

96-0178: ATM Switch Performance Testing Experiences

Gojko Babic, Arjan Duresi, Raj Jain

The Ohio State University

Contact: Jain@CIS.Ohio-State.Edu

<http://www.cis.ohio-state.edu/~Jain/>



- ❑ Measuring current metrics using current monitors
- ❑ Effect of monitor finite accuracy
- ❑ Background traffic
- ❑ Test configurations
- ❑ Experiences with Delay, Throughput, Frame Loss Rate, Fairness

Facts about Monitors

- ❑ Generators/Analyzers are expensive.
⇒ Test configurations should use as few of them as possible even for switches with large number of ports.
- ❑ Monitors have finite resolution.
Ours had a resolution of $0.5 \mu\text{s}$.
Averaging eliminates the effect of resolution.
- ❑ Monitors have internal path delays that must be subtracted from measured delay. For example, on an OC-3 link of 10 m, the measured inter-arrival cell time at full load was $2.83 \mu\text{s}$ but the CTD was $3.33 \mu\text{s}$.
- ❑ The measured inter-arrival times are accurate but the CTD

Measuring Frame Latency: Issues

- ❑ Most current monitors measure “Cell Transfer Delays (CTD).” How to compute MIMO frame latency from CTD?
- ❑ Monitors have limited clock accuracy. How does it affect latency measurements?
- ❑ What background traffic is appropriate?

CTD and MIMO Frame Latency


- MIMO = Min{LILO, FILO-NFOT}
or equivalently,

$$\text{MIMO} = \begin{cases} \text{LILO} & \text{if input rate} \leq \text{output rate} \\ \text{FILO - NFOT} & \text{if input rate} \geq \text{output rate} \end{cases}$$

Here, NFOT = Normalized frame output time
= Frame input time \times Input Rate/output rate

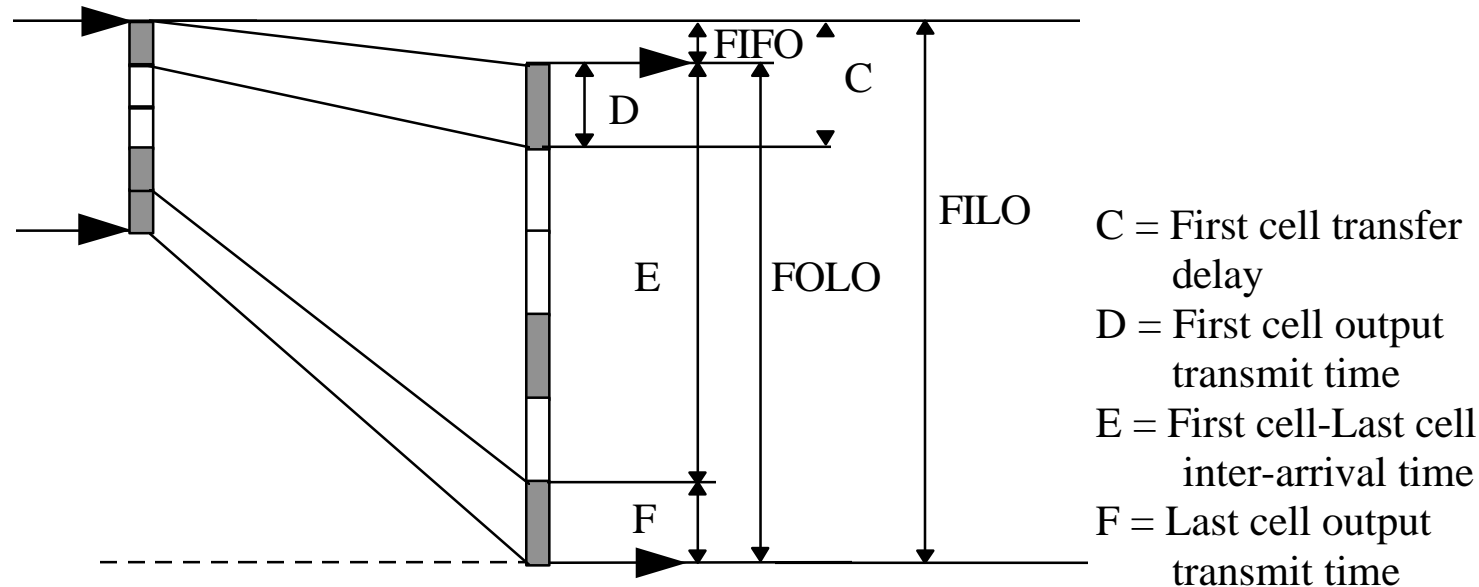
- CTD = First bit in to last bit out
= FILO cell latency

Input Rate \leq Output Rate



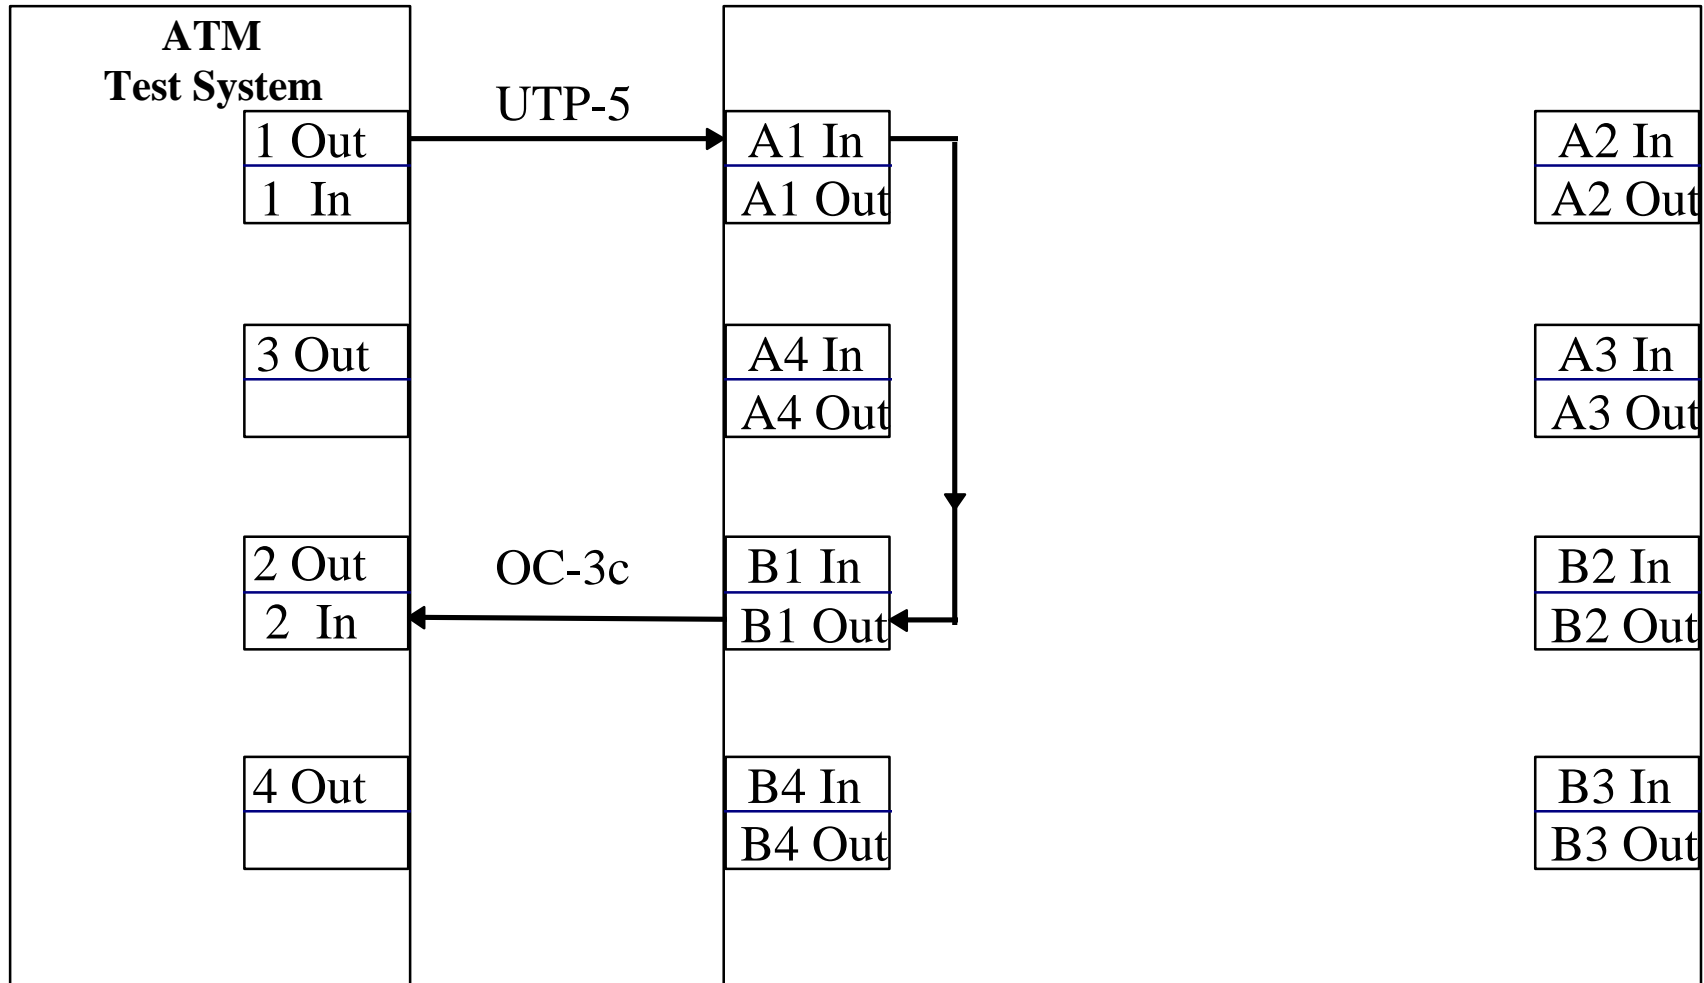
- ❑ MIMO Frame Latency
= LILO Frame latency
= Last cell's CTD - Last cell's input transmit time
- Test system's error
- ❑ Need to measure only last cell's CTD in this case.

Input Rate \geq Output Rate



- $MIMO = FILO - NFOT = FIFO + FOLO - NFOT$
 $= (\text{First cell's CTD} - \text{First Cells output transmit time}) +$
 $(\text{First cell to last cell inter-arrival time} + \text{Last cell output}$
 $\text{transmit time}) - NFOT - \text{Test system's error}$
- Need to measure First cell's CTD and first cell to last cell inter-arrival time.

Test Configuration for Latency (No background Traffic)



Test Method

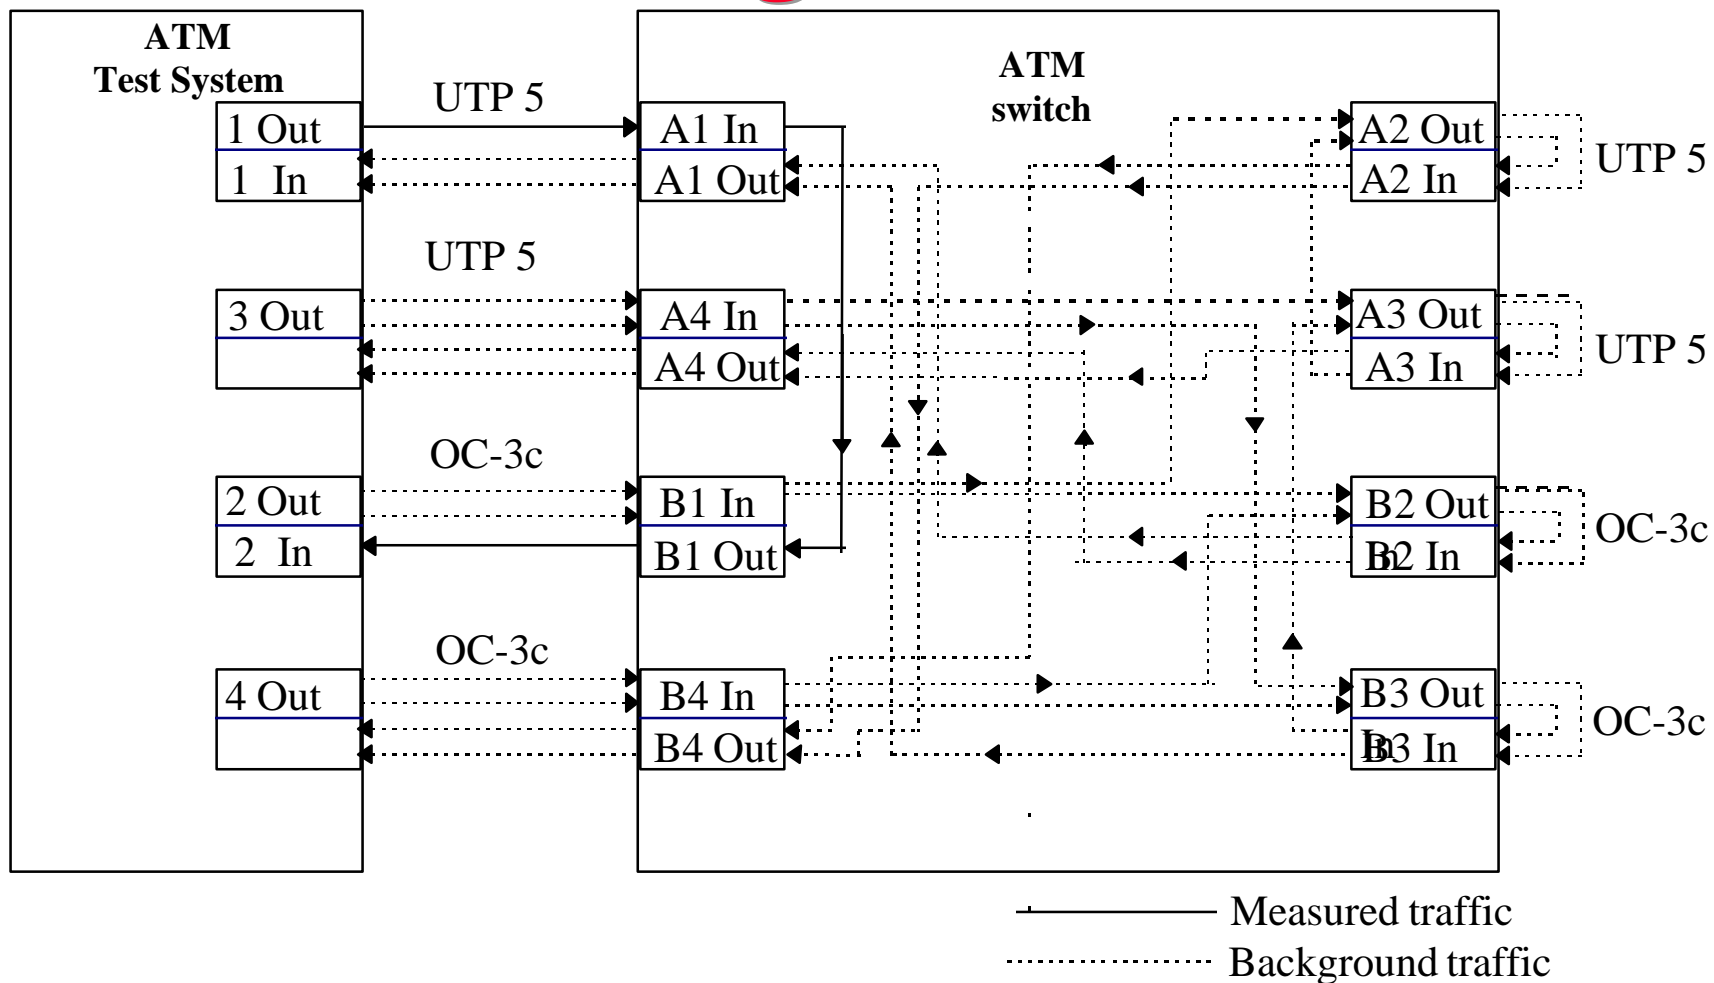
- ❑ Can measure MIMO latency using case I or II. Same result with either method.
- ❑ Foreground Traffic:
 - ❑ 4.63 Frames/sec
⇒ Inter-frame time of 0.216 sec
 - ❑ 192 cells/frame
 - ❑ Total 1000 frames
- ❑ Background Traffic: None
- ❑ Record average CTD of different cells

Test Results and Lessons

Run #	1st cell	2nd cell	2-191 cells	3-96 cells	97th cell	3-190 cells	98-190 cells	191st cell	192nd cell
1	19.02		20.52						20.77
2	19.06		20.54						20.78
3	19.04	19.21				20.53		20.79	20.77
4	19.07	19.21		20.31	20.75		20.76		20.78
5	19.07	19.19		20.32	20.73		20.76		20.78
6	19.14		20.58						20.83
7	19.13		20.58						20.81


- ❑ Cell CTD depends upon the cell's position in the frame.
- ❑ Cells later in the frame cells have larger CTD than those earlier in the frame
- ❑ Run 4: MIMO Latency = $20.78 - 3.33 = 17.45 \mu\text{s}$

Test Configuration with Background Traffic




Lessons

4-Port Switch



Foreground



Background

- ❑ Baseline configurations require all ports to be fully loaded. Need n generators and n analyzers.
- ❑ Can test with 2 generators and one analyzer by using “wrap arounds”
- ❑ The foreground traffic should not share any generator/analyzer logic (i.e., in the same direction on the same port) with background traffic to avoid distortions caused by the monitor.

Background Traffic

- ❑ Background traffic has a load intensity, arrival pattern, and service class
- ❑ Load Intensity:
 - ❑ For an n port switch:
Maximum background load (MBL)
 $= (n-1) * \text{Port rate}$
 - ❑ Other background intensities are expressed as a percentage of MBL
- ❑ Arrival Pattern: Equally spaced frames
- ❑ Service Class: CBR (higher priority than foreground)
UBR (Same priority as foreground)

Measurement Results

- With UBR background with frames of 2004 cells, Foreground frame size = 1000 cells

Load %	1st cell	2nd -999th cells	Last cell	MIMO	FIFO	Dif %
25	19.2	20.96	21.08	17.75	15.87	11.85
40	19.24	21.13	21.32	17.99	15.91	13.07
50	19.28	21.20	21.36	18.03	15.95	13.04
60	19.38	21.54	21.95	18.62	16.05	16.01
65	19.4	21.38	21.6	18.27	16.07	13.69
70	19.52	21.43	21.65	18.32	16.19	13.16
75	19.47	28.81	36.31	32.98	16.14	194.34
80*	19.32	57.54	94.17	90.84	15.99	468.11
90*	19.4	98	168.08	164.75	16.07	925.2
97*	19.49	122.42	207.54	204.21	16.16	1163.7

- * Foreground traffic is lost
- See contribution for other cases measured

Lessons

- ❑ FIFO frame latency does not change with background load \Rightarrow Does not reflect performance degradation caused by background.
- ❑ FIFO measures only first cell's latency \Rightarrow Not a good frame level metric
- ❑ MIMO latency depends upon the background intensity.
- ❑ Foreground traffic is lost even though it does not share any port in the same direction with background. It is important to know the background intensity at which this happens.


Lessons (Cont)

- Near MBL, background intensity points should be closely spaced while near low load the points can be far apart.
 - Proposal: Measure at 0, 0.5, 0.75, 0.875, 0.9375, 0.9687, ...
 $1-2^{-k}$, $k=0, 1, 2, \dots$
- Latency depends upon the background frame size and the foreground frame size.
 - Proposal: Measure at AAL payload sizes of 64B, 1518B, 9188B, and 64kB

Lessons (Cont)

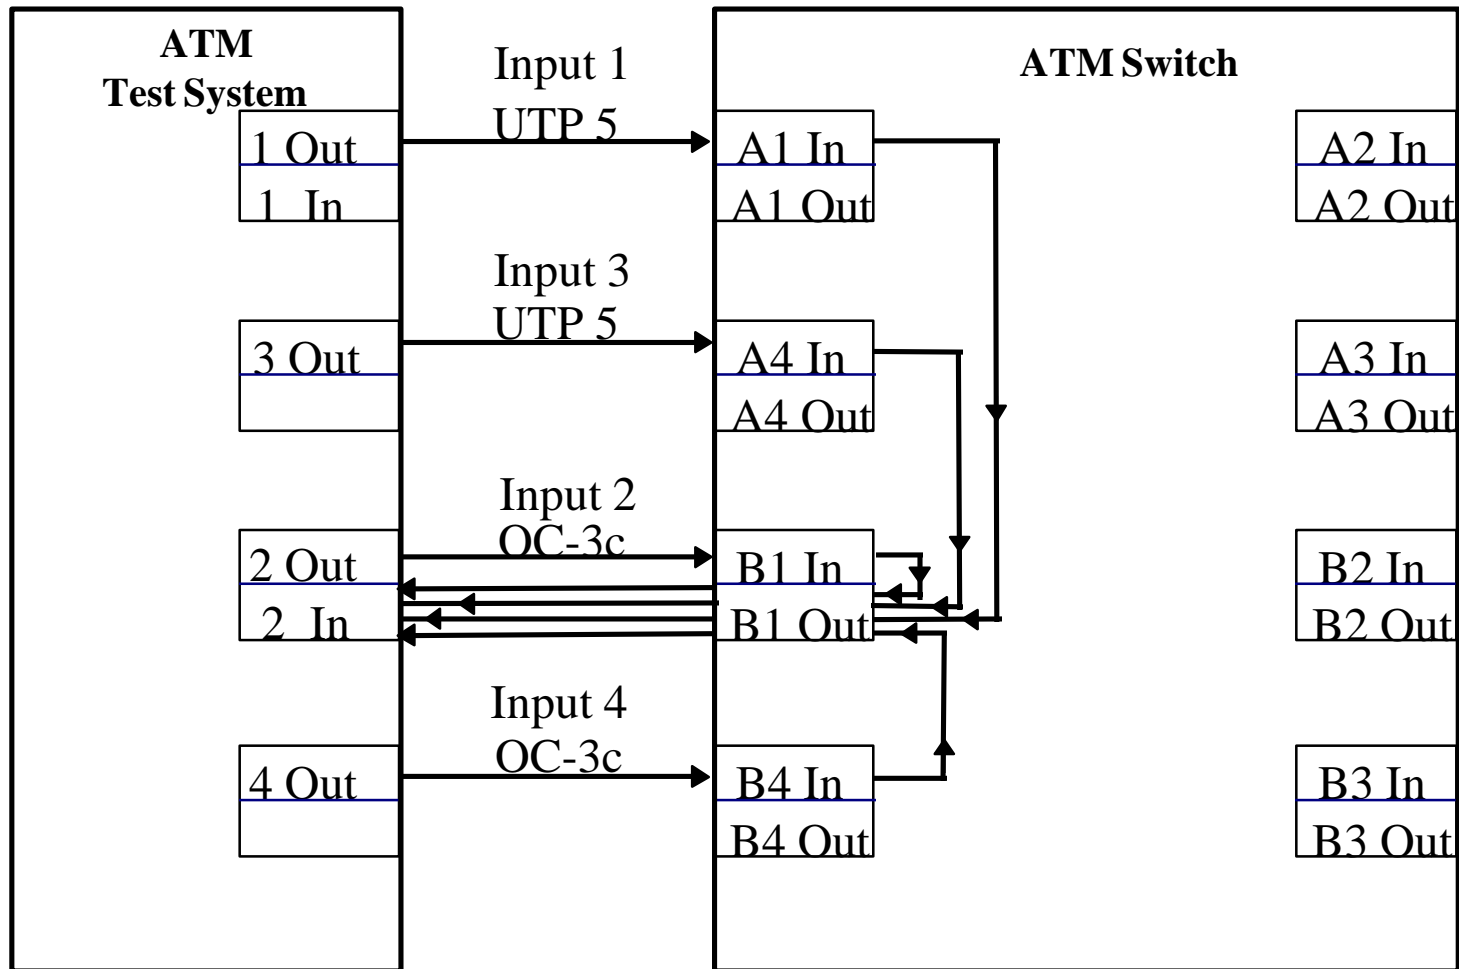
- Latency depends upon the background traffic class and its priority relative to foreground traffic's priority.
 - Proposal: Measure with highest, lowest and same priority background traffic.

Throughput




- ❑ Baseline specifies four configurations: n-to-n cross, n-to-n, n-to-1, and 1-to-n (multicast)
- ❑ Measured throughput for n-to-1 case with $n=2, 3, 4$
- ❑ Our monitor limited to one AAL5 VC/port
⇒ Used AAL 3/4
(but results will not be different for AAL5)

Test Configuration



Lessons



- ❑ Throughput increases and suddenly drops to zero at 100% output link rate
- ❑ Lossless throughput = Peak throughput
⇒ Remove peak throughput metric
- ❑ Full-load throughput = 0 for n-to-1

Lessons (Cont)

- Little variability in results
⇒ No need to report std. error.
- No unfairness when underloaded
But frame loss is different for different streams
⇒ Replace “Throughput Fairness” by “Loss Fairness”

Results

- Load = 100.32% of output link rate


Metric	Input 1	Input 2	Input 3	Input 4
Cell Loss Ratio	0.0036	0.0022	0.0033	0.0026
Frame Loss Ratio	0.2620	0.2050	0.2890	0.2260
Cell Miss-ins. Rate [cell/sec]	0.0000	0.0000	0.0000	0.0000

- Load = 120% of output link rate

Metric	Input 1	Input 2	Input 3	Input 4
Cell Loss Ratio	0.0637	0.0520	0.0630	0.0771
Frame Loss Ratio	0.7340	0.7350	0.6310	0.8760
Cell Miss-ins. Rate[cell/sec]	146	117	137	181

Lessons

- ❑ Cell loss ratios may be small but frame loss ratios are large
- ❑ There is unfairness in frame loss ratio
- ❑ At 100.32% load, 22-26% of the frames are lost
- ❑ At 120% load, 63-87% of the frames are lost
- ❑ At 400% load, almost all frames are lost
- ❑ At 120% and higher load, there is cell misinsertion



Summary

- ❑ MIMO latency can be measured even with current cell-level monitors.
- ❑ CTD of a cell depends upon a cell's position in the frame and, therefore, varies widely. Mean CTD is statistically not meaningful.
- ❑ Frame transfer delay depends upon foreground intensity, service class, and frame size and upon background intensity, service class, and frame size.
- ❑ Loopbacks can be used to fully load an n-port switch using just one generator/analyzer.

Summary (Cont)

- ❑ Peak load is equal to lossless throughput, we may remove one of the two metrics
- ❑ Variance in throughput is negligible. Remove standard error of throughput.
- ❑ For n-to-1 configurations, full load throughput is zero. Remove full load throughput.
- ❑ Replace “Throughput fairness” by “Loss fairness”