****

ATM Forum Document Number: ATM_Forum/97-0834

****

Title: Modifications to Appendix A of Performance Testing Baseline Text on MIMO Latency

****

Abstract: A new definition of MIMO latency was adopted in the July'97 ATM Forum meeting. The appendix of the baseline text needs to be modified to reflect this change. This contribution proposes the revised text.

****

Source:
Gojko Babic, Raj Jain, Arjan Durresi, Justin Dolske.
The Ohio State University

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

The presentation of this contribution at the ATM Forum is sponsored by NASA.

****

Date: September 1997

****

Distribution: ATM Forum Technical Working Group Members (AF-TEST, AF-TM)

****

****

Appendix A: Defining Frame Latency on ATM Networks

A.1. Introduction

This appendix discusses delays, and the performance metrics characterizing them, that an ATM network introduces to its frames. We are concerned with delays caused by node processing, such as switching and routing, as well as queuing delays that may be introduced by the background traffic and inter-network link transmission delays. On the other hand, transmission delays introduced by input and output links of a network component should not be attributed to the component. Also, note that characteristics of traffic generators (e.g., host speeds) should not affect network performance metrics. The discussion in this Appendix applies to any network element (including switches, multiplexors, inverse-multiplexors, wires) or any combination of such network elements. Although we frequently use the term "switch," the discussion applies equally well to other network elements, whole networks, or parts of networks.

In the case of a single bit, the switch (network) delay is generally defined as the time between the instant the bit enters the system and the instant the bit exits from the system. Figure A.1 illustrates the single-bit latency.

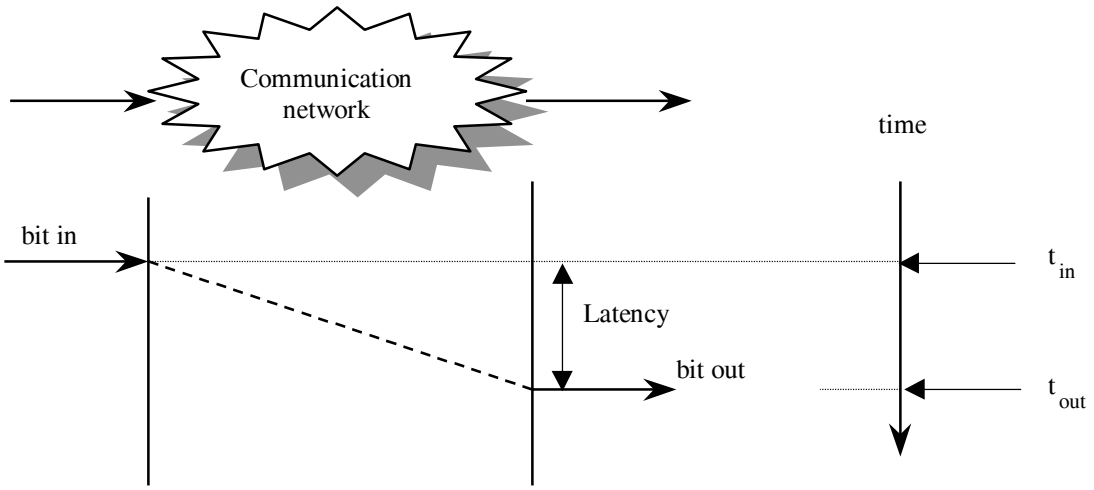

Figure A.1: Latency for a single bit

For multi-bit frames, the usual way to define the frame latency introduced by a switching device is to apply one of the following four definitions:

- FIFO latency: Time between the first-bit entry and the first-bit exit
- LILO latency: Time between the last-bit entry and the last-bit exit
- FILO latency: Time between the first-bit entry and the last-bit exit
- LIFO latency: Time between the last-bit entry and the first-bit exit

Figure A.2 illustrates the usual frame latencies (FIFO, LILO, FILO and LIFO) in a scenario with a contiguous frame on both input and output, passing through the given communication network which has an input link rate lower than the output link rate.

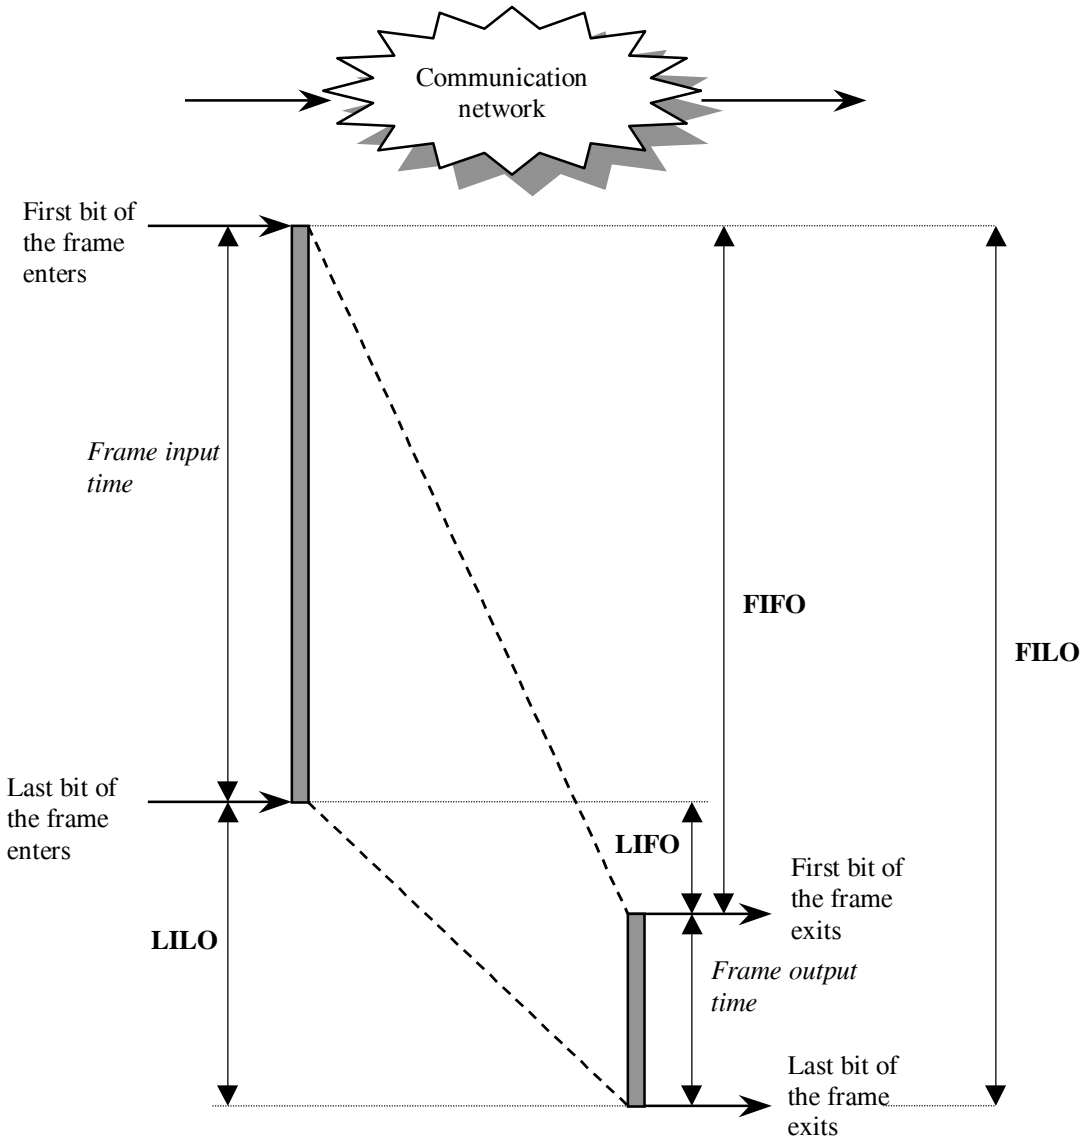

Figure A.2: Usual frame latencies

Unfortunately, as it will be shown later, none of the four above metrics is appropriate for an ATM network. In this appendix, we introduce and justify a new latency metric called "MIMO" latency. This new latency metric applies to any type of network where the frames may be contiguous or discontinuous, although our primary interest is an ATM environment. To define the MIMO latency, we introduce the concept of a "zero-delay" switch, which is in some sense the best a switch can do. The delay of any other switch is defined as the latency over and above the delay of a zero-delay switch.

This appendix is organized as follows. In the next section, we analyze why the usual frame latencies are not appropriate in an ATM environment. We introduce the MIMO latency in Section A.3. In Section A.4, we introduce the concept of a zero-delay switch and its processing of individual cells and contiguous frames. We discuss delays introduced to discontinuous frames passing through a zero-delay switch in Section A.5. Section A.6 presents the method for calculating the FILO latency of frames passing through a zero-delay switch. An equivalent, but easier to use, definition of MIMO latency is developed in Section A.7. Section A.8 of this appendix presents derivations of expressions for MIMO latency calculation based on cell-level data. The last section discusses the user perceived delay in data communication networks.

A.2. Usual Frame Latencies as Metrics for ATM Switch Delay

An ATM switch has to deal with both contiguous and discontinuous frames. This is because ATM switches do cell-switching, i.e., an ATM switch may transmit a received cell of any frame without first waiting for other cells of that frame to arrive. Thus, frames sent and received in an ATM environment are not always contiguous. Even if the input frame is contiguous, the ATM switch may transmit discontinuous frames, i.e., it may introduce idle periods, unassigned cells and/or cells of other frames between cells of the frame.

The above factors make the usual frame latency metrics inappropriate for ATM switches. In this section, we show why LIFO, FIFO and FILO latencies are not appropriate metrics for an ATM switch. Later in this appendix, we shall show that LILO latency is an appropriate metric only in certain cases.

LIFO Latency

In [1], the delay in a packet-switching network is defined as the time between a "packet entry event" and a "packet exit event." A packet entry event is defined to occur at the time when the last bit of the frame enters a network, while a packet exit event is defined to occur when the first bit of the frame exits a network. This is equivalent to LIFO latency, which is considered as an appropriate metric for store-and-forward packet-switching networks because:

- packets are contiguous on both input and output and
- it is accepted that the transmission delay during packet input is an intrinsic delay for a store and forward device, for which the switch should not to be penalized.

Newer networking devices are not necessarily store-and-forward. Some of them are cut-through devices that start emitting the frame before it is received completely. Figure A.3 illustrates the case of a frame passing through a cut-through switching device with three of the four usual latencies indicated. LIFO latency is not shown because the first bit of the frame exits before the last bit of the frame enters and the LIFO latency is negative. This is a common case with cut-through devices. Thus, LIFO latency is not a good indicator of the switch delay for any cut-through type device, and as such it is

inappropriate for an ATM environment, where cut-through forwarding of frames is the normal mode of operation.

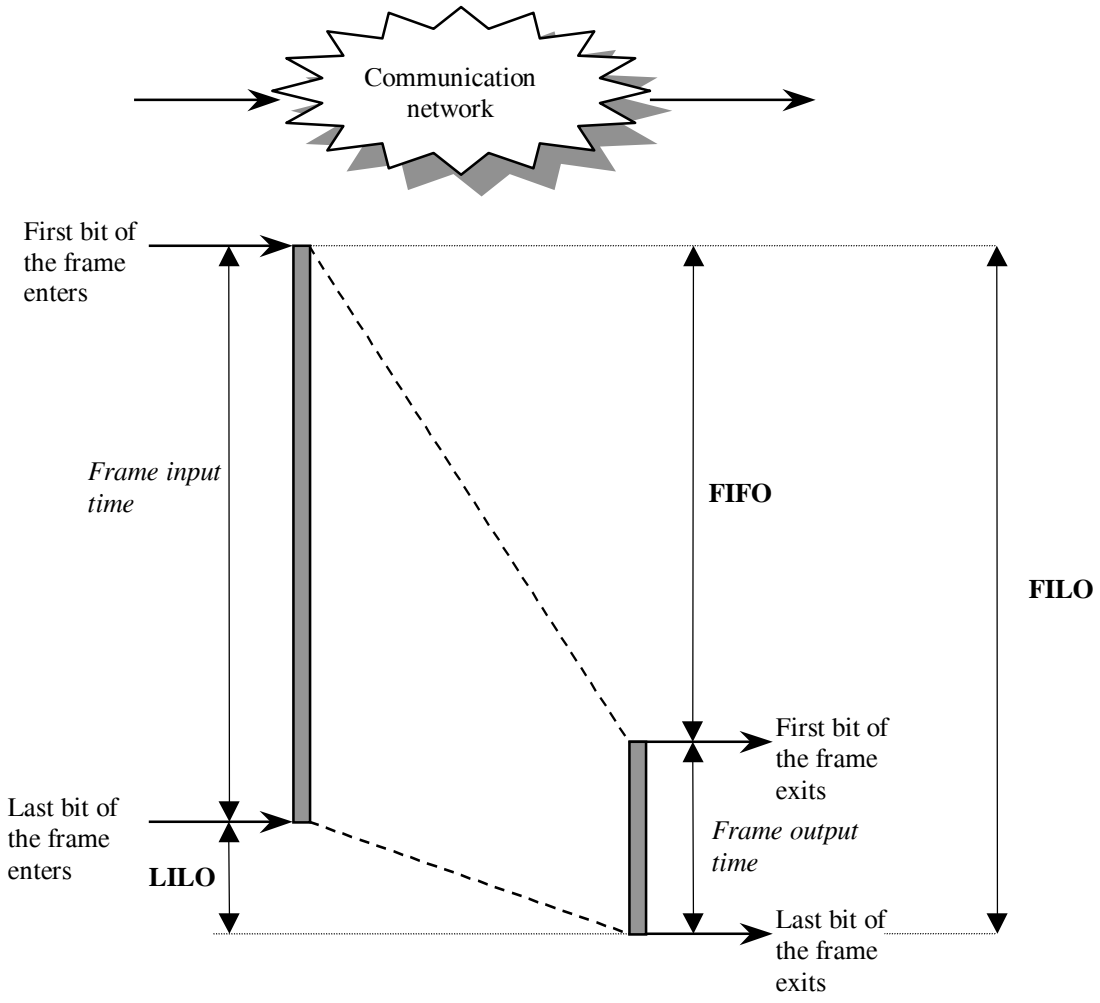

Figure A.3: Latencies of a frame passing through a cut-through switching device

FIFO Latency

It is interesting to note that [2] provides a LIFO latency definition as the delay metric for store and forward switching devices, as well as a FIFO latency definition for bit forwarding devices (i.e. cut-through switching devices). The introduction of FIFO latency as a delay metric is an attempt to avoid negative values for the delay through cut-through devices.

While FIFO latency may provide meaningful results if the frames are continuous, it may provide useless results if the frames are discontinuous. It is possible to have a very low FIFO delay while delays for the other parts of the frames are high. Again, since frames on ATM networks are generally discontinuous, FIFO latency is not a meaningful measure of frame latency. Figure A.4 illustrates this point.

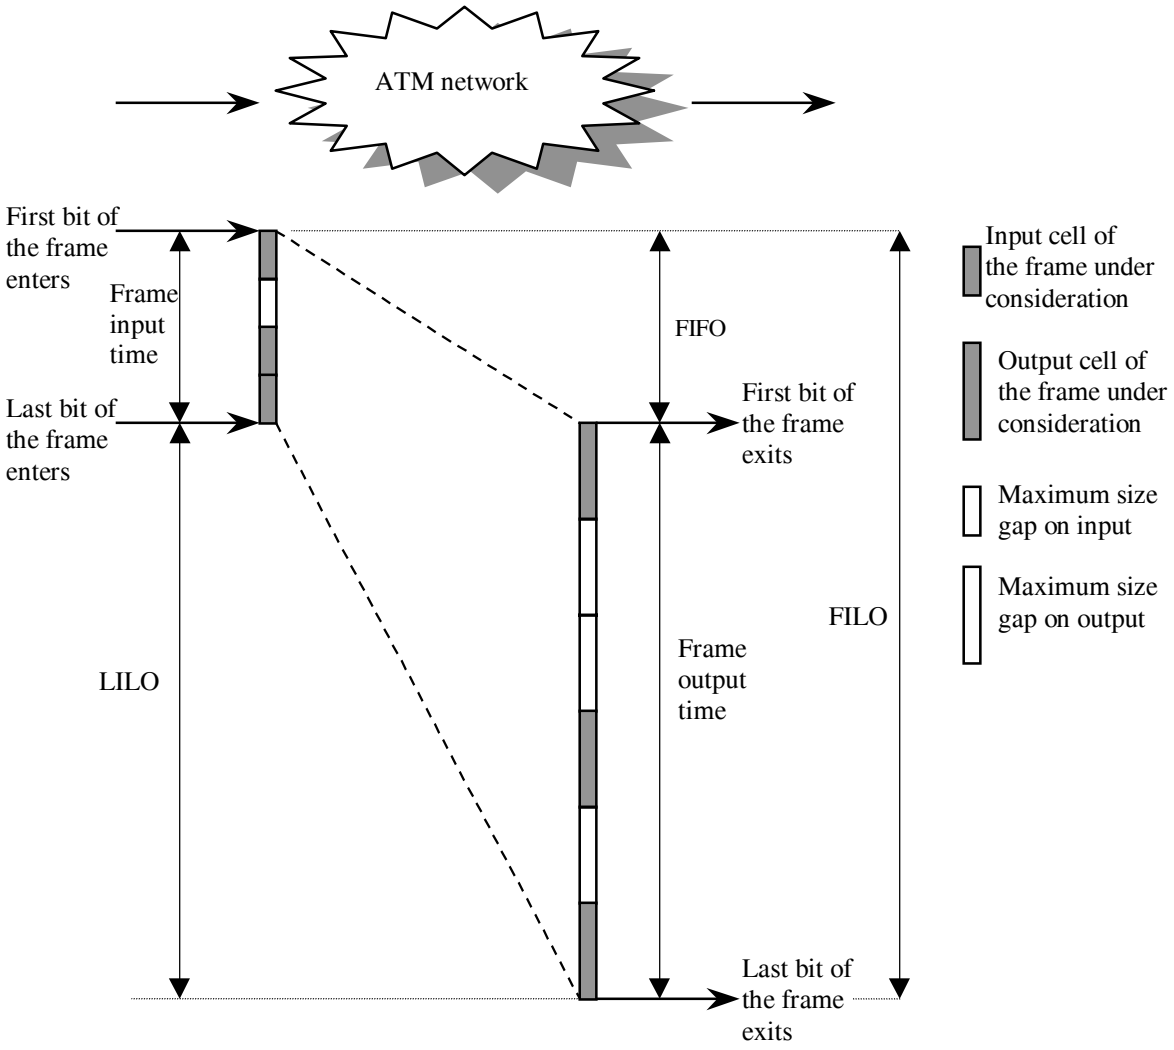

Figure A.4: Usual Latencies in an ATM Environment

In this case, the frame consists of 3 cells passing through an ATM switch with the input link rate higher than the output link rate. The frame is discontinuous on both input and output. The last cell is delayed considerably more than what FIFO latency would indicate.

It is possible to have one pattern of idle periods or unassigned cells (positions and a number of them) on the input of a given frame, and a completely different pattern on the output of the same frame. Note that it is also possible for a switch to remove idle periods or unassigned cells from the input, "transmitting" fewer of them on output, as we shall illustrate later.

In Figure A.4, as well as in the rest of this appendix, an unassigned cell, an idle period or a cell of another frame between cells of a given frame is indicated as a gap. In Figure A.4

the frame on input has a one-cell gap after the first cell of the frame, followed by the two remaining cells of the frame. On output, there is a two-cell gap after the first cell and then a one-cell gap between the second and the third cell of the frame.

From Figure A.4, it can be observed that it is possible for a switch to have a small FIFO latency if the first cell of a frame is transmitted quickly. However, if the later cells are delayed considerably, the receiver is not able to assemble the frame. FIFO latency does not reflect the expansion and compression of gaps on output. This is why FIFO latency is not an appropriate delay metric for switches in the ATM environment.

FILO Latency

From any of the previous three figures it can be noted that the relationship between FILO and LILO latency is as follows:

FILO latency = LILO latency + Frame Input Time

Although FILO and LILO latencies are related (one can compute one given the other), LILO latency is a preferred metric since it is independent of frame input time. FILO latency is different for different frame input patterns. Suitability of LILO and FILO metrics under various circumstances is discussed after introducing MIMO latency in the next section.

A.3. MIMO Latency Definition

MIMO latency (Message-In Message-Out) is a performance metric that defines the delay introduced upon a frame passing through a switch (or any other network component). When applied to a single switch, the MIMO latency accounts only for delays introduced by the switch (because of switching and other processing) and is independent of the frame input time, output transmission time, and other physical layer delays introduced on the input and output links.

Succinctly, MIMO latency is defined as follows:

MIMO latency = FILO latency – NFOT

where

- NFOT (Nominal Frame Output Time) is equal to the FILO latency of a given frame passing through *a zero-delay switch*.

We define a zero-delay switch as a switch that handles incoming frames in such way that they are transmitted on the output link without any unnecessary time consuming processing.

The above definition implies that MIMO latency is the difference between the measured FILO latency of a frame passing through the given switch and the FILO latency of the

same frame passing through a zero-delay switch. As defined, MIMO latency has the desired property of always being positive (or zero for a zero-delay switch).

The MIMO latency is not limited to switches. It applies to all types of communication devices, including repeaters, multiplexers, (store-and-forward or cut-through) bridges, routers, ATM switches, wires, or any combination of these. MIMO latency also accounts for discontinuous frames on the input and/or output. For discontinuous frames on input, gaps may include idle periods, unassigned cells and/or cells from other frames. For discontinuous frames on output, it is assumed that there are no cells from other frames inserted between the cells of the given frame, but idle periods or unassigned cells are allowed. It should be realized that the last assumption does not present a limitation for measurements in benchmarking environments.

In the following two sections, we explore the concept of a zero-delay switch in depth.

A.4. Cell and Contiguous Frame Latency Through A Zero-Delay Switch

Figure A.5 illustrates the latency that one-bit frame would experience while passing through a zero-delay switch. As expected, a zero-delay switch should start transmission on the output link as soon as the bit arrives on the input link. Thus, the latency of a single bit through a zero-delay switch is equal to zero. A wire of a zero length is one example of a zero-delay switch.

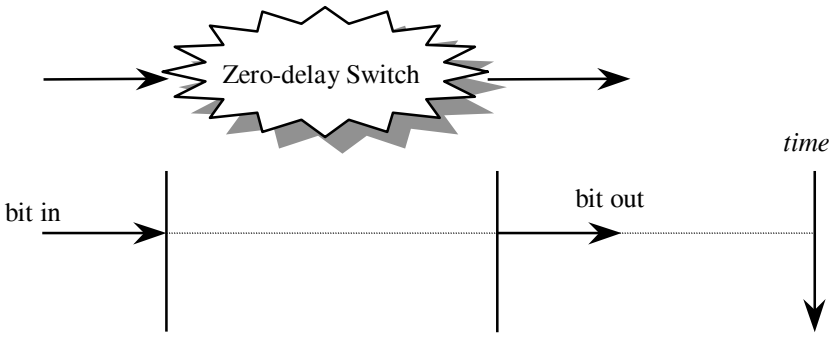

Figure A.5: Latency of one bit passing through the zero delay switch

Figure A.6 illustrates how a zero-delay switch would handle a cell consisting of multiple bits. The desired performance depends upon the relationship between the input and output link rates. In the case when the input link rate is equal to the output link rate, as presented in Figure A.6a, a zero-delay switch transmits each bit as soon as it arrives. Thus, each bit of the cell experiences zero latency in a zero-delay switch. A zero-length wire is one example of a zero-delay device.

Figure A.6b illustrates the case when the input link rate is higher than the output link rate. In this case, outputting (transmitting) a bit takes longer than inputting it. The zero-delay switch can transmit only the first bit as soon as it is received. The other bits of the cell

can not be transmitted immediately as they arrive, because the transmission of all previously received bits has not yet finished. Bits at the end of the cell wait longer then bits at the beginning. Thus, a zero-delay switch in this situation should be intelligent enough to do appropriate buffering of incoming bits. A zero-length wire with a FIFO buffer is an example of a zero-delay device that can handle inputs faster than the output.

Figure 6c illustrates the case when the input link rate is lower than the output link rate. A zero-delay switch does not start transmission of the first bit immediately after it is received, but after an appropriate delay. Bits at the beginning of the cell are delayed more than bits at the end, with larger delays for slower output link rates. Only the last bit of a cell has no delay and it is transmitted immediately upon its arrival. Thus, a zero-delay switch would be intelligent enough to avoid under-runs by appropriately delaying the transmission of incoming bits. A zero-length wire with an "intelligent" FIFO buffer is an example of such a zero-delay device.

It should be realized that the illustrations in Figure A.6 apply not only to cells, but also to contiguous frames passing through a zero-delay switch.

Note that a repeater can be considered as a zero-delay switch with input link rate equal to output link rate. Thus, Figure A.6a illustrates how a repeater handles incoming frames.

Also, note that a multiplexer, with n links on input and the output link capacity equal to the sum of input link capacities, can be considered as a zero-delay switch with input link rate lower than output link rate. For a multiplexer with two input links of rates equal to one half of the output link rate, Figure A.6c illustrates how the multiplexer would handle incoming frames. Similarly, a demultiplexer can be considered as a zero-delay switch with an input-link rate higher than the output-link rate. Figure A.6b illustrates operation of a two-output demultiplexer.

Based on Figure A.6, Table 1 provides (qualitative) indications for the four usual frame latency metrics applied to a zero-delay switch. None of the latencies has a zero value in all three cases, as it should be for the latency of a frame passing through a zero-delay switch.

Table 1: Usual Latencies Applied to a Zero-Delay Switch

	FIFO	LILO	LIFO	FILO
Input rate = Output rate	0	0	negative	positive
Input rate > Output rate	0	positive	negative	positive
Input rate < Output rate	positive	0	negative	positive

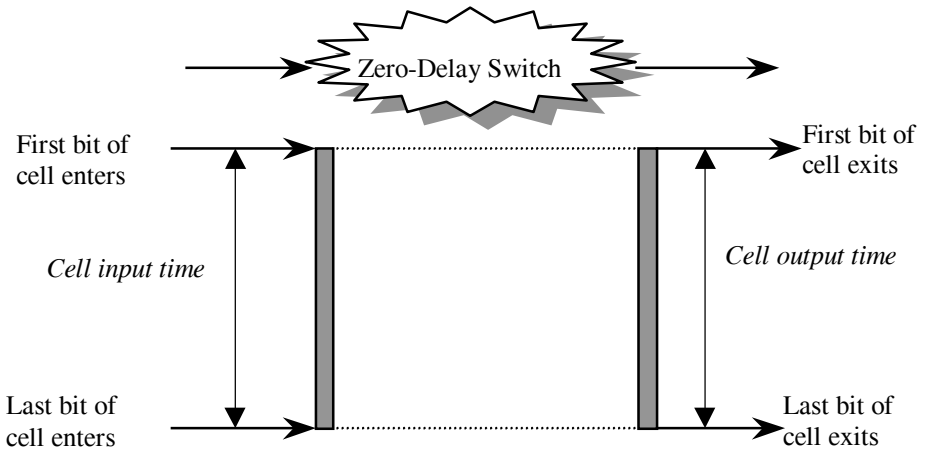

a. Input Link Rate = Output Link Rate

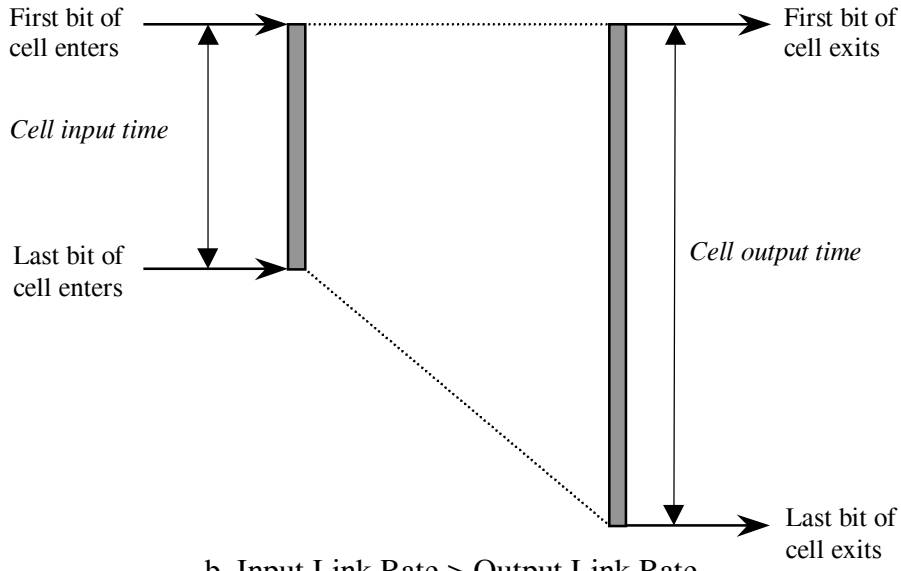

b. Input Link Rate > Output Link Rate

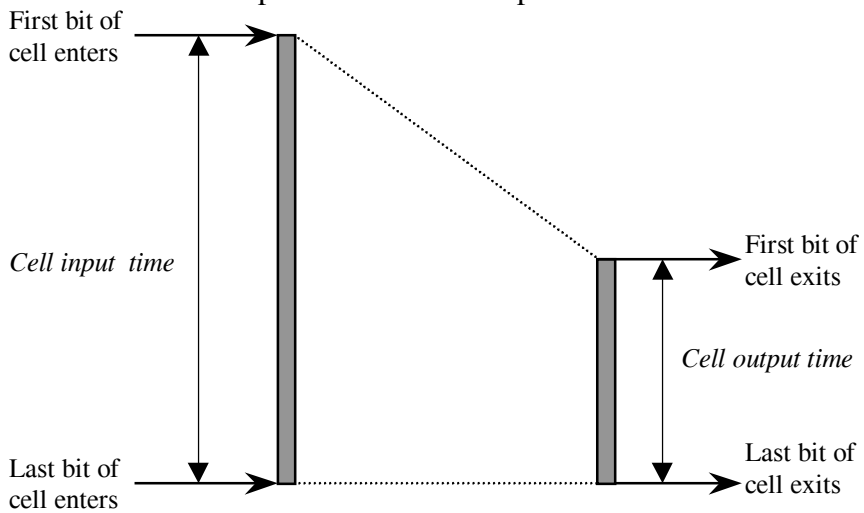

c. Input Link Rate < Output Link Rate

Figure A.6*:* Latency of one cell passing through a zero-delay switch

A.5. Latency of Discontinuous Frames Passing Through a Zero-Delay Switch

In this section, we consider how a zero-delay switch handles discontinuous frames in an ATM environment. In particular, we are interested in FILO latency, since it is used in the MIMO latency definition.

Figure A.7 illustrates one of two possible cases of a frame passing through a zero-delay switch with an input link rate higher than the output link rate. The frame includes two cells and the input link rate is 4 times the output link rate. The two cells start arriving at time t = 0 and t = 5, respectively. A zero-delay switch will start transmitting the first cell at time t = 0 and finish at time t = 4. The second cell can be transmitted without waiting and it is finished at t = 9. This is how long a zero-delay switch will take to transmit this frame. Hence, FILO latency of a zero-delay switch for this frame is 9. This is the normalized frame output time (NFOT) for this input pattern. No device can transmit this frame any faster. If a device takes longer, the difference between the FILO latency of the device and NFOT is considered as the delay introduced by the device.

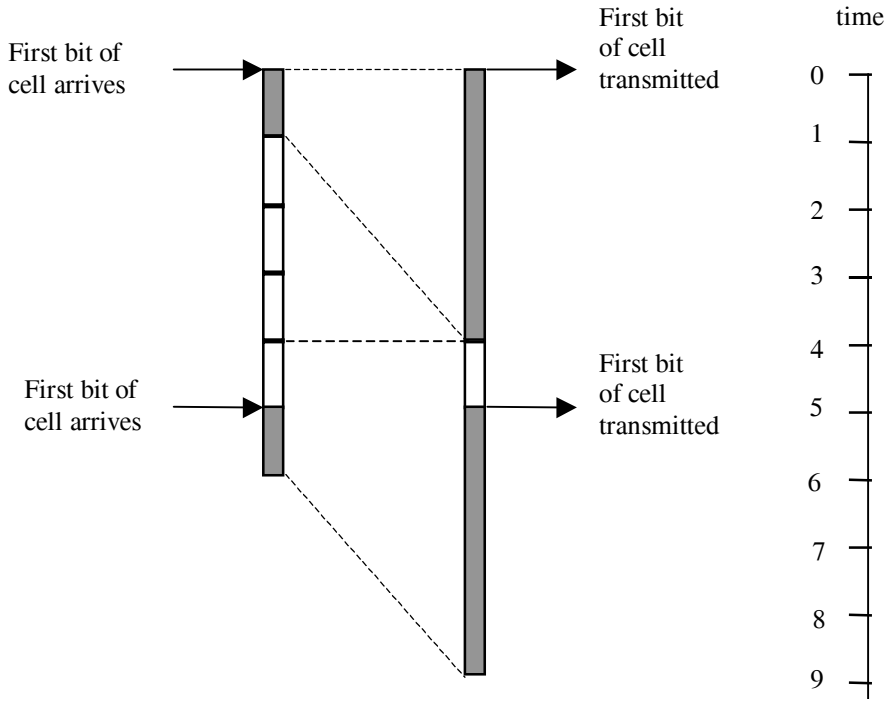

Figure A.7: Zero Delay Switch Operations, no Cell Waiting Case (Input rate > Output Rate)

Figure A.8 shows the other possible case of a frame passing through a zero-delay switch with an input link rate higher than the output link rate. As in Figure A.7, the frame has two cells and the input link rate is 4 times the output link rate. However, the frame has a different gap pattern. The second cell arrives at time t = 2 and thus has to wait. A zero-delay switch will start transmitting the first cell at time t = 0 and finish at time t = 4. The

second cell can be transmitted at t = 4 and finished at t = 8. Hence, FILO latency of a zero-delay switch for this frame is 8.

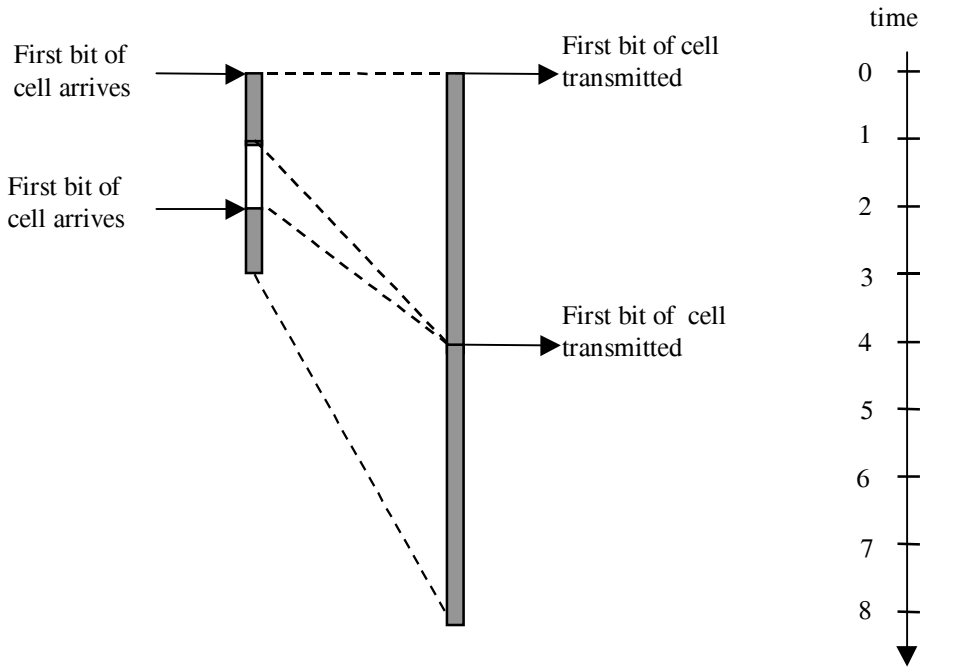

Figure A.8: Zero Delay Switch Operation, Cell Waiting Case (Input Rate > Output Rate)

Thus, in the case when the input link rate is higher than the output link rate, it is possible that:

- an incoming cell can be transmitted immediately (no cell waiting case) or
- an incoming cell has to wait for previously received cells of the same frame to be transmitted (cell waiting case).

Thus, for a given discontinuous frame, it is possible that some cells have to wait on previously received cells of the same frame, while some cells can be transmitted without waiting. Also, notice that a zero-delay switch is decreasing the size of each gap from input, with some gaps being completely removed.

Figure A.9 illustrates the only possible case of a frame passing through a zero-delay switch with an input rate lower than the output rate. Again, the frame includes two cells but the output link rate is now four times the input link rate. The two cells arrive at time t = 0 and t = 5, respectively. A zero-delay switch will start transmitting the first cell at time t = 3 (not at t = 0, in order to avoid an underrun), and finish at time t = 4. The second cell starts at t = 8 and finishes at t = 9. This is how long a zero-delay switch will take to transmit this frame. Hence, the FILO latency of a zero-delay switch for this frame is 9.

Note that in the case when the input rate is lower than the output rate, a cell never has to wait for completion of transmissions of previously received cells. Also, notice that in this case, a zero-delay switch does not eliminate any gaps from the input, although each gap

is enlarged on output. Additionally, when back-to-back cells are received on the input, new gaps are introduced between cells on the output.

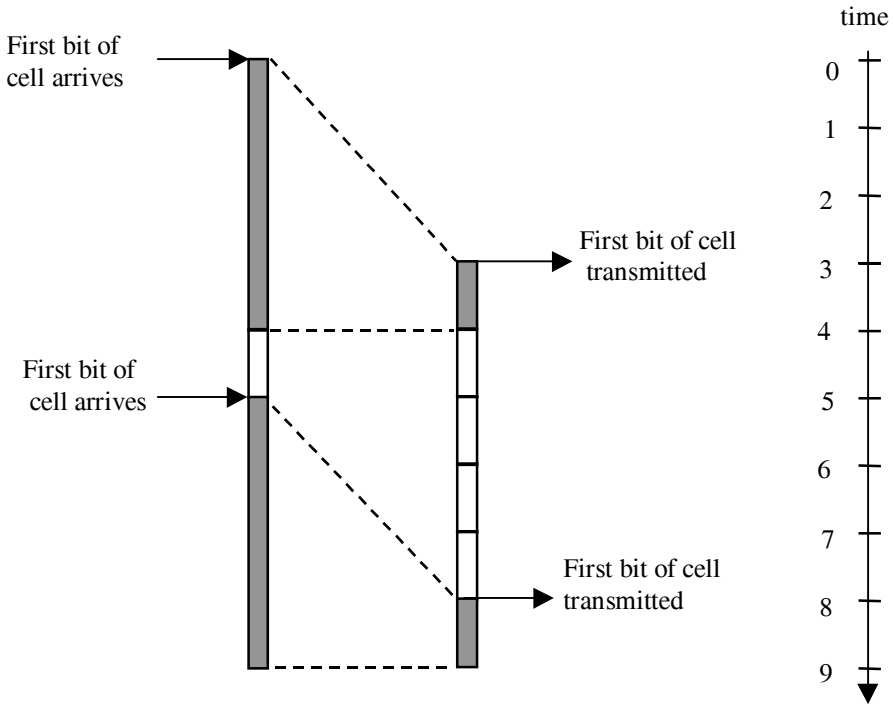

Figure A.9: Zero-Delay Switch Operation
(Input Rate < Output Rate)

A.6. Calculation of FILO Latency for a Zero-Delay Switch

The MIMO definition introduces NFOT as the FILO latency of a frame passing through a zero-delay switch. In this section, we explain how to obtain NFOT "on the fly," i.e., when a frame pattern is not known in advance, but cell arrival times can be obtained in real time. We define the following parameters:

- CIT = cell input time = 424[bits] / Input Link Rate [bits/sec]
- COT = cell output time = 424[bits] / Output Link Rate [bits/sec]

The procedure for NFOT calculation is as follows:

a. Initially NFOT = 0 and time t is measured from the arrival of the first bit of the first cell in a zero-delay switch.
b. For each cell with its first bit arriving at time t, update NFOT as follows:

$$NFOT = \max\{t, NFOT\} + CT$$

where:

$$CT = \begin{cases} CIT & \text{if input link rate} \le \text{output link rate} \\ COT & \text{if input link rate} \ge \text{output link rate} \end{cases}$$

A.7. Equivalent MIMO Latency Definition

An equivalent MIMO latency definition, which is more convenient for use in frame latency measurements and calculations when the input link rate is lower than or equal to the output link rate, can be derived as follows.

Input link rate ≤ output link rate, implies that CIT ≥ COT. A zero-delay switch will transmit the last bit of each cell of the frame as soon as it is received. In particular, the last bit of the frame is transmitted as soon as it is received. Thus, NFOT in these cases is equal to the frame input time:

$$\text{NFOT} = \text{Frame Input Time}$$

and,

$$\begin{aligned}\text{MIMO latency} &= \text{FILO latency} - \text{NFOT} \\ &= \text{FILO latency} - \text{Frame Input Time} \\ &= \text{LILO latency}\end{aligned}$$

Then the equivalent MIMO latency definition is:

$$\text{MIMO latency} = \begin{cases} \text{LILO latency} & \text{if Input Link Rate} \le \text{Output Link Rate} \\ \text{FILO latency} - \text{NFOT} & \text{otherwise} \end{cases}$$

Throughout this discussion, we assume that the link rates are used in latency computation. If other rates are used, there is the potential for strange results. For example, it is possible that a carrier may offer a lower rate contract to a customer on a higher rate link. If the peak cell rate for the traffic contract is less than the link rate, and this peak cell rate is used for MIMO calculations, then the MIMO value may be negative, depending on the scheduling of cells on the link and the traffic contract. Using the link rate in MIMO calculations avoids this potential problem.

A.8. Measuring MIMO Latency

To measure MIMO latency for a frame passing through the System Under Test (SUT), the times of occurrence for the following two events need to be recorded:

- the first-bit of the frame enters into the SUT,
- the last-bit of the frame exits from the SUT.

The time between these two events is the FILO latency. NFOT can be obtained from the cell pattern of the test frame on input as explained in Section A.6. Substituting FILO latency and NFOT into the MIMO latency formula would give the SUT's delay for a given frame.

If the input link rate is lower than or equal to the output link rate, it is easier to calculate MIMO latency. In this case, the times of occurrence for the following two events need to be recorded:

- the last-bit of the frame enters into the SUT,
- the last-bit of the frame exits from the SUT.

The time between these two events is the LILO latency, which is equal to the MIMO latency for the frame. Note that the cell arrival pattern does not matter in this case.

Contemporary ATM monitors provide measurement data at the cell level. Considering that the definition of MIMO latency uses bit level data, we now describe how to calculate MIMO latency using measurements at the cell level. Standard definitions of two cell level performance metrics, which are of importance for MIMO latency calculation are:

- cell transfer delay (CTD) , defined as the time between the first bit of the cell entering the switch and the last bit of the cell leaving the switch,
- cell inter-arrival time, defined as the time between arrival of the last bit of the first cell and arrival of the last bit of the second cell.

In cases where input link rate is higher than output link rate, according to the MIMO latency definition, FILO latency has to be measured. From Figure A.10, it can be observed that:

FILO latency = First cell's transfer delay + First cell to last cell inter-arrival time

Thus, to calculate MIMO latency when the input link rate is higher than or equal to the output link rate, it is necessary to measure the transfer delay of the first cell of a frame and the inter-arrival time between the first cell and the last cell of a frame.

In cases when input link rate is lower than or equal to output link rate, it is sufficient to measure LILO latency. From Figure A.11, it can be observed that:

LILO latency = Last cell's transfer delay – CIT

Thus, to calculate MIMO latency when the input link rate is lower than or equal to the output link rate, it is necessary to measure the transfer delay of the last cell of a frame.

7. User Perceived Delay

It should be pointed out that MIMO latency measures only the SUT's contribution to the delay. It does not include the delay caused by components not in the SUT's control. In particular, it does not include the frame input time. However, a user using the system does have to wait while the frame is being sent to the SUT. A user typically assembles the frame and gives it to the network. The user starts waiting as soon as the first bit starts entering the system and cannot do any meaningful work until the last bit exits the network. Thus, user perceived performance is reflected by FILO latency.

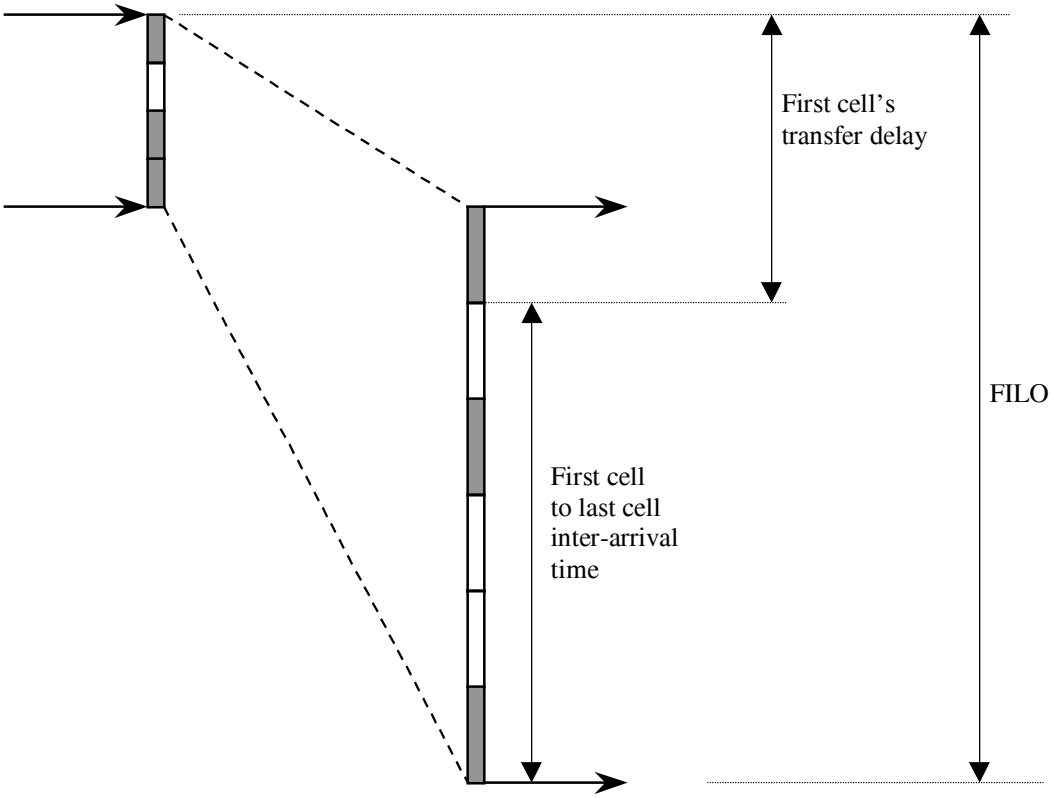

Figure A.10: FILO Latency Calculation
(Input Rate > Output Rate)

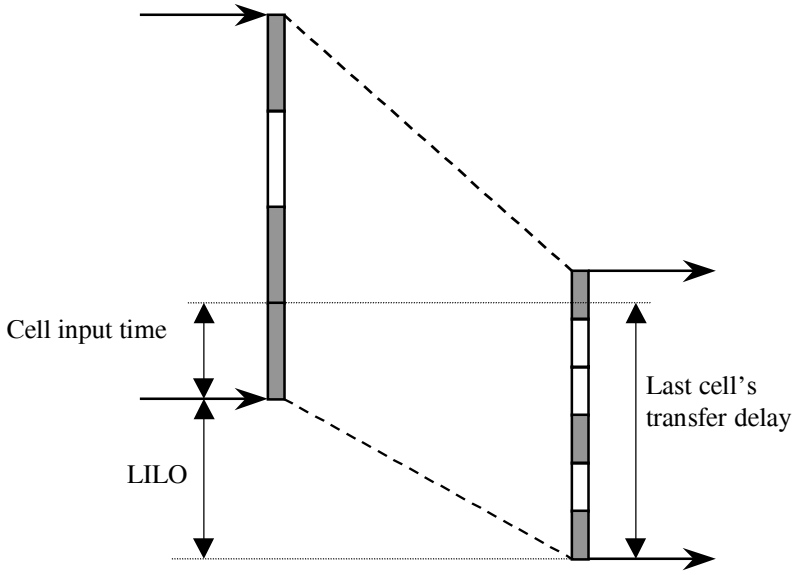

Figure A.11: LILO Latency Calculation
(Input rate ≤ Output Rate)

Figure A.12 illustrates the relationships between the user perceived performance and MIMO latency in two scenarios with continuous frames. In the first scenario, the input link rate is same as the output link rate. In the second scenario, the output is slower. The switch delay, as given by MIMO latency, is same in both cases; but the user perceived delay, as given by FILO latency, is different. For the case in Figure A.12b, FILO latency is worse. It can be observed that the user perceived delay depends upon input/output link speeds. On the other hand, network delay measured by MIMO latency is independent of link speeds. The difference between those two delays is the frame latency through a zero-delay switch.

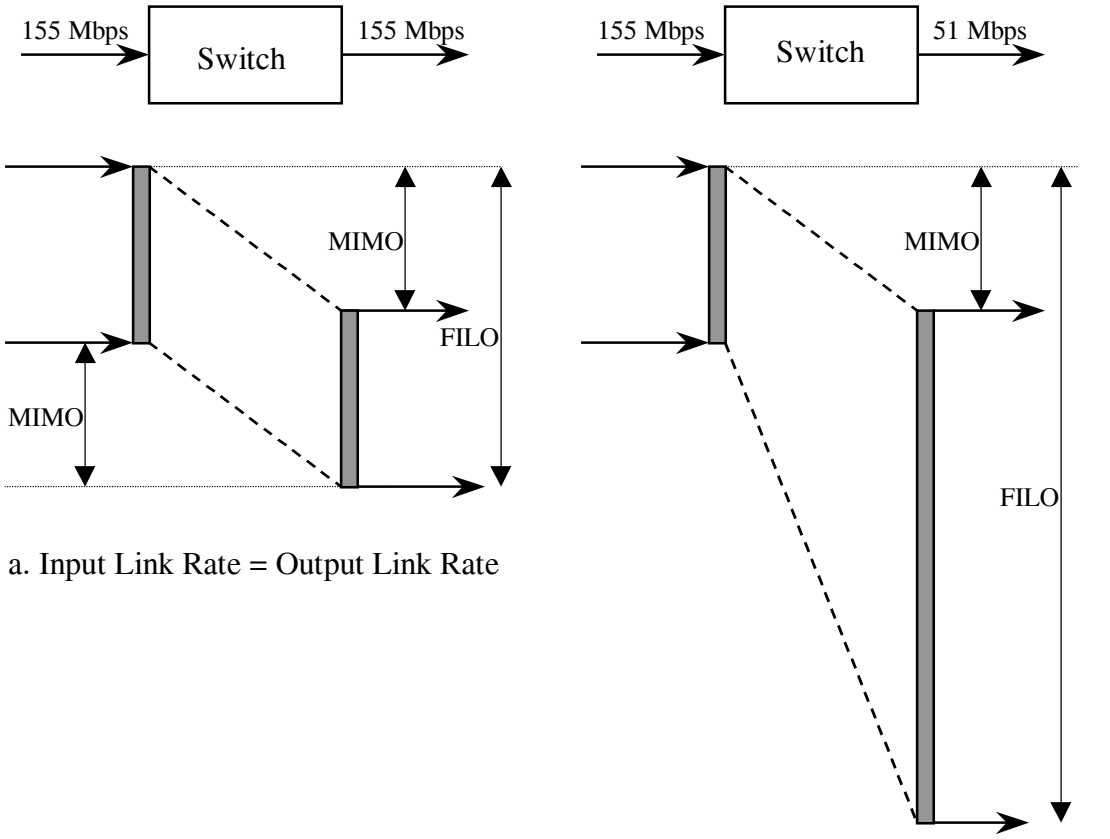

a. Input Link Rate = Output Link Rate

b. Input Link Rate > Output Link Rate

Figure A.12: FILO Latency as User Perceived Delay

References:

[1] CCITT Recommendation X.135: " Speed of Service (Delay and Throughput) Performance Values for Public Data Networks when Providing International Packet Switched Service", 1992

[2] S. Bradner, "Benchmarking Terminology for Network Interconnection Devices", RFC 1242

[3] ITU-T Recommendation I.356, "B-ISDN ATM Layer Specification," ITU-Study Group 13, Geneva, 1995