96-0811: Frame-Level Throughput and Latency Metrics - Proposed Text

Raj Jain, Gojko Babic, Bhavana Nagendra

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

- Throughput
- Latency

Goals

- Frame-level only.
- Definitions applicable at AAL layer
  Extendible to higher layers.
- Applicable to switch or group of switches

Throughput

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

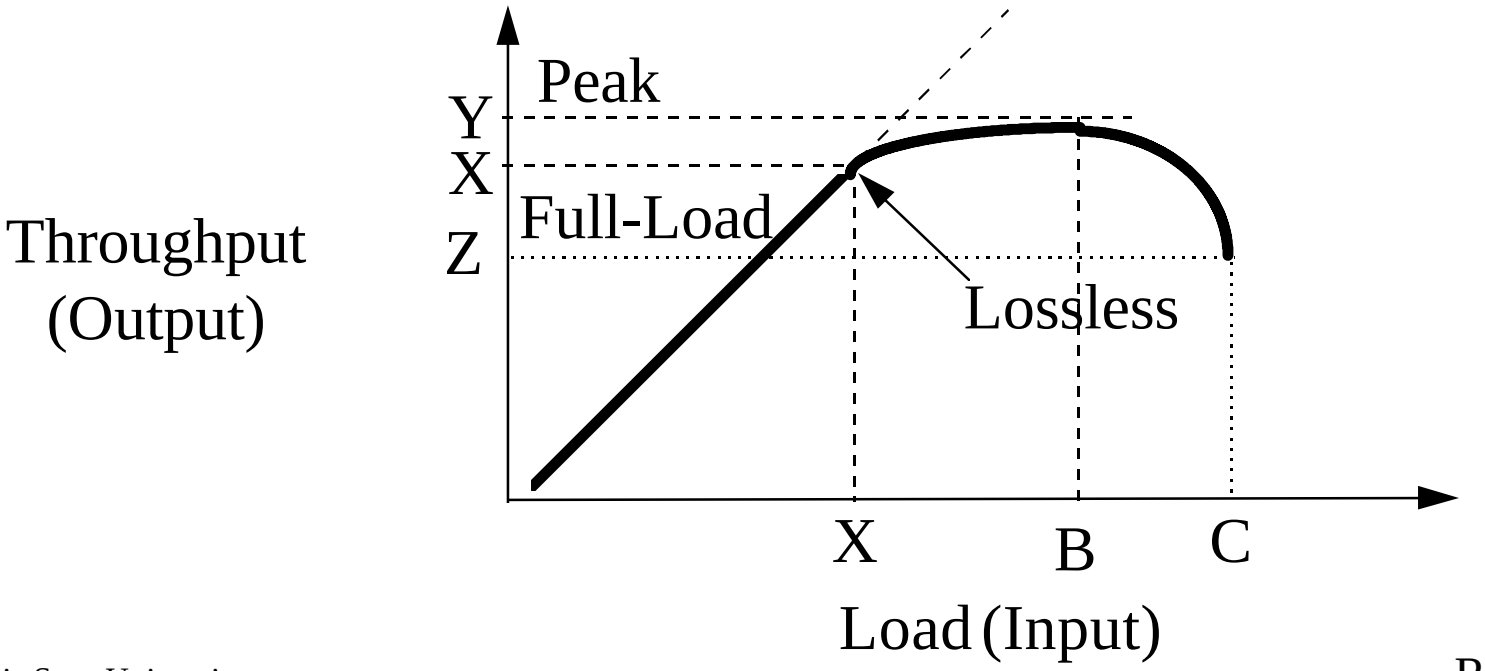

Statistical Variations

- Repeat NRT times for TRT seconds each
  Default NRT = 30, TRT = 60 seconds
- Sample = $\{T_1, T_2, T_3, ..., T_n\}$
- Sample mean $T = (1/n)\Sigma T_\iota$
- Sample Standard Deviation $\sigma_T = (\Sigma(T\iota - T)^2)/(n-1)$
- Standard Error = Standard Deviation/ $\sqrt{n}$
- $100(1-\alpha)\%$ confidence interval
  $= (T - z_{[1-\alpha/2]}\text{Std\_Err} , T - z_{[1+\alpha/2]}\text{Std\_Err})$
Confidence	α	Z
90%	0.1	1.615
99%	0.01	2.346
99.9%	0.001	3.291

Traffic Pattern

- n-to-n Straight: n Vcs
  i to i+1 mod n

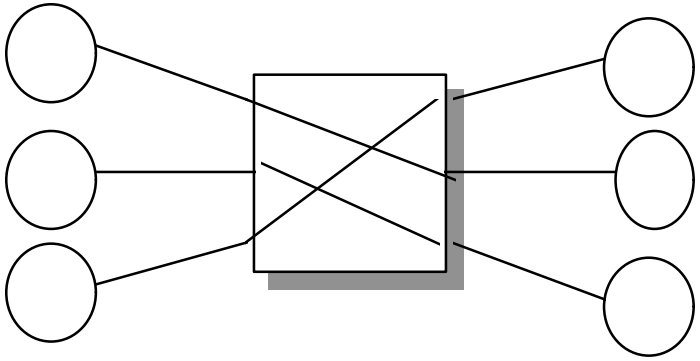

- n-to-n Cross: n^2 Vcs

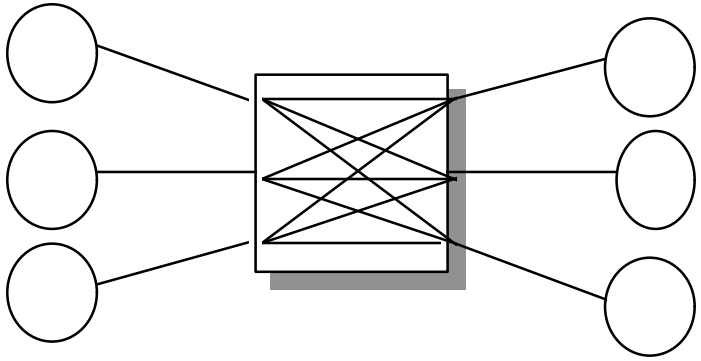

- n-to-1: n Vcs

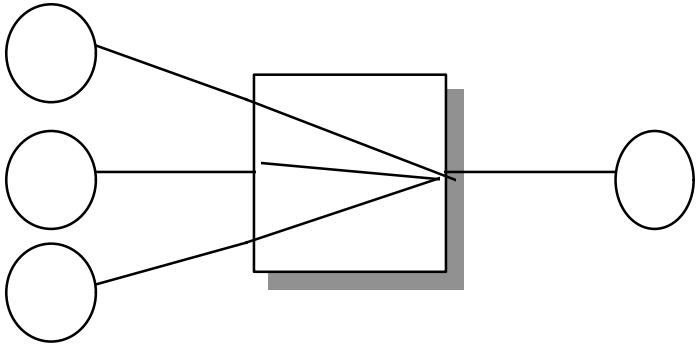

- 1-to-n Straight: 1 Vc

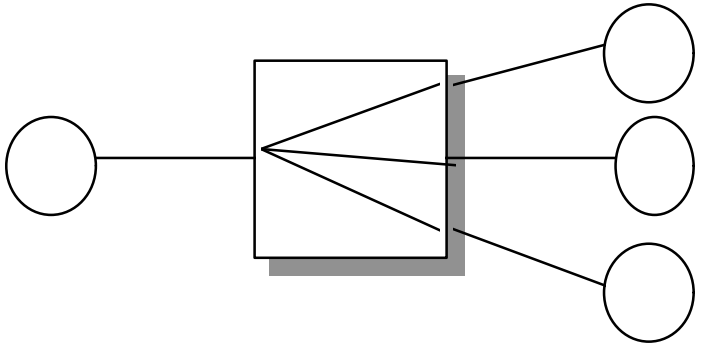

Background Traffic

- With and without background traffic
- To be defined
- Without background traffic until then

Frame Latency

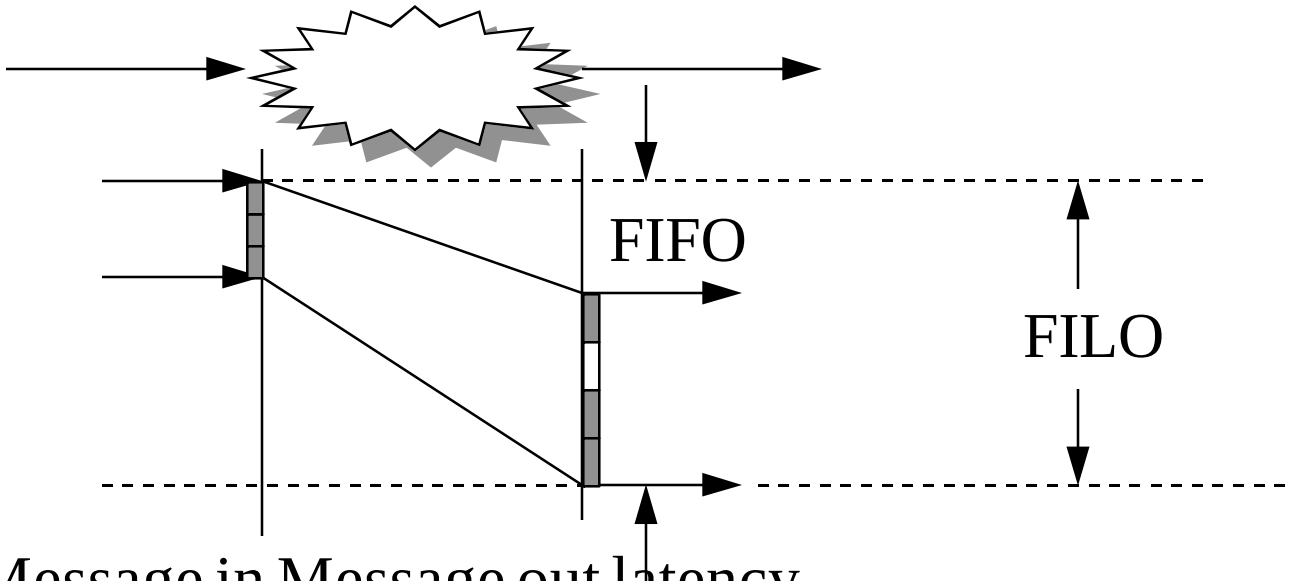

- MIMO = Message in Message out latency
  = min{LILO, FILO-Normalized Frame Output Time}
- NFOT = Frame size/output link rate
- Applies even when: Input rate <> Output rate
  Even when frames are not contiguous
- Unit: μs

Statistical Variations

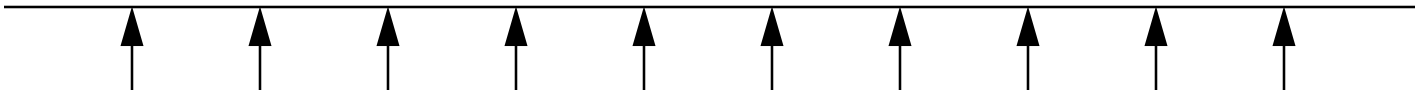

- Send NML cells at TTL/(NML + 1) intervals
- NML = number of marked cells for latency measurement
- TTL = Total Time for Latency measurement
- Default: NML = 30 TTL = 31 seconds
- Calculate mean and standard error
  (Same way as for throughput)

Background Traffic

- With and without background traffic
- Background traffic to be specified
- Without background traffic until then

Reporting Results

	Throughput							
Traffic	Lossless		Peak		Full-load		Latency	
Pattern	Mean	Std Err	Mean	Std Err	Mean	Std Err	Mean	Std Err
n-to-n Straight								
n-to-n Cross								
n-to-1								
1-to-n								

Default Parameter Values

Parameter	Meaning	Default
NRT	Number of repetitions of throughput experiments	30
TRT	Time of each repetition of throughput experiment	60 sec
FSA	Frame Size for AAL performance experiments	9188 Bytes
NML	Number of marked frames sent in latency experiments	30
TTL	Total time of latency experiments	31 sec

Summary

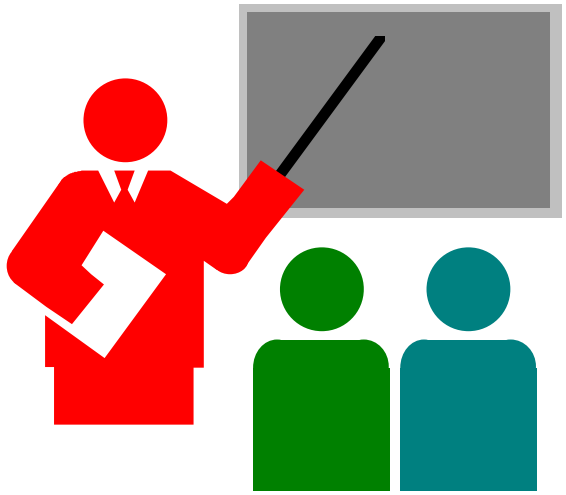

- Throughput: Lossless, peak, full-load
- Latency = Min{LILO, FILO- NFOT} = MIMO

Motion

- Include the text of 96-08011 in the baseline draft.

Latency

- For a single bit:
  Total latency = Bit in to bit out = Switch latency

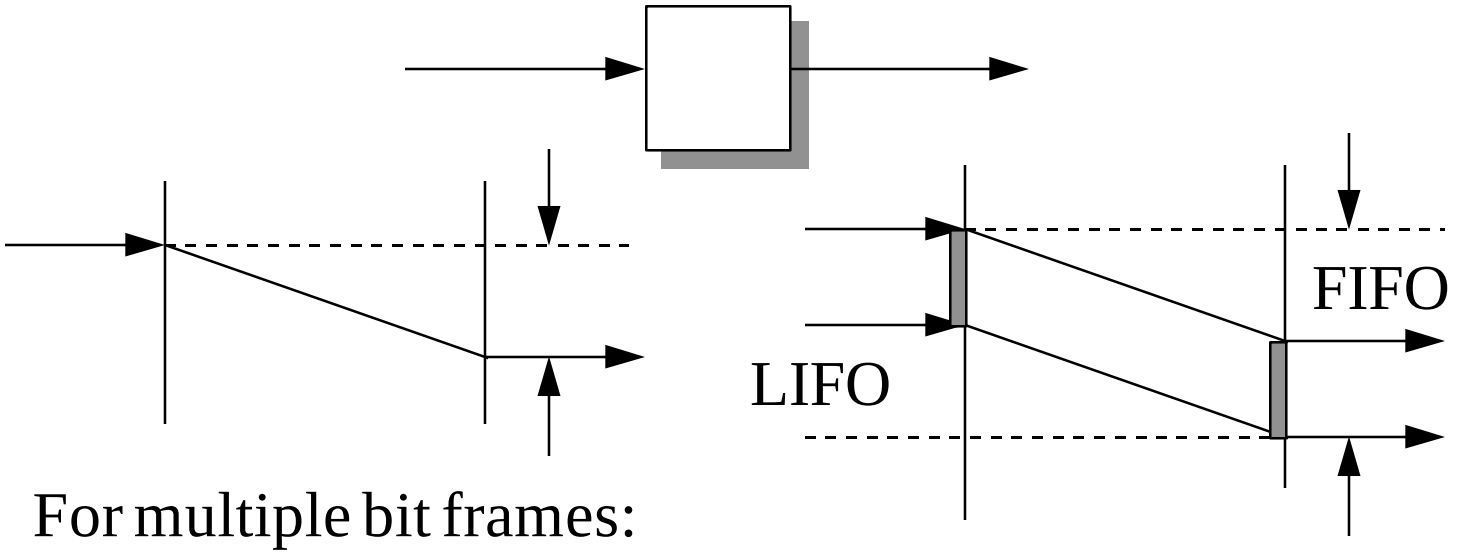

- For multiple bit frames:
  - FIFO = First bit In to First bit Out
  - LILO = Last bit In to Last bit Out
  - FILO = First bit In to Last bit Out
  - LIFO = Last bit In to First bit Out

Latency: Multiple Bit Frames

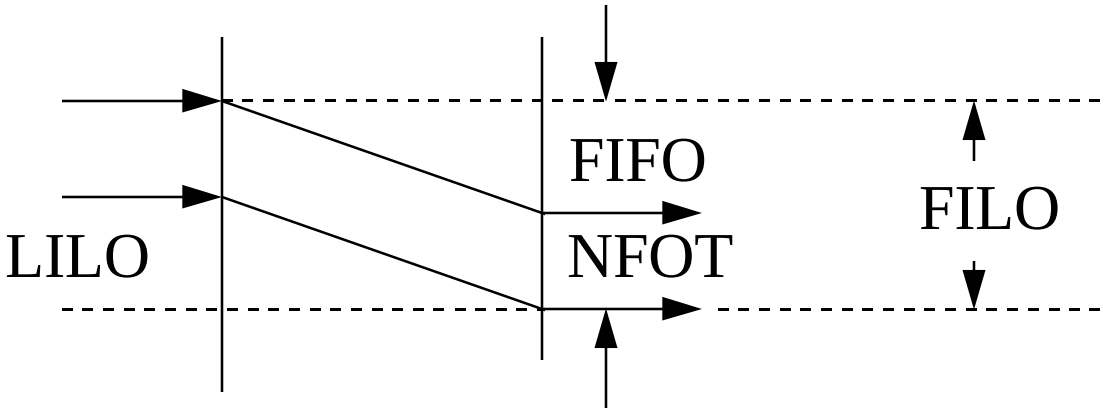

- FIFO = LILO
- FILO = FIFO + Frame time
- LIFO = LILO - Frame time
- Nominal Frame output time NFOT = Frame size/output speed
- Total Delay = FILO = Switch latency + Frame time
- Switch Latency = FILO - NFOT= FIFO = LILO

Multiple Bit Frames (Cont)

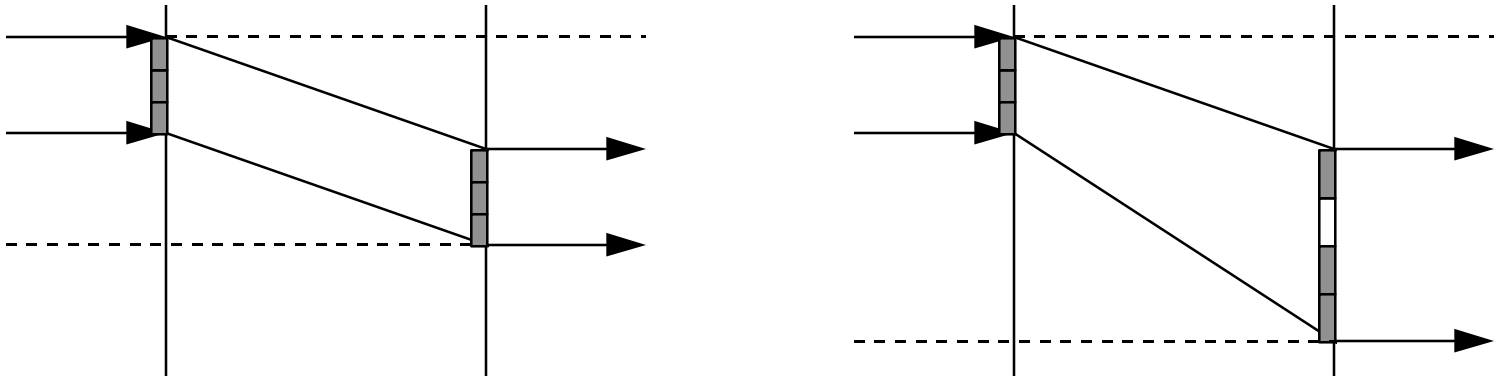

- Switch Latency = FILO - Frame Time = FIFO = LILO
- This assumes contiguous frames ⇒ No idle cells intermingled
- Also assumes input and output lines are of same speed.
- FIFO does not reflect the degradation caused by gaps
- LILO does not reflect the degradation caused by output speed.
- FILO- NFOT is similarly incorrect if input < output speed
- MIMO = Min{FILO- NFOT, LILO} is the correct measure.

Latency: Comparison

No.	Case	FIFO	LILO	FILO-NFOT	MIMO
1a	Input = output, contiguous frame, zero-delay switch	√	√	√	√
1b	Input = output, contiguous frame, nonzero-delay switch	√	√	√	√
1c	Input = output, non-contiguous frame, zero-delay switch	This case is not possible			
1d	Input = output, non-contiguous frame, nonzero-delay switch	×	√	√	√
2a	Input > output, contiguous frame, zero-delay switch	√	×	√	√
2b	Input > output, contiguous frame, nonzero-delay switch	√	×	√	√

Latency: Comparison (Cont)

No.	Case	FIFO	LILO	FILO-NFOT	MIMO
2c	Input > output, non-contiguous frame, zero-delay switch	This case is not possible			
2d	Input > output, non-contiguous frame, nonzero-delay switch	×	×	√	√
3a	Input < output, contiguous frame, zero-delay switch	×	√	×	√
3b	Input < output, contiguous frame, nonzero-delay switch	×	√	×	√
3c	Input < output, non-contiguous frame, zero-delay switch	×	√	×	√
3d	Input < output, non-contiguous frame, nonzero-delay switch	×	√	×	√

Case 1a: Input = Output Speed, Contiguous Frame, Zero-delay Switch

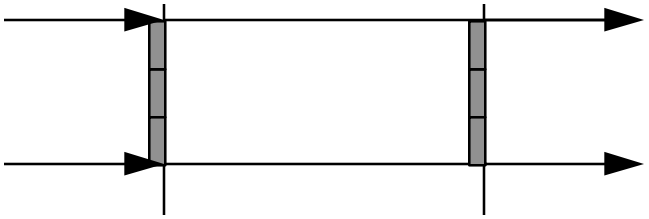

- Zero-delay switch = wire (or a cut-through switch)
- Switch Latency = FILO - NFOT= FIFO = LILO = MIMO =0
- All four alternatives give correct answer

Case 1b: Input = Output Speed, Contiguous Frame, Nonzero Delay Switch

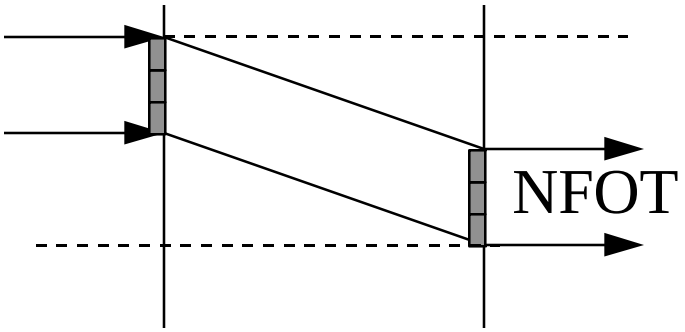

- Example: A very long wire
- Switch Latency = FILO - NFOT= FIFO = LILO = MIMO
- All four metric give the same answer.

Case 1c: Input = Output Speed, Non-Contiguous Frame, Zero Delay Switch

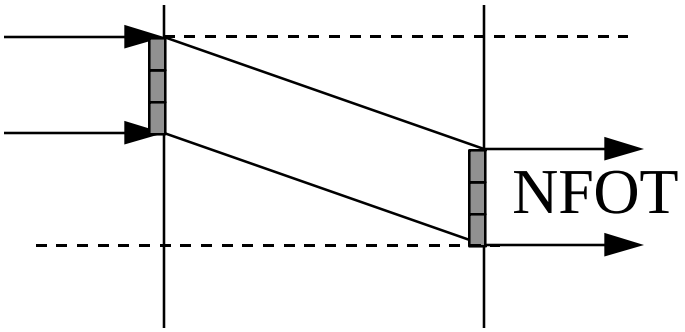

- Switch = A short wire
- All bits exit as soon as the enter.
  $\Rightarrow$ Non-contiguous frame is not possible.

Case 1d: Input = Output Speed, Non-Contiguous Frame, Nonzero Delay Switch

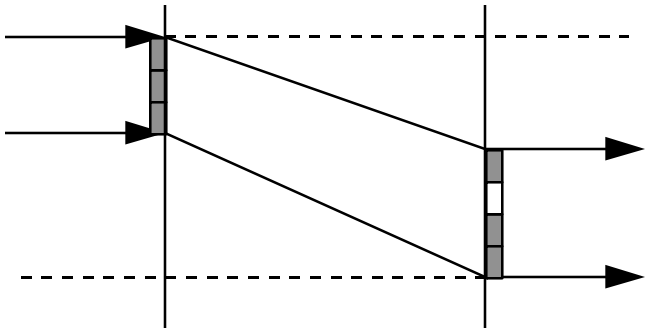

- FIFO does not reflect the degradation caused by gaps. FIFO is not a correct measure of switch latency.
- Switch Latency = total delay - frame time
- FILO - NFOT= LILO = MIMO
- Other three metrics give the same answer and are correct.

Case 2a: Input > Output Speed Contiguous Frame, Zero-Delay Switch

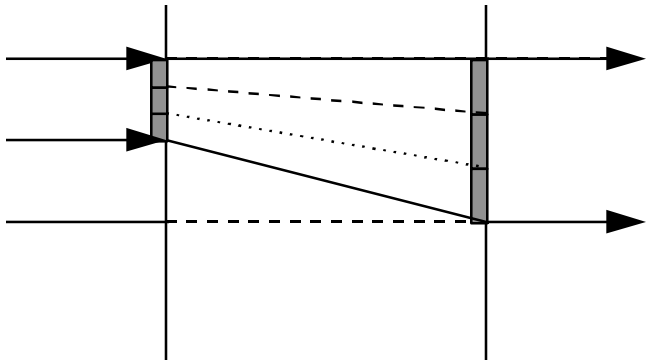

- Zero delay switch = Two port memory
- Total Delay = FILO
  = Switch Latency + Nominal Frame Output Time
- Switch Latency = FILO - NFOT= FIFO =0
- In this case, LILO will give wrong answer.
- LILO > FILO - NFOT $\Rightarrow$ MIMO = FILO- NFOT = 0
- MIMO is also correct

Case 2b: Input > Output Speed Contiguous Frame, Nonzero-Delay Switch

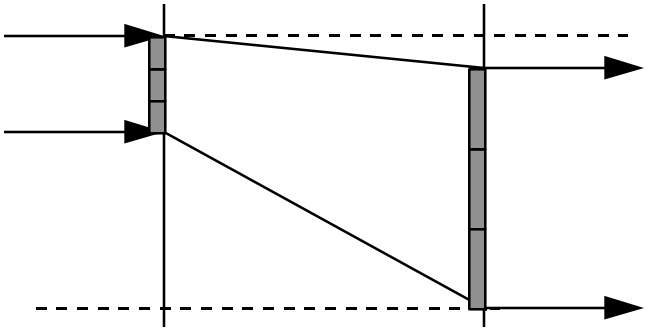

- Total Delay = FILO = Switch Latency + Frame Time
- Switch Latency = FILO - NFOT= FIFO
- In this case, LILO is affected by the output speed and gives wrong answer.
- LILO > FILO - NFOT $\Rightarrow$ MIMO = FILO - NFOT
- MIMO is also correct

Case 2c: Input > Output Speed
Non-contiguous Frame , Zero-delay Switch

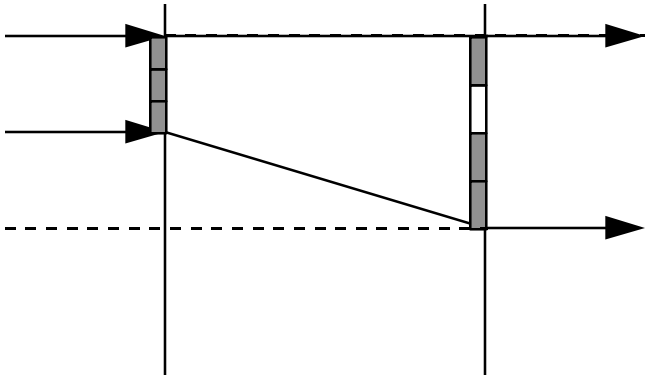

- Zero-delay switch = Two port memory
- A zero-delay switch will not introduce any gaps
  ⇒ This case is not possible.

Case 2d: Input > Output Speed, Non-contiguous Frame, Nonzero delay Switch

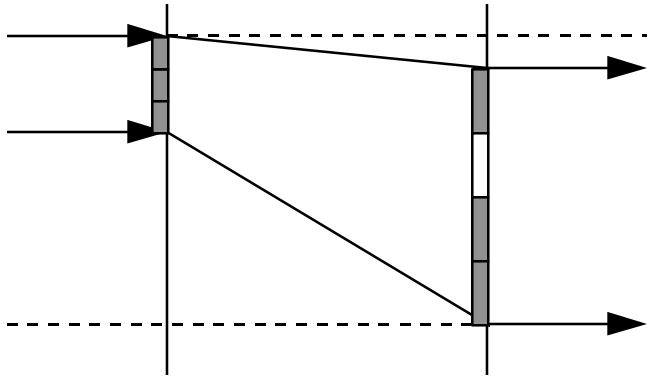

- Switch Latency = total delay - frame time = FILO - NFOT
- FIFO does not reflect gaps ⇒ FIFO is wrong
- LILO is affected by output speed and is wrong.
- LILO > FILO - NFOT
  ⇒ MIMO = min{LILO, FILO- NFOT} = FILO- NFOT
  ⇒ MIMO is also correct.

Case 3a: Input < Output Speed, Contiguous Frame, Zero-delay Switch

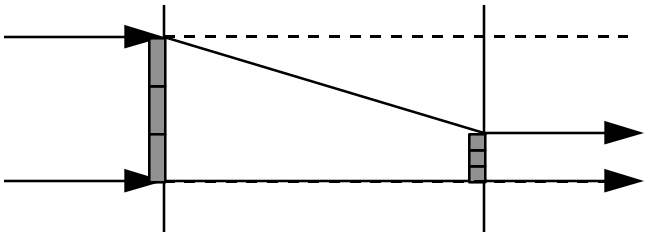

- Zero-delay Switch = 2-port memory
- Contiguous frames are possible only if the transmission of the first bit is timed based on frame size.
- FIFO is non-zero ⇒ Not a correct measure of switch latency
- FILO - NFOT = FIFO = non-zero. Also incorrect.
- LILO is zero. So it is correct.
- MIMO = Min{LILO, FILO - NFOT} is zero. It is also correct.

Case 3b: Input < Output Speed, Contiguous Frame, Nonzero-delay Switch

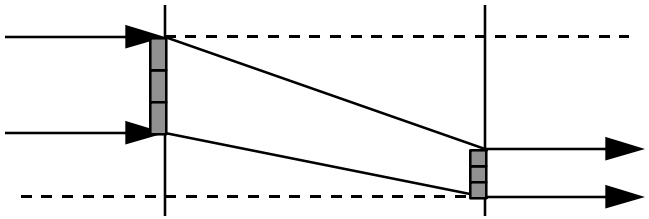

- To maintain frame contiguity, the departure of first bit has to be scheduled depending upon the output speed.
- FIFO can be made arbitrarily large by increasing the output link speed (and not changing the switch at all).
  FILO - NFOT= FIFO is similarly incorrect.
- LILO is the only metric that can be argued to be correct.
- LILO < FILO - NFOT
  MIMO = Min{LILO, FILO - NFOT} = LILO
  MIMO is also a correct measure.

Case 3c: Input < Output Speed, Non-Contiguous Frame, Zero-delay Switch

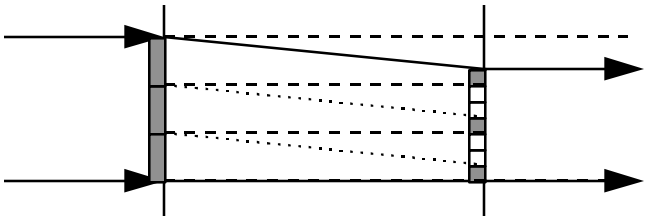

- Zero-delay switch = two port memory
- To obtain contiguous cell, the first bit of the cell is sent such that the last bit can be sent immediately upon arrival.
- FIFO is non-zero. So it is incorrect.
- FILO - NFOT is non-zero. So it is incorrect.
- Gaps are caused by speed difference. Not attributable to switch.
- LILO is zero. So it is correct
- MIMO = Min{LILO, FILO - NFOT} is zero.
  So it is also correct.

Case 3d: Input < Output Speed, Non-Contiguous Frame, Nonzero-delay Switch

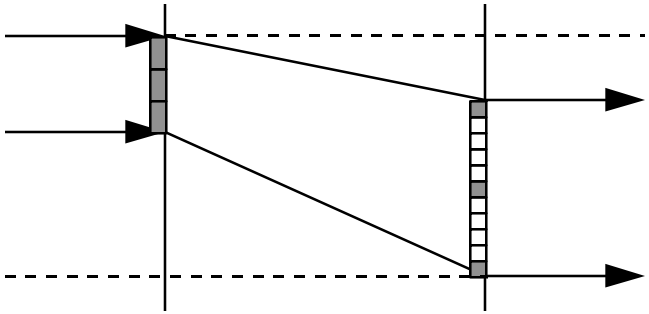

- To maintain cell contiguity, first bit transmission time depends upon the output speed. FIFO can be made arbitrarily large by increasing the output link speed (and not changing the switch).
- FIFO can also be made small by sending the first cell fast but introducing idle cells later $\Rightarrow$ FIFO is not correct.
- FILO - NFOT > FIFO is similarly incorrect.
- LILO is the only metric that can be argued to be correct.
- LILO < FILO - NFOT $\Rightarrow$ MIMO = LILO