96-1268: MIMO Latency - Revised Definition

Raj Jain, Gojko Babic

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

Edu

Jain/

Overview

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

Overview: Cases To Consider

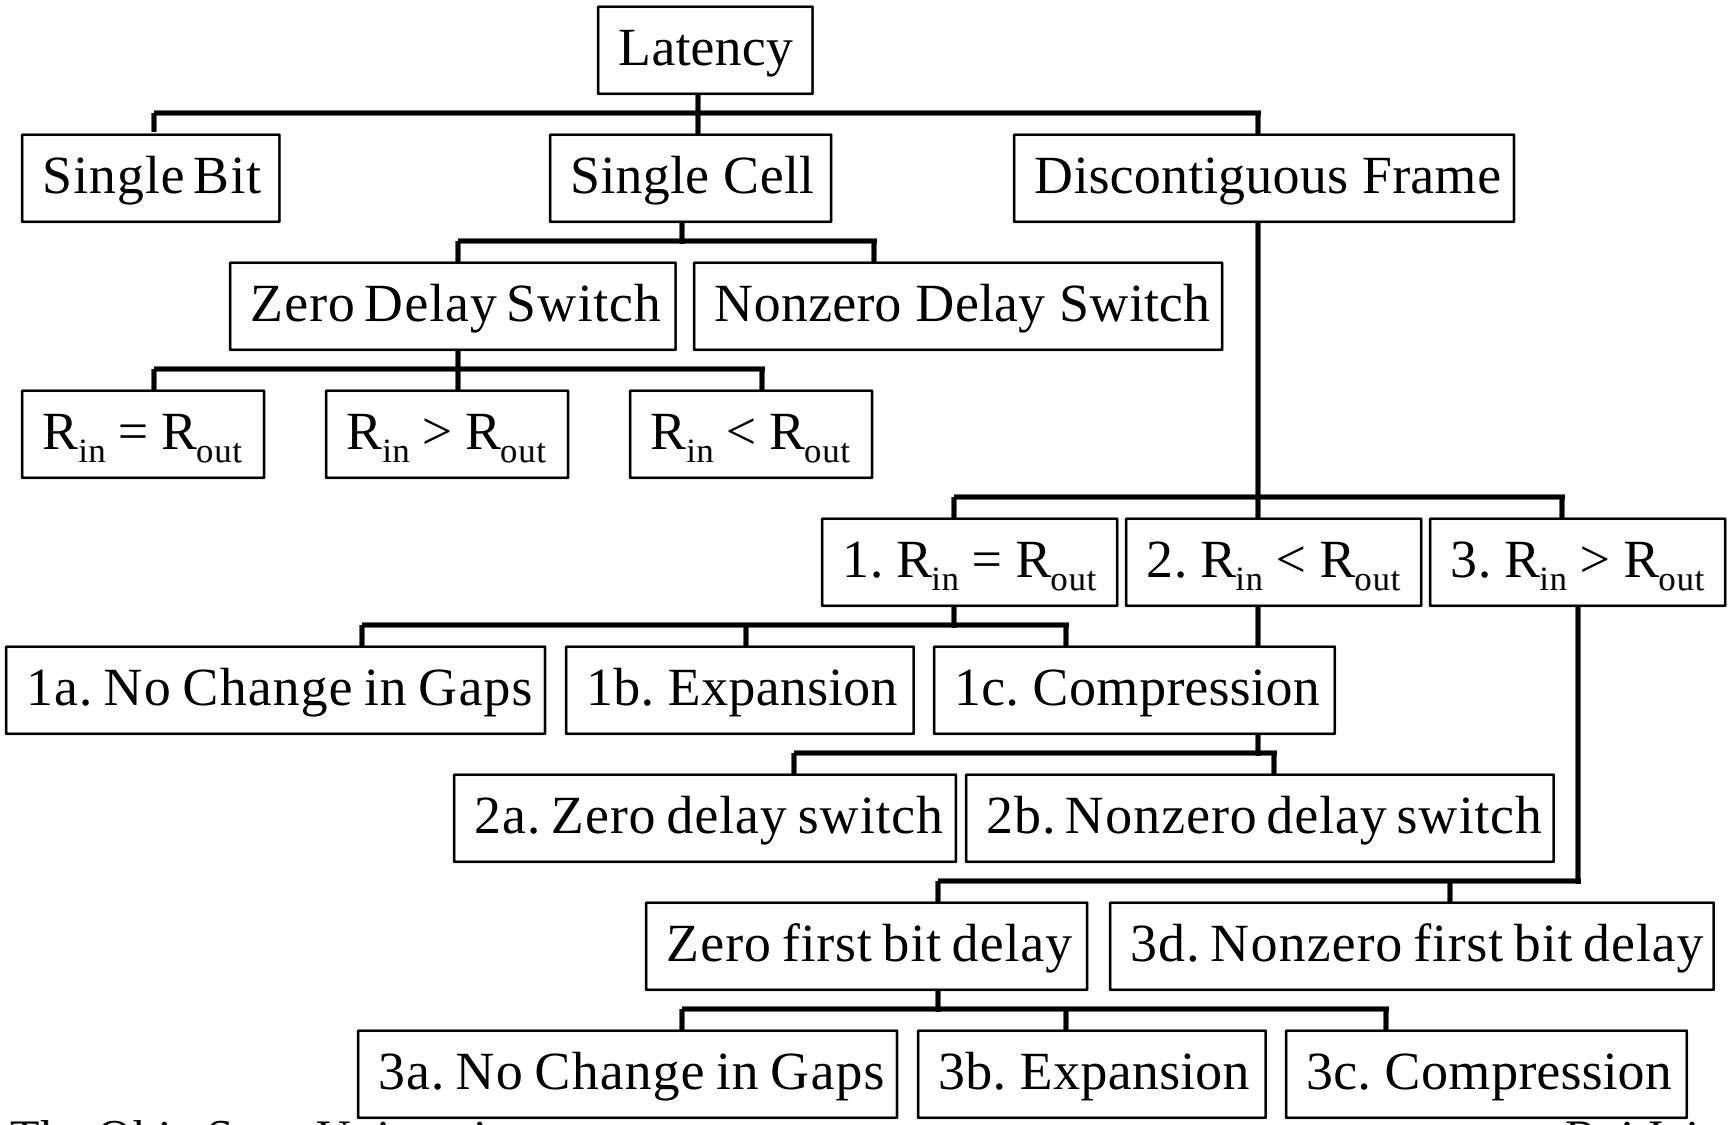

Latency: Single Bit

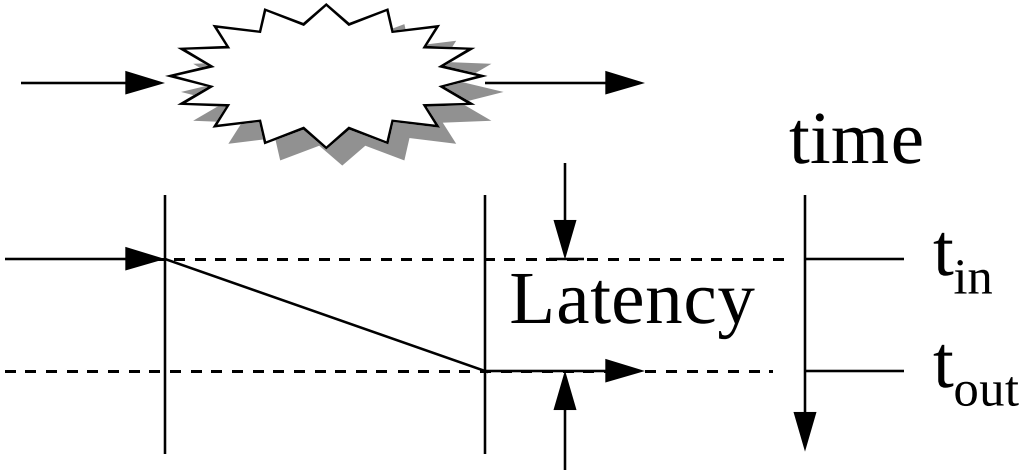

- Latency = $t_{out} - t_{in}$

Latency: Single Cell

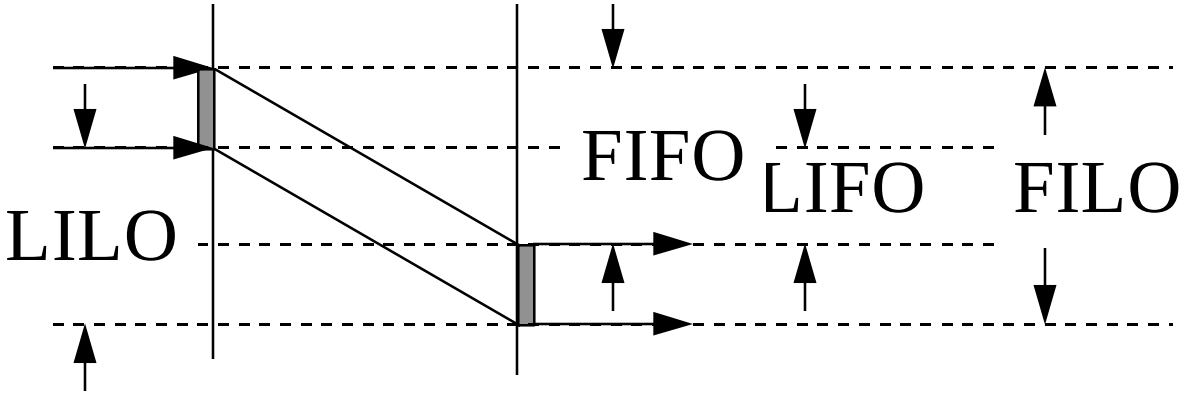

- FIFO = First-bit in to first-bit out
- LIFO = Last-bit in to last-bit out
- FILO = First-bit in to last-bit out
- LILO = Last-bit in to last-bit out
- FIFO = LILO = FILO - Cell time = LIFO + Cell time
- Assumes input speed = output speed

Measurement Points

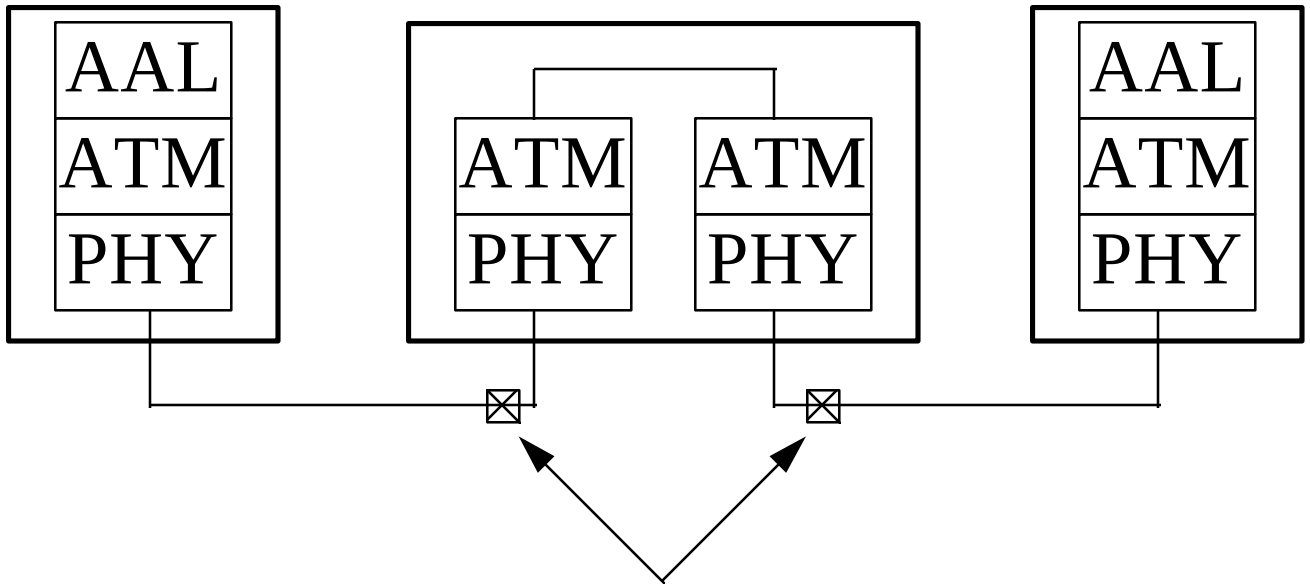

- Host speed should not affect the measured switch performance
- Delay caused by input/output link speeds should not be attributed to switch latency.

Latency: Single Cell, Input Speed ≠ Output Speed

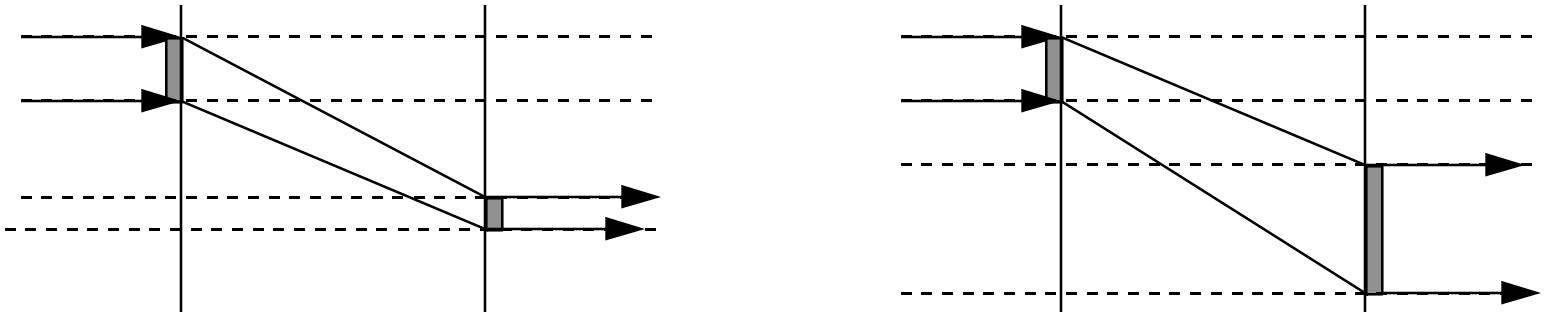

(a) Input speed < Output speed (b) Input speed > Output speed

- FIFO = FILO - Cell output time
- LILO = FILO - Cell input time
- LIFO = FILO - (Cell input time + Cell output time)
- We will show that MIMO is the correct switch latency
  MIMO = Min{LILO, FILO - Cell Input Time*R_{in}/R_{out}}

Simple Switches

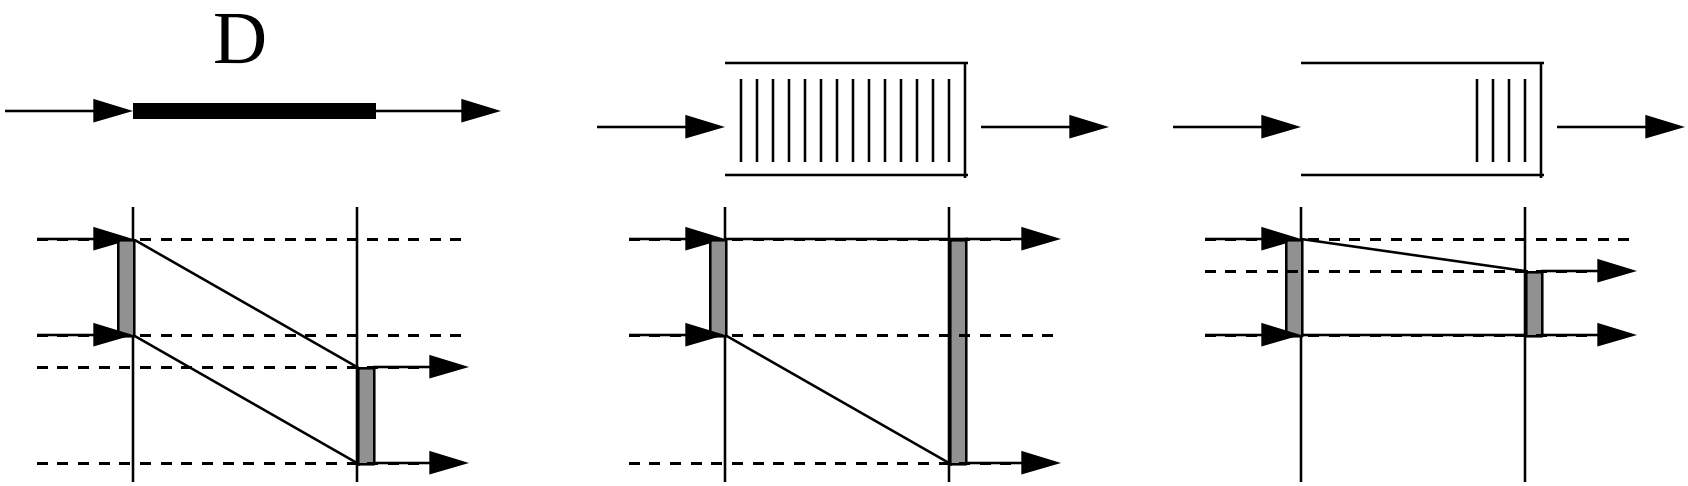

(a) Wire: $R_{in} = R_{out}$ (b) FIFO: $R_{in} > R_{out}$ (c) Intelligent FIFO: $R_{in} < R_{out}$

- 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

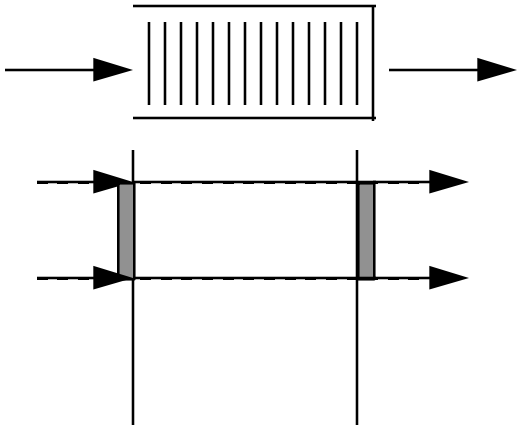

- FIFO is zero $\Rightarrow$ Correct
- LILO is zero $\Rightarrow$ Correct
- MIMO = Min{LILO, FILO - Cell input time*R_{in}/R_{out}}
  = 0 $\Rightarrow$ Correct

Zero-Delay Switch
Input Speed > Output Speed

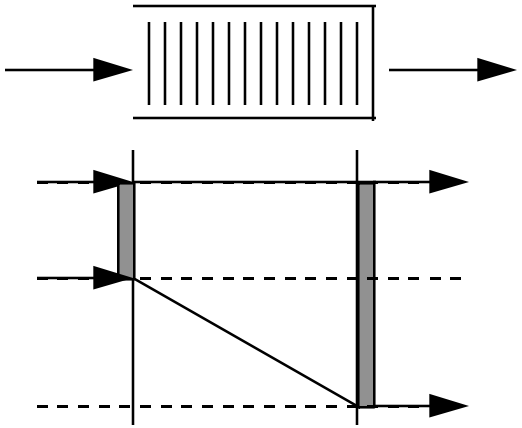

- FIFO is zero $\Rightarrow$ Correct
- LILO is non-zero $\Rightarrow$ Incorrect
- MIMO = Min{LILO, FILO - Cell input time*R_{in}/R_{out}}
  = 0 $\Rightarrow$ Correct

Zero-Delay Switch
Input Speed < Output Speed

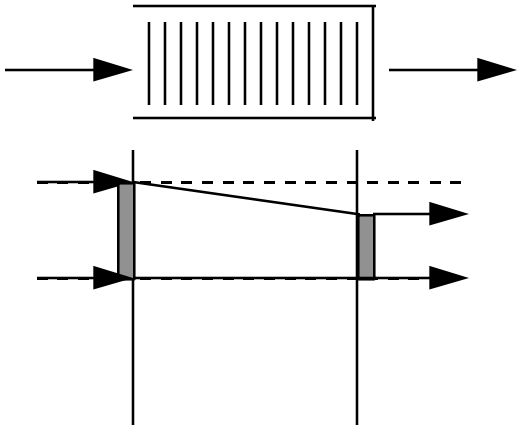

- FIFO is non-zero⇒ Incorrect
- LILO is zero⇒ Correct
- MIMO = Min{LILO, FILO - Cell input time*R_{in}/R_{out}}
  = 0 ⇒ Correct

Non-Zero Delay Switches

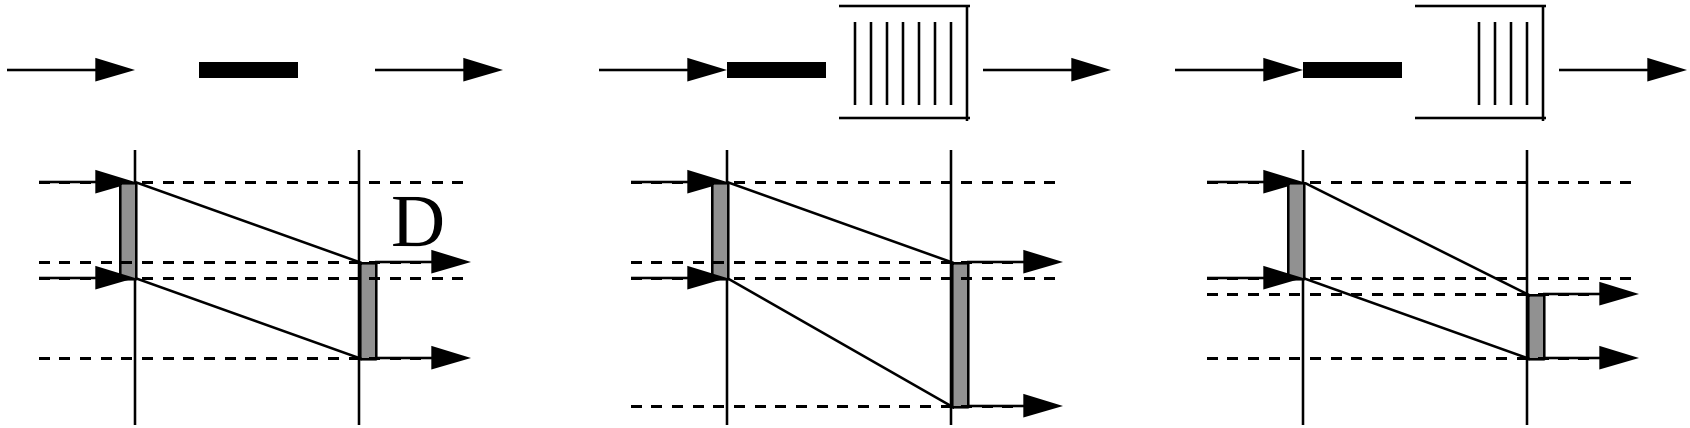

(a) $R_{in} = R_{out}$ (b) $R_{in} \geq R_{out}$ (c) $R_{in} < R_{out}$

(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

Discontiguous Frames

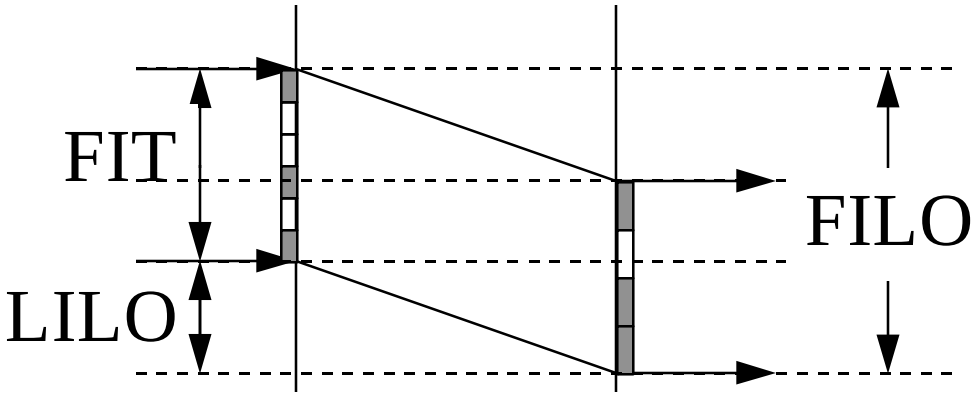

- Frames coming in and out of a switch are not contiguous
- Position and number of gaps may change
- MIMO = Min{LILO, FIT*R_{in}/R_{out}}
  Where, FIT = Frame input time
  = First-bit in to last-bit in

1. Input Speed = Output Speed

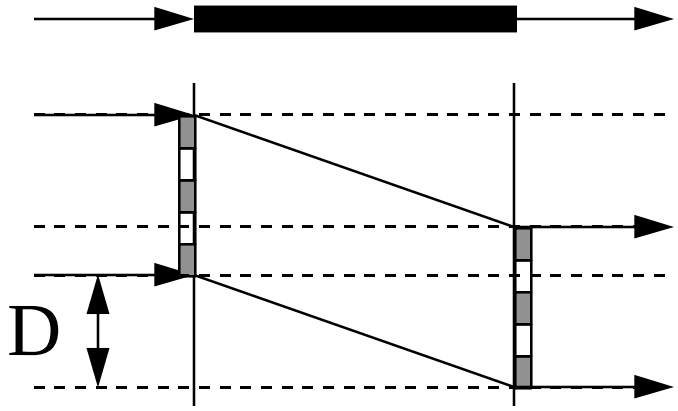

(a) No change

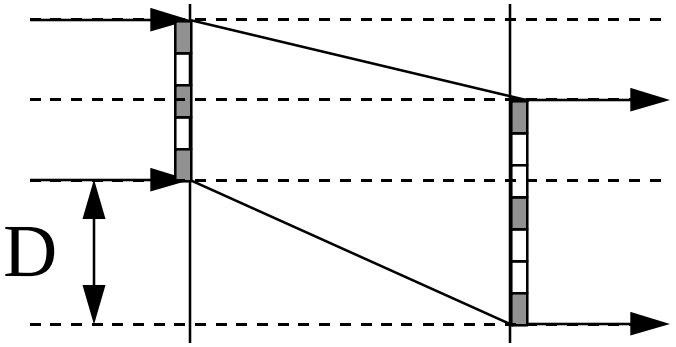

(b) Expansion

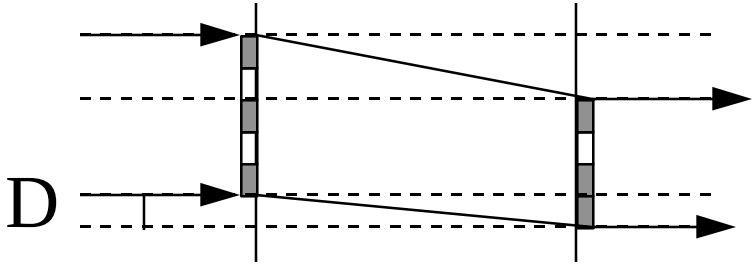

(c) Compression

1a. Input Speed = Output Speed No Change in Gaps

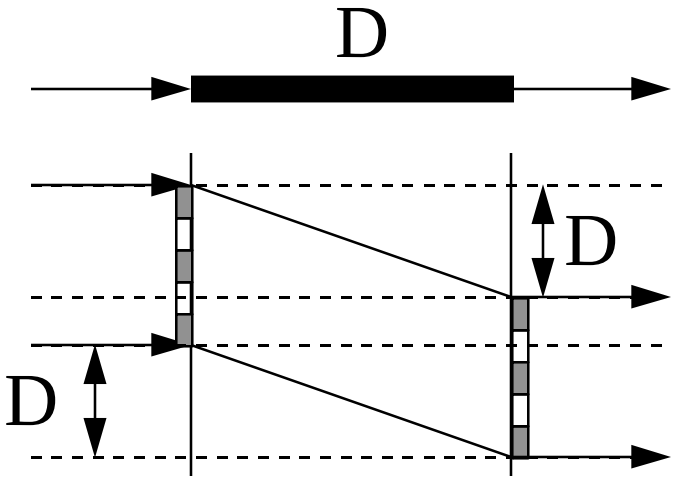

- FIFO = D
- LILO = D
- FILO - FIT = FIFO = D
  $\Rightarrow$ MIMO = Min{LILO, FILO-FIT} = D
- All three metrics are correct

1b. Input Speed = Output Speed Expansion in Gaps

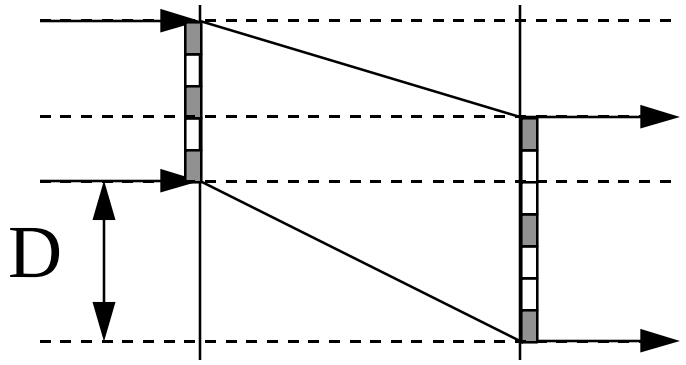

- FIFO does not reflect expansion.
  $\Rightarrow$ FIFO is incorrect.
- LILO = FILO - FIT = FIFO + expansion = D
  $\Rightarrow$ MIMO = Min{LILO, FILO-FIT} = D
- LILO and MIMO are correct.

1c. Input Speed = Output Speed Compression in Gaps

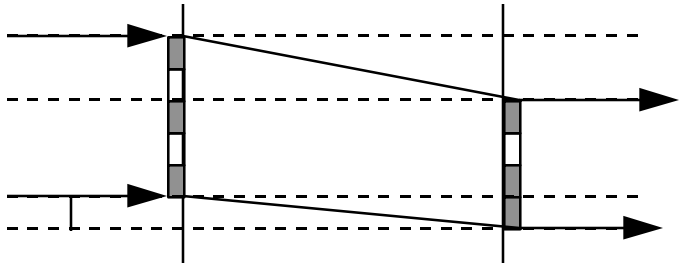

- FIFO does not reflect compression.
  $\Rightarrow$ FIFO is incorrect.
- LILO = FILO - FIT = FIFO - compression = D
  $\Rightarrow$ MIMO = Min{LILO, FILO-FIT} = D
- LILO and MIMO are correct.

Summary: Discontiguous Frames Input Speed = Output Speed

No.	Case	FIFO	LILO	MIMO
1a	No Change in gaps	√	√	√
1b	Expansion	×	√	√
1c	Compression	×	√	√

- MIMO is the only metric that applies to all cases so far.

2a. Input Speed < Output Speed Zero-Delay Switch

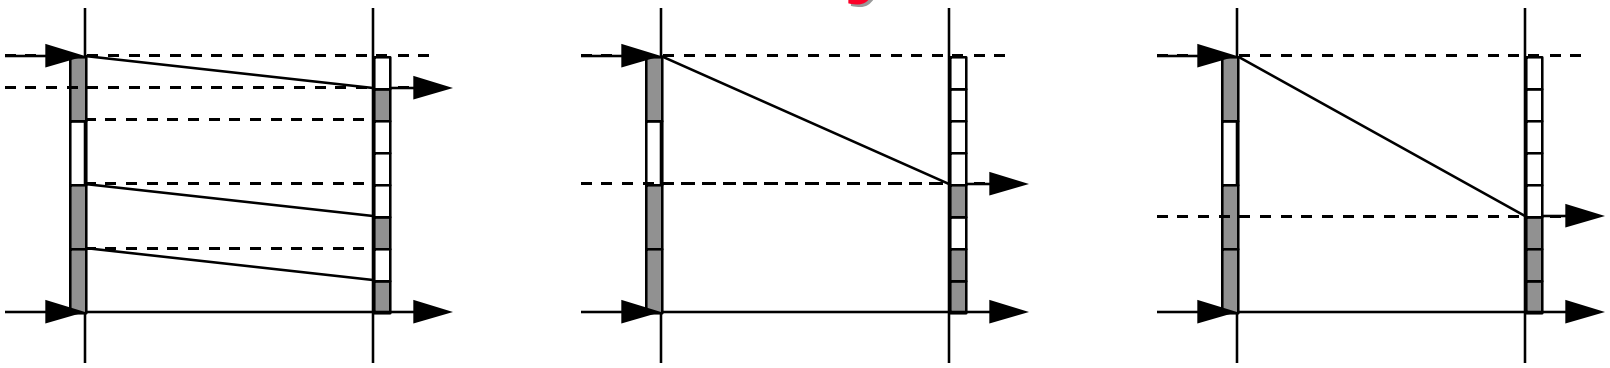

- Need to store bits to avoid underrun
- In the best case, last bits of each cell has no delay
  - FIFO is non-zero $\Rightarrow$ FIFO is incorrect.
  - LILO = 0 $\Rightarrow$ LILO is correct
  - FILO = FIT $\Rightarrow$ FILO - FIT*R_{in}/R_{out} = Positive
    $\Rightarrow$ MIMO = Min{LILO, FILO-FIT *R_{in}/R_{out} } = 0

2b. Input Speed < Output Speed Nonzero-Delay Switch

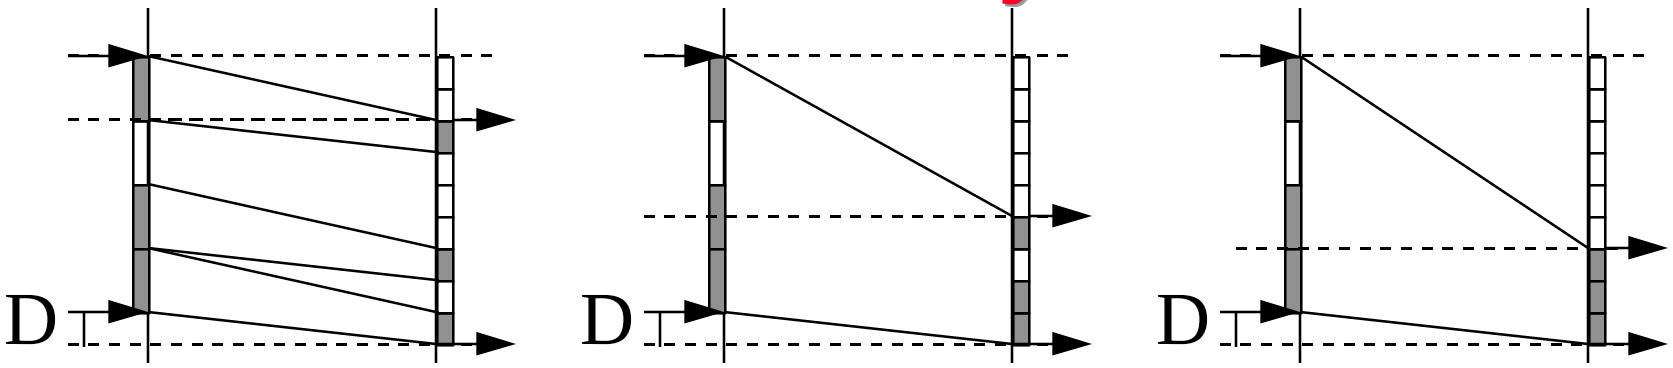

- In each case:
  - FIFO > D $\Rightarrow$ FIFO is incorrect.
  - LILO = D $\Rightarrow$ LILO is correct
  - FILO = FIT + D $\Rightarrow$ FILO - FIT*R_{in}/R_{out} > D
    $\Rightarrow$ MIMO = Min{LILO, FILO-FIT *R_{in}/R_{out} } = D

3a. Input Speed > Output Speed No Change in Gaps

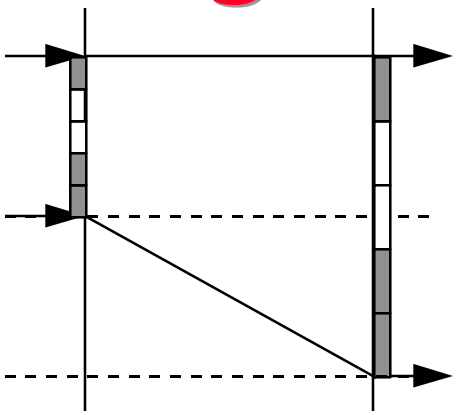

- Switching delay = 0
- FIFO is zero $\Rightarrow$ FIFO is correct.
- LILO is nonzero $\Rightarrow$ LILO is incorrect
- FILO = FIT*R_{in}/R_{out} $\Rightarrow$ FILO - FIT*R_{in}/R_{out} = 0
  $\Rightarrow$ MIMO = Min{LILO, FILO-FIT *R_{in}/R_{out} } = 0
  MIMO is correct

3b. Input Speed > Output Speed Expansion in Gaps

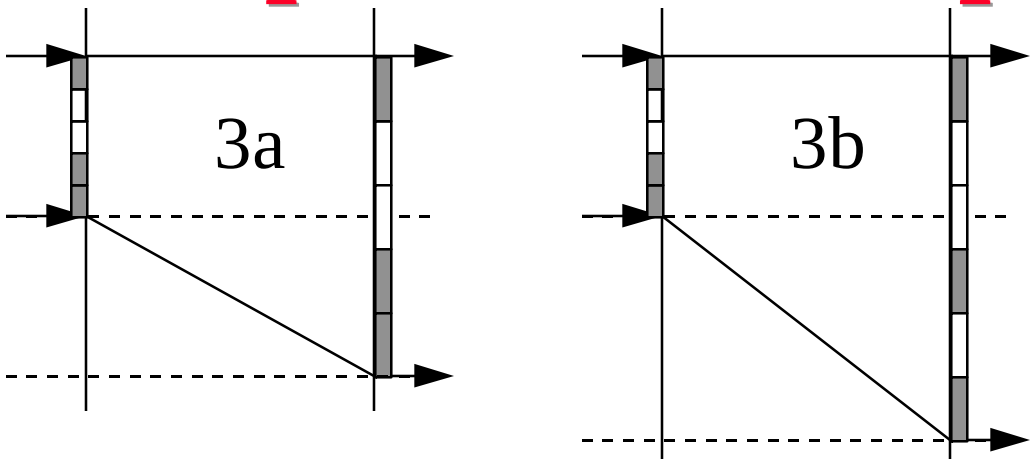

- Worse than 3a. Switching delay = Two cell input time
- FIFO is zero $\Rightarrow$ FIFO is incorrect.
- LILO > Two cell input time $\Rightarrow$ LILO is incorrect
- FILO > FIT*R_{in}/R_{out} $\Rightarrow$ FILO - FIT*R_{in}/R_{out} = 2
  $\Rightarrow$ MIMO = Min{LILO, FILO-FIT *R_{in}/R_{out} } = 2
  MIMO is correct

3c. Input Speed > Output Speed Compression in Gaps

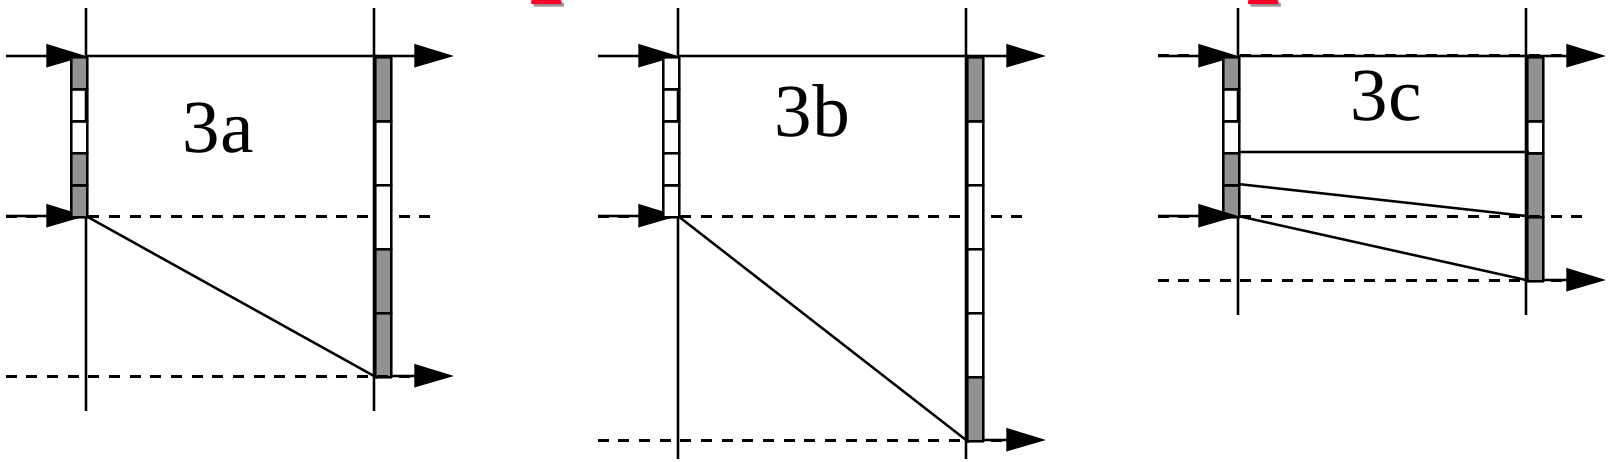

- This case is better than 3a by 3 cells. Switch latency = -3
- FIFO is zero ⇒ FIFO is incorrect.
- LILO = 2 ⇒ LILO is incorrect
- FILO = 7, FIT = 5,
  ⇒ MIMO = Min{LILO, FILO-FIT $*R_{in}/R_{out}$ } = -3
  MIMO is correct

Summary: Discontiguous Frames Input Speed > Output Speed Zero First Bit Delay

No.	Case	FIFO	LILO	MIMO
3a	No Change in gaps	√	×	√
3b	Expansion	×	×	√
3c	Compression	×	×	√

- MIMO is the only metric that applies to all cases so far.

3d. Input Speed > Output Speed Nonzero First Bit Delay

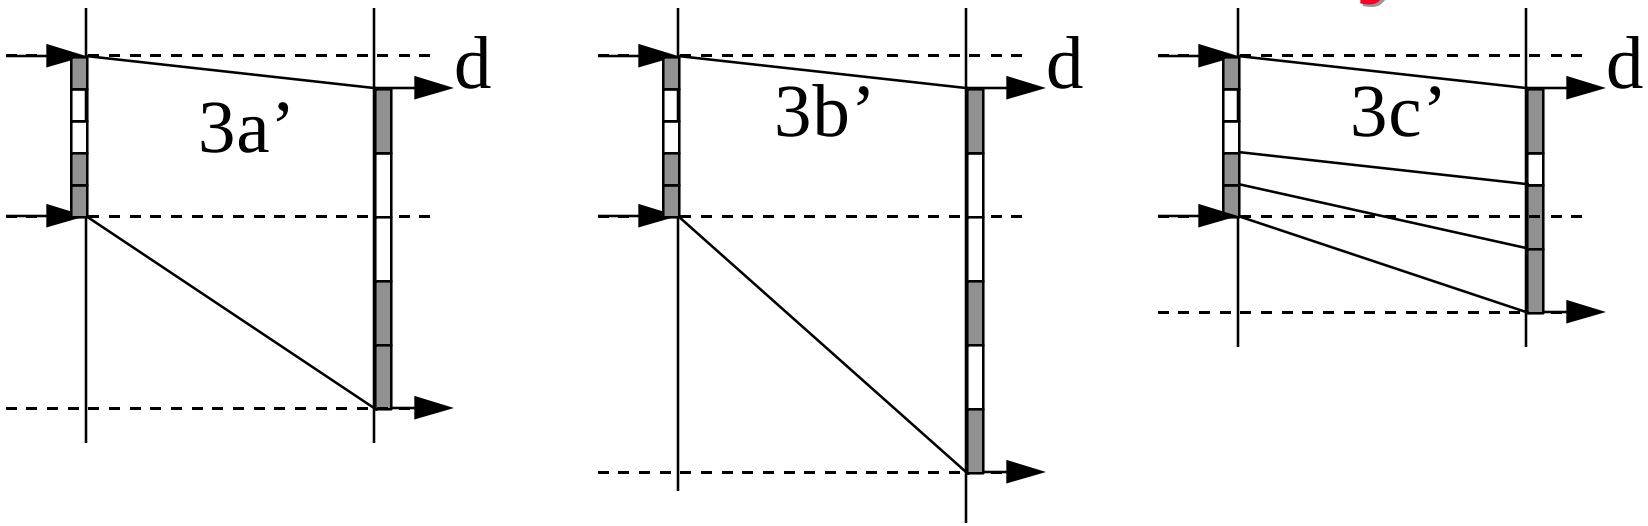

- These cases are similar to 3a, 3b, 3c except that the switch delay has increased by d, where d is the first bit delay.
- It can be shown that the previous summary table applies here also.

Summary: Discontiguous Frames Input Speed > Output Speed

No.	Case	FIFO	LILO	MIMO
3a+3a'	No Change in gaps	√	×	√
3b+3b'	Expansion	×	×	√
3c+3c'	Compression	×	×	√

- MIMO is the only metric that applies to all cases.

User Perceived Delay

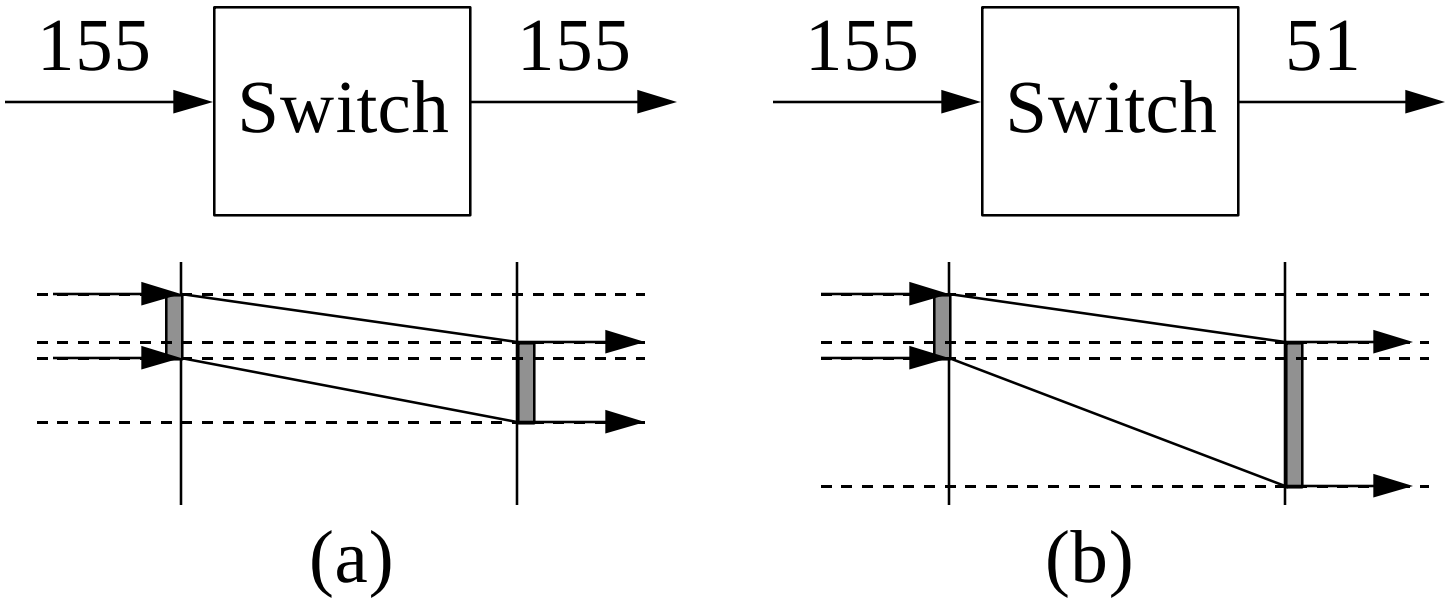

- The switch latency is same in both cases.
- The user perceived performance in case b is worse than that in case a.

User Perceived Performance (Cont)

- The user perceived performance depends upon the link speeds
- The switch latency measured by MIMO is independent of the input/output link speeds.
- User can not begin its work until the last bit has been received and so user perceived performance is reflected by LILO.
- The difference LILO-MIMO is due to link speed differential. This is called buffering delay and is nonzero only if the input has to be buffered (input speed > output speed).

User Perceived Performance (Cont)

- In all other cases, buffering delay is zero (Input speed $\geq$ output speed). In these cases,

  $$\text{MIMO} = \text{LILO}.$$

- In general:
  User Perceived Delay = LILO
  = MIMO + Buffering Delay
- Buffering delay and user perceived delay is never negative.
- Switches are smart and can compress a frame resulting sometimes in negative delay (when input speed < output speed)

Summary

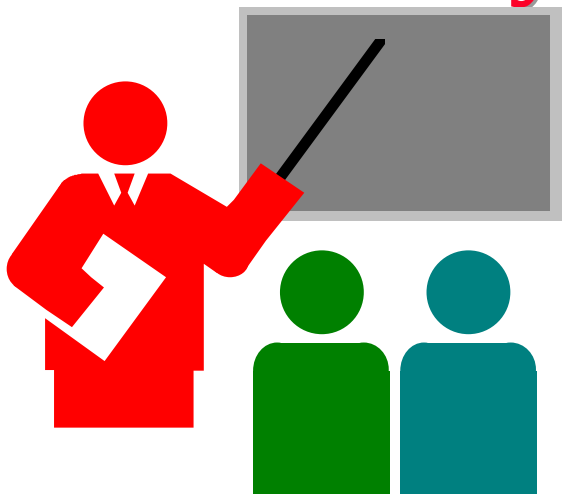

- Switch Latency:
  MIMO = Min{LILO, FILO - FIT*R_{in}/R_{out}}
- Switch latency is not affected by host or link speeds
- User perceived delay = LILO
- Buffering delay = LILO - MIMO