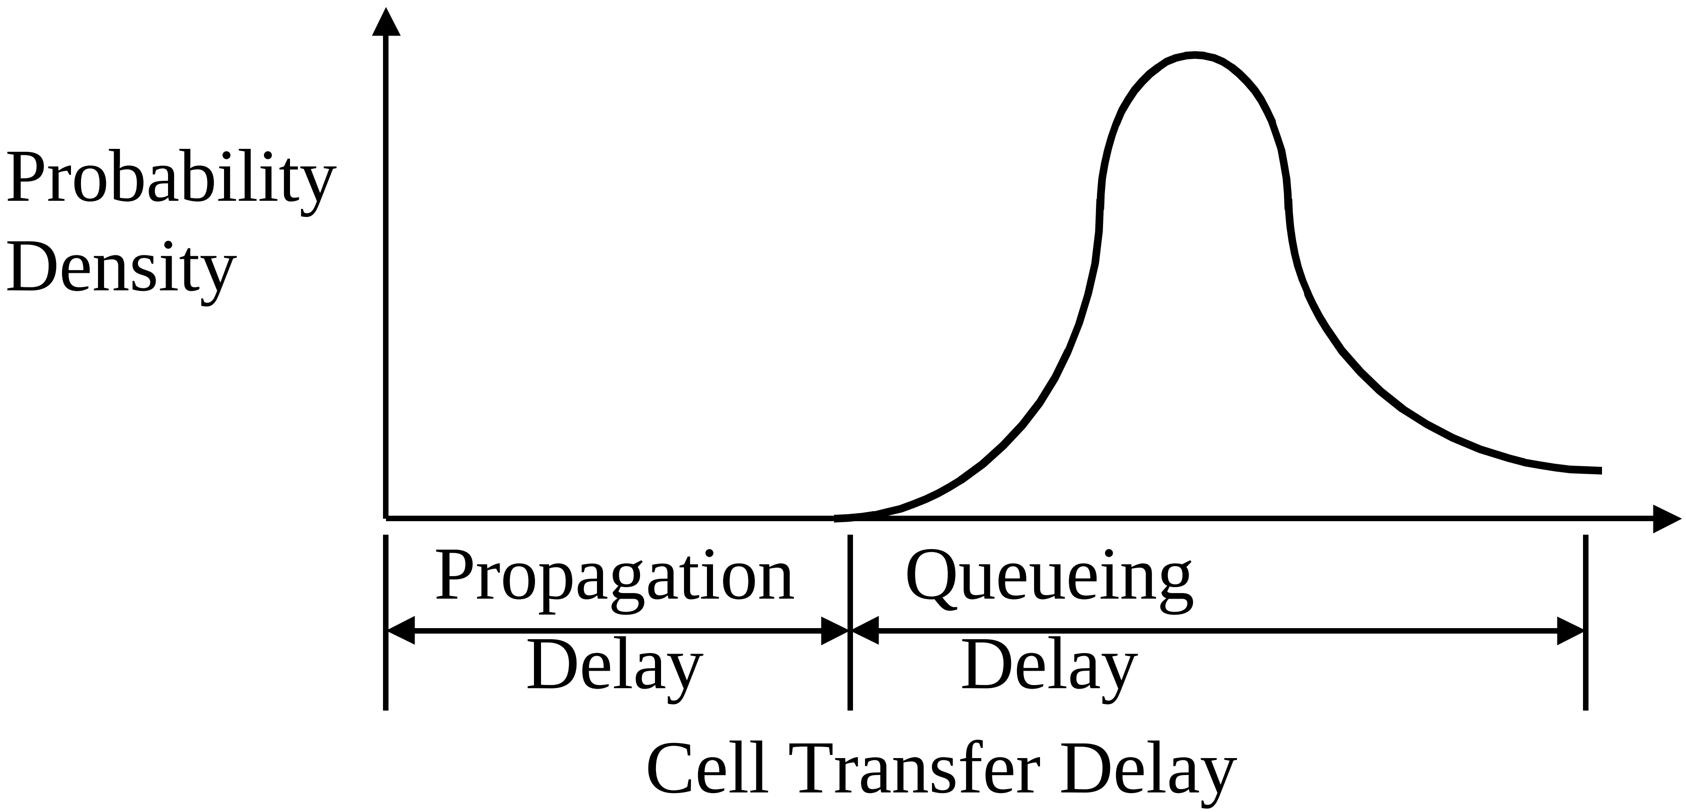

- For VBR voice, we need to specify Max CTD

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%

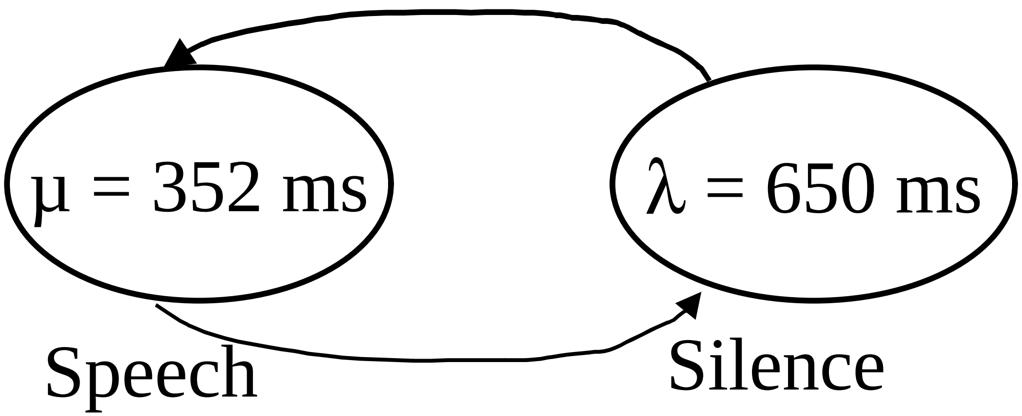

Performance Metrics

- 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
	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

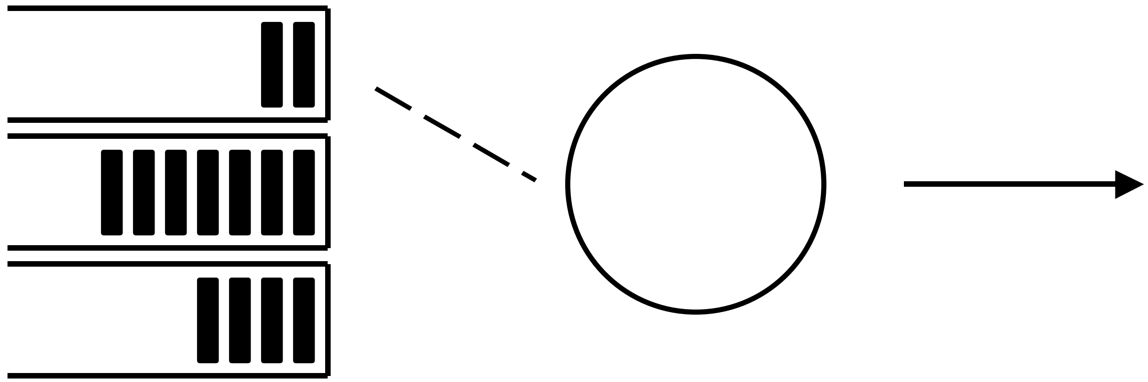

- Round Robin (RR)
- Earliest Deadline First (EDF)
- Longest Queue First (LQF)

Scheduling Results: 1 Buf/VC

NS	Buf	Sched	CLR	DVQ	Fairn
20	1	rr	0.0000	0.0000	1.0000
20	1	lqf	0.0000	0.0000	1.0000
20	1	edf	0.0000	0.0000	1.0000
24	1	rr	0.0000	0.0000	1.0000
24	1	lqf	0.0000	0.0000	1.0000
24	1	edf	0.0000	0.0000	1.0000
30	1	rr	0.1126	0.0011	1.0000
30	1	lqf	0.1126	0.0013	1.0000
30	1	edf	0.1126	0.0011	1.0000
35	1	rr	0.2400	0.0024	1.0000
35	1	lqf	0.2418	0.0027	1.0000
35	1	edf	0.2400	0.0024	1.0000

Scheduling Policies: Results I

- With more than 24 users, the cell loss rate is more than 10^{-3}
  - Compression does not allow overbooking
  - 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
20	2	rr	0.0000	0.0000	1.0000
20	2	lqf	0.0000	0.0000	1.0000
20	2	edf	0.0000	0.0000	1.0000
24	2	rr	0.0000	0.0000	1.0000
24	2	lqf	0.0000	0.0000	1.0000
24	2	edf	0.0000	0.0000	1.0000
30	2	rr	0.0616	0.0006	1.0000
30	2	lqf	0.0488	0.0010	1.0000
30	2	edf	0.0616	0.0006	1.0000
35	2	rr	0.1964	0.0031	1.0000
35	2	lqf	0.1764	0.0025	1.0000
35	2	edf	0.1964	0.0031	1.0000

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 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	Buf	Sched	CLR	DVQ	Fairness
40	2	rr	0.3865	0.0074	1.0000
40	2	lqf	0.3579	0.0047	1.0000
40	2	edf	0.3865	0.0073	1.0000
48	2	rr	0.6423	0.0132	1.0000
48	2	lqf	0.6161	0.0078	0.9999
48	2	edf	0.6371	0.0130	1.0000
60	2	rr	2.5959	0.0384	0.9999
60	2	lqf	2.4932	0.0354	0.9971
60	2	edf	2.5353	0.0357	0.9999
65	2	rr	4.9184	0.0693	0.9997
65	2	lqf	4.6462	0.0636	0.9899
65	2	edf	4.8210	0.0648	0.9998

Scheduling Results: Heavy Load

NS	Buf	Sched	CLR	DVQ	Fairness
70	2	rr	8.2518	0.1235	0.9994
70	2	lqf	7.9017	0.1027	0.9732
70	2	edf	8.1647	0.1075	0.9996
75	2	rr	12.7650	0.2079	0.9987
75	2	lqf	12.4222	0.1546	0.9363
75	2	edf	12.7535	0.1882	0.9990

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)

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 ⇒ No scheduling required
- Buffer threshold: 80% (for selective drop)

Drop Policies Results

NS	Drop	CLR	DVQ	Fairness
20	tail	0.0000	0.0000	1.0000
20	sel	0.0000	0.0000	1.0000
24	tail	0.0000	0.0000	1.0000
24	sel	0.0000	0.0000	1.0000
30	tail	0.0361	0.0011	1.0000
30	sel	0.0361	0.0011	1.0000
35	tail	0.1746	0.0027	1.0000
35	sel	0.1746	0.0027	1.0000
40	tail	0.3611	0.0049	1.0000
40	sel	0.3611	0.0049	1.0000

Drop Polices Results: Heavy Load

NS	Drop	CLR	DVQ	Fairness
48	tail	0.5938	0.0075	1.0000
48	sel	0.5938	0.0075	1.0000
60	tail	2.3042	0.0772	0.9990
60	sel	2.3042	0.0772	0.9990
65	tail	4.4562	0.1901	0.9971
65	sel	4.6682	0.0484	0.9998
70	tail	7.8797	0.3257	0.9861
70	sel	8.0486	0.0826	0.9994
75	tail	12.4850	0.4631	0.9636
75	sel	12.6091	0.1315	0.9991

Drop Policies: Results

- At low loads (up to 60%) both schemes behave identically.
- At higher loads, selective drop is better over plain FIFO drop.
- Fairness of selective discard is very close to 1.

Summary

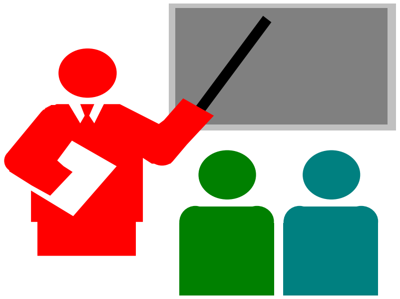

- Overbooking VBR voice causes queueing and performance becomes unacceptable.
- Instead of overbooking, it is better to fill the left-over bandwidth by ABR or UBR.
- Small buffering (1 or 2 cells ok). Larger buffering makes delay unacceptable.
- Scheduling or drop policies are important at higher loads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