Performance Benchmarking of ATM Devices

Raj Jain

Professor of Computer and Information Science

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

Our Team

- Dr. Gojko Babic
- Dr. Arjan Durresi
- Justin Dolske
- Shabbir Shahpurwala

- What is Performance Benchmarking?
- Goals of ATM Forum Performance Testing WG
- Current definitions: Throughput, latency, fairness, frame loss rate, maximum burst size, connection establishment latency
- MIMO Frame Latency
- Measurement experiences

Dictionary Definition

- **Benchmark** *v. trans*. To subject (a system) to a series of tests in order to obtain prearranged results not available on competitive systems.

From: The Devil's DP Dictionary
S. Kelly-Bootle

Other Networking Benchmarks

- Benchmarking Methodology Working group (BMWG) formed in January 1990
- RFC 1242 "Benchmarking Methodology for Network Interconnection Devices" written in July 1991.
- RFC 1944, "Benchmarking Methodology for Network Interconnect Devices", May 17, 1996, 30 pp.
- Defined a number of terms that are commonly (mis)used by vendors

Why Do This at ATM Forum?

- ATM Forum has the most interest of making ATM successful (compared to IETF)
- Confusion caused by differing terminology and differing benchmarks will eventually lead to customer dis-satisfaction
- Better customer information will contribute to more customer satisfaction and more sales and hence success of ATM.

Cell Level vs Frame Level

- Performance benchmarking
  = Performance seen by the user ≠ Cell level QoS
  For example, CLR = 0.1% may mean a frame loss rate of 0.1% in one switch or 0.001% in another.
- Data applications care for frame loss rate and not CLR.
- Video applications care for
  - Frame loss rate
  - Frame delay variation
  - Frame transfer delay

SCOPE: Goals

- Define metrics that help the customer compare various ATM (and possibly non-ATM) equipment.
- The metrics should be independent of switch architectures.
  They should apply to all architectures.
- Develop precise methodologies for measuring these metrics.
  Methodology = Procedure + Configuration + Traffic Pattern
  $\Rightarrow$ Anyone (user or vendor) can conduct it and come up with the same result.

Non-Goals

- ATM Forum will not do any measurements.
- Forum will not certify any measurements.
- Will not set any performance thresholds
  - Setting thresholds can kill the performance-cost tradeoffs
  - Example 1: Frame loss rate should be no more than 1%
  - Example 2: Switch delay should be less than 1 ms.

OSU National ATM Benchmarking Lab

- Performance Benchmarking
- Presentations at NetWorld+Interop Atlanta (Sep 1995)
- Presentation at ATM Forum (Oct 1995)
- Defining metrics and measurement methodology
- Benchmarks run in our lab
- The benchmark scripts can be run by any manufacturer
- Modeled after Harvard benchmarking lab for routers

Metrics

- Throughput: Lossless, Peak, Full Load
- Latency: MIMO
- Frame Loss Rate
- Throughput Fairness
- Maximum Frame Burst Size
- Connection establishment latency

Throughput

- Lossless, Peak, Full-load
- Unit = bits/sec

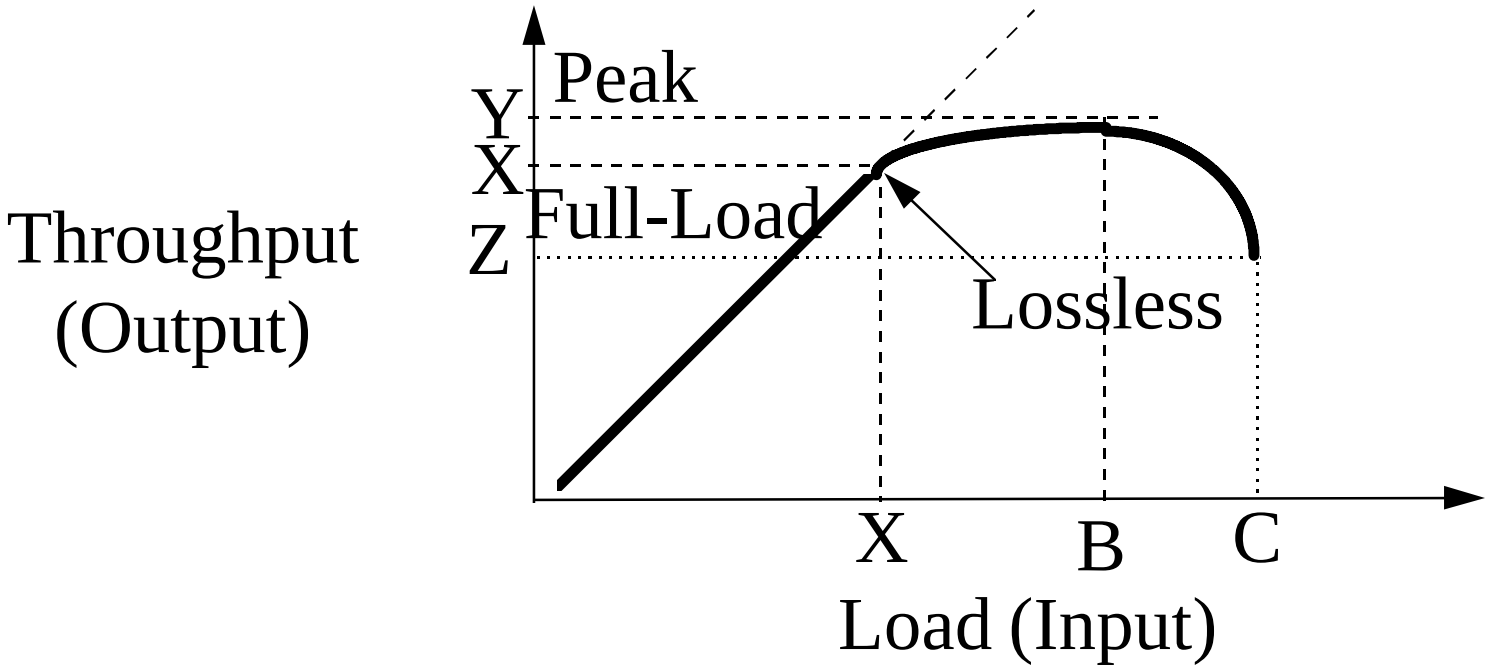

Connection Configurations for Foreground Traffic

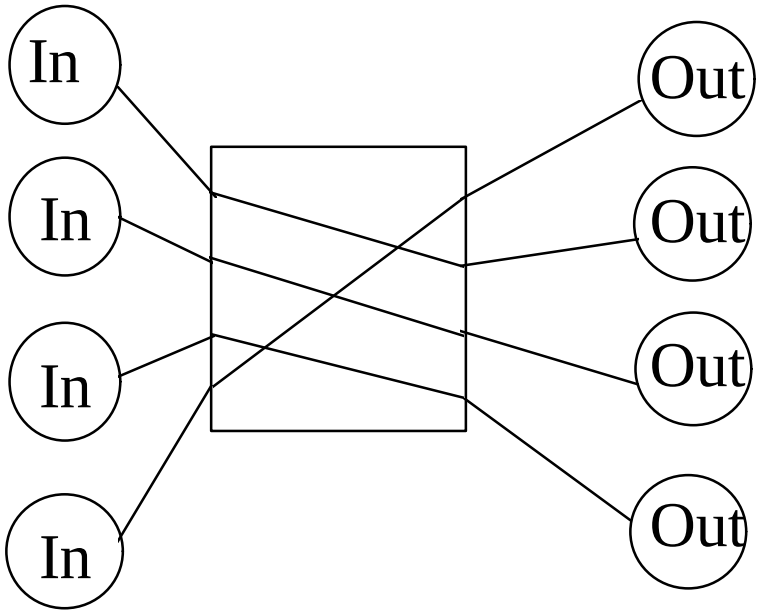

a. n-to-n straight: *n* VCCs; *n*=4

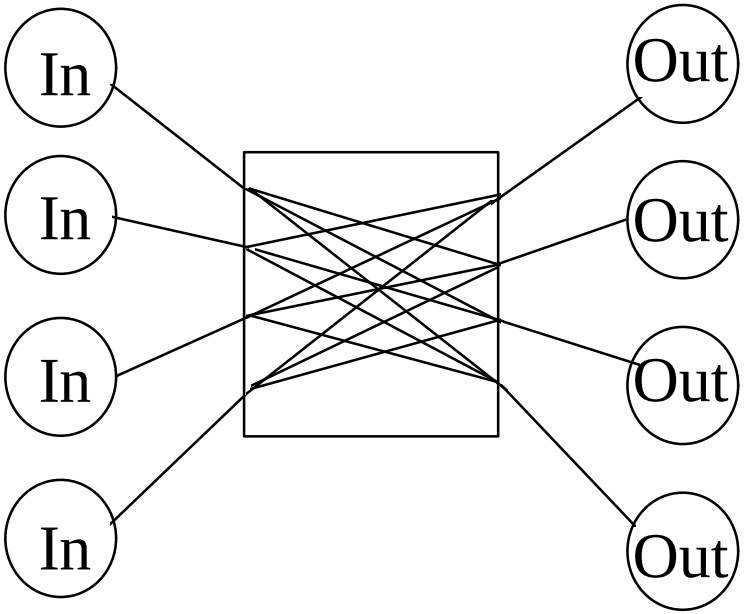

b. n-to-(n–1) full cross: *n*× (*n*–1) VCCs; *n*=4

Note: Inputs are shown on the left. Outputs on the right.

Configurations (Cont)

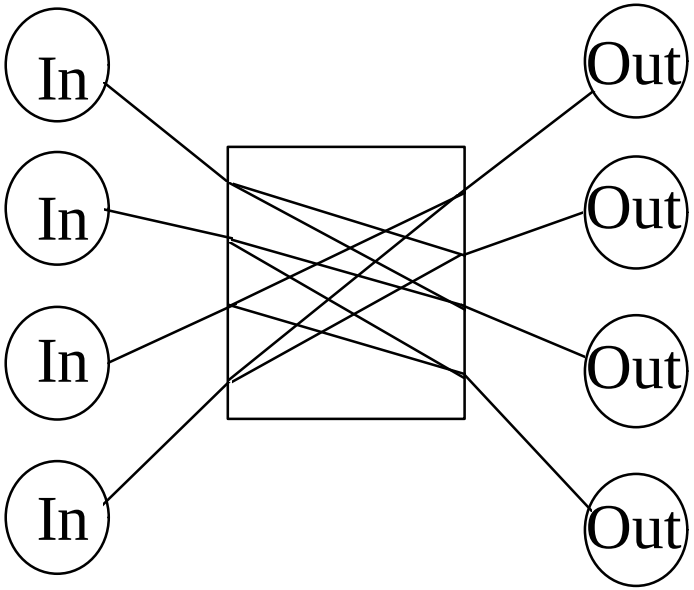

c. n-to-m partial cross: $n \times m$ VCCs; n=4, m=2

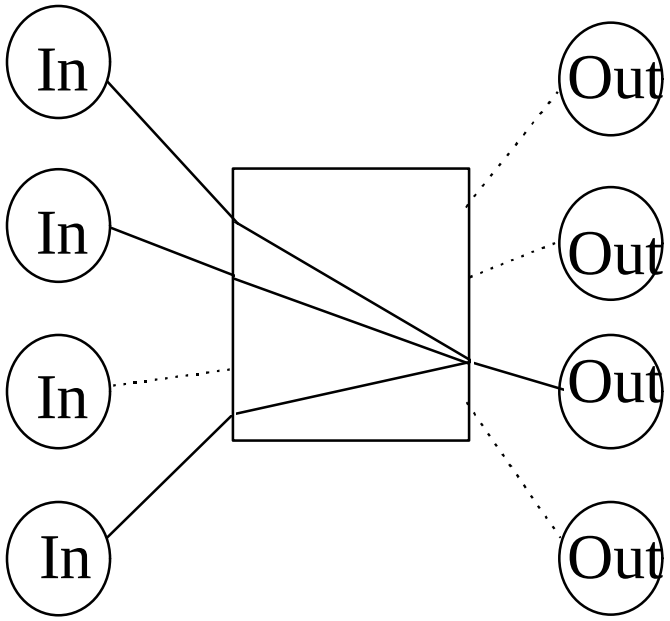

d. k-to-1: k VCCs; k=3

Configurations (Cont)

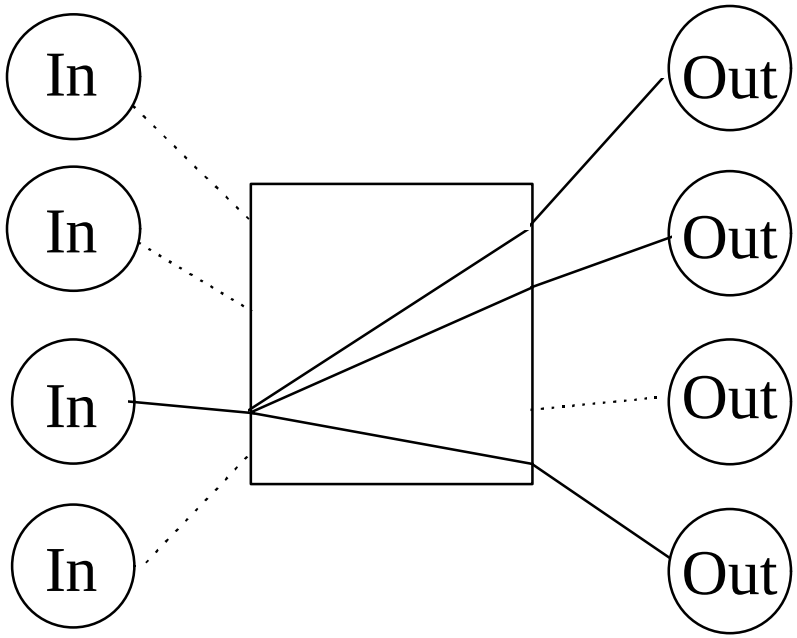

e. 1-to-(n-1): one (multicast) VCC

Scaleable Test Configurations

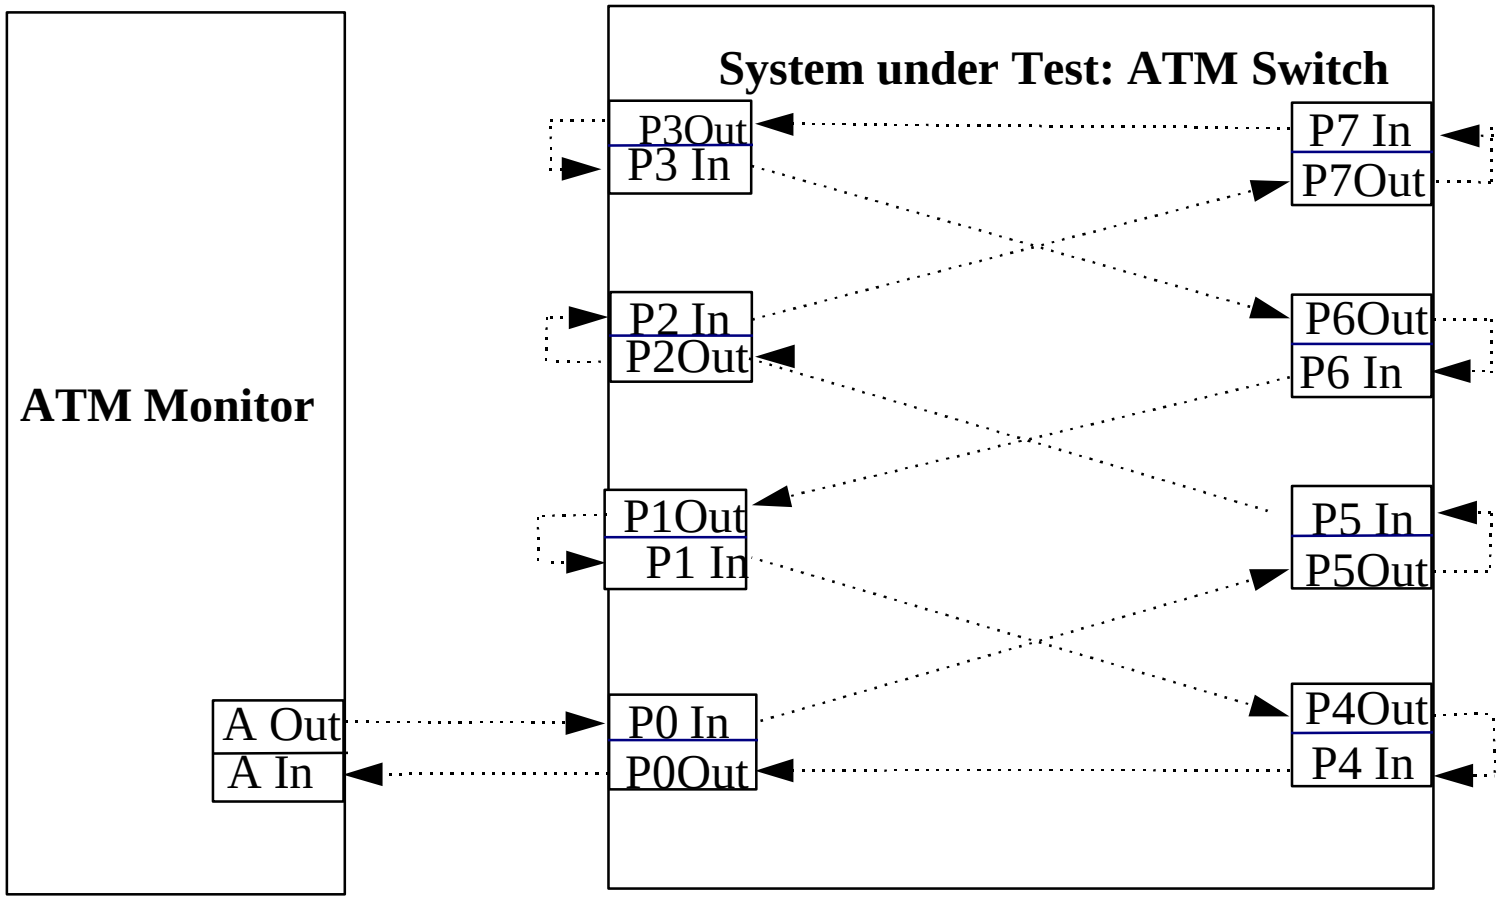

n-to-n Straight

Scaleable Test Configurations II

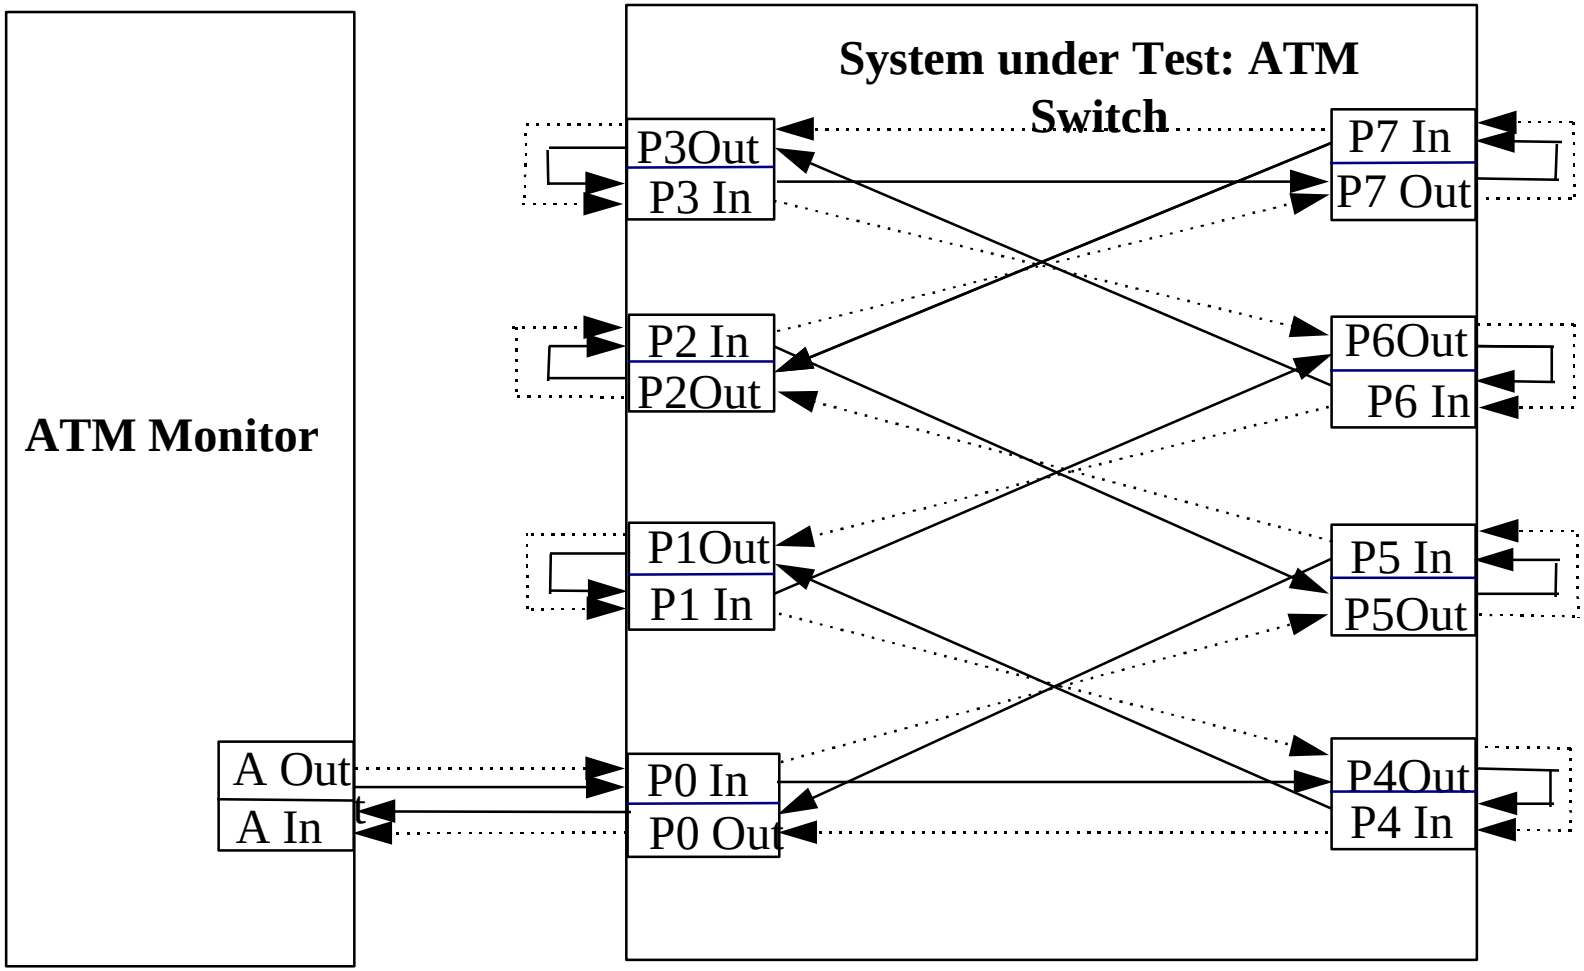

8-to-2 partial cross

Latency Measurement: Overview

- Input frame not contiguous
- Output frame not contiguous
- Input Speed ≠ Output Speed
- Single Cell
- Single Bit

Latency: Single Bit

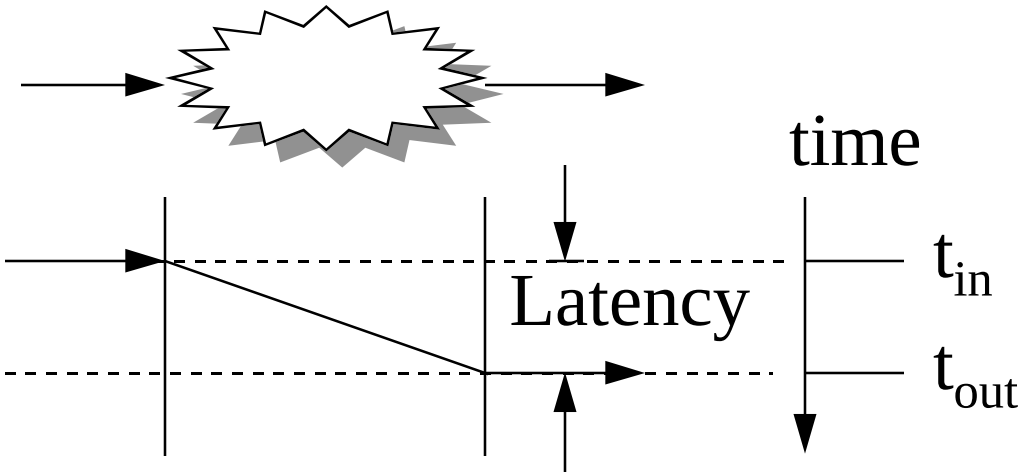

- Latency = t_{out} - t_{in}

Latency: Single Cell

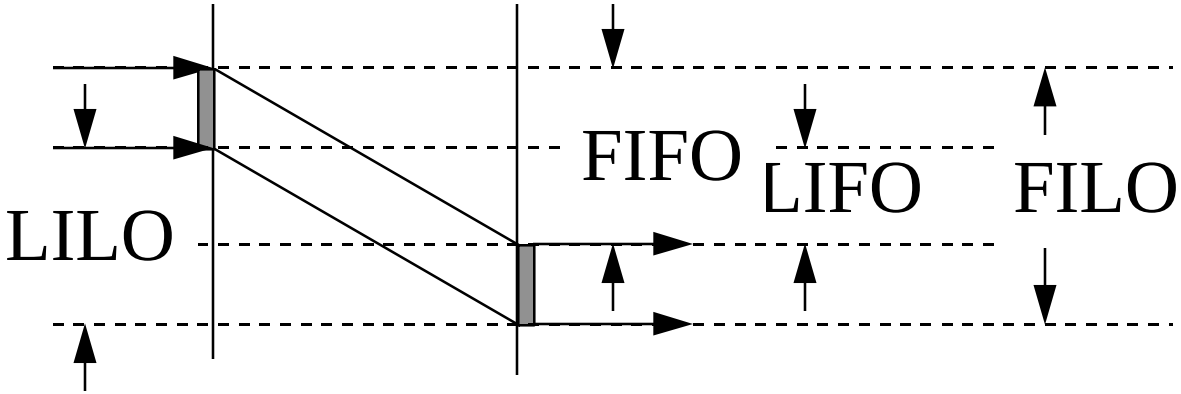

- FIFO = First-bit in to first-bit out
- LIFO = Last-bit in to last-bit out
- FILO = First-bit in to last-bit out
- LILO = Last-bit in to last-bit out
- FIFO = LILO = FILO - Cell time = LIFO + Cell time
- Assumes input speed = output speed

Frame Latency

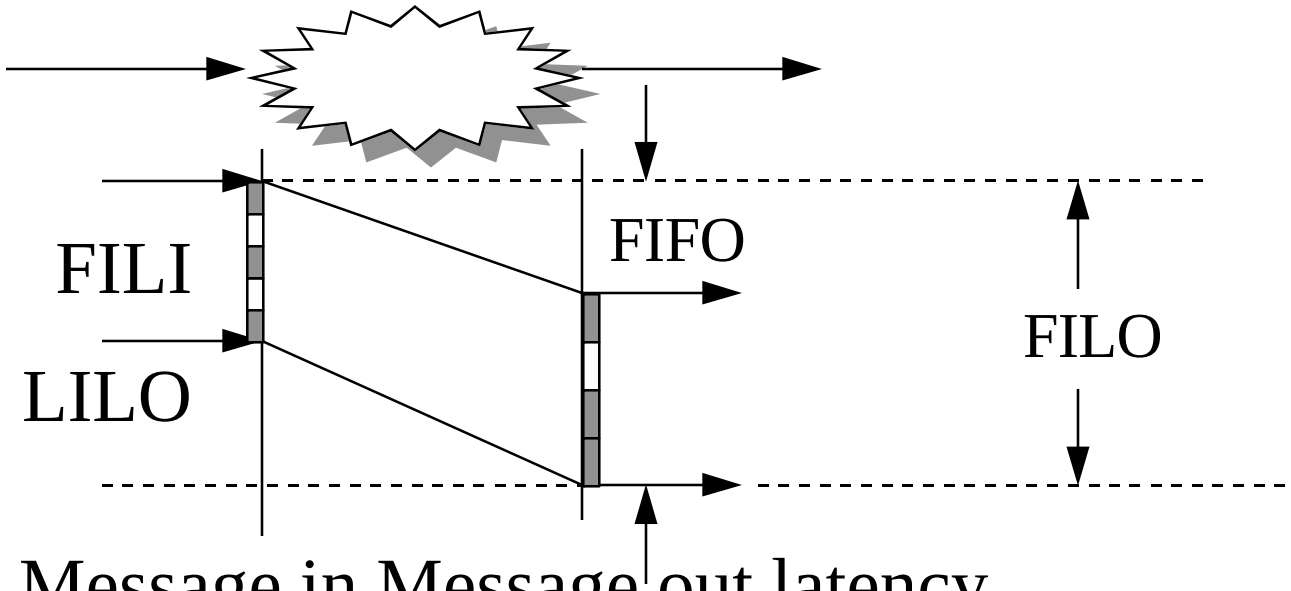

- MIMO = Message in Message out latency
  = min{LILO, FILO-Normalized Frame Output Time}
- NFOT = FILI*Input link rate/output link rate
- Applies even when frames are not contiguous
  Even when: Input rate <> Output rate

Latency Definition: Requirements

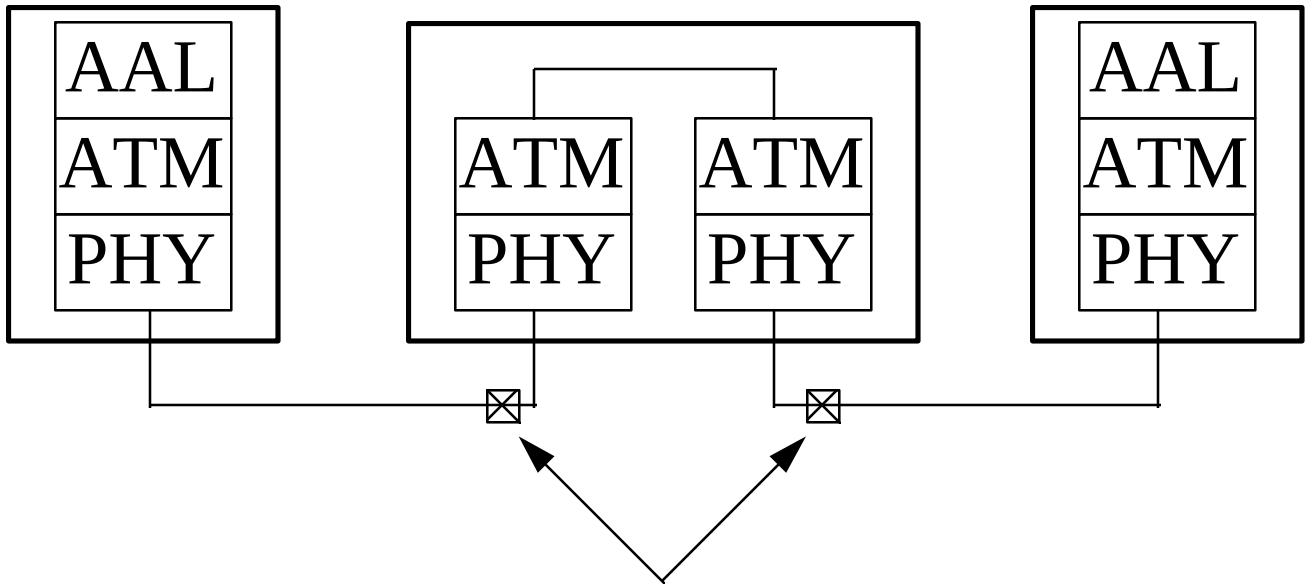

- Host speed should not affect the measured switch performance
- Delay caused by input/output link speeds should not be attributed to switch latency.

Simple Switches

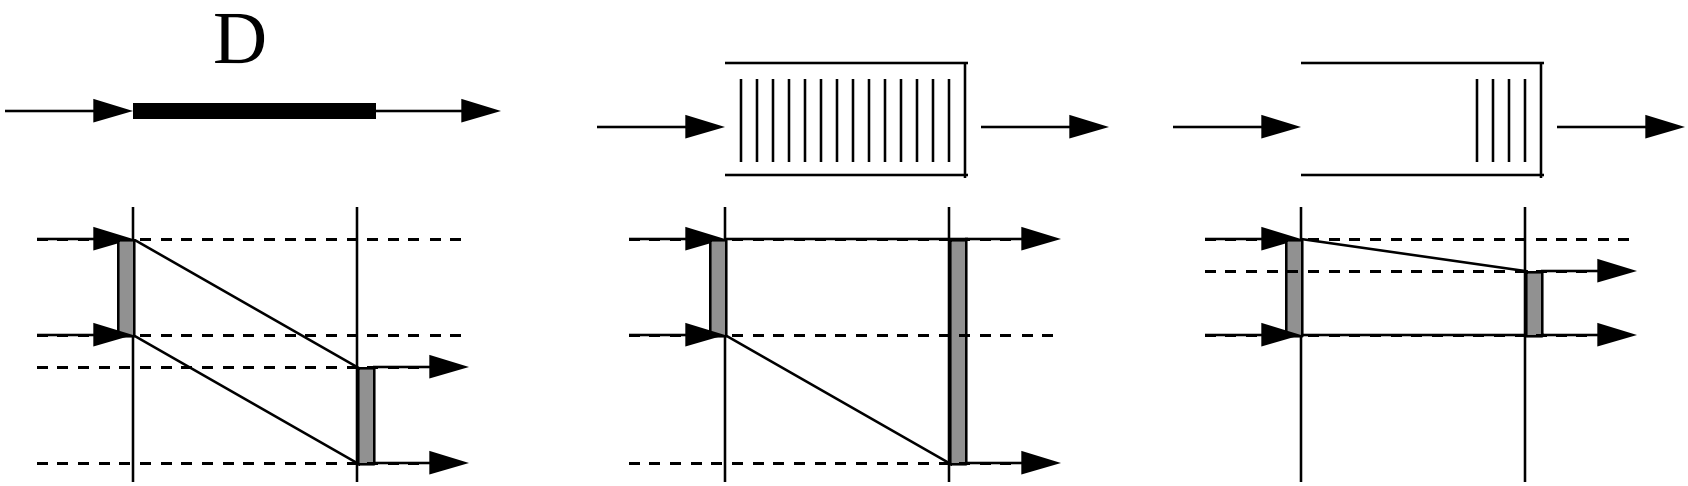

(a) Wire: $R_{in} = R_{out}$ (b) FIFO: $R_{in} > R_{out}$ (c) Intelligent FIFO: $R_{in} < R_{out}$

- Intelligent FIFO:
  Knows cell size. Holds first bit to avoid underrun.
- These components have known switching latencies
- Combinations of these ⇒ Known latency switches

Zero-Delay Switch
Input Speed = Output Speed

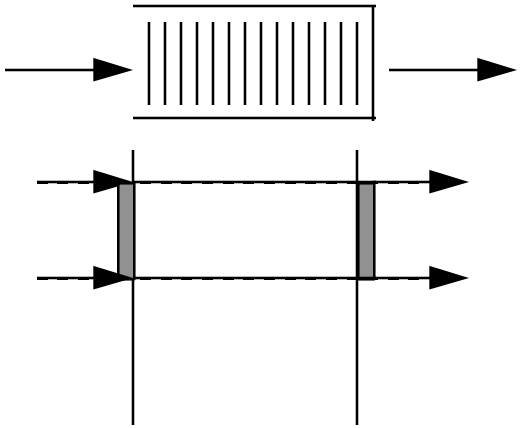

- FIFO is zero ⇒ Correct
- LILO is zero ⇒ Correct
- MIMO = Min{LILO, FILO - Cell input time*R_{in}/R_{out}}
  = 0 ⇒ Correct

Zero-Delay Switch
Input Speed > Output Speed

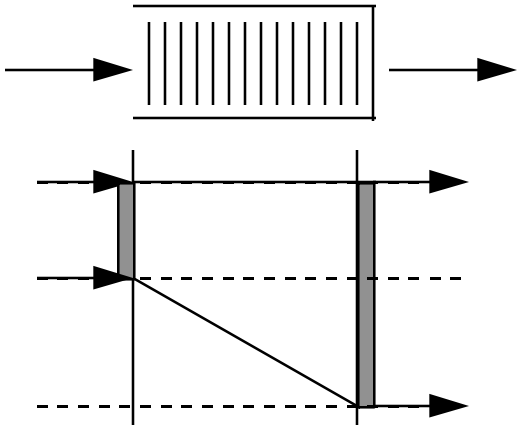

- FIFO is zero ⇒ Correct
- LILO is non-zero ⇒ Incorrect
- MIMO = Min{LILO, FILO - Cell input time*R_{in}/R_{out}}
  = 0 ⇒ Correct

Zero-Delay Switch
Input Speed < Output Speed

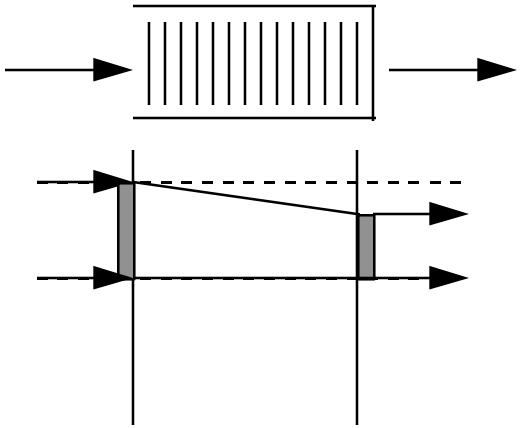

- FIFO is non-zero ⇒ Incorrect
- LILO is zero ⇒ Correct
- MIMO = Min{LILO, FILO - Cell input time*R_{in}/R_{out}}
  = 0 ⇒ Correct

Non-Zero Delay Switches

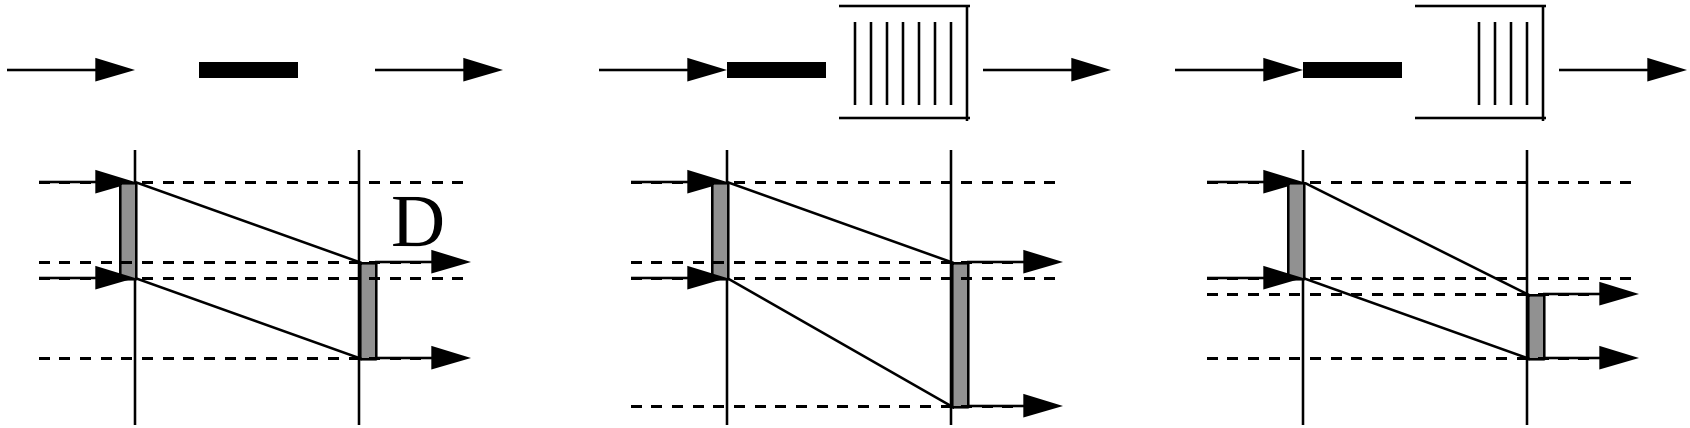

(a) $R_{in} = R_{out}$ (b) $R_{in} \geq R_{out}$ (c) $R_{in} < R_{out}$

(a) FIFO = LILO = MIMO = D

(b) FIFO = MIMO = D, LILO is incorrect

(c) LILO = MIMO = D, FIFO is incorrect

Summary: Single Cell

No.	Case	FIFO	LILO	MIMO
1	Input speed = output speed	√	√	√
2	Input speed ≥ output speed	√	×	√
3	Input speed < output speed	×	√	√

- MIMO is the only metric that applies to all cases.
- These results also apply to contiguous frames

Discontigous Frames: Summary

No.	Case	FIFO	LILO	MIMO
1aD	Input rate = Output rate, no change in gaps	√	√	√
1bD	Input rate = Output rate, expansion of gaps	×	√	√
1cD	Input rate = Output rate, compression of gaps	×	√	√
2aD	Input rate < Output rate, no change in gaps	×	√	√
2bD	Input rate < Output rate, expansion of gaps	×	√	√
2cD	Input rate < Output rate, compression of gaps	×	√	√
3aD	Input rate > Output rate, no change in gaps	√	×	√
3bD	Input rate > Output rate, expansion of gaps	×	×	√
3cD	Input rate > Output rate, compression of gaps	×	×	√

Frame Latency from Cell-level Data

- If Input rate ≤ Output Rate, MIMO Latency = LILO
  ⇒ MIMO can computed from last cell's CTD
- If Input rate ≥ Output Rate
  MIMO Latency = FIFO + Frame output time - Normalized Frame Output Time
  ⇒ MIMO can be computed from first cell's CTD and first cell to last cell inter-arrival times
- If Input rate = output rate, either one can be used

1st cell delay	last cell delay	1st cell-last cell inter-arrival time	MIMO latency [1]	FIFO latency	FOLO time	FILO latency	MIMO latency [2]
21.5	21.5	541.0	18.2	18.2	543.83	562.03	18.44
18.5	21.0	543.5	17.7	15.2	546.33	561.53	17.94

Scaleable Test Configuration

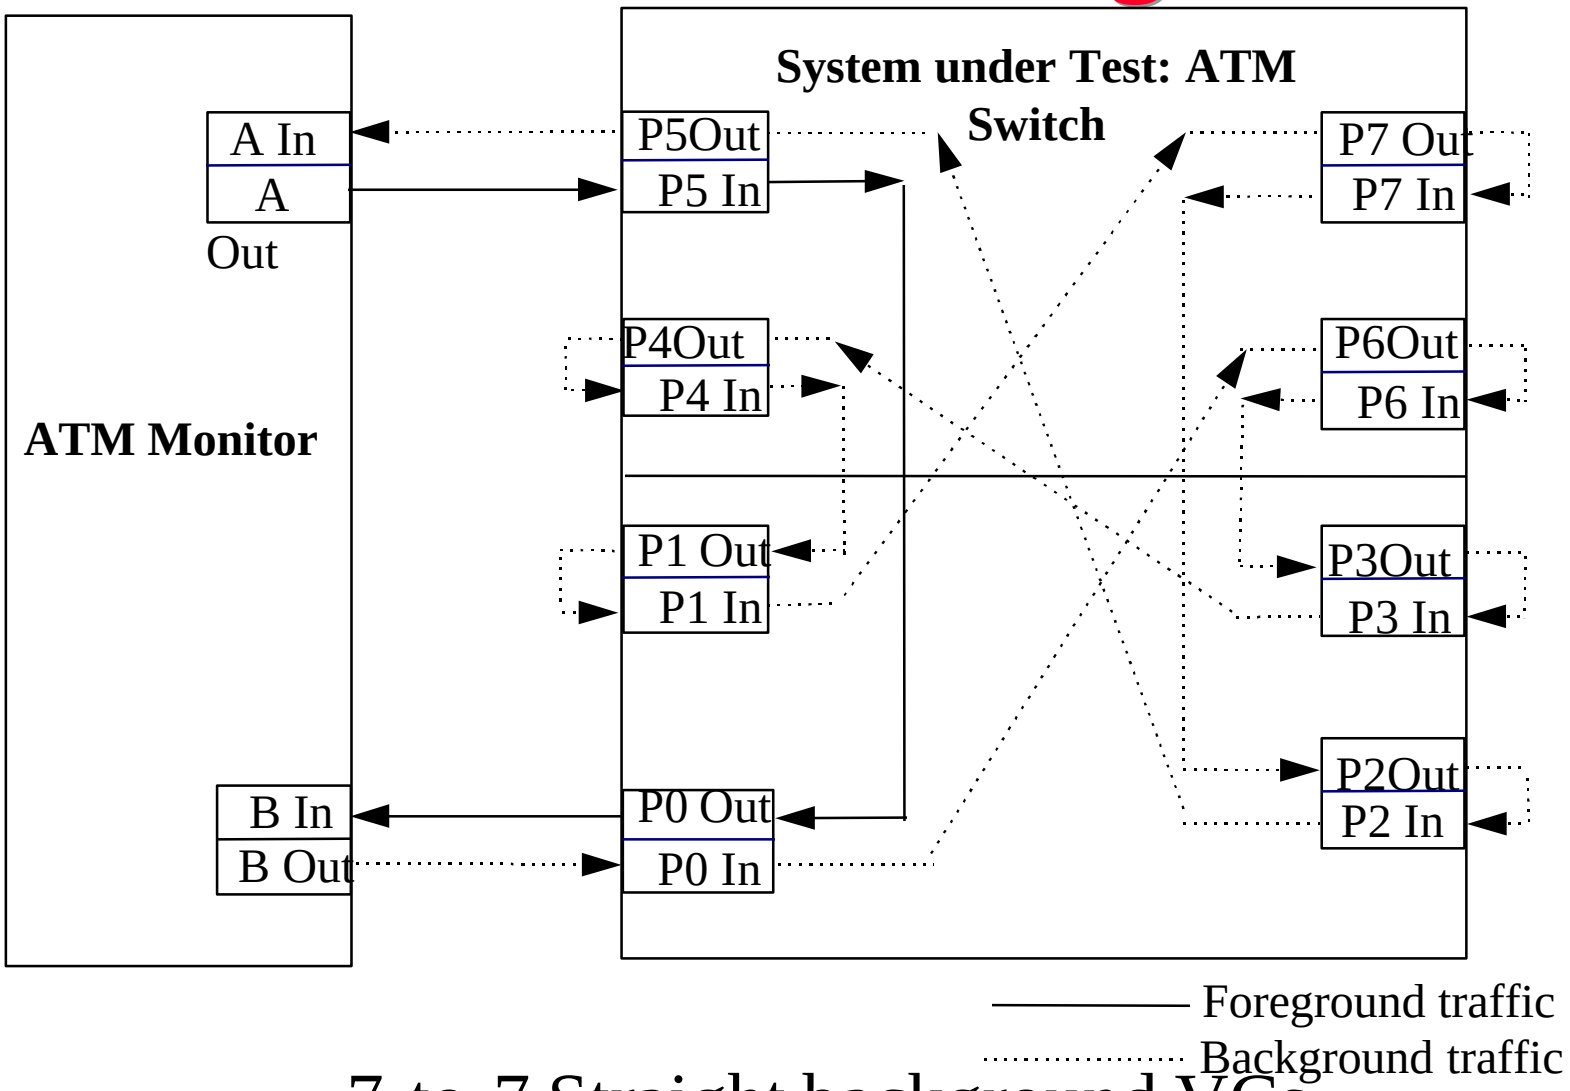

7-to-7 Straight background VCs

Performance Metrics

1. Throughput: Lossless, peak, full-load
2. Latency = Min{LILO, FILO- NFOT} = MIMO
3. Throughput fairness = $(\Sigma x_i)^2/(n\Sigma x_i^2)$
4. Frame loss rate = (Input Rate - Throughput)/Input Rate
5. Maximum frame burst size = # of back to back frames
6. Call establishment latency
   = MIMO latency of Setup + MIMO latency of Connect

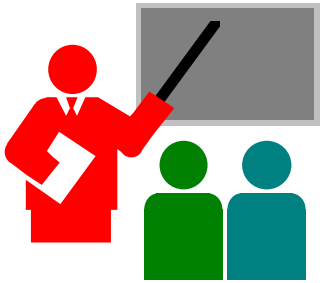

Summary

- Performance benchmarking is important to avoid customer confusion
- Frame-level not cell-level metrics
- Test configurations: n-to-1, n-to-n, n-to-k partial cross, n-to-(n-1) full cross, ...
- Scalable test configurations require only one or two generator and monitor regardless of number of ports.
- MIMO latency can be measured even with current cell-level monitors.

References

- See http://www.cis.ohio-state.edu/~jain/
- Raj Jain and Gojko Babic, "Performance Testing Effort at the ATM Forum: An Overview," To appear in IEEE Communications Magazine, Version April 10, 1997, 11 pp., http://www.cis.ohio-state.edu/~jain/papers/perf_com.ps