The ATM Forum
Technical Committee

Inverse Multiplexing
for ATM (IMA) Specification
Version 1.0

AF-PHY-0086.000

July, 1997

The ATM Forum
Worldwide Headquarters
2570 West El Camino Real, Suite 304
Mountain View, CA 94040-1313
Tel:+1-650-949-6700
Fax:+1-650-949-6705

Acknowledgments

The following people participated in the development of the IMA specification.

Richard Cam
Brian Carlton
Marian Chan
Jim Frimmel
Andre Green
John Guerrero
Amar Gupta
Frank Gurnee
Thomas W. Hill
Andreas Klug
Ian Lewis
Chris Martin
David McDysan
Michael Mroz
Jim Millar
Richard Proctor
Richard Vallée
Gerhard Vogt
David Wetle
Greg Wetzel

Richard Vallée, Editor
Rick Townsend, Chairman of the Physical Layer Working Group

Table of Contents

List of Figures

List of Tables

1. Introduction

This is the Inverse Multiplexing for ATM (IMA) specification, the purpose of which is to provide inverse multiplexing of an ATM cell stream over multiple physical links and to retrieve the original stream at the far-end from these links. Multiplexing of the ATM cell stream is performed on a cell by cell basis across the multiple links.

This specification is offered for private and public User-to-Network Interfaces (UNI's), Network-to-Network Interfaces (NNI's) and Broadband Inter-Carrier Interfaces (BICI's). It defines a new sublayer located between the Physical Layer Interface Specific Transmission Convergence (TC) sublayer and the ATM layer: the IMA sublayer. This specification also defines modifications to the Interface Specific TC sublayer on which IMA is implemented. It also defines the IMA Management Information Base (MIB) objects.

This specification does not deal with ATM layer functions such as UNI/NNI/BICI implementation requirements for signaling, traffic management and OAM.

1.1 Overview of Inverse Multiplexing for ATM

A methodology is described which provides a modular bandwidth for user access to ATM networks and for connection between ATM network elements at rates between the traditional order multiplex levels, for instance, between the DS1/E1 and DS3/E3 levels in the asynchronous digital hierarchies. DS3/E3 links are not necessarily readily available throughout a given network and therefore the introduction of ATM Inverse Multiplexers provides an effective method of combining the transport bandwidths of multiple links (e.g., DS1/E1 links) grouped to collectively provide higher intermediate rates.

The ATM Inverse Multiplexing technique involves inverse multiplexing and de-multiplexing of ATM cells in a cyclical fashion among links grouped to form a higher bandwidth logical link whose rate is approximately the sum of the link rates. This is referred to as an IMA group.

Figure 1 on page 12 provides a simple illustration of the ATM Inverse Multiplexing technique in one direction. The same technique applies in the opposite direction.

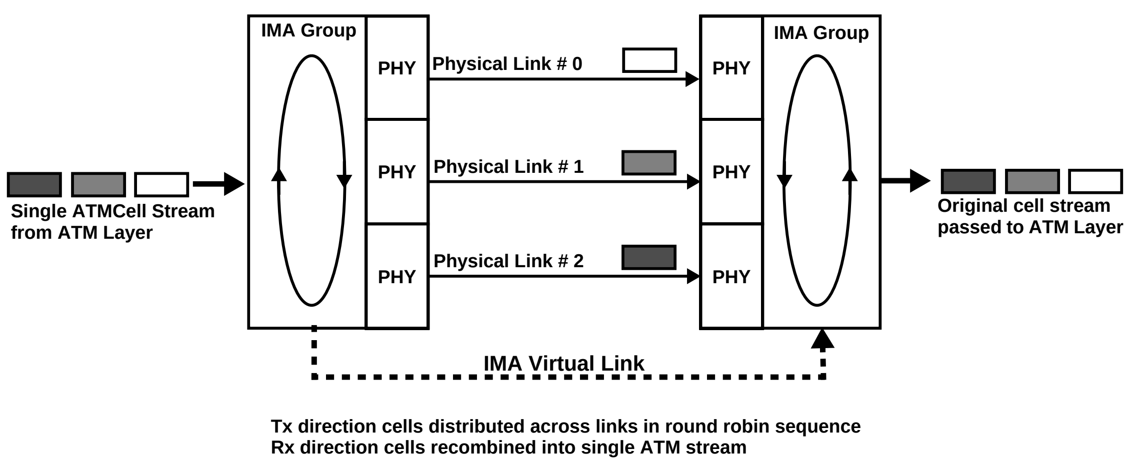

Figure 1 Inverse Multiplexing and De-multiplexing of ATM Cells via IMA Groups

IMA groups terminate at each end of the IMA virtual link. In the transmit direction, the ATM cell stream received from the ATM layer is distributed on a cell by cell basis, across the multiple links within the IMA group. At the far-end, the receiving IMA recombines the cells from each link, on a cell by cell basis, recreating the original ATM cell stream. The aggregate cell stream is then passed to the ATM layer.

The transmit IMA periodically transmits special cells that contain information that permit reconstruction of the ATM cell stream at the receiving IMA end after accounting for the link differential delays, smoothing CDV introduced by the control cells, etc. These cells, defined as IMA Control Protocol (ICP) cells, provide the definition of an IMA frame. The transmitter must align the transmission of IMA frames on all links (shown in Figure 2). This allows the receiver to adjust for differential link delays among the constituent physical links. Based on this required behavior, the receiver can detect the differential delays by measuring the arrival times of the IMA frames on each link.

At the transmitting end, the cells are transmitted continuously. If there are no ATM layer cells to be sent between ICP cells within an IMA frame, then the transmit IMA sends Filler cells to maintain a continuous stream of cells at the physical layer. The insertion of Filler cells provide cell rate decoupling at the IMA sublayer. The Filler cells should be discarded by the receive IMA.

A new OAM cell is defined for use by the IMA protocol. This cell has codes that define it as an ICP or Filler cell.

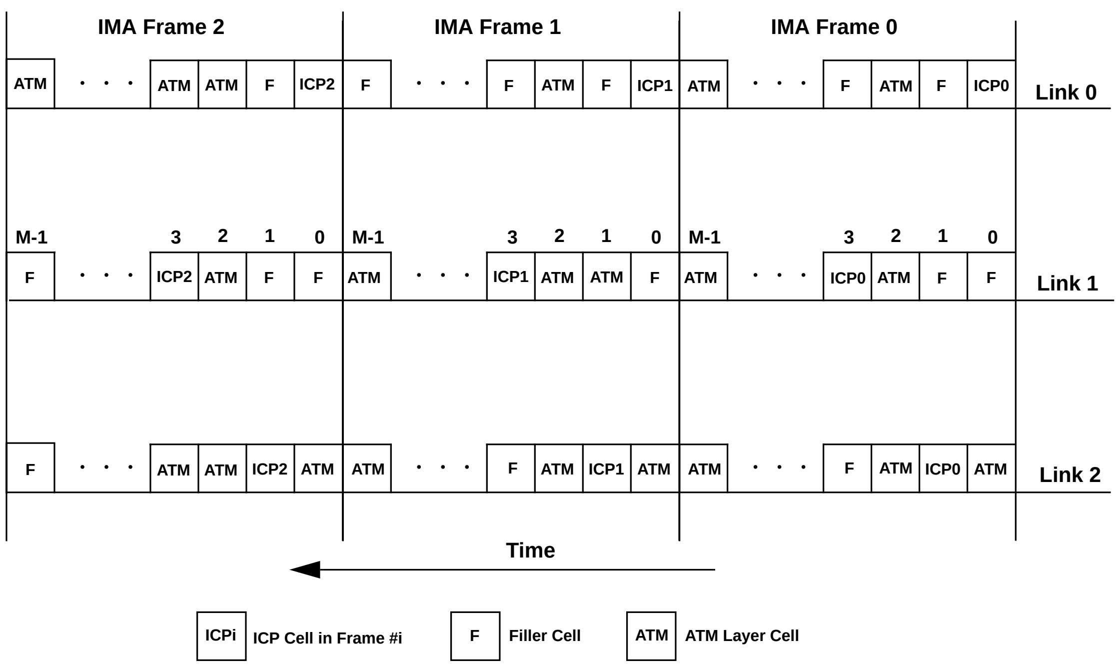

Figure 2 Illustration of IMA Frames

1.2 Terminology

This specification uses three levels for indicating the degree of compliance necessary for specific functions, procedures and coding associated with the IMA specification:

- Requirement **(R-x)**: functions, procedures and coding necessary for operational compatibility.

- Conditional Requirement **(CR-y)**: functions, procedures and coding necessary for operational compatibility if the specified optional function is implemented.
- Option **(O-z)**: functions, procedures and coding that may be useful, but are not necessary for operational compatibility.

A **(R-x)** at the beginning of the paragraph shall apply to all the requirements in that paragraph, unless otherwise specified. The x, y, and z represent the order of the requirement, conditional requirement and option as they appear in this specification.

2. Terms and Definitions

AAL	ATM Adaptation Layer Examples of AALs are AAL1, AAL3, AAL4 and AAL5.
Active	This is a link state indicating the link is capable to pass ATM layer cell in the specified direction.
Anomaly	Discrepancy between the actual and desired characteristic of an item. An anomaly may or may not affect the ability of an item to perform a required function (see [3]).
Asymmetric Configuration	Configuration where the IMA group is not required to configure all links in both the transmit and receive directions.
Asymmetric Operation	Mode where the IMA group is allowed to carry ATM traffic over links that are not Active in both transmit and receive directions.
ATM	Asynchronous Transfer Mode
ATM Layer Cell	Cells that are exchanged between the ATM layer and IMA sublayer.
AVCR	AVailable Cell Rate
B-ISDN	Broadband ISDN
BICI	Broadband Inter-Carrier Interface
Blocked	This a group state indicating than the group has been inhibited.
Blocking	This is a transitional state that allows graceful transition into the Unusable without loss of ATM layer cells.
CAC	Connection Admission Control
CBR	Constant Bit Rate
CDV	Cell Delay Variation
CID	Cell Identification (ICP CID)
CLR	Cell Loss Ratio
Config-Abort	This a group state indicating than the group has rejected the group parameters proposed the FE IMA group.
CPE	Customer Premises Equipment
CRC	Cyclic Redundancy Check
CTC	Common Transmit Clock configuration This is a configuration where the transmit clocks of all the links within the IMA group are derived from the same clock source.
Data Round-Robin	This represents the scheduler transmitting or receiving cells in a round-robin fashion.
Defect	A defect is a limited interruption in the ability of an item to perform a required function. It may or may not lead to maintenance action depending on the results of additional analysis. Successive anomalies of an item to perform a required function are considered as a defect (see [3]).
DSU	Data Service Unit
DXI	Digital eXchange Interface
ES	Errored Seconds Such as defined in ANSI T1.231[3]
Failure	A failure is the termination of the ability of an item to perform a required function (see [3]).

Fault	A Fault is an implementation specific NE event triggered, e.g., due to wrong LSM behavior at the FE. However, Fault does not cover the detection of invalid ICP cell content.
FC	Failure Count
FE	Far-End
Filler cell	This cell is used to fill in the IMA frame when no cells are available at the ATM layer. It is used for performing cell rate decoupling at the IMA sublayer (like Idle cell for single-link interface).
GSM	Group State Machine State machine defining the behavior of the IMA group.
GTSM	Group Traffic State Machine This is defined to indicate when traffic is exchanged with the ATM layer.
HEC	Header Error Check
HSSI	High Speed Serial Interface
ICP cell	IMA Control Protocol cell
ICO	ICP Cell Offset
IEC	Inter-Exchange Carrier
IETF	Internet Engineering Task Force
IDCC	IMA Data Cell Clock
IDCR	IMA Data Cell Rate This represents the rate at which IMA data cell should be exchanged between the IMA sub-layer and the ATM layer.
IFSN	IMA Frame Sequence Number
IMA	Inverse Multiplexing for ATM
IMA frame	The IMA frame is used as the unit of control in the IMA protocol. It is defined as M consecutive cells, numbered 0 to M-1 on each link, across the N links in an IMA group. One of the M cells on each of the N links is an ICP cell that occurs within the frame at the ICP cell offset position (the offset may be different on different links). The IMA frame is aligned on all links. “Aligned” means that it is transmitted “simultaneously”. Differential link delay can cause the reception to “misaligned” in time; the alignment is recovered by the link delay synchronization mechanism. M is defined during start-up; N is determined by the UM and the IMA link start-up procedure; and the ICP “stuff” mechanism is a controlled violation of the IMA consecutive frame definition.
IMA group	Group of links at one end used to establish an IMA virtual link to another end.
IMA RR	IMA Round Robin This is the distribution or retrieval of ATM layer and Filler on a cell by cell basis, when Tx=On or Rx=On respectively.
IMA sublayer	Sublayer part of the Physical layer and located between the interface specific Transmission Convergence (TC) sublayer and the ATM layer.
IMA virtual link	Virtual link established between 2 IMA ends over a number of physical links (IMA group).
In Group	This is an event indicating that a link has configured within an IMA group.
Inhibiting	This represents the action to voluntary disable the capacity of the group or the link to carry ATM layer cells reason other than reported problems.
Insufficient-Links	This a group state indicating than the group has not sufficient links to be in the Operational state.

ISDN	Integrated Services Digital Network
ITC	Independent Transmit Clock configuration This is a configuration where there is a transmit clock of at least one link within the IMA group that is not derived from the same clock source as the clock source of any other transmit links.
ITU-T	International Telecommunication Union - Telecommunication Standardization Sector
IV	ICP cell Violations
LASR	Link Addition and Slow Recovery procedure
Layer Management Functions	The Layer Management functions relate to processing of actions such configuration, fault monitoring, performance monitoring within the layer.
LCD	Loss of Cell Delineation defect The LCD defect is reported when the OCD anomaly has persisted for the period of time specified in ITU-T Recommendation I.432[41]. The LCD defect is cleared when the OCD anomaly has not been detected for the period of time specified in ITU-T Recommendation I.432.
LCD-RDI	Loss of Cell Delineation Remote Defect Indicator The LCD-RDI defect is reported to the FE when a link defect is locally detected. The LCD-RDI defect is not always required on link interface.
LDS	Link Delay Synchronization The LDS is a event indicating that the link is synchronized with the other links within the IMA group with respect to differential link delay.
LCR	Link Cell Rate Cell rate (cell/sec) over the portion of the link bandwidth available for carrying cells.
LID	Link IDentifier
LIF	Loss of IMA Frame defect The LIF defect is the occurrence of persistent OIF anomalies for at least 2 IMA frames.
Link Management	This function provides direct management of the physical links (e.g., DS1/E1 links).
LODS	Link Out of Delay Synchronization defect The LODS is a link event indicating that the link is not synchronized with the other links within the IMA group.
LOF	Loss Of Frame
LOS	Loss Of Signal
LSB	Least Significant Bit
LSI	Link Stuff Indication
LSM	Link State Machine
M	IMA Frame Size
MIB	Management Information Base
MSB	Most Significant Bit
N	Number of links configured in an IMA group
NE	Near-End (local end)
N_{on}	Number of links passing ATM layer cells in the transmit or receive direction.
NMS	Network Management System
NNI	Network-Node Interface

Not Configured	This is a group state indicating that the group does not exist yet.
Not In Group	This is used as an event or a state indicating that a link is no longer configured within an IMA group.
OAM	Operations And Maintenance
OCD	Out of Cell Delineation anomaly As specified in ITU-T Recommendation I.432[41], an OCD anomaly is reported upon the occurrence of ALPHA consecutive cells with incorrect HEC, and is no longer reported after detecting DELTA consecutive cells with correct HEC.
OIF	Out of IMA Frame anomaly The OIF is the occurrence of IMA anomalies listed later in this specification.
OOF	Out Of Frame
Operational	This a group state indicating than the group has sufficient links on both Tx and Rx directions to carry ATM layer cells.
OSI	Open System Interconnection
PMD	Physical Medium Dependent
ppm	part per million
PRC	Primary Reference Clock This term is from ITU-T Recommendation G.811[34]. This is a synonym of PRS.
PRS	Primary Reference Source This term is from ANSI T1.101[1]. This is a synonym of PRC.
QoS	Quality of Service
P_{Tx}	Minimum number of links required to be Active in the transmit direction for the IMA group to move into the Operational state.
P_{Rx}	Minimum number of links required to be Active in the receive direction for the IMA group to move into the Operational state.
RDI	Remote Defect Indicator
RFC notes	Requests For Comments notes
RFI	Remote Failure Indicator
RM cell	Resource Management cell
PRI	Primary Rate Interface
Rx	Receive (side)
Rx=On	The receive end is passing received ATM cells to the ATM layer.
Rx Failed	The Rx Failed condition is entered after a specified persistence of a link defect (e.g., LCD) or an IMA defect (LIF or LODS). It causes the LSM to enter into the Unusable state. The mechanism is implementation specific and could be separated from the failure definition due to a general used Network element-wide mechanism/parameter for failure detection.
SAR	Segmentation And Reassembly
SCCI field	Status and Control Change Indication field
SES	Severely Errored Seconds
SICP cell	Stuff ICP cell One of the 2 ICP cell comprising a stuff event.

Stuff event	Repetition of an ICP cell over one IMA link to compensate for timing difference with other links within the IMA group.
SMI	Structure of Management Information Defined in RFC 1442[19], it is divided into three parts: module definitions, object definitions, and trap definitions.
SNMP	Simple Network Management Protocol
Start-up	This a group state indicating than the group is waiting to see the FE in Start-up.
Start-up-Ack	This a group transitional state, when both groups start-up and the FE group parameters have been accepted.
Symmetric Configuration	Configuration where the IMA group is required to configure all links in both the transmit and receive directions.
Symmetric Operation	Mode where the IMA group is required to only carry ATM layer cells on links that are Active in both transmit and receive directions.
TC	Transmission Convergence
TRL	Timing Reference Link This is the link selected as the reference to derive the IDCR.
TRLCR	Timing Reference Link Cell Rate
Tx	Transmit (side)
Tx=On	The transmit end is sending ATM cells coming from the ATM layer.
UNI	User Network Interface
Unit Management	Entity managing a unit such as an IMA unit.
UAS	UnAvailable Seconds Such as defined in ANSI T1.231[3].
Unusable	This is a link state indicating the link is not in use due to fault, inhibition, etc.
Usable	This is a link state indicating the link is ready to operate in the specified direction, but it is waiting to move to Active.
UUS	UnUsable Seconds Seconds where the link state is Unusable.
VCI	Virtual Connection Identifier
VPI	Virtual Path Identifier

3. IMA Objectives

The main objectives that guided the development of the IMA Specification are listed below. Only the performance objectives applicable to the case where IMA is used over DS1/E1 links have been identified.

- Available for private/public UNI's and NNI's and BICI's.
- Provide transport of a single ATM cell stream at rates between existing link rates (e.g., between DS1/E1 and DS3/E3 rates)
- Applicable over DS1/E1 or other interfaces carrying ATM cells.
- Operating on the interfaces operating at the same nominal link cell rate (LCR).
- Provide for the provisioning of bandwidth in physical link increments (e.g., DS1/E1 increments).
- Provide transparent transport of the ATM layer and higher layer signals (including the preservation of cell order, control of cell delay variation, cell format, etc.).
- Use a cell based multiplexing technique for converting a single ATM stream into multiple lower speed ATM streams to be sent over independent links and retrieving the initial ATM cell stream from the links at the far-end.
- Maintain compatibility with current ATM Forum Physical Layer specifications (i.e., same transmission convergence sublayer including the relevant transmission MIB).
- Detect and reject lines with delay greater than the provisioned maximum differential delay tolerance.
- Add/delete links while maintaining the IMA group in service.
- Handle link failure and automatic link recovery.
- Support the use of leased and dial lines (e.g., ISDN lines) providing the same LCR, including links running in configurations where the transmit clocks of the links in the group are not all derived from the same source.
- Support symmetric and asymmetric IMA configurations, and symmetric and asymmetric IMA operations.
- Support of all ATM traffic/QoS classes including CBR.
- When running over DS1/E1 links, the IMA interface should compensate for link differential delays of at least 25 milliseconds.
- When running over DS1/E1 links, the IMA interface should contribute minimal CDV so as to support an end-to-end service objective of 1 millisecond CDV.
- When running over DS1/E1 links, the IMA interface should exhibit cell level error performance comparable to an individual DS1/E1 link. The effects of both random and bursty errors should be considered.
- Support system easily integrable with current network management system.
- Provide IMA status MIB - Management Information Base.
- Allow re-use of existing DS1/E1 Physical Layer devices.

4. IMA Reference Models

This section provides reference models of different aspects of IMA.

4.1 IMA Sublayer in Layer Reference Model

Figure 3 on page 21 illustrates the layer reference model including the Inverse Multiplexing for ATM (IMA) sublayer. It is based on the B-ISDN protocol reference model defined in ITU-T Recommendation I.321[38].

The IMA sublayer is part of the physical layer. It is located between the traditional Transmission Convergence sublayer and the ATM layer.

Figure 3 shows the main functions of the IMA sublayer.

		User Plane Functions	Layer Management Functions	Plane Management Functions
ATM Layer				
Physical Layer	IMA Transmission Convergence Sublayer	• ATM Cell stream splitting and reconstruction • ICP cells insertion/removal • Cell rate decoupling • IMA Frame Synchronization • Stuffing • Discarding of cells with bad HEC	• IMA Connectivity • ICP Cell Errors (OIF, LIF) • LIF/LODS/RDI-IMA defect processing • RDI-IMA alarm generation • Tx/Rx IMA Link State Report	• IMA group configuration • Link addition/deletion • ATM cell rate change • IMA group failure notification • IMA statistics
	Interface Specific Transmission Convergence Sublayer	• Cell scrambling/descrambling (if required) • No cell discarding • Cell delineation • Header error correction (if required) • HEC generation/verification	• HEC Error Indication • LCD-RDI alarm generation (if required)	• LCD failure notification • TC stats
	Physical Medium Dependent Sublayer	• Bit timing • Line coding • Physical Medium	• Local alarm processing • RDI alarms generation	• Link failure notification • PMD stats

Figure 3 IMA Sublayer in Layer Reference Model

4.2 Reference Model of Unit Operating IMA

This section describes the functions that characterize an unit that supports an IMA interface. Figure 4 on page 22 shows the functional blocks of such unit. This specification only covers the IMA functions and portions of the Link and Unit Management functions required for managing the IMA interface. All others are implementation dependent or defined in other ATM Forum specifications.

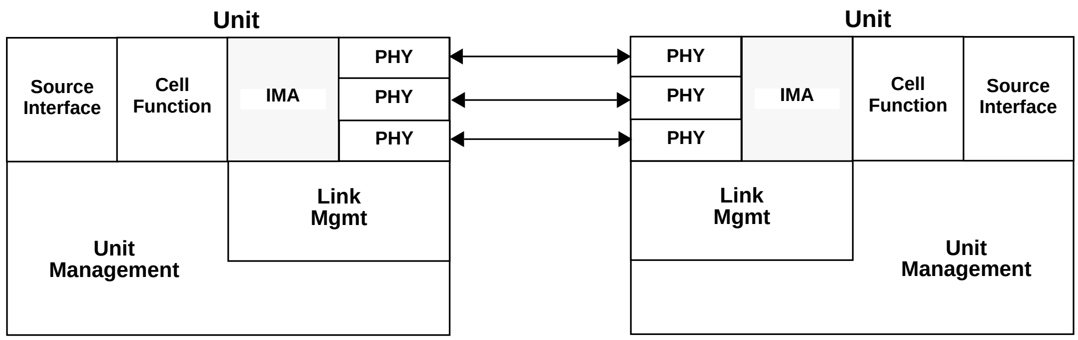

Figure 4 Functional Blocks of Units Hosting an IMA Interface

The following sections describe in more detail the functional blocks of the IMA unit.

4.2.1 Source Interface

The Source Interface provides the connection (typically proprietary) to an internal data bus (e.g., an ATM switch, router, computer). The Source Interface might also be a standard interface (e.g., DXI over HSSI). In the first case the IMA is "integrated" into a network element supporting other functions and may therefore provide such (non-standardized) functions as integrated management. In the second case, the IMA is a stand-alone unit.

There is no requirement that a particular Source Interface must be supported. The Source Interface choice is vendor specific. There is also no requirement that the Source Interface on one side match the Source Interface on the other (e.g., one side could be a direct interface to a data bus in an ATM switch and the other could be an ATM DSU supporting the IMA function).

4.2.2 Cell Function

The Cell Function is dependent upon the Source Interface:

- If the Source Interface emits ATM cells, the Cell Function is null (traffic shaping, if any, may be accomplished outside the unit). In this case, OAM cells, RM cells, etc., must be passed transparently through the source interface.
- If the Source Interface does not emit ATM cells, the Cell Function must arrange for the output of the source interface to be converted into ATM cells. There is no limitation on the Cell Function other than it must emit ATM cells. This requirement does not preclude the Cell Function implementing traffic shaping, buffering for frames/cells, etc. Support within the cell function is an implementation issue (e.g., OAM cell flows, RM cells, VPI/VCI to frame-level mapping).
- There is no requirement that the Cell Function on one side match the Cell Function on the other (e.g., one side could be null, as would be the case with a direct interface to a data bus in an ATM switch, and the other could support full AAL1, AAL3/4, AAL5 SAR functions, as might be the case for an ATM DSU supporting the IMA function).

4.2.3 Inverse Multiplexing

The Inverse Multiplexing function controls the distribution of cells onto the group of links made available to the IMA and handles differential delays and actions to be taken when links are added/dropped or when they are failed/restored. In the receive direction, the IMA performs differential delay compensation and recombines the cells into the original cell stream with the original inter-cell spacing or something. In essence, the Inverse Multiplexing function emulates a single UNI/NNI/BICI physical link; the IMA process of splitting and recombining streams is as transparent to the layer above as a traditional single-link physical layer interface.

4.2.4 Link Management

This function provides direct management of the physical links (e.g., DS1/E1 links). This may include such functions as establishing those links, for example DS1/E1 links via ISDN PRI signaling over a D channel. Other functions in Link Management include the notification of network management operations upon detection of link defects, collecting Facility Data Link information, instantiating the physical link MIB objects, etc.

Link Management functions also relate to the management and establishment of the links available to the IMA function.

- When implementing an IMA virtual link over a specific type of links (e.g., DS1/E1 links), the IMA MIB objects should inter-work with the link MIB if specified in an RFC (e.g., RFC 1406 [18] for DS1/E1).
- The Link Management function should not preclude the option of using dial-up services (e.g., ISDN service). If such are supported, the Link Management function should perform the connection management (establishment and release). Unit management should determine the establishment and release policies (i.e., when to initiate a new DS1/E1 link, when to release it, etc.).

4.2.5 Unit Management

The Unit Management (UM) function provides for the management of all functions of the whole unit. For example, capabilities required by users would include support of MIBs indicating the status of the IMA function, integration of alarms in the IMA, provision of a configuration interface, SNMP access to the unit, etc.

The IMA MIBs which allows management of the IMA functions of the unit are specified in Section 15 on page 80 and shall be implemented if the UM is SNMP based. Other unit management functions are for further study.

4.3 Timing Reference Model

On a link basis, the unit clocking should comply with the appropriate requirements for the links (e.g., ATM Forum DS1 and E1 specifications[8][9]).

On a group basis, the timing may be based on the timing configuration guideline presented in Figure 5 on page 24. Other timing configurations are feasible but are not discussed here.

Switch Sw and switches Sw(0) to Sw(N-1) appearing in Figure 5 indicate the possible configurations.

The “Tx Clock Unit” may derive its timing from an internal source, an external source or from any incoming link.

If loop timing (also referred as slave clocking[33]) is required, Sw(0) to Sw(N-1) could be set to position “1” on the associated link. Loop timing may be used when synchronizing to a digital network on a per link basis. Another example of loop timing would be using one of the incoming links as the source of the “Tx Clock Unit” and have all the link switches set to “0”.

The term Independent Transmit Clock (ITC) is used for operation where the transmit clock on each link is independently derived from a clock source. An example is illustrated in Figure 5 on page 24 by having at least one of switches Sw(0) to Sw(N-1) being set to position “1”.

The term Common Transmit Clock (CTC) is used when the same clock is used for all links. An example of this is illustrated in Figure 5 on page 24 by having all switches Sw(0) to Sw(N-1) being set to position “0”.

Appendix G on page 127 shows application scenarios intended to provide guidance in selecting a timing configuration.

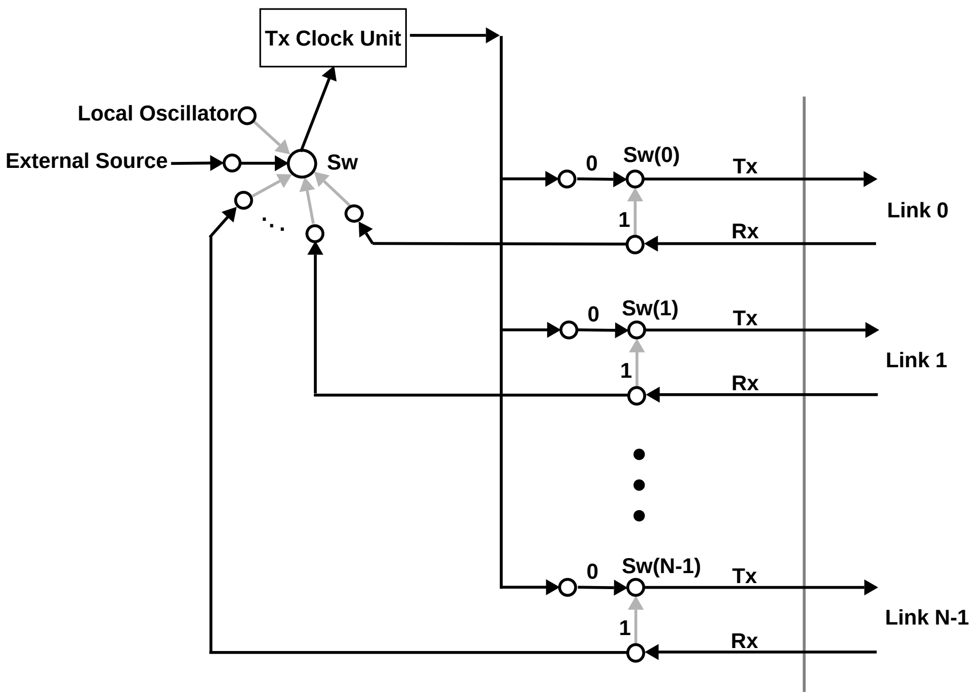

Figure 5 IMA Timing Configuration Reference Model

5. Basic IMA Protocol Definitions

5.1 IMA Physical Link Characteristics

This section specifies the basic architecture of the IMA interface.

(R-1) The IMA interface shall utilize a number (N) of transmission links designated as an IMA group by the unit management (UM). These links shall all operate at the same nominal link cell rate (LCR).

(R-2) The IMA interface shall be connected to another IMA interface over clear channel facilities. This implies that all cells generated by the transmit IMA unit shall only be terminated at the receive IMA unit.

5.2 Transmission Convergence Sublayer Specification

This section defines:

- exceptions to the interface specific transmission convergence functions when IMA sublayer is implemented, and
- basic IMA sublayer functions

5.2.1 Exceptions to the Interface Specific Transmission Convergence Sublayer

The physical layer HEC generation/verification, cell header correction (if required), cell delineation and cell payload scrambling/descrambling (if required) should be implemented as specified in the appropriate physical layer specification for the links to be used in the IMA group.

The following exceptions to the interface specific Transmission Convergence (TC) sublayer shall be met:

(R-3) The interface specific TC sublayer shall pass all cells to the IMA sublayer or provide an indication that a cell was received (this includes HEC errored cells).

(R-4) Idle cells shall not be used at any time to perform cell rate decoupling.

Cell rate decoupling is performed at the IMA sublayer.

5.2.2 IMA Specific Transmission Convergence Sublayer Part

5.2.2.1 IMA General Characteristics

(R-5) The transmit IMA shall assign a unique Link Identifier (LID) to each physical link within the IMA group. The LID shall, not be changed while the link is a member of an IMA group.

(R-6) The selected LID value shall be between 0 and 31 inclusive.

(R-7) The transmit IMA shall distribute ATM cells arriving from the ATM layer (including any unassigned cells) over the N links in a cyclic round robin fashion, and on a cell-by-cell basis.

(R-8) The transmit IMA shall distribute the ATM layer cells over the links using an ascending order based on the LID assigned to each link within the IMA group.

(R-9) Each end of the IMA virtual link shall use the IMA Control Protocol (ICP) cell format defined in Table 2 on page 29 to convey IMA configuration, synchronization, status and defect information to the far-end (FE).

(R-10) The transmitting IMA sublayer shall perform cell rate decoupling by inserting IMA Filler cells between ICP cells when there is no cell available at the ATM layer.

(R-11) The receive IMA shall:

- accept ATM cells from the N links according to ascending order based on the Link ID (LID) received in the ICP cells on the incoming link,
- compensate for differential link delays and re-build the original ATM cell stream,
- discard Filler cells,
- discard cells with bad HEC,
- process and discard ICP cells,
- pass the aggregate ATM cell stream to the ATM layer (this includes unassigned cells), and
- not mis-order cells.

(R-12) The receive IMA shall use the ICP cell to maintain link delay and protocol synchronization and to determine the differential delay between the links in the IMA group.

5.2.2.2 IMA OAM Cell Definition

Two IMA OAM cells are defined:

- Filler cell
- IMA Control Protocol (ICP) cell

These cells are defined in Table 1 on page 27 and Table 2 on page 29, respectively. The transmission of the octets is row-by-row, from top to bottom. In each octet, the most-significant bit is transmitted first (MSB = bit 7, LSB = bit 0).

The following fields are common to both Filler and ICP cells:

- cell header,
- OAM label, and
- cell identification

(R-13) The Filler and ICP cell header is the same and shall contain a fixed combination as specified in Table 1 on page 27 and Table 2 on page 29[1].

(R-14) Octet 6 of Filler and ICP cells shall contain the OAM Label and be set to “0x01”, to indicate IMA Version 1.0.

(R-15) The Cell Identification (CID) bit, bit 7 of octet 7 of the Filler and ICP cells, shall be used to identify the IMA OAM cell as an ICP or Filler cell.

(R-16) Octets 52-53 are defined as specified in ITU-T Recommendation I.610[42], Section 7.1, for octets 52-53 of the F1 and F3 flow OAM cells.

1. The value of octets 1 to 4 defined in this specification will be proposed to ITU-T, Working Group 13.

The following subsections gives details of the Filler and ICP cells respectively.

5.2.2.2.1 Filler Cell Definition

This section defines the Filler cell. This cell is used for performing cell rate decoupling at the IMA sublayer.

(R-17) The format of the Filler cell shown in Table 1 on page 27 shall be used.

Table 1 Filler Cell Format

Octet	Label	Comments
1 - 5	ATM cell header	Octet 1 = 0000 0000, Octet 2 = 0000 0000, Octet 3 = 0000 0000, Octet 4 = 0000 1011, Octet 5 = 0110 0100 (valid HEC)
6	OAM Label	bits 7-0: OAM Label value 00000001: IMA Version 1.0
7	Cell ID Link ID	Bit 7: IMA OAM Cell Type (0: Filler cell) Bits 6-0: Unused and set to 0
8-51	Unused	Set to 0x6A as defined in ITU Recommendation I.432 [41] for unused bytes.
52 - 53	CRC Error Control	Bits 15-10: Reserved field for future use - default value coded all zero. Bits 9-0: CRC-10 as specified in ITU Recommendation I.610[42].

The Filler cell can be validated by verifying the value of the header, OAM Label, CID and CRC fields.

5.2.2.2.2 ICP Cell Definition

The ICP cell structure is detailed in Table 2 on page 29.

The content of the ICP cell has been divided in five classes:

- A: Link specific information, transmitted only over the specific link.
- B: Group specific information, transmitted over all links in the group.
- C: Link specific information, however transmitted over all links in the group.
- D: Unused octet
- E: End-to-end channel

(R-18) The same content of fields appearing under classes B and C shall be transmitted over all links within an IMA group.

(R-19) Bits 4-0 of octet 7 shall be used to identify the Link ID (LID) on which the ICP cell is sent. Its range shall be from 0 to 31.

The IMA Frame Sequence Number (IFSN) field (octet 8) is used to indicate the sequence number of the IMA frame. The ICP Cell Offset field (octet 9) is used to indicate the location of the ICP cell within the IMA Frame. The use of these is defined in more detail in Section 5.2.2.3 on page 30.

The Link Stuff Indication (LSI) fields, bits 2-0 in octet 10, is used to indicate the occurrence of the next stuff event. The use of the LSI field is specified in more detail in Section 7 on page 36.

The Status and Control Change Indication (SCCI) field (octet 11) is used to indicate changes over octets 12 to 49 (classes B and C). Its is specified in more detail in Section 5.2.2.4 on page 32.

The IMA ID field (octet 12) is used to indicate the ID of the IMA group. The field is set to a value between 0 and 255 inclusive (8 bit value). Its use is specified in more detail in Section 5.2.2.5 on page 33.

The Group Status and Control field (octet 13) is used to indicate group status and control information. The use of the Group Status and Control field is explained in more detail in Section 5.2.2.6 on page 33 and Section 10.2 on page 53.

The Transmit Timing Information field (octet 14) is used to indicate the synchronization and the timing reference link (TRL) used by the transmit end. The use of the Transmit Timing Information field is explained in more detail in Section 7 on page 36.

The Tx Test Control, Tx Test Pattern and Rx Test Pattern fields (octets 15-16-17) are used to implement the test pattern procedure as specified in Section 13 on page 77.

The Link Information fields (octets 18 to 49) are used to exchange IMA link specific information between both ends of the IMA virtual link. The detail of the Link Information field is shown in Table 3 on page 30.

(R-20) The Link 0 Information field shall be used to exchange information on the transmit link having LID = 0, the Link 1 Information field for transmit link having LID = 1, etc.

The use of the Link Information is described in more detail in Section 10.1.6 on page 52 and Section 12.1.2 on page 69.

(R-21) Octet 50 within the ICP cell is unused and shall be set to 0x6A as defined for unused octets in ITU-T Recommendation I.432[41].

(R-22) The End-to-End Channel field (octet 51) shall be reserved as a proprietary channel. The field shall be set to “0” if unused.

(R-23) The IMA shall not rely on the processing of the End-to-End Channel field for any IMA functionality defined within this specification.

(R-24) At the receive end, the HEC and CRC-10 of the ICP cell shall be checked before relying on the fields covered by these two error checks.

Table 2 ICP Cell Format

Octet	Class	Label	Comments
1 - 5		ATM cell header	Octet 1 = 0000 0000, Octet 2 = 0000 0000, Octet 3 = 0000 0000, Octet 4 = 0000 1011, Octet 5 = 0110 0100 (valid HEC)
6		OAM Label	bits 7-0: OAM Label value 00000001: IMA Version 1.0
7	A	Cell ID and Link ID	Bit 7: IMA OAM Cell Type (1: ICP cell) Bits 6-5: Unused and set to 0 Bits 4-0: Logical ID for physical link range (0... 31)
8	A	IMA Frame Sequence Number	IMA Frame Sequence Number: from 0 to 255 and cycling.
9	A	ICP Cell Offset	Range (0... M-1): indicates position of ICP cell within the IMA frame.
10	A	Link Stuff Indication	Bits 7 - 3: Unused and set to 0 Bits 2 - 0: 111 = No imminent stuff event, 100 = Stuff event in 4 ICP cell locations (optional), 011 = Stuff event in 3 ICP cell locations (optional), 010 = Stuff event in 2 ICP cell locations (optional), 001 = Stuff event at the next ICP cell location (mandatory), 000 = This is one out of the 2 ICP cells comprising the stuff event (mandatory).
11	B	Status & Control Change Indication	Bits 7-0: Status change indication: 0 to 255 and cycling (count to be incremented every change of octets 12-49)
12	B	IMA ID	Bits 7-0: IMA ID
13	B	Group Status & Control	Bits 7-4: Group State 0000 = Start-up, 0001 = Start-up-Ack, 0010 = Config-Abort - Unsupported M, 0011 = Config-Abort - Incompatible Symmetry, 01xx = Reserved for other Config Abort reasons, 1000 = Insufficient-Links, 1001 = Blocked, 1010 = Operational, Others: Reserved Bits 3-2: Symmetry of Group 00 = Symmetrical configuration and operation, 01 = Symmetrical configuration and asymmetrical operation (optional), 10 = Asymmetrical configuration and asymmetrical operation (optional), 11 = Reserved Bits 1-0: IMA Frame Length (00: M=32, 01: M=64, 10: M=128, 11: M=256)
14	B	Transmit Timing Information	Bits 7-6: Unused and set to 0 Bit 5: Transmit Clock Mode: (0: ITC mode, 1: CTC mode) Bits 4-0: Tx LID of the timing reference link (0 to 31)
15	B	Tx Test Control	Bits 7-6: Unused and set to 0 Bit 5: Test Link Command (0:inactive, 1: active) Bits 4-0: Tx LID of test link (0 to 31)
16	B	Tx Test Pattern	Bit 7-0: Tx Test Pattern (value from 0 to 255).
17	B	Rx test Pattern	Bit 7-0: Rx Test Pattern (value from 0 to 255).
18	C	Link 0 Information	Bits 7-5: Transmit State (see Table 3 on page 30) Bits 4-2: Receive State (see Table 3 on page 30) Bits 1-0: Rx Defect Indicators (see Table 3 on page 30)
19-49	C	Link 1-31 Info	Status and control of link with LID in the range 1-31.
50	D	Unused	Set to 0x6A as defined in ITU I.432 for unused bytes.
51	E	End-to-end channel	Proprietary channel (set to 0 if unused).
52 - 53		CRC Error Control	Bits 15-10: Reserved field for future use - default value coded all zero. Bits 9-0: CRC-10 as specified in ITU Recommendation I.610[42].

Table 3 Encoding of the Link Information Fields

Bits	Encoding		
7-5	**Tx State**	**State**	**Additional information (see note 1)**
	000	Not In Group	
	001	Unusable	No reason given
	100	Unusable	Inhibited (vendor specific)
	101	Unusable	Failed (not currently defined)
	010	Unusable	Fault (vendor specific)
	011	Unusable	Mis-connected
	110	Usable	
	111	Active	
4-2	**Rx State**		
	000	Not In Group	
	001	Unusable	No reason given
	100	Unusable	Inhibited (vendor specific)
	101	Unusable	Failed
	010	Unusable	Fault (vendor specific)
	011	Unusable	Mis-connected
	110	Usable	
	111	Active	
1-0	**Rx Defect**		
	00	No defect	
	01	Physical Link defect (e.g. LOS, OOF/LOF, LCD)	
	10	LIF	
	11	LODS	

Note 1. The Unusable encoding is optional. As a minimum, the “No reason given” one of them shall be used to report the Unusable state.

5.2.2.3 IMA Frame Definition

The IMA frame is used as the unit of control in the IMA protocol. It is defined as M consecutive cells, numbered 0 to M-1 on each link, across the N links in an IMA group. One of the M cells on each of the N links is an ICP cell that occurs within the frame at the ICP cell offset position (the offset may be different on different links). The IMA frame is aligned on all links. “Aligned” means that it is transmitted “simultaneously”. Differential link delay can cause the reception to “misaligned” in time; the alignment is recovered by the link delay synchronization mechanism. M is defined during start-up; N is determined by the UM and the IMA link start-up procedure; and the ICP stuff mechanism, as defined in Section 7 on page 36, is a controlled violation of the IMA consecutive frame definition.

(R-25) An IMA Frame shall be composed of M consecutive cells transmitted on each link within the IMA group.

(R-26) The ICP cell shall be sent on each link once per IMA frame, hence every M cells.

(R-27) The IFSN field in the ICP cell shall be used to indicate the sequence number of the IMA frame. The field shall increment from 0 to 255 and repeat the sequence.

(R-28) The IMA Frame Sequence number in the ICP cell shall increment with each IMA frame on a per-link basis, but within an IMA frame, the sequence number contained in the ICP cell of each link is identical.

(R-29) The transmit IMA shall align the transmission of IMA frame on all links within an IMA group.

Figure 2 on page 13 shows an example of the transmission of IMA frames over three links. On link 0, the ICP cells have their cell offset set to zero (i.e., they are the first cell in each IMA frame). On link 1, the ICP cells have the ICP cell offset set to 3 and on link 2, the ICP cells have their ICP cell offset set to 1. In practice, these ICP cells should be distributed more evenly over the IMA frame but are shown closer for ease of illustration. It should also be noted that within an IMA frame, the ICP cells on all links have the same IMA Frame Sequence number. The ICP cell offset is described in more detail in the next section.

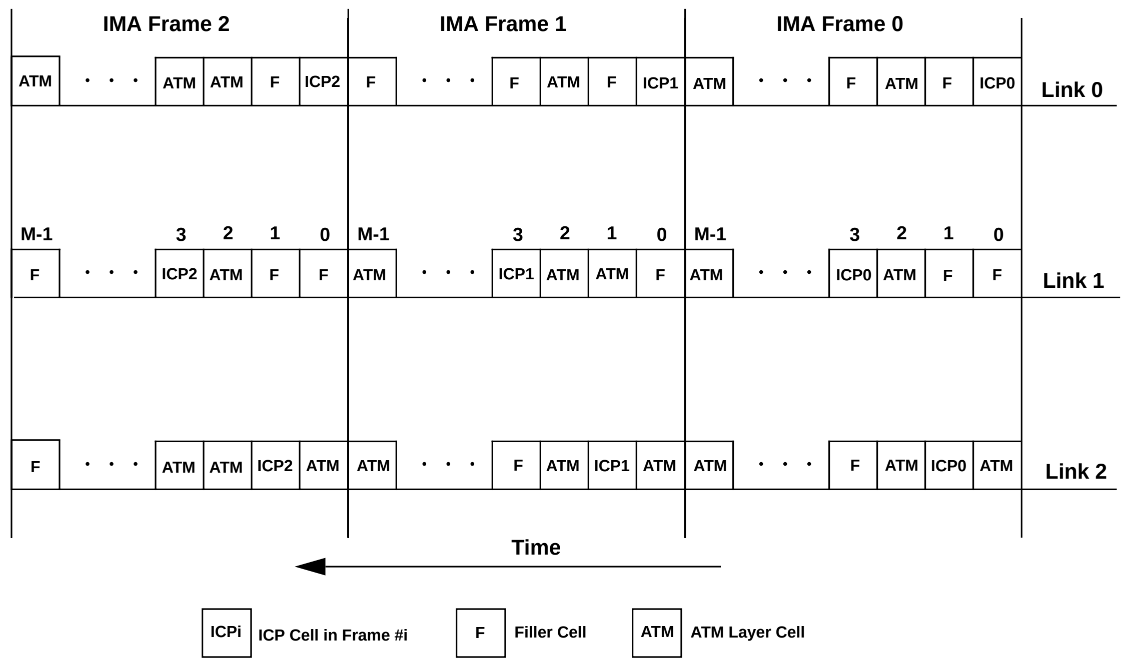

Figure 6 Illustration of IMA Frames

5.2.2.3.1 ICP Cell Offset

The ICP cell may be located anywhere within the IMA frame (e.g., the ICP cell offset may be any value from 0 to 127 if M = 128).

An ICP cell offset of 0 means that the ICP cell is the first cell of the IMA frame on that particular link. An ICP cell offset of M-1 means that the ICP cell is the last cell of the IMA frame on that link.

(R-30) The ICP Cell Offset field (octet 9) shall be used to indicate the location of the ICP cell within the IMA Frame. The field shall be set to a value between 0 to M-1 where M is the IMA frame length in cells.

(O-1) The transmit IMA may fill the following locations to spread ICP cells: 0, (1/2)M, (1/4)M, (3/4)M, (1/8)M, (3/8)M, (5/8)M, (7/8)M, etc. The transmit end could select the first unused location appearing in the list when adding a new link within the IMA group.

(R-31) The offset of the ICP cell sent over any link shall be fixed when the link is assigned a LID and this offset shall remain the same until the link is no longer part of the group.

Spreading the ICP cells in the IMA frame from link to link allows faster notification of configuration where the links exhibit the same propagation delay and minimum delay introduction related to the handling of ICP cells.

5.2.2.3.2 IMA Frame Length (M)

The IMA Frame Length (M) shall be processed as follows:

(R-32) The value of M used by an IMA shall always be carried in the ICP cell within the IMA Frame Length field.

(R-33) The IMA interface shall support M = 128.

(O-2) The support of M equals to 32, 64, 256 is optional; no other values are supported.

(R-34) The value M shall be configured at group start-up time, shall remain fixed thereafter and shall have a default value of 128.

Section 10.2 on page 53 describes the group start-up procedure in more detail.

Changing the transmitted M value at any other time will impact the IMA Frame Synchronization Mechanism (see Section 11 on page 65).

(CR-1) If (O-2) is used, the value of M used by the transmit end shall be as the one configured by the UM.

(CR-2) If (O-2) is used, it is not required that both directions of the IMA virtual link use the same M.

(CR-3) If (O-2) is used, the receive end shall synchronize its incoming links using the received M value for IMA frame synchronization.

(R-35) If the receive end does not support the received M, it shall abort the start-up procedure using the corresponding code defined in the “Group Status and Control” field (see Section 10.2.1.2 on page 55).

The value M may be configurable.

5.2.2.4 Processing of the SCCI Field

(R-36) The SCCI field shall be set to the previous transmitted SCCI value, incremented modulo 256, to indicate a change on at least one of the fields appearing in octets 12 to 49 in the transmitted ICP cell (see Table 2 on page 29).

The monitoring of the ICP cell at the receiving end is implementation dependent. The receiver may select only one or more links to monitor the value of the SCCI field.

(R-37) The SCCI field shall be used to identify received ICP cells for processing when ICP cells are monitored on more than one link, or when the monitored link has changed. Fields in octets 12 to 49 shall be processed if the SCCI field has advanced beyond the SCCI value of the last processed ICP cell.

5.2.2.5 IMA ID

The IMA ID is used to identify the IMA ID group.

(R-38) The IMA ID used by the transmit end shall be carried in the IMA ID field.

(R-39) The IMA ID shall be set at group start-up time.

Changing the transmitted IMA ID value at any other time will impact the IMA Frame Synchronization Mechanism (see Section 11 on page 65).

Both ends of the IMA virtual link can independently select an IMA ID. There is no need to ensure a unique IMA ID on each end.

The IMA ID may be configurable.

5.2.2.6 Group Symmetry

The IMA protocol is defined to allow symmetric or asymmetric cell rate transfer over the IMA virtual link. The IMA interface can be configured in three modes:

- *Symmetrical Configuration and Operation*: this the default and mandatory mode; in this mode, the IMA interface is required to configure each IMA link in both transmit and receive directions; it is also enforced to only transmit and receive ATM layer cells over links that are Active in both directions.
- *Symmetrical Configuration and Asymmetrical Operation*: in this case, the IMA interface is required to configure each IMA links in both transmit and receive directions; it is permitted to keep transmitting ATM layer cells over a link in the transmit direction while the link is not Active in the receive direction or vice versa. This behavior is described in Appendix C on page 117.
- *Asymmetrical Configuration and Operation*: in this case, the IMA interface is not required to configure all IMA links in both transmit and receive directions; it is also permitted to keep transmitting ATM layer cells over a link in the transmit direction while the link is not Active in the receive direction or vice versa. This behavior is described in Appendix C on page 117.

The “Symmetry” field, specified in Table 2 on page 29, is used to indicate the symmetry of the IMA group.

(R-40) The symmetry of the group shall be only established or changed at group start-up time.

(R-41) The “Symmetrical Configuration and Operation” mode shall be supported by all IMA implementations.

(O-3) The use of the “Symmetrical Configuration and Asymmetrical Operation” mode is optional.

(O-4) The use of the “Asymmetrical Configuration and Operation” mode is optional.

(R-42) If the NE does not support the symmetry mode proposed by the FE, or the proposed symmetry proposed by the FE and the configured symmetry of the NE do not match, the NE shall abort the start-up procedure using the appropriate code defined in the “Group Status and Control” field (Table 2 on page 29).

In order to allow a fast recovery when (O-3) or (O-4) is used at the NE and when the FE IMA can only be configured to the “Symmetrical Configuration and Operation” mode, the NE may adjust to the symmetric operation.

(R-43) Only the valid combinations of symmetrical configurations specified in Table 4 on page 34 shall be used at each end of the IMA virtual link.

The symmetry of the group may be configurable.

Table 4 Valid and Invalid Combinations of Different Symmetrical Configurations

			End B		
			Symmetrical Configuration		Asymmetrical Configuration
			Symmetrical Operation	Asymmetrical Operation	Asymmetrical Operation
End A	Symmetrical Configuration	Symmetrical Operation	Valid	Invalid	Invalid
		Asymmetrical Operation	Invalid	Valid	Invalid
	Asymmetrical Configuration	Asymmetrical Operation	Invalid	Invalid	Valid

6. Quality of Service (QoS) Requirements

(R-44) The IMA interface shall support all ATM traffic/QoS classes supported at the ATM layer.

7. Support for Common and Independent Transmit Clock Operations

In order to accommodate the use of Common Transmit Clock (CTC) and Independent Transmit Clock (ITC) configurations as defined in Section 4 on page 21, this specification defines the insertion of a stuff event to compensate for timing difference between the links within an IMA group.

The IMA link transmit interfaces may be locked to one clock source or may be plesiochronous with respect to the other links. When plesiochronous, one of the IMA transmit buffers may become depleted or overflow due to timing differences. The IMA cell stuff mechanism is provided to allow error-free ITC operation.

In the CTC mode, stuffing is also inserted on all links at nominal rate specified in Section 8 on page 38. This allows inter-operation between transmitter running the CTC and ITC modes and fixes a single nominal rate for the IMA Data Cell Rate (IDCR) defined in Section 8 on page 38.

(R-45) The transmit IMA shall indicate to the far-end (FE) in which transmit clock mode it is running in the “Transmit Clock Mode” field in the ICP cell (see Table 2 on page 29).

The transmit IMA is allowed to indicate that it is in the ITC mode even if all the transmit clocks of the links in the group are derived from the same source.

(R-46) The transmit IMA shall support the CTC mode.

(R-47) The transmit IMA shall only indicate to the FE that it is in the CTC mode when all the transmit clocks of the links in the group are derived from the same source.

(O-5) The transmit IMA may also support the ITC mode.

(CR-4) If (O-5) is used, the cell stuffing mechanism procedure shall also be invoked to prevent link transmit buffer under-run or over-run.

In all cases, the link transmit buffers should not deplete.

(R-48) The transmit IMA shall indicate a stuff action in the ICP cell preceding a stuff event using the mandatory Link Stuffing Indication (LSI) codes specified in Table 2 on page 29.

(R-49) The transmit IMA shall perform stuffing by repeating the ICP cell containing the Link Stuffing Indication (LSI) code (thus both ICP cells indicate that “this cell is 1 out of the 2 ICP cells comprising the stuff event”).

One of the two ICP cells comprising the stuff event is identified as a Stuff ICP (SICP) cell. Either one may be treated as a SICP cell.

The insertion of a stuff event is illustrated in Figure 7 on page 37.

The three additional indications (010, 011 and 100) may provide greater confidence in detecting a SICP occurrence.

(O-6) The transmit IMA may also indicate an incoming stuff event in the forth, third and second ICP cells preceding the stuff event using the optional LSI codes defined in Table 2 on page 29.

(R-50) The transmit IMA shall not introduce a stuff event more than once every 5M ICP, Filler and ATM layer cells over any link.

Appendix B.4 on page 112 explains why 5M cells has been selected.

(R-51) The receiving IMA shall remove one of any two consecutive ICP cells with LSI code indicating that "this cell is 1 out of the 2 ICP cells comprising the stuff event". The removed SICP cell shall not be counted as a cell for the purposes of determining the IMA round-robin sequence.

(R-52) The receive IMA shall operate with CTC and ITC transmission modes.

(R-53) When the FE IMA transmit mode does not match the NE IMA transmit mode, the NE IMA shall inform the UM of the mismatch. This shall not cause a restart of the group.

It is up to the implementer to decide how to process the incoming stuff cell indication codes. It is recommended that the receive IMA rely on at least one ICP cell with a correct CRC-10.

Figure 7 on page 37 shows the ICP cells before, during and after the stuff event occurrence.

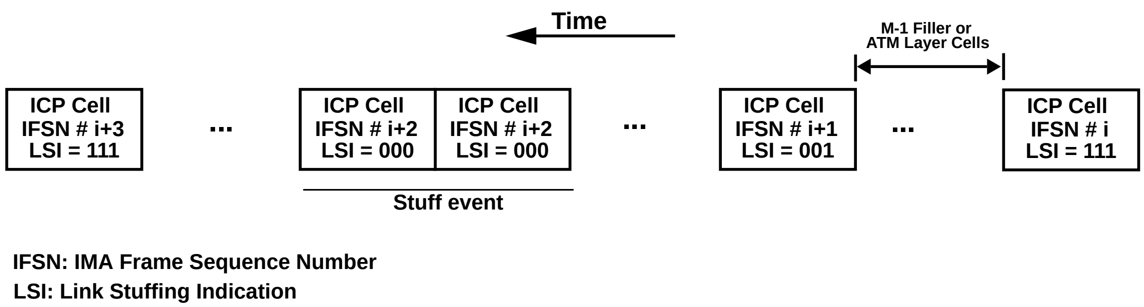

Figure 7 Insertion of a Stuff Event over One Link (Using only Mandatory Advanced Indication)

8. IMA Data Cell Rate Implementation

This section defines the requirements of the IMA Data Cell Rate (IDCR) implementation to minimize the introduction of CDV by the IMA virtual link.

The IDCR is defined as the rate at which ATM cells should be passed from the ATM layer to the IMA sublayer on transmit and passed from the IMA sublayer to the ATM layer on receive. In order to emulate single-link physical layer interfaces, the IDCR is nominally a constant for any given configuration of links in the Active state.

The IDCR is derived from one of the physical links identified as the timing reference link (TRL). The TRL is used to pass synchronization from the transmit to the receive end.

The IMA IDCR is derived, at both transmit and receive, according to the following equation:

$$IDCR = N_{on} \times TRLCR \times \left(\frac{M-1}{M}\right) \times \left(\frac{2048}{2049}\right) \qquad \text{(EQ. 1)}$$

In the transmit direction, N_{on} represents the number of links which currently transmit cells passed from the ATM layer (Tx=On as specified in Section 10.1.3 on page 46). In the receive direction, N_{on} represents the number of links which currently receive ATM layer cells to be passed to the ATM layer (Rx=On as specified in Section 10.1.3 on page 46).

TRLCR is the TRL cell rate providing by the TRL in the transmit and receive direction, respectively.

M is the IMA frame length in cell units. The “(M-1)/M” factor account for an ICP cell every M cells.

The TRLCR is scaled down by “2048/2049” to account for the insertion of one stuff event on the TRL after every 2048 ICP, Filler and ATM layer cells (e.g., every 16 IMA frames if M equals 128). This implies that in the case of ITC mode, the links other than the TRL could be faster or slower than the TRL which would require more or less frequent insertion of a stuff event to prevent transmit buffer underruns/overruns on these links. Appendix B.4.2 on page 113 explains the choice of 2048.

8.1 Behavior of the Transmit End

Figure 8 on page 38 shows the behavior of the transmit end.

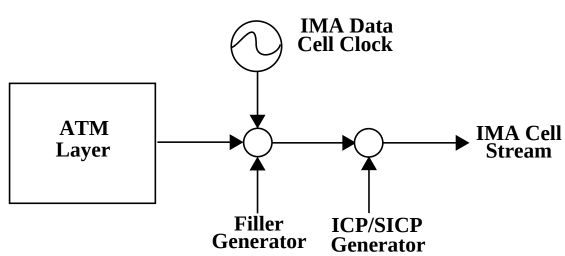

Figure 8 Behavior of the Transmit End

(R-54) The actions taken for cell rate decoupling at the IMA transmitter shall not inject a Filler cell if an ATM layer cell is available for scheduling. The IMA transmitter shall perform a check that an

ATM Layer cell is available and accept that cell when the Tx IMA Data Cell Clock (IDCC) ticks and only when the Tx IDCC ticks.

This means that the IMA interface should emulate the behavior of one physical interface while interacting with the ATM layer.

It should also be noted that the IDCR is independently determined by the transmitter at each end of the IMA virtual link. The Rx IDCR is recovered by the receive IMA using the indicated Tx TRL.

8.1.1 Derivation of the Transmit IDCR from the TRL

(R-55) The transmit IMA shall derive the Tx IDCR from the selected TRL according to (EQ. 1).

8.1.2 Selection of the TRL

(R-56) A link shall be selectable as the TRL if its transmit link state is set to Active (see Section 10 on page 43).

(R-57) If there is no link in the Active state, the IMA shall select one of the links in Usable state if any, else one of the links in the Unusable state shall be selected.

(R-58) The TRL shall be selected only during IMA start-up by the transmit IMA and changed only when the selected TRL is deleted or no longer satisfies requirement (R-56).

(R-59) Once selected, the TRL shall be indicated to the FE over the “Transmit Timing Information” field in the ICP cell (see Table 2 on page 29).

As long as requirements (R-56) and (R-58) are met, the criteria for selecting the TRL at the transmit end is vendor specific.

8.1.3 Stuffing on the TRL and Other Links

(R-60) When the transmit IMA is operating in the CTC mode, the IMA shall introduce a stuff event every 2048 ICP, Filler and ATM layer cells on all links.

(CR-5) If (O-5) is used, the transmit IMA shall introduce a stuff event every 2048 ICP, Filler and ATM layer cells on the TRL.

(CR-6) If (O-5) is used, the transmit IMA shall introduce stuff events on links other than the TRL in order to compensate the timing difference between the TRL and the other links.

The previous requirements implies that the receive IMA can use a modulo-2049 counter as an additional mean to predict the occurrence of stuff event on the TRL, in the case the transmit is in the ITC mode, or on all links in the case the transmit in the CTC mode, once a stuff event has been detected.

8.2 Behavior of the Receive End

This section defines the behavior of the receive IMA through requirements (R-60), (R-61) and (R-62). These requirements may be regarded as not applicable to an implementation if and only if the following conditions both apply:

- the IMA receiver is directly built into end equipment that directly terminates the ATM layer (i.e., terminates all ATM connections), and

- the system is only capable of carrying services that either do not require CDV control (e.g., some data services), or where the CDV in handled in some other way (e.g., absorbed in a play-out buffer at the ATM layer connection termination).

An implementation that satisfies both of the above conditions may conform to these requirements if desired.

The IMA receiver should remove the CDV introduced by a transmitter behaving as described in Section 8.1 on page 38. The CDV to be removed is in the form of ICP cells, including Stuff ICP (SICP) cells. Figure 9 on page 40 shows the behavior of the IMA receive end. The CDV introduced due to the emulated single physical interface behavior, as mentioned in Section 8.1 on page 38, cannot be removed. The Rx IDCR should be derived from the incoming IMA link identified as the TRL by the FE IMA transmitter over the "Transmit Timing Information" field in the ICP cell.

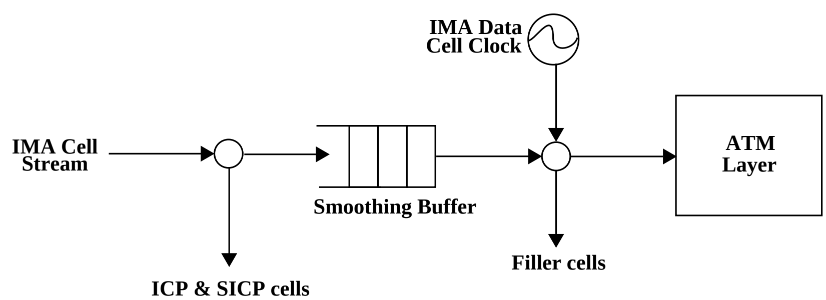

Figure 9 Behavior of the IMA Receive End

(R-61) The CDV attributed to the presence of ICP cells shall be removed by a behavior equivalent to providing a small smoothing buffer into which cells are placed after reordering and after removing ICP cells (including SICP cells), but not Filler cells.

If the smoothing buffer is physically realized, the handling of overflow and underflow conditions, if any, are implementation dependent.

(R-62) If the TRL is in the Working state (as defined in Figure 20 on page 70), the Rx IDCR made available to the ATM layer of the receiver shall be derived as specified in equation (EQ. 1) using the incoming link indicated by the FE IMA transmitter as the TRL in the "Transmit Timing Information" field in the ICP cell (see Table 2 on page 29). However, the receiver behavior is implementation specific during a short period (up to 100 milliseconds) after receiving the far-end indication of a change of TRL to a different link.

If the IMA Frame Synchronization mechanism is not in the Working state, the recovery of the transmit IDCR is implementation specific until the Working state of that link is recovered or another TRL is selected by the transmitter.

Appendix B.5 on page 114 gives an example of the recovery of the IDCR at the receive end.

(R-63) Zero or one cell only shall be made available to the ATM layer at an IMA data cell clock tick. The behavior of the IMA receiver shall be equivalent to following: when the IMA data cell clock at the receiver ticks, one cell shall be removed from the smoothing buffer. If the cell is a Filler cell, then the Filler cell shall be discarded and nothing passed to the ATM layer. If the cell is not a Filler cell, then it shall be passed to the ATM layer.

If there are no cells available (e.g., in the smoothing buffer), then the behavior is implementation dependent. Note that this condition is an anomaly since there should always be a cell available even if it is only a Filler cell.

Note that the IMA cell clock rate will change and be re-adjusted to the new TRL if the current TRL fails or is deleted from the IMA group.

Note that the abstract behavior described on the above requirements assumes that the cell stream, in the order required by the ATM layer, has been reconstructed in the smoothing buffer. During the addition or deletion of a link into/from the IMA group, the IDCR must be recalculated and the new rate applied. There are dependencies on the time at which the IMA cell rate should actually change. The play-out of data that had been received at the old cell rate and that might not have been transmitted to the ATM layer (in particular the ATM layer and Filler cells that are in the delay compensation buffers) should proceed at the data rate at which they were received. This would affect the CDV during transient phases in a link addition/recovery or deletion procedure.

9. Differential Link Delay on IMA

9.1 Differential Link Delay Requirement on Transmitter

(R-64) The transmit IMA shall not introduce more than 2.5 cell times at the physical link rate of differential delay among the constituent links.

The choice of 2.5 cell times at the underlying link rate should be sufficient to accommodate the following:

- a gap of 1 cell time among the inputs to the constituent physical link interface,
- plus 1 cell time to transmit a SICP cell, and
- plus 0.5 cell time for processing.

The IMA stuffing mechanism should be sufficient to ensure that the frames do not wander from alignment by more than the tolerance.

Delay compensation at the receiving end may be performed using the definition of the IMA frame along with ICP cells.

9.2 Differential Link Delay Compensation at Receiver

(R-65) The amount of link differential delay tolerated by an IMA implementation shall be up to at least 25 milliseconds when used over DS1/E1 links.

(O-7) The amount of link differential delay tolerance may be configured up to the maximum value supported by the IMA implementation.

It should be noted that both ends of the IMA virtual link may be configured with different amounts of tolerable differential delay (i.e., say one up to 25 ms and one greater than 25 ms).

Appendix A on page 107 provides a detailed example of link differential delay compensation.

Appendix B.3 on page 111 describes the limitation of the IMA protocol with respect to link differential delay.

10. IMA Interface Operation

This section defines the operation of the IMA individual link and group.

10.1 IMA Link Operation

The IMA link operates independently for both the transmit and receive directions at each end of the IMA virtual link. It allows for a smooth introduction of each link in the round-robin. It also allows graceful handle of error conditions and removal of a link.

A Link State Machine (LSM) is defined for the transmit and receive directions of each IMA link.

10.1.1 Overview of the Link State Machine

Figure 10 on page 44 and Figure 11 on page 45 provide an overview of the link state machine (LSM) in both the transmit and receive directions respectively. A simplified Harel state machine notation [12] is used in Figure 10 and Figure 11. In this notation, the top field of a block identifies the state, while the bottom specifies actions. Arrows denote state transition paths with the event triggering the transition shown as overlaid text. Local events or FE state transitions are used as events. Conditions such as Not In Group and Unusable are locally defined and controlled, no indications received from the FE.

One out of 4 states in each direction can be viewed by the FE over the “Link Information” fields defined in Table 3 on page 30:

- Not In Group: the link is not configured within an IMA group.
- Unusable: link is configured, but not in use due to fault, inhibition, etc.
- Usable: link ready to operate, but waiting for the other end to be Usable or Active
- Active: link actively passing ATM layer cells from/to the ATM layer

The LSM starts in the Not In Group state until the link is configured by the UM. Once configured, the LSM moves out to the Unusable state. From the Unusable state, the LSM provides an operating cycle between the Unusable, Usable and Active states.

The Unusable state provides a synchronization point with any vendor-dependent or group level control of link usability. Link error conditions (e.g., physical or IMA errors) and vendor-dependent conditions (e.g., BER over certain limits) can delay the transition of the LSM into the Usable state or may bring the LSM back to the Unusable state from the Usable or Active state.

The Unusable state also allows the IMA to voluntary delay the restoration of the link for other reason than detected problems. Later sections will refer to this as the inhibition of the link.

The Usable state is an extra state between the Unusable and Active states that allows to coordinate the NE and FE when bringing up the link. It also provides a clear synchronization point before activating the links ready to be set Active.

The Unusable and Usable states are also used as synchronization points between the local Tx and Rx LSMs when the IMA group is operated symmetrically.

Once the Rx LSM is in the Usable state, it can move into the Active only once the FE is reporting that the Tx LSM is Usable. This transition is also subjected to a group-wide synchronization of link activation.

Once the Tx LSM is in the Usable state, it can move into the Active only once the FE is reporting that the Rx LSM is Active. This transition is also subjected to a group-wide synchronization of link activation.

The Tx and Rx LSMs also have the Deleted state to provide a graceful deletion of the link without loss of ATM layer cells.

The Rx LSM also has the Blocking sub-state which is a transitional state toward the Unusable state to provide a graceful de-activation of a link without loss of ATM layer cells.

For symmetrical operation, many of the local events are handled in both Tx and Rx LSMs at the same time.

The next section provides more detail on the Tx and Rx LSM requirements.

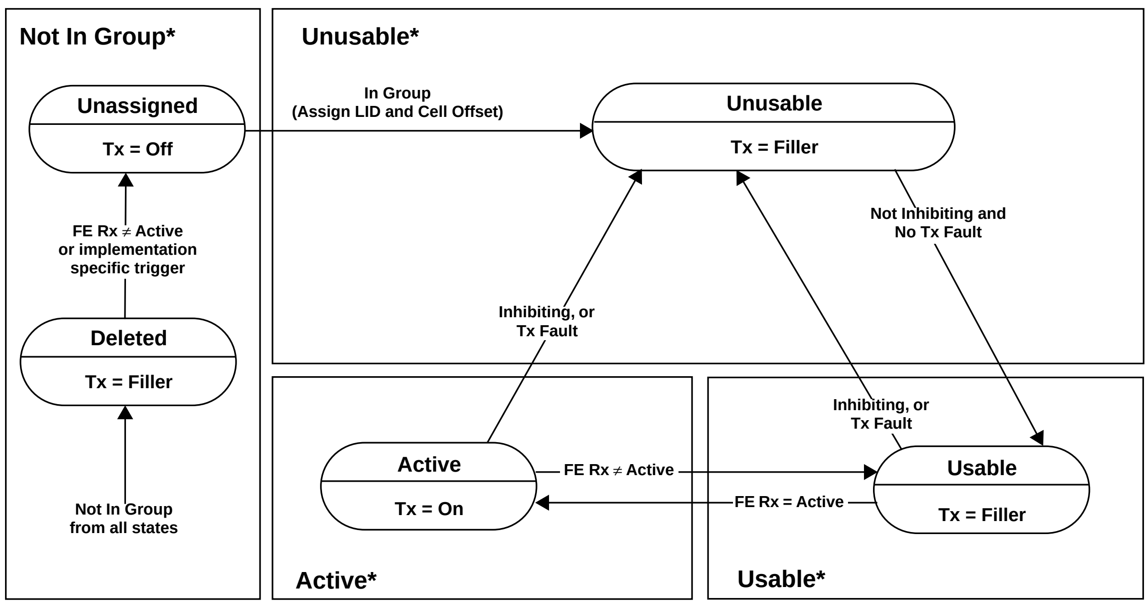

Notes:

- ***: State required to be signalled to the FE.**
- **The Not In Group state is the starting state.**
- **Tx = Off indicates that the transmit end is not expected to be sending IMA frames.**
- **Tx = Filler indicates that the transmit end is sending IMA frames containing only Filler cells.**
- **Tx = On indicates that the transmit ATM layer cell transfer from the ATM layer has been enabled for that link.**
- **FE Rx = Active is a FE state signalled via ICP cells indicating that FE Rx is Active.**
- **FE Rx ≠ Active is a FE state signalled via ICP cells indicating that FE Rx is not Active.**
- **The transition from Unusable to Usable may be delayed due to an ongoing group-wide synchronized start-up or link addition/recovery of links (refer to the group start-up procedure requirement (R-82) and the LASR procedure requirement (R-88)).**
- **The transition from Usable to Active may be delayed due to an ongoing group-wide synchronized start-up or link addition/recovery of links (refer to the group start-up procedure requirement (R-84) and the LASR procedure requirement (R-90)).**
- **Fault: the report of a fault is implementation specific.**
- **See Table 7 on page 47 for detailed definition of transition triggers.**

Figure 10 Transmit Link State Machine Overview

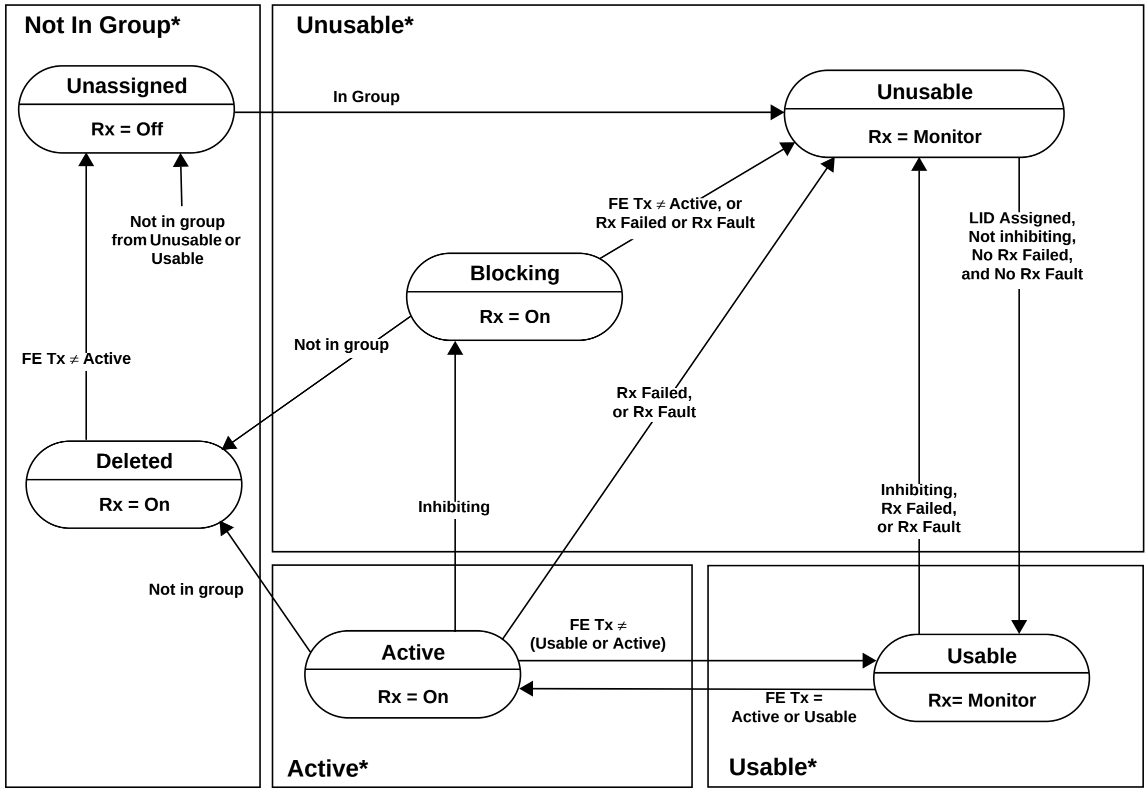

<u>Notes:</u>

- *: State required to be signalled to the FE.
- Not In Group is the starting state.
- Rx = Off indicates that no cells are processed by the receive IMA on that link.
- Rx = Monitor indicates that the receiver is only processing incoming ICP cells on that link.
- Rx = On indicates that the receive ATM layer cell transfer to the ATM layer has been enabled for that link.
- FE Tx = Active is a FE state signalled via ICP cells indicating that FE transmit is Active.
- The transition from Usable to Active may be delayed due to an ongoing group-wide synchronized start-up or link addition/recovery of links (refer to the group start-up procedure requirement (R-83) and the LASR procedure requirement (R-89)).
- FE Tx ≠ Active is a FE state signalled via ICP cells indicating that FE transmit is not Active.
- Fault: a fault in the link or in the protocol; the meaning of this may be implementation specific.
- See Table 7 on page 47 for detailed definition of transition triggers.

Figure 11 Receive Link State Machine Overview

10.1.2 Definitions of the Transmit/Receive States and Events

10.1.2.1 Transmit Link States

The following four transmit link states are defined:

- Not In Group: the link is not configured, there are two sub-states:
  - — Unassigned: no information about the link exists.
  - — Deleted: the link has been removed from the group; this transitional state is to ensure that the other end is still not receiving ATM layer cells before it moves to the Unassigned state.

- Unusable: the link is configured but cannot be used. There are many reasons for this state including:
  - — Test failed: mis-connectivity has been found as a result of a test. The Test Pattern procedure defined in Section 13 on page 77 is an example of test that can be used.
  - — Fault: a fault has been detected both on the link and in the link protocol.
  - — Inhibited: operation of the link is blocked for some locally defined application or implementation dependent reason. The link may otherwise be used.
- Usable: the link is ready to be used, it is awaiting the FE to activate its receiver before sending any ATM layer cells; IMA frames containing only Filler cells are being transmitted; but the link is not in the data round robin.
- Active: the link is transmitting ATM layer cells; it is part of the data round robin.

10.1.2.2 Receive Link States

The following four receive link states are defined:

- Not In Group: the link is not configured; there are two sub-states:
  - — Unassigned: no information about the link exists.
  - — Deleted: the link has been removed from the group; this transitional state is to ensure that the other end is no longer Active (transmitting ATM layer cells) before it moves to Unassigned sub-state.
- Unusable: the link is configured, but cannot be used; there are many reasons for this including:
  - — Idle: insufficient information to start link establishment
  - — Fault: a fault has been detected both on the link and in the link protocol.
  - — Failed: the receiver has failed due to the persistence of a defined defect. Examples of defects are LCD, LIF, and LODS. These are described in more details in Section 12.1 on page 69.
  - — Inhibited: operation of the link is blocked for some locally defined application or implementation dependent reason. The link may otherwise be used.

  The sub-state “Blocking” has been defined under the Usable to allow graceful transition into the Unusable state without loss of ATM layer cells.
- Usable: the link is ready to be used, and it is awaiting the FE Tx to be Usable or Active.
- Active: the receive IMA is receiving ATM layer cells on this link; the receive IMA considers this link as part of the data round robin.

10.1.3 Definitions of the Transmit/Receive Link Actions

Table 5 on page 46 and Table 6 on page 47 shows the actions that should be performed by the LSM in the transmit and receive directions respectively.

Table 5 Actions at the Transmit End

Action	Meaning
Tx = Off	Indicates that the transmit end is not expected to be sending IMA frames on that link.
Tx = Filler	Indicates that the transmit end is sending IMA frames containing only Filler cells on that link.
Tx = On	Indicates that the transmit ATM layer cell transfer from the ATM layer has been enabled for that link.

Table 6   Actions at the Receive End

Action	Meaning
Rx = Off	Indicates that no cells are processed by the receive IMA on that link.
Rx = Monitor	Indicates that the receiver is only processing incoming ICP cells on that link.
Rx = On	Indicates that the receive ATM layer cell transfer to the ATM layer has been enabled for that link.

10.1.4   Link Events Driving the Link State Machines

Table 7 on page 47 defines the local events driving the Tx and Rx LSMs. Depending on the symmetry of the configuration, some events would be input to both Tx and Rx state transitions.

Table 7   Local Events Driving the Transmit and Receive State Machines

Event	Meaning	Fully Symmetric	Asymmetric Operation	Asymmetric Configuration
In Group	The assignment of a link in an IMA group by the UM. This corresponds to the moment the link is configured.	Both	Both	One only
Not in Group	The removal of an IMA link from a group by the UM.	Both	Both	One only
LID Assigned	Sufficient information for the link to begin communication.	Rx only	Rx only	Rx only
Rx Failed	Condition entered after the persistent detection of a defect at the receiver (see Section 12.1.3 on page 69 and Section B.2 on page 110). The criteria for entering the condition is implementation specific.	One only	One only	One only
No longer Rx Failed	No longer in the Rx Failed condition.	One only	One only	One only
Fault	Entering the Fault condition. A fault in the link or in the protocol. The meaning of this may be implementation dependent.	Vendor dependent	Vendor dependent	One only
No longer Fault	No longer in the Fault condition.	Vendor dependent	Vendor dependent	One only
Inhibiting	The link is not to be used for some reason (implementation and application specific). An inhibiting action on a link in the Active state is performed without loss of ATM layer cells.	Both	One only	One only
No longer inhibiting	Inhibiting reason removed.	Both	One only	One only
Test related	These are possible ways that the testing could be performed.	Probably Tx only (see note 1)	Probably Tx only (see note 1)	Probably Tx only (see note 1)

Note 1.   Only the transmit end is expected to perform the connectivity test.

The FE LSM states reported over the incoming ICP cells are used to drive the NE LSM.

10.1.5   Link State Transition Tables

Table 8 on page 49 and Table 9 on page 51 show the transition from the current to the new LSM state and the corresponding actions to be executed upon the occurrence of any possible events.

(R-66) The Tx LSM shall be controlled using the state machine transition illustrated in Table 8 on page 49.

(R-67) The Rx LSM shall be controlled using the state machine transition illustrated in Table 9 on page 51.

Table 8 on page 49 and Table 9 on page 51 shall be interpreted according to the following rules:

- The first row presents the current state of the LSM. The current state shall be signalled to the FE via the ICP cells.
- The second row presents the sub-state of the LSM. The sub-states are not signalled, although the corresponding actions shall apply.
- The first column presents the events that trigger state transition. Four group of events are defined. Events may appear simultaneously and shall be treated sequentially. The order is implementation specific.

Table 8 Detailed Transmit Link State Machine

Primary NE Tx State	**Not In Group**		**Unusable**	**Usable**	**Active**
Sub-State	**Unassigned (Tx=Off)**	**Deleted (Tx=Filler)**	**Unusable (Tx=Filler)**	**Usable (Tx=Filler)**	**Active (Tx=On)**
NE Tx Events (see note 1)					
In group	Unusable: Tx = Filler (see note 2)	--	--	--	--
Not in Group	-- (see notes 3,4)	--	Deleted: Tx = Filler	Deleted: Tx = Filler	Deleted: Tx = Filler
Fault	N/A (see note 5)	--	--	Unusable: Tx = Filler	Unusable: Tx = Filler
No longer Fault	N/A	--	Unusable or Usable (see note 6)	N/A	N/A
Inhibiting	--	--	--	Unusable: Tx = Filler	Unusable Tx = Filler
No longer inhibiting	--	--	Unusable or Usable (see note 6)	--	--
FE Rx States Received in ICP Cell					
Not In Group	N/A	Unassigned: Tx = Off	--	--	Usable: Tx = Filler
Unusable	N/A	Unassigned: Tx = Off	--	--	Usable: Tx = Filler
Usable	N/A	Unassigned: Tx = Off	--	--	Usable: Tx = Filler
Active	N/A	--	--	Active or Usable (see note 7)	--
NE Rx State (see note 8)					
Not in Group	--	--	N/A*	N/A*	N/A*
Unusable	N/A* (see note 9)	N/A*	--	--	Usable: Tx = Filler
Usable	N/A*	N/A*	--	--	--
Active	N/A*	N/A*	--	--	--
FE Tx State Received in ICP Cell (see note 7)					
Not in Group	--	--	N/A*	N/A*	N/A*
Unusable	N/A*	N/A*	--	--	Usable: Tx = Filler
Usable	N/A*	N/A*	--	--	--
Active	N/A*	N/A*	--	--	--

Note 1. If NE Tx and Rx events and FE Tx and Rx events appear simultaneously they shall be treated sequentially. The order is implementation specific.

Note 2. The LID and ICP Cell Offset shall also be assigned at this time.

Note 3. "--" means the event is ignored in this state.

Note 4. The LID is assigned, the link goes to Unusable; it may subsequently be tested, or go to Usable, or stay in Unusable because of a fault, or some inhibiting actions.

Note 5. N/A indicates that this event should not happen, but may be ignored otherwise.

Note 6. Proceed to Usable with Tx=Filler, only if there are no reasons to remain in Unusable, i.e., no Tx fault and not inhibiting.The transition to Usable may be delayed due to an ongoing group-wide synchronized start-up or link addition/recovery of links (refer to the group start-up procedure requirement (R-82) and the LASR procedure requirement (R-88)).

Note 7. For symmetric configuration and operation, proceed to Active with Tx = On, only if NE Rx state is Active or Usable and the FE Tx is Active or Usable. This ensures that this direction is only used if the other direction is also Usable. Under asymmetric operation (see Appendix C on page 117) proceed to Active with Tx = On. The transition to Active with Tx=On may be delayed due to an ongoing group-wide synchronized start-up or link addition/recovery of links (refer to the group start-up procedure requirement (R-84) and the LASR procedure requirement (R-90)).

Note 8. These states are interpreted only if the group is set in the “Symmetric Configuration and Operation” mode.

Note 9. N/A* indicates that for a symmetrically operated group this should not happen, but may be ignored otherwise.

Table 9 Detailed Receive Link State Machine

Primary NE Rx State	Not In Group		Unusable		Usable	Active
Sub-State	**Unassigned (Rx = Off)**	**Deleted (Rx = On)**	**Unusable (Rx = Monitor)**	**Blocking (Rx = On)**	**Usable (Rx = Monitor)**	**Active (Rx = On)**
NE Rx Events (see note 1)						
In Group	Unusable: Rx = Monitor (see note 2)	--	--	--	--	--
Not In Group	-- (See note 3)	--	Unassigned Rx = Off	Deleted: Rx = On	Unassigned: Rx = Off	Deleted: Rx = On
Failed Condition	N/A (see note 4)	--	--	Unusable: Rx = Monitor	Unusable: Rx = Monitor	Unusable: Rx = Monitor
No longer Failed	N/A	--	Unusable or Usable (see note 5)	--	--	--
Fault Condition	N/A	--	--	Unusable: Rx = Monitor	Unusable: Rx = Monitor	Unusable: Rx = Monitor
No longer Fault	N/A	--	Unusable or Usable (see note 5)	--	--	--
Inhibiting	--	--	--	--	Unusable: Rx = Monitor	Blocking: Rx = On
No longer inhibiting	--	--	Unusable or Usable (see note 5)	--	--	--
FE Tx States received in ICP cell						
Not In Group	N/A	Unassigned: Rx = Off	--	Unusable: Rx=Monitor	--	Usable: Rx=Monitor
Unusable	N/A	Unassigned: Rx = Off	--	Unusable: Rx=Monitor	--	Usable: Rx=Monitor
Usable	N/A	Unassigned: Rx = Off	--	Unusable: Rx=Monitor	Usable or Active (see note 6)	Active or Usable (see note 7)
Active	N/A	--	--	--	N/A	--
NE Tx State (see note 8)						
Not in Group	--	--	N/A*	N/A*	N/A*	N/A*
Unusable	N/A* (see note 9)	N/A*	--	--	--	Usable
Usable	N/A*	N/A*	--	--	--	--
Active	N/A*	N/A*	--	--	--	--
FE Rx States received in ICP cell (see note 8)						
Not in Group	--	--	N/A*	N/A*	N/A*	N/A*
Unusable	N/A*	N/A*	--	--	--	--
Usable	N/A*	N/A*	--	--	--	--
Active	N/A*	N/A*	--	--	--	--

Note 1. If NE Tx and Rx events and FE Tx and Rx State events appear simultaneously they shall be treated sequentially. The order is implementation specific.

Note 2. The initial state will be Unusable, until the LID has been assigned, after which if there are no reasons to remain Unusable the link will proceed to Usable.

Note 3. "--" means the event is ignored in this state.
Note 4. N/A indicates that this event should not happen, but may be ignored otherwise.
Note 5. Process to Usable with Rx = Monitor, only if there are no reasons to remain in Unusable, i.e., the LID has been assigned, there are no failures, no fault and not inhibited.
Note 6. For symmetrical configuration and operation, proceed to Active with Rx = On, only if NE Tx state is Active or Usable. Under asymmetric operation (see Appendix E on page 116) proceed to Active with Rx = On. The transition to Active with Rx=On may be delayed due to an ongoing group-wide synchronized start-up or link addition/recovery of links (refer to the group start-up procedure requirement (R-83) and the LASR procedure requirement (R-89)).
Note 7. If the IMA group is operating symmetrically then; if the NE Tx state is Unusable, proceed to "Usable" with Rx = Monitor, otherwise remains Active with Rx = On.
Note 8. These are interpreted only if the group is set in the "Symmetrical Configuration and Operation" mode.
Note 9. N/A* indicates only that for groups set in the "Symmetrical Configuration and Operation" mode this should not happen, but may be ignored otherwise.

10.1.6 Processing of the Tx and Rx State Fields

(R-68) The encoding of the "Tx State" and "Rx State" fields of the Link Information field (Table 3 on page 30) shall be transmitted according to the link state at the time the ICP cell is constructed. Any change in any link operation status shall be reflected in the ICP cell within the next 2M cells on that link and the SCCI field shall be set for the first ICP cell to reflect a change.

Ongoing group-wide synchronized start-up or link addition/recovery of links may enforce to the report of multiple LSMs at the same time. This is covered in more details in Section 10.2.1 on page 53 and Section 10.2.2 on page 62.

(R-69) The IMA shall use one of the Unusable encoding when reporting the Unusable state. At a minimum, the "No reason given" shall be used.

(O-8) The IMA may optionally use "Inhibited", "Failed", "Fault" or "Mis-connected" as a reason when reporting the Unusable state.

(R-70) Each processed incoming ICP cell with new state indication shall cause re-evaluation of the Tx and Rx LSMs operation state.

Table 10 on page 52 shows the valid combinations of the Tx and Rx state values for all symmetrical and asymmetrical configurations and operations.

(R-71) Each end shall only use the valid combination of Tx and Rx states as specified in Table 10 on page 52.

Table 10 Valid Combinations of the Tx and Rx states

Tx state	Rx state	Symmetrical configuration and operation	Symmetrical configuration, asymmetrical operation	Asymmetrical configuration and operation
Not In Group	**Not In Group**	Valid	Valid	Valid
Not In Group	**Unusable**	Not Valid	Not Valid	Valid
Not In Group	**Usable**	Not Valid	Not Valid	Valid
Not In Group	**Active**	Not Valid	Not Valid	Valid
Unusable	**Not In Group**	Not Valid	Not Valid	Valid
Unusable	**Unusable**	Valid	Valid	Valid
Unusable	**Usable**	Valid	Valid	Valid

Table 10 Valid Combinations of the Tx and Rx states

Tx state	Rx state	Symmetrical configuration and operation	Symmetrical configuration, asymmetrical operation	Asymmetrical configuration and operation
Unusable	**Active**	Not Valid	Valid	Valid
Usable	**Not In Group**	Not Valid	Not Valid	Valid
Usable	**Unusable**	Valid	Valid	Valid
Usable	**Usable**	Valid	Valid	Valid
Usable	**Active**	Valid	Valid	Valid
Active	**Not In Group**	Not Valid	Not Valid	Valid
Active	**Unusable**	Not Valid	Valid	Valid
Active	**Usable**	Not Valid	Valid	Valid
Active	**Active**	Valid	Valid	Valid

Section D on page 119 presents examples of link state exchanges between both ends.

10.2 IMA Group Operation

This section defines the operation of the IMA group. The operation of the IMA group is governed by the Group State Machine (GSM), the Group Traffic State Machine (GTSM), and by the Link Addition and Slow Recovery (LASR) procedure. These three items are used to ensure reliable transmission and reception of ATM layer cells across all links in the Active state. This includes the negotiation of group parameters (i.e., symmetry and M values), the bringing up of the IMA group, and the graceful addition/recovery and deletion of links to and from the group.

10.2.1 IMA Group State Machine Definition

10.2.1.1 Overview of the Group State Machine

An overview of the GSM is shown in Figure 12 on page 54. The GSM starts in the Not-Configured state from which it goes into the Start-up state. This transition corresponds to the time the UM configures the IMA group. In the Start-up state, the IMA proposes group parameters to the FE and waits to determine if the proposed parameters are accepted by the FE. This is done by monitoring the “Group Status and Control” field in the incoming ICP cells. The FE will indicate if it accepts the proposed group parameters by returning Start-up-Ack. The FE will return Config-Aborted if it rejects the parameters.

While in the Start-up state, the local IMA also has to determine if it can accept the group parameters proposed by the FE. If it accepts the proposed parameters it goes into the “Start-up-Ack state. If it rejects the proposed parameters, it goes into the Config-Aborted state.

Once in the Start-up-Ack state, the local IMA now waits for one of the signals specified in Figure 12 to move to the “Insufficient-Links” state. At this point, the GSM starts paying attention to the individual LSMs. It stays in the Insufficient-Links state until it gets the amount of links required to be Active. This number is specified in term of P_{Tx} and P_{Rx}. P_{Tx} represents the number of links required to be Active in the transmit direction. P_{Rx} represents the number of links required to be Active in the receive direction. P_{Tx} and P_{Rx} are equal if the group is operating symmetrically. It is required that the GSM receives indications from the individual LSMs about the status of the links.

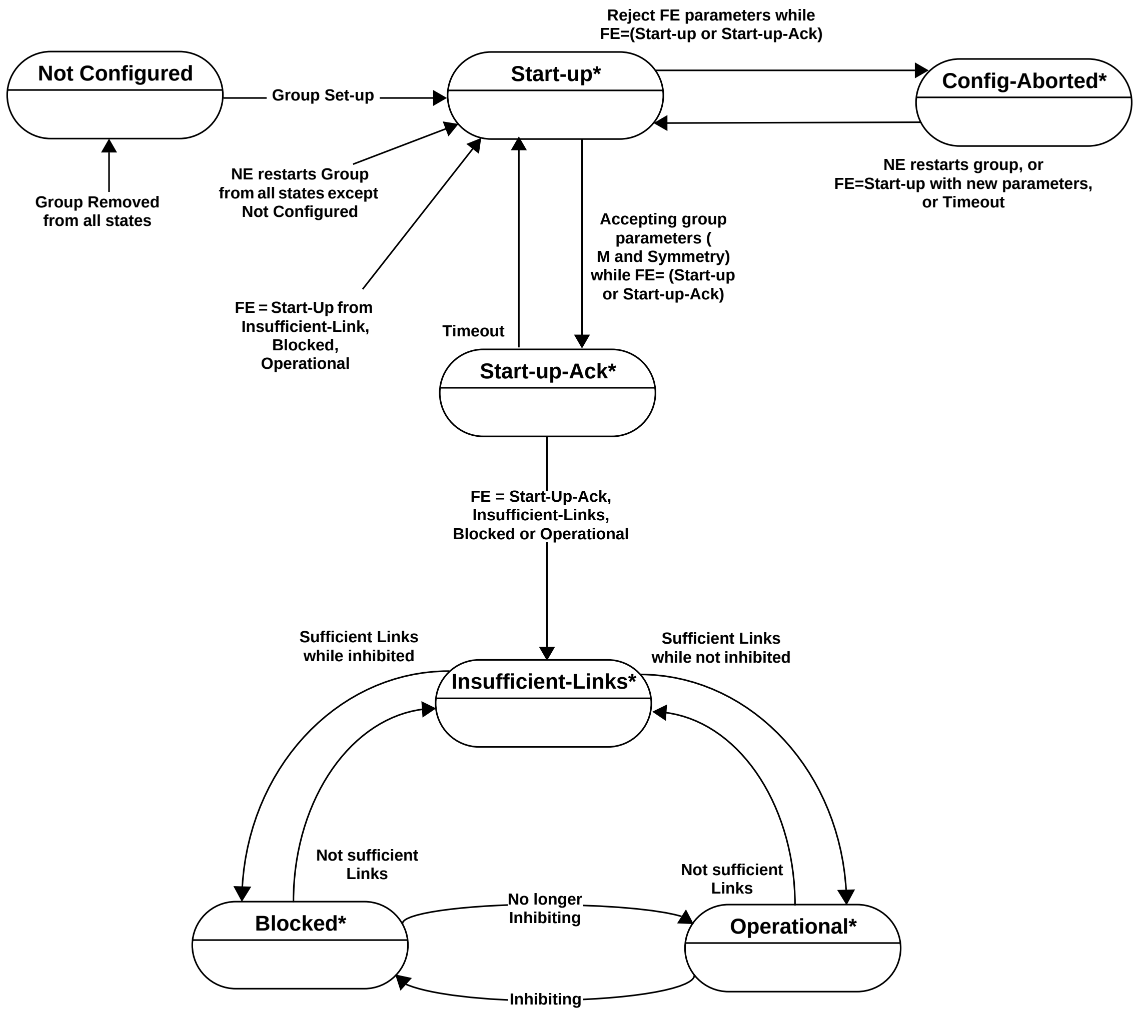

Notes:

- *: Group state required to be signalled.
- Config Aborted State: To avoid a race hazard, the NE GSM should remain in this state for at least one second; this is to allow for at least one round trip of communication between both ends on the IMA virtual links.
- Start-up-Ack: The NE GSM should wait for at least one second before moving to Start-up. This return to Start-up occurs if one of the expected signals which would cause the move to Insufficient-Links is not received from the FE GSM.
- Sufficient Links transition is defined as:
  For symmetric operation: at least P_{Tx} links Active in both directions, where P_{Tx} and P_{Rx} are equal.
  For asymmetric operation: at least P_{Tx} transmit links Active and P_{Rx} receive links Active.

Figure 12 Overview of the Group State Machine

Once the group reaches the amount of links required to be Active, it moves into the Operational state if it has not been inhibited. Once in the Operational state, the group is allowed to receive ATM cells and passes them to the ATM layer. It also requires to receive the Operational state signal from the FE before starting transmitting cells passed from the ATM layer. Figure 13 on page 55 illustrates the Group Traffic State Machine

(GTSM) that indicates the capability of the group to transmit cells from the ATM layer. ATM layer cells can be only transmitted by the NE when both the NE and FE IMA groups are in the Operational state.

The GTSM allows the NE to ensure that both ends have sufficient links before starting transmitting cells to the FE.

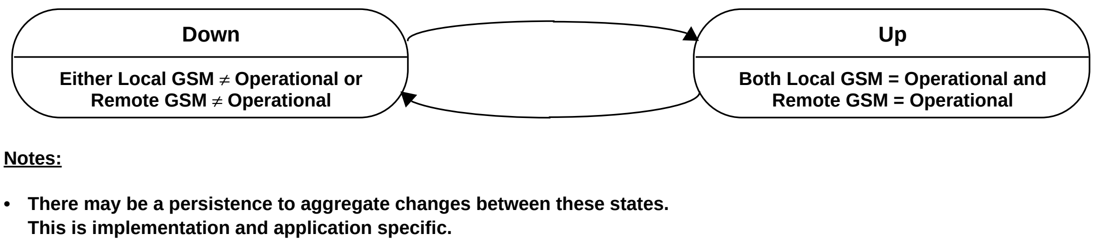

Figure 13 Group Traffic State Machine

Figure 12 on page 54 also shows the presence of 2 time-outs. One time-out is used to allow the GSM to return from the Config-Aborted to the Start-up state to indicate to the FE that it is ready to propose and evaluate new group parameters. A second time-out is used to return from Start-up-Ack to Start-up when the NE is not receiving one of the FE signals required to transition to Insufficient-Links.

10.2.1.2 Detailed Group State Machine Definition

This section defines the characteristics of the GSM. Table 11 on page 55 describes the state of the GSM.

Table 11 Description of GSM States

State	Description
Not Configured	The group does not exist
Start-up	This end is in start-up, and is waiting to see the FE in start-up. When sufficient communication with the FE is achieved, the Group Parameters (M, Symmetry) are recorded and the group moves to Start-up-Ack.
Start-up-Ack	This is a transitional state, when both groups start-up, they then move through this state to the Insufficient-Links state. While in this state it ignores an indication from the FE that it is in Start-up.
Config-Aborted	This state is entered when the FE tries to use unacceptable configuration parameters.
Insufficient-Links	This state implies that this end does not have “Sufficient Links”. There are two definitions of “Sufficient Links” depending on group symmetry. Sufficient Links is defined as: For symmetric operation: at least P_{Tx} links are Active in both directions (P_{Tx}=P_{Rx}). For asymmetrical operation: at least P_{Tx} links are Active in the transmit direction and at least P_{Rx} links are Active in the receive direction.
Blocked	The group is blocked (e.g., inhibited by the UM). The group can be blocked for maintenance purposes while sufficient links are Active in both directions.
Operational	The group is not inhibited and has sufficient links in both the Tx and Rx directions. The IMA has now the capability to receive ATM layer cells and passes to the ATM layer from the IMA sublayer.

(R-72) The states listed in Table 11 on page 55, with the exception of the Not Configured state, shall be reported to the FE group using the value defined in the “Group Status and Control” field in Table 2 on page 29

The values P_{Tx}, and P_{Rx}, mentioned in Table 11 on page 55, are application specific and may be configurable parameters.

(R-73) P_{Tx}, and P_{Rx} shall be greater than 0.

(R-74) The transmit end shall send over each link the same value in the “Group Status and Control” field (see Table 2 on page 29) for at least 2 consecutive IMA frames.

This is required in order to ensure that the FE is properly informed.

(R-75) The GSM shall validate the Rx M and the Rx IMA ID over at least one link before moving the Start-up-Ack state. The validated Rx M and IMA ID shall be used by the IMA Frame Synchronization mechanism defined in Section 11 on page 65.

Other group fields in the ICP cells do not have to be validated over all links once the received M and IMA ID have been validated.

Table 12 on page 56 lists and describes the local events causing GSM state transitions.

Table 12 Local Events Driving the GSM

Event	Description
Group Setup	The establishment of the IMA Group by the UM.
Group Removed	Removal of the IMA Group by the UM.
Acceptance of FE Group Parameters	All the FE group parameters are acceptable while the FE GSM is either in the Start-up or Start-up-Ack state.
Rejection of FE Group Parameters	One or more of the FE group parameters are unacceptable. For example the proposed M value is not supported or the requested group symmetries are different.
Inhibiting	An implementation specific shut down of the group for a reason other than insufficient links (e.g., by the UM).
No longer inhibiting	Removal of the inhibiting condition.
Sufficient Links	See the description of the Insufficient-Links state in Table 11 on page 55.
Not Sufficient Links	See the description of the Insufficient-Links state in Table 11 on page 55.
Events that cause a new start-up attempt from Config-Aborted	This can be driven by many causes, a new local configuration, a realization that the FE is attempting start-up with a different configuration, or simply a timeout to retry. These are implementation specific.

(R-76) If the requested configuration is unacceptable, the IMA shall indicate this using the Config Aborted state for at least 1 second. This is to allow for at least one round trip of communication between the ends (4M cells plus round trip time).

(R-77) The IMA shall support the GSM state transitions as defined in Table 13 on page 57.

Table 13 shows the state transition of the GSM. The first row lists the current state, and the first column indicates the event triggering state change. Two categories of events (NE events and FE states) are defined. If events occur simultaneously, they shall be treated sequentially. The order is implementation specific.

Table 13 GSM State Transition

NE State	Not Configured (see note 1)	Start-up	Start-up-Ack	Config-Aborted	Insufficient-Links	Blocked	Operational
NE Events							
Group Setup	Start-up	--	--	--	--	--	--
Group Removed (see note 2)	--	Not Configured	Not Configured	Not Configured	Not Configured	Not Configured	Not Configured
Acceptance of FE Group Parameters	--	Start-up or Start-up-Ack (see note 3)	N/A (see note 4)	N/A	N/A	N/A	N/A
Rejection of FE Group Parameters	--	Config-Aborted	Config-Aborted	--	N/A	N/A	N/A
Inhibiting	--	--	--	--	--	--	Blocked
Not Inhibiting	--	--	--	--	--	Operational	--
Sufficient links	--	--	--	--	Blocked or Operational (see note 5)	--	--
Not sufficient links	--	--	--	--	--	Insufficient-Links	Insufficient-Links
Local Restart (see note 6)	--	Start-up	Start-up	Start-up	Start-up	Start-up	Start-up
Timeout (see note 7)	--	--	Start-up (see note 8)	Start-up (see note 9)	--	--	--
FE State signalled via ICP cell							
Start-up	--	Start-up-Ack or Config-Abort (see note 10)	--	Start-up or Config-Abort (see note 11)	Start-up	Start-up	Start-up
Start-up-Ack	--	Start-up-Ack or Config-Abort (see note 10)	Not Configured	--	--	--	--
Config-Aborted	--	--	--	--	--	--	--
Insufficient-Links	--	--	Insufficient-Links	--	--	--	--
Blocked	--	--	Insufficient-Links	--	--	--	Operational
Operational	--	--	Insufficient-Links	--	--	--	Operational (see note 12)

Note 1. Starting state.
Note 2. The Group Removed events shall be used only when the GSM is in the Start-up, Config-Abort and Insufficient-Links state.
Note 3. Proceed to Start-up-Ack if FE=(Start-up-Ack or Start-up), else remain Start-up.
Note 4. N/A (Not Applicable) indicates that this event should not happen, but may be ignored otherwise.
Note 5. If inhibiting go to Blocked state, otherwise to Operational state.
Note 6. The local restart shall be used only when the GSM is in the Start-up, Config-Abort and Insufficient-Links state.
Note 7. Time-outs are only meaningful for Start-up-Ack (see note 8) and Config-Aborted (see note 9).
Note 8. This time-out shall be at least one second. This will provide enough time for the FE GSM to notice that the NE GSM is in Start-up-Ack and to change to Start-up-Ack itself.
Note 9. When the GSM aborts due to unacceptable configuration, it shall remain in the Config-Aborted state for at least one second (at least one round delay between both ends).
Note 10. If the new parameters are acceptable proceed to Start-up-Ack, otherwise proceed to Config-Abort.
Note 11. If receiving FE=Start-up with new group parameters, then proceed to Start-up, else remain Config-Abort.
Note 12. The GTSM is up in this combination of states, down otherwise.

(R-78) Each IMA interface shall independently determine whether the group is up or down with respect to carrying traffic as in Figure 13 on page 55 and report the result according to Table 14 on page 58.

(R-79) When the GTSM is in the Down state, it shall be reported to the UM and ATM layer management. This shall be the only notification to the ATM layer (management) about Physical Layer defect or failure (other notification to the UM are possible).

Table 14 GTSM States and Corresponding Report Priority

GTSM State	Reasons	Report Priority
Up		
Down	FE Start-up	Most important
Down	NE Aborted (with given reason)	
Down	FE Aborted (with given reason)	
Down	NE Insufficient Links	
Down	FE Insufficient Links	
Down	NE Blocked	
Down	FE Blocked	Least important

(R-80) The IMA interface shall not drop any ATM layer cells when adding or recovering links while the GSM is maintained in the Operational state.

(R-81) The IMA interface shall not drop any ATM layer cells when deleting or inhibiting links while the GSM is maintained in the Operational state.

When links are being deactivating due to link failure or link deletion, the bandwidth of the IMA group decreases. Section B.1 on page 110 explains the possible actions to be performed in this case.

10.2.1.3 Synchronized Link Activation during Group Start-up Procedure

This section defines extra requirements to synchronize the activation of the links during the Group Start-up procedure.

Once the NE IMA has determined which links exhibit acceptable differential link delay, it is now ready to activate the links in both directions. It is important to activate the links in the receive direction at the same time and then later activate the links in the transmit direction at the same time to minimize the number of bandwidth changes on the IMA virtual link.

(R-82) The group start-up procedure shall ensure that all accepted links have their states changed to Tx=Usable over the same advancing of the SCCI field.

(R-83) The group start-up procedure shall ensure that all accepted links have their states changed to Rx=Active over the same advancing of the SCCI field.

(R-84) The group start-up procedure shall ensure that all accepted links have their states changed to Tx=Active over the same advancing of the SCCI field.

(R-85) The GSM shall wait a minimum of one second, unless all the configured links are being reported Tx=Usable by FE, before reporting links Rx=Active.

(R-86) The GSM shall wait a minimum of one second, unless all the configured links are being reported Rx=Active by FE, before reporting links Tx=Active.

Figure 14 on page 60 illustrates the detailed behavior at each end of the IMA virtual link during group start-up.

The flow of the start-up procedure over time is presented in Figure 15 on page 61. The figure highlights the key check points required to be passed before moving further in the process. The links become part of the sequencing round robin (RR) when the transmitting end has allocated them a LID and the receiving end has started synchronizing them (delay compensation for accommodating longest propagation delay links).

Appendix C on page 117 details the start-up procedure when using a different value for P_{Tx} and P_{Rx} in the case of optional asymmetrical operations.

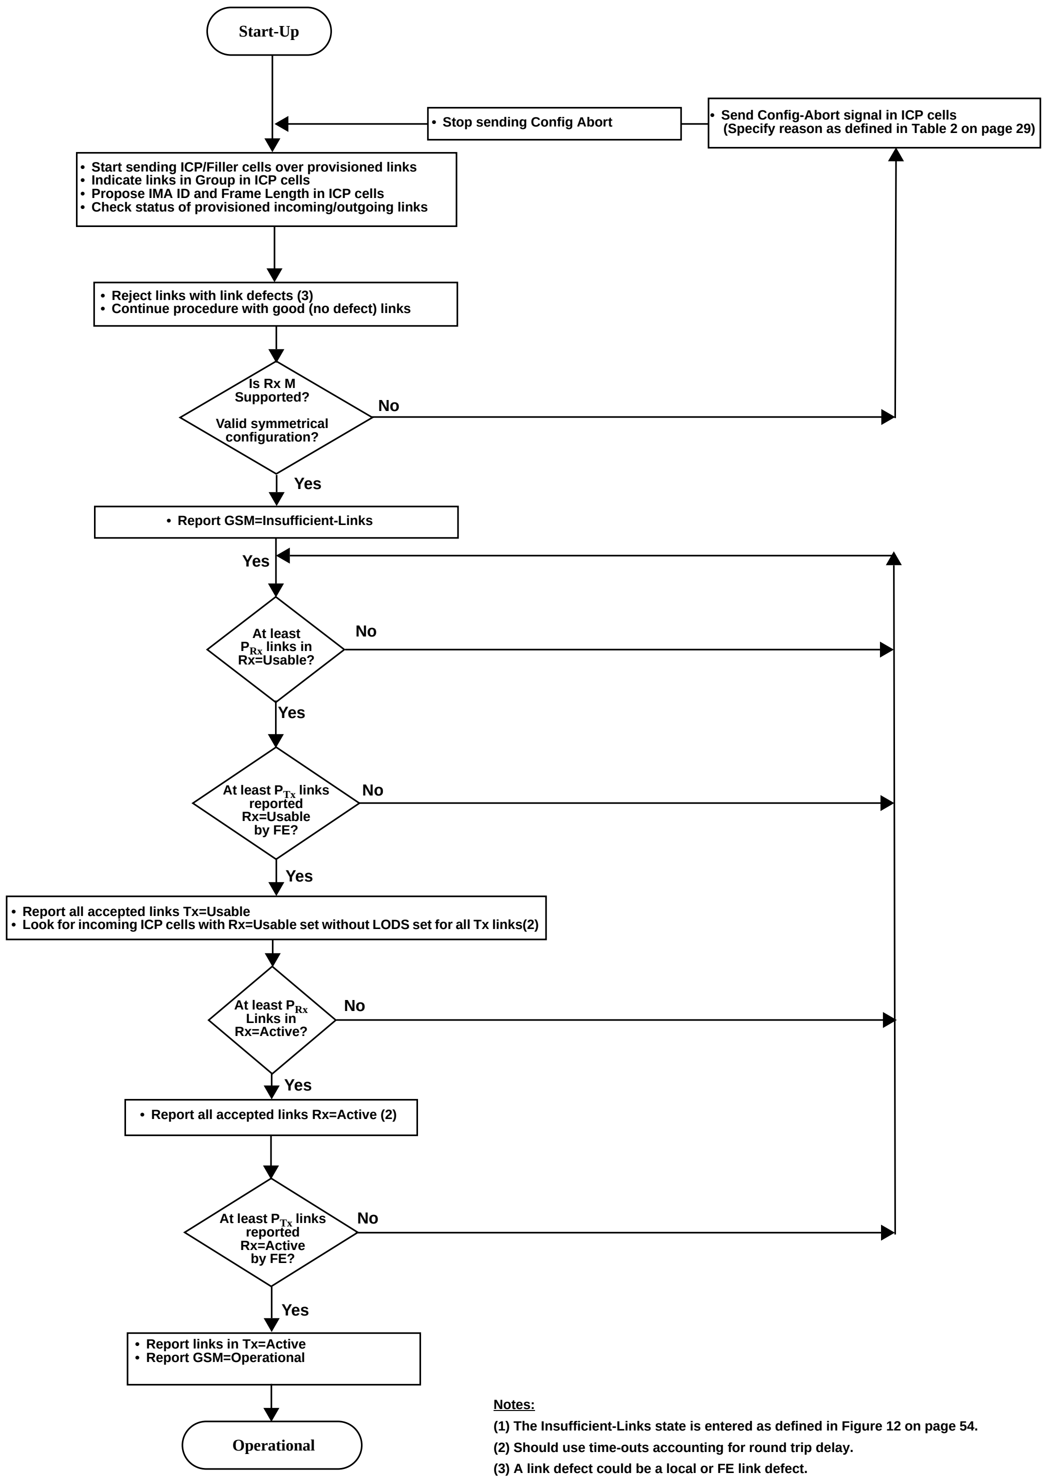

Figure 14 Example of Synchronized Link Activation during Group Start-Up

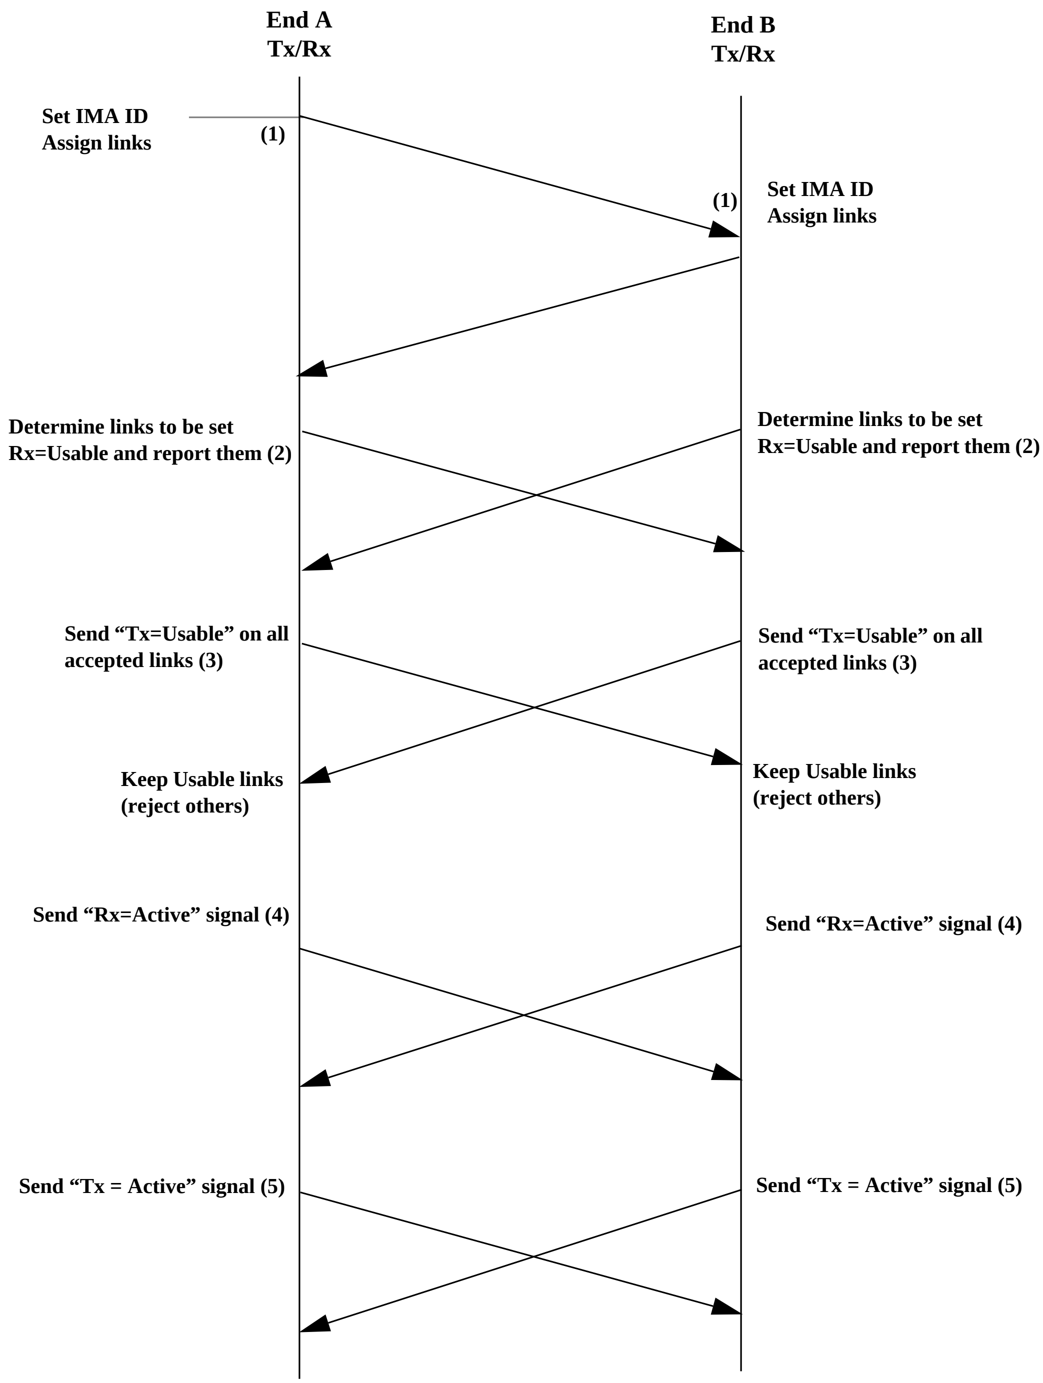

(1) Each link is assigned a LID
(2) At this point, links reported Rx=Usable are those that have been recognized within the IMA group, that do not exhibit any defects and exhibit proper link differential delay.
(3) All links are reported Tx=Usable at the same time.
(4) All links are reported Rx=Active at the same time.
(5) The transmit links can be activated once they all have been recognized Active in the receive direction or after some time-out has expired and P out of N of the receive links have been recognized Active. Links are reported Tx=Active at the same time.

Figure 15 Example of Synchronized Link Activation during Group Start-up over Time

10.2.2 Link Addition and Slow Recovery (LASR) Procedure

The Link Addition and Slow Recovery (LASR) procedure is defined to synchronize the insertion of new or recovered links within the IMA RR. Figure 16 gives an example of the group and link actions at each end of the IMA virtual link. It is required that each end of the IMA virtual link process the LASR procedure in order to complete the addition and activation of the new link.

(R-87) Only one LASR procedure per IMA group shall be executed at any time. The link addition shall be used to add one, or more than one link at the same time.

(R-88) The LASR procedure shall ensure that all new or recovered links have their states changed to Tx=Usable over the same advancing of the SCCI field.

(R-89) The LASR procedure shall ensure that all new or recovered links have their states changed to Rx=Active over the same advancing of the SCCI field.

(R-90) The LASR procedure shall ensure that all new or recovered links have their states changed to Tx=Active over the same advancing of the SCCI field.

(R-91) The GSM shall wait a minimum of one second, unless all the new or recovered links are being reported Tx=Usable by FE, before reporting links Rx=Active.

(R-92) The GSM shall wait a minimum of one second, unless all the new or recovered links are being reported Rx=Active by FE, before reporting links Tx=Active.

The LASR procedure also applies when re-activating a link.

(R-93) If the LASR procedure is in progress, the addition of one or more new links triggered by the UM or a possible slow link recovery (as defined in Section 12.1.3 on page 69) shall be delayed until the link addition procedure is completed or aborted (link addition capability is enabled once again as shown in Figure 16 on page 63).

(R-94) The execution of the LASR procedure over time is shown in Figure 17 on page 64.

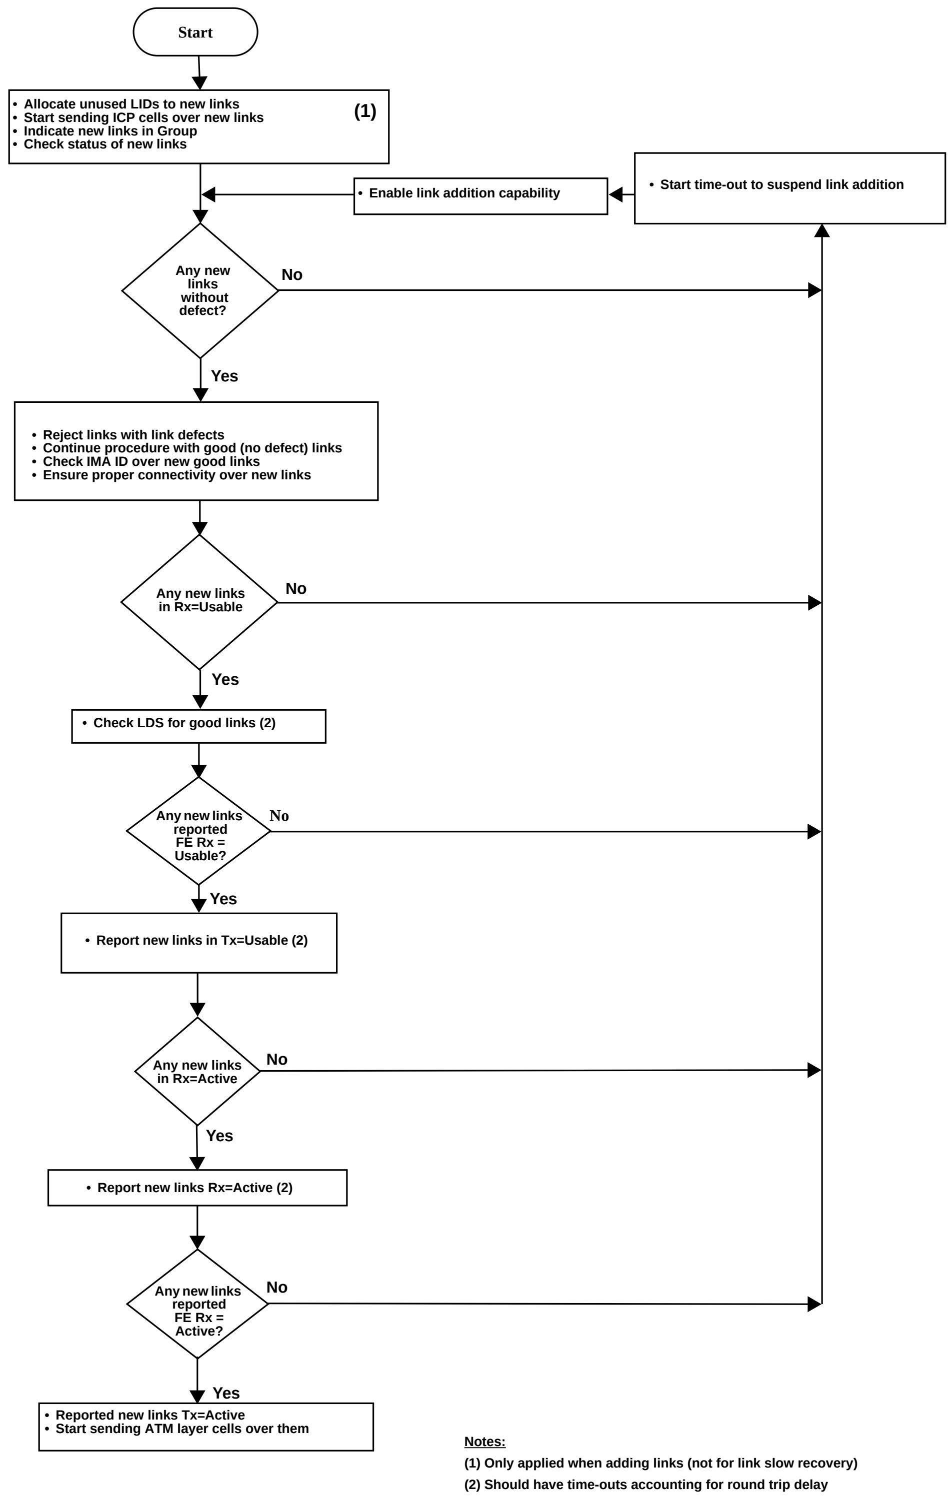

Figure 16 Link Addition and Slow Recovery Procedure Example

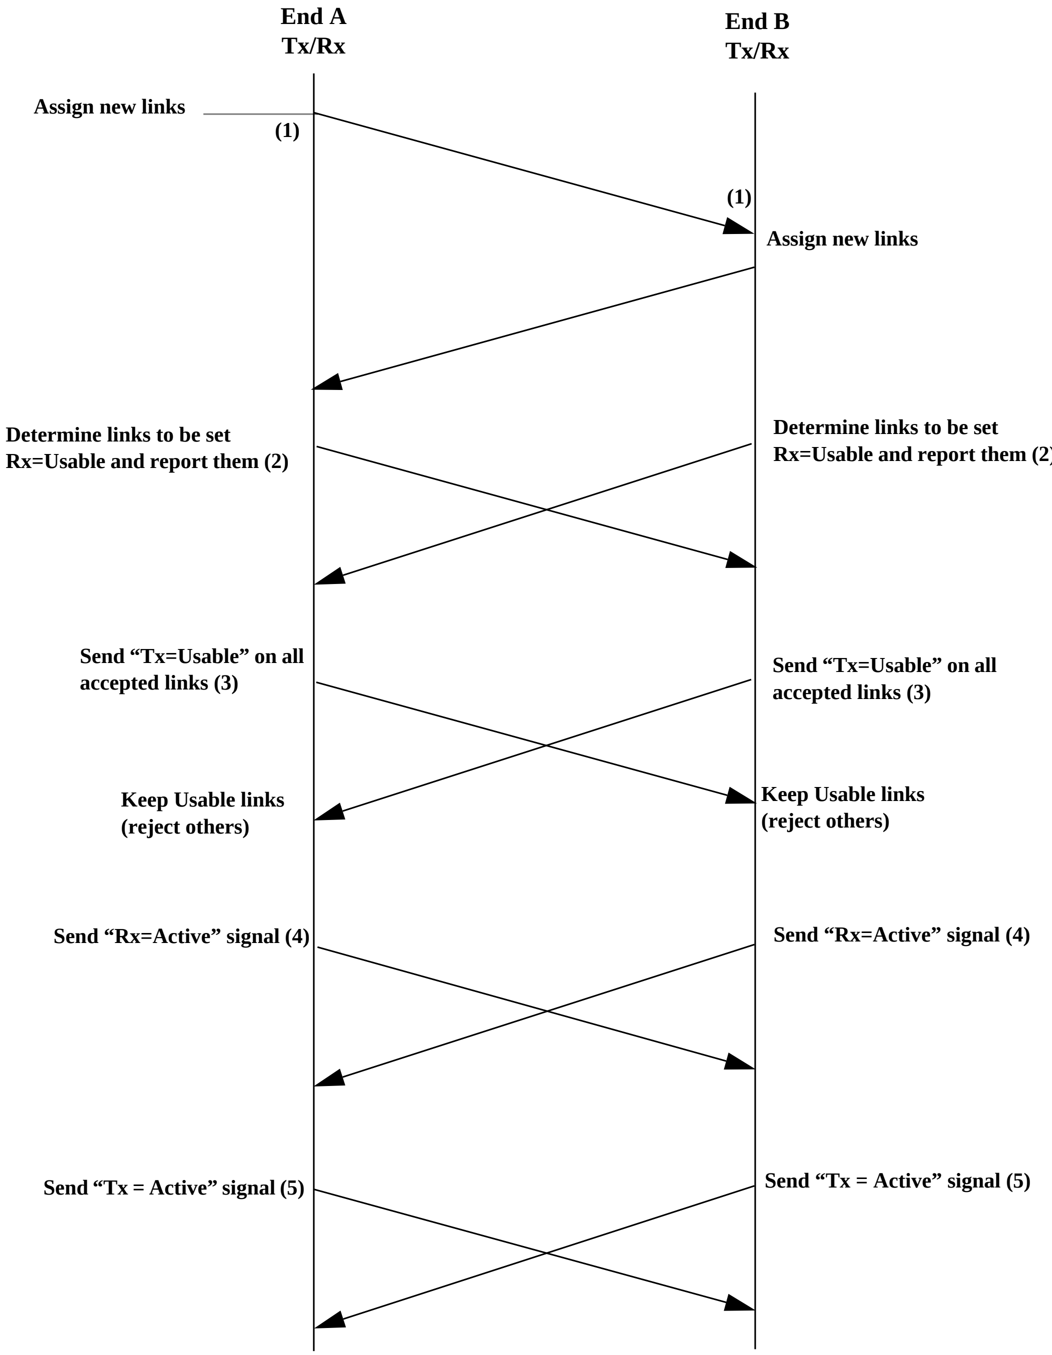

(1) Each link is being assigned a LID.
(2) At this point, new links are those that are Usable on both ends.
(3) All new links should be reported Tx=Usable at the same time.
(4) All new links should be reported Rx=Active at the same time.
(5) The new links can be activated in the transmit direction once they all have been recognized Active in the receive direction or after some time-out has expired and some of the links have been recognized Active in the receive direction. All new links should be reported Tx=Active at the same time.

Figure 17 Example of Link Addition Procedure over Time

11. IMA Frame Synchronization Mechanism

This section defines the IMA frame synchronization mechanism. It is based on the cell delineation mechanism in ITU-T Recommendation I.432[41] and ANSI T1.646[6].

(R-95) Each IMA link in the group shall provide the IMA frame synchronization mechanism as shown in Figure 18 on page 66.

In the standards for the Physical Layer specification like ITU-T Recommendation I.432 or ANSI T1.646 the synchronization mechanism is specified independent of the error/maintenance handling. The synchronization states form a basis for the different error and maintenance states. The same applies for this specification.

The IMA frame synchronization mechanism consists of the IMA HUNT, IMA PRESYNC and IMA SYNC states. By analogy with ITU-T Recommendation I.432 and ANSI T1.646 the checking scheme should be Cell-by-Cell in the IMA HUNT state and Frame-by-Frame in the IMA SYNC and IMA PRESYNC states.

(R-96) The synchronization mechanism (implemented for each link separately) shall work independently of the other defect conditions and the link delay compensation.

(R-97) The default value of Alpha(α), Beta(β) and Gamma(γ), variables appearing in Figure 18 on page 66, shall be supported as specified in Table 15 on page 65.

(O-9) The support of values of Alpha(α), Beta(β) and Gamma(γ) other than the default values, such as specified in Table 15 on page 65, is optional.

Table 15 Alpha, Beta and Gamma Values

Setting	Range	Default (Required support)
Alpha (α) (consecutive invalid ICP cells)	1-2	2
Beta (β) (consecutive errored ICP cells)	1-5	2
Gamma (γ) (consecutive valid ICP cells)	1-5	1

(R-98) Upon the occurrence of HEC/CRC errored cell in the ICP position, the receiving end shall ignore the cell content but shall assume that it was an ICP cell.

(O-10) The IMA Frame Synchronization mechanism may also go into the Hunt state when no longer getting cells from the physical layer.

Table 16 on page 66 provides the definitions of the terms appearing in Figure 18 on page 66.

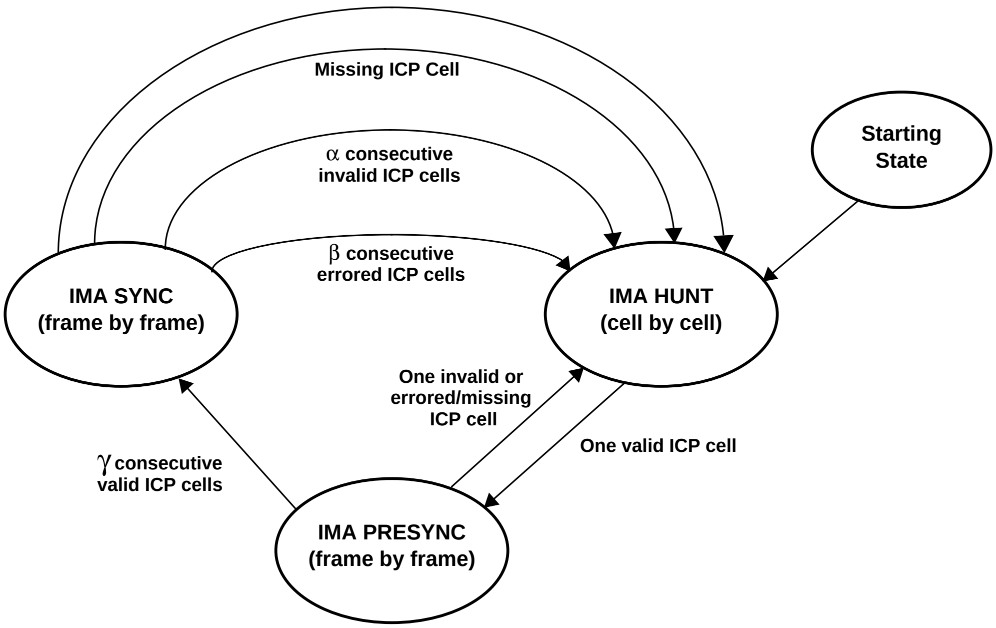

Notes:

- See both (R-97) and (O-9) for the value of α, β and γ.
- In brackets: checking scheme in each state (cell by cell or frame by frame checking)
- See Table 16 on page 66 for definitions of terms.

Figure 18 IMA Frame Synchronization Mechanism

Table 16 Definitions for IMA Frame Synchronization Mechanism

Term	Definition
Invalid ICP Cells (see note 1)	Cell with good HEC&CRC and CID=ICP at expected frame position with (unexpected IMA label) or (unexpected LID) or (unexpected IMA ID) or (received M not equal to expected M of the IMA group) or (unexpected IMA frame number) or (unexpected ICP cell offset)
Errored ICP Cell (see note 1)	A cell with a HEC or CRC error at expected ICP cell position if it is not a missing cell.
Valid ICP Cell	(Cell with CID equal to ICP) and (No IMA OAM header, IMA label, HEC or CRC error) and (expected LID) and (expected IMA ID) and (received M equal to expected M of the IMA group) and (expected IMA frame number) and (expected ICP cell offset)

Table 16 Definitions for IMA Frame Synchronization Mechanism

Term	Definition
Expected/unexpected	LID, IMA ID, M, ICP cell offset: expected/unexpected related to value fixed during group start-up or link addition; when in Start-up, any value is to be expected. IMA frame number, stuff indication: related to separate reference counting at Rx side (see note 1).
Missing ICP Cell	(Cell located at ICP cell location without HEC error and without IMA OAM cell header) or (without HEC error and with IMA OAM cell header and without CID equal to ICP)
Consecutive ICP Cells (see note 1)	ICP cells sent on the same links in consecutive IMA frames (excepting the SICP cell in a stuff event, and regardless of offset). Thus one consecutive ICP cell is an ICP cell on one link in one IMA frame, two consecutive ICP cells are the cells on the same link in two consecutive IMA frames, etc.

Note 1. Only applicable to ‘frame by frame’ checking

11.1 IMA Frame Synchronization with Stuff Events

The IMA should preserve the IMA frame synchronization while receiving stuff events subjected to HEC or CRC errors.

Figure 19 on page 68 shows cases when one or more of the ICP cells preceding and comprising the stuff event are HEC or CRC errored.

(R-99) The receive IMA shall maintain synchronization for cases 1, 2, 3 and 6 represented in Figure 19 on page 68.

(O-11) The receive IMA may optionally maintain synchronization for cases 4, and 5 in Figure 19 on page 68. Case 4 requires the use of cell-by-cell hunting as described in Section 11 on page 65.

(O-12) The receive IMA may optionally maintain synchronization for case 7 in Figure 19 on page 68, when passing stuff indication over more than one of the previous ICP cells and when β, as defined in Section 11 on page 65, is greater than 2.

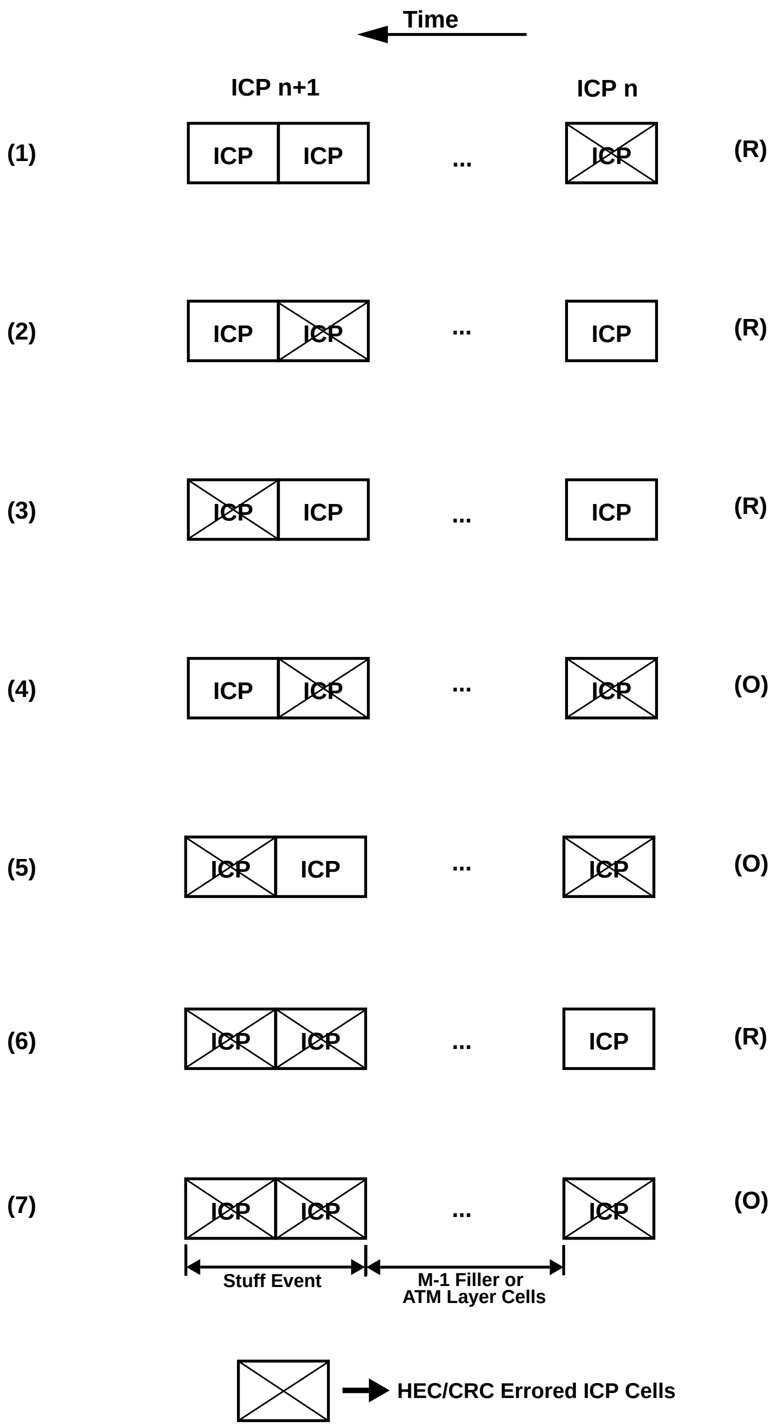

Figure 19 Error on ICP Cells Preceding and Comprising a Stuff Event

12. IMA Interface OAM Implementation

This section specifies the physical layer OAM functions and procedures required on the IMA interface.

12.1 IMA OAM Functions

12.1.1 IMA Group Maintenance Signals

Section 10.2.1 on page 53 specifies the signals to be used for reporting the actual state of the GSM and GTSM.

12.1.2 IMA Link Maintenance Signals

Section 10.1.6 on page 52 specifies the signals to be used for reporting the actual transmit and receive link states to the FE.

(R-100) The following link remote defect indicators shall be used to report local IMA defects:

- Link defects: indicates that a link defect has been detected at the local end (i.e., LOS, OOF/LOF, AIS or LCD).
- LIF: indicates that LIF defect has been detected at the local end.
- LODS: indicates that LODS defect has been detected at the local end.

(R-101) If several defects are detected at the same time, the defect with the highest priority, as listed in Table 17 on page 69, shall be reported.

Table 17 Priorities of Defect Reports

Priority	Rx Defect
1 (highest)	Link defect (e.g., LOS, OOF/LOF, AIS or LCD)
2	LIF
3	LODS

The IMA defects (LIF or LODS) are defined in Table 18 on page 74.

(R-102) The receive IMA shall report Rx defects no later than 2M cells after detection.

This is no immediate relationship between any defect reporting and the transmit and receive link states.

12.1.3 IMA Link Error Handling

This section defines the procedures to handle IMA link errors.

(R-103) Error handling of IMA links shall be as specified in Figure 20 on page 70 and Figure 21 on page 71.

Error handling regarding the anomalies, defects and failures are specified based on the IMA frame synchronization mechanism defined in Figure 18 on page 66. Figure 20 on page 70 shows the relationship between states of the IMA frame synchronization mechanism and the error/maintenance states also derived from the cell delineation procedure. The IMA Working state corresponds to the case where the IMA frame synchronization mechanism has been in the IMA SYNC state. The IMA OIF anomaly corresponds to the cases when the IMA frame synchronization mechanism has exited the IMA SYNC state. The IMA LIF defect corresponds to the case where the OIF anomaly has persisted for at least 2 IMA frames. The IMA LIF defect state is exited when the IMA frame synchronization mechanism has been in the IMA SYNC for at least 2 IMA frames.

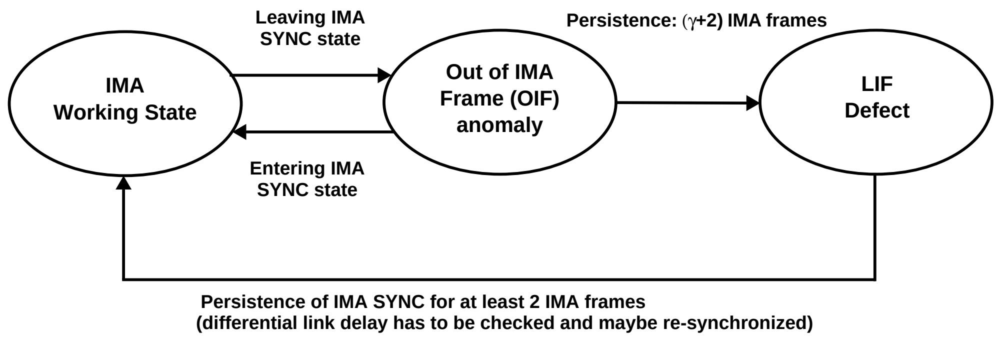

Figure 20 IMA Error/Maintenance State Diagram

(R-104) During OCD or OIF anomaly conditions on one link, no ATM layer cells shall be passed to the ATM layer for that link by the receive IMA. Any ATM layer cells accumulated before the occurrence of the OCD or OIF anomaly shall be passed to the ATM layer and only the ATM layer cells received after the anomaly condition has been detected shall be replaced with Filler cells.

(R-105) On IMA sublayer backward reporting of an Rx defect shall not be generated until LIF (or LODS) defect state is entered. The reporting shall be as specified in Section 12.1.2 on page 69.

The two major states shown as largest boxes in Figure 21 on page 71 are distinguished by whether or not the FE is informed during the error handling. In the case where the FE is not informed, the sub-states are distinguished by whether cells received in the sub-state are passed to the ATM layer or discarded.

In general, the detection of failures is based on the persistence of defects and received “FE defect reporting” and is implementation specific. The same applies for the detection of the Rx Failed condition defined in Section 10.1.4 on page 47.

The action of maintaining synchronization also means that the IMA receiver passes valid incoming cells to the ATM layer as if the bad link was simply receiving Filler cells.

It should be noted that there is only one state in which incoming ATM layer cells can be passed to the ATM layer.

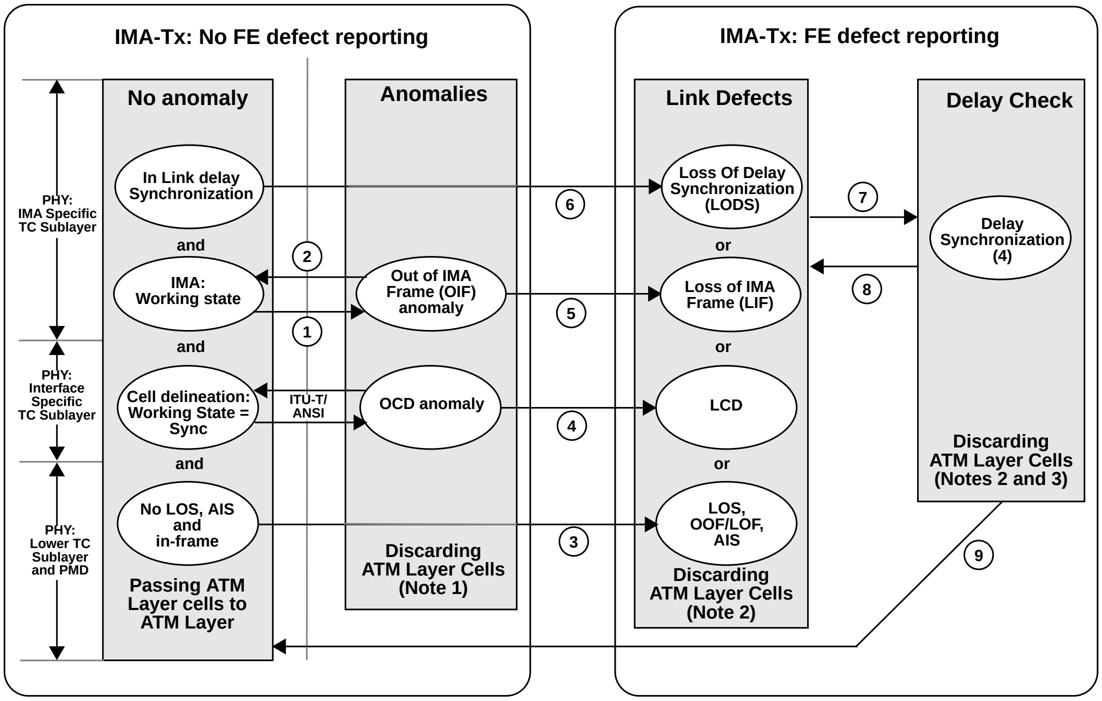

State transitions:

1. Transition according to Figure 18 on page 66 for leaving IMA SYNC
2. Entering SYNC state according to Figure 18 on page 66. Entering the IMA SYNC state triggers a new delay synchronization for that link.
3. LOS, OOF(DS1)/LOF(E1) or AIS or other lower layer notification is reported at the physical layer.
4. LCD notification is reported at the physical level.
5. Persistence: 2 IMA frames.
6. Loss Of Delay Synchronization (LODS)
7. No LOS, OOF/LOF, AIS, LCD or LIF condition detected
8. LOS or LOF or AIS or LCD or LIF or LODS
9. No error reporting by IMA protocol, recovery of the link as specified in Section 12.1.3.1 on page 71.

Notes:

(1) Fast recovery: ATM layer cells that are received are discarded. When Rx=On, the IMA shall emulate the reception of ATM layer cells by substituting Filler cells.

(2) Medium recovery: the receive IMA is waiting for persistent defect. The Rx=On actions is performed until the persistence leads to the entrance of the Rx Failed state. When Rx=On, the IMA shall emulate the reception of ATM layer cells by substituting Filler cells.

(3) Slow recovery: Rx Failed occurs. When Rx=On, the IMA shall emulate the reception of ATM layer cells by substituting Filler cells.

(4) Differential delay synchronization should be checked and re-synchronized.

Figure 21 IMA Error Handling Overview

12.1.3.1 IMA Link Recovery Mechanisms

According to Figure 21 on page 71 three types of link recovery are possible:

1. Fast recovery without entering the defect zone (e.g., going directly from “IMA Out of Frame” anomaly to the IMA Working state as described in Figure 20 on page 70); the FE is not aware of the anomaly and the fast recovery.
2. Medium recovery: a defect is reported and the FE defect reporting is initiated; the FE is aware of the defect and the medium recovery.
3. Slow recovery: a Rx Failed condition is entered due to the persistence of a defect; this triggers the deactivation of that link and the recovery by the IMA protocol as specified in Section 10.1.5 on page 47; the FE is aware of the deactivation of the link and the slow recovery. Examples of the slow recovery are shown in Figure 42 on page 123 and Figure 43 on page 124.

The slow recovery mechanism using the IMA protocol is dependent on how the local end processes the defect and declares the Rx Failed condition. The occurrence of the Rx Failed condition will cause the LSM to move to the Unusable state (see Section 10.1.5 on page 47).

12.2 IMA Performance and Failure Alarm Monitoring

This section describes the requirements for IMA performance and failure alarm monitoring.

Figure 22 on page 72 illustrates the way monitored parameters are used by both the layer management entity, for driving the LSM, and for failure and performance monitoring. NE events (in terms of anomalies, signals, and defects) and FE reports (in terms of anomalies, signals, defects and error conditions) are defined to provide performance monitoring parameters. NE events and (processed) failures are also defined to make up FE reports for the transmitted signal.

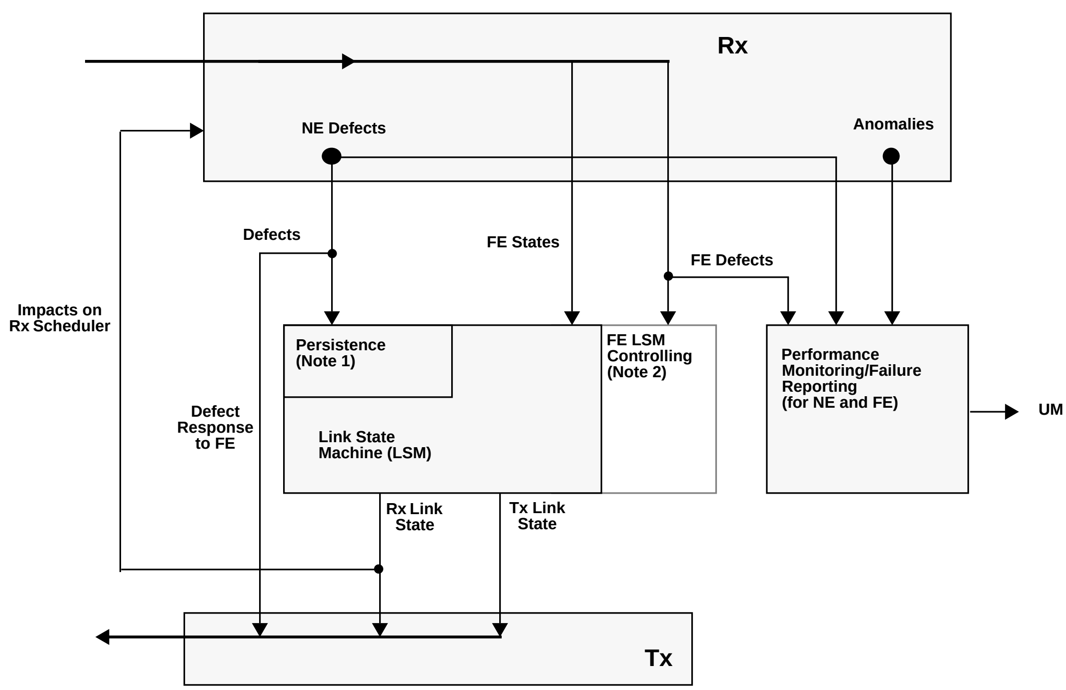

Note 1. Results in Rx Failed as a result of the persistence check (according to Appendix B.2 on page 110) and triggers the LSM (to go to Rx Unusable (Failed) state).

Note 2 Optional: the FE can observe the NE persistence behavior and force the NE behavior by changing Tx state to Unusable, thereby forcing the NE out of Active. The FE may do this if, in its estimation, the NE persistence processing is unacceptable.

Figure 22 Structure of anomaly and defect processing

12.2.1 Performance Monitoring Objectives

Performance monitoring is the process of continuous collection, analysis and reporting of performance data associated with a transmission entity (such as described in ANSI T1.231[3]). Similar to ITU-T Recommendation G.826[37], the following process is defined:

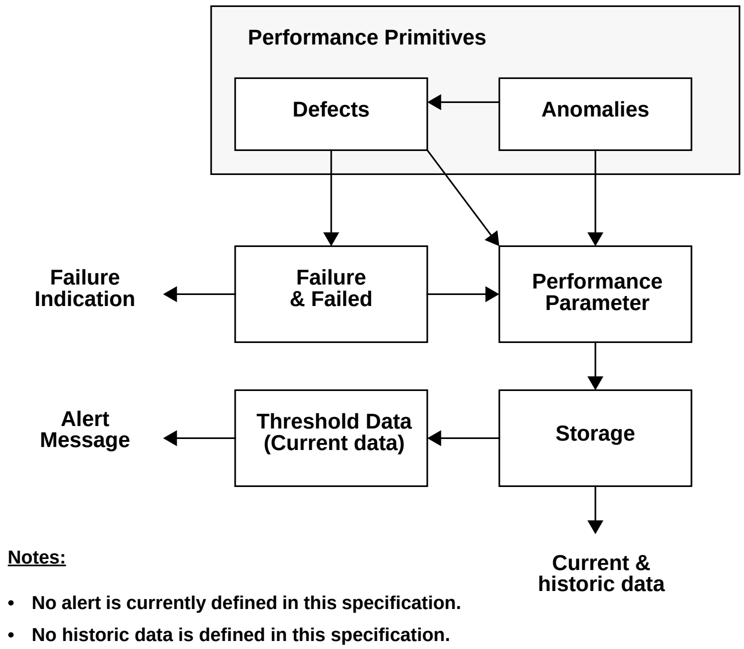

Figure 23 Performance Monitoring Process (based on ANSI T1.231, Figure 4)

The IMA Performance Primitives consist of

- *Anomalies* with
  a. Bit error related events (e.g., errored ICP cells)
  b. Synchronization related events (e.g., OIF)
- *Defects* with
  a. IMA Frame error events (e.g., LIF)
  b. Other defect events (e.g., LODS)
- *Failure alarm*: a failure alarm is the termination of the ability of an item to perform a required function (ANSI T1.231). A failure alarm is declared when the defect (or response FE defect) persists for x seconds (default: 2.5 ± 0.5 seconds). The failure alarm is different than the Failed condition defined in Section 10.1.4 on page 47.
- *Performance Parameters* specified for each sublayer or OAM level separately are based on performance primitives and failures. IMA performance parameters are accumulated in a continuous fashion. Performance parameters are inhibited during seconds when the unavailability of the link or the group has been reported and when the counting of anomaly events is inhibited during defect detection[1].

1. The IMA performance parameters could optionally be accumulated over 15 minutes and 24 hours accumulation periods.

- The Intermittent defect handling like the hit integration (refer to Appendix B.2 on page 110) is only used for special applications and is not required in ANSI T1.231.

12.2.2 Performance Monitoring Parameters

12.2.2.1 IMA Performance Primitives

(R-106) Table 18 on page 74 indicates the performance primitives that shall be detected by the IMA interface.

Table 18 IMA Link and Group Anomalies and Defects

Performance Primitives	Events	Definition
Anomalies (IMA Specific)	Errored ICP cell	Same definition as in Figure 18 on page 66.
	Invalid ICP cell	Same definition as in Figure 18 on page 66.
	Missing ICP cell	Same definition as in Figure 18 on page 66.
	OIF events	Leaving the IMA SYNC state.
Defects (IMA Specific)	LIF	Definition as described in Section 12.1.3 on page 69.
	LODS	A LODS defect is reported when the differential link delay between the link and the other links in the group is over the tolerable differential link delay.
	RDI-IMA	One of the available remote defect indicators (including IMA link specific defect) is indicated in the link related “Link Information” field.

12.2.2.2 IMA Performance Parameters

The IMA performance parameters are based on principles for the sub-layers of DS1 and E1, e.g. ANSI T.231, ANSI T1.646 and ITU-T Recommendation G.826.

Table 19 on page 74 indicates the required and optional IMA performance parameters.

Table 19 IMA Performance Parameters

(R)/(O)	Link/ Group	Performance Parameter	Definition
(R-107)	Link	IV-IMA	ICP Violations: count of errored, invalid or missing ICP cells during non-SES-IMA or non-UAS-IMA condition.
(R-108)	Link	SES-IMA	Count of one second intervals containing ≥ 30% of the ICP cells counted as IV-IMAs (see note 1), or one or more link defects (e.g., LOS, OOF/LOF, AIS or LCD), LIF, LODS defects during non-UAS-IMA condition.
(R-109)	Link	SES-IMA-FE	Count of one second intervals containing one or more RDI-IMA defects during non-UAS-IMA-FE condition.
(R-110)	Link	UAS-IMA	Unavailable seconds: unavailability begins at the onset of 10 contiguous SES-IMA and ends at the onset of 10 contiguous seconds with no SES-IMA.
(R-111)	Link	UAS-IMA-FE	Unavailable seconds at FE: unavailability begins at the onset of 10 contiguous SES-IMA-FE and ends at the onset of 10 contiguous seconds with no SES-IMA-FE.

Table 19 IMA Performance Parameters

(R)/(O)	Link/ Group	Performance Parameter	Definition
(R-112)	Link	Tx-UUS-IMA	Tx Unusable seconds: count of Tx Unusable seconds at the NE LSM.
(R-113)	Link	Rx-UUS-IMA	Rx Unusable seconds: count of Rx Unusable seconds at the NE LSM.
(R-114)	Link	Tx-UUS-IMA-FE	Tx Unusable seconds at FE: count of seconds with Tx Unusable indications from the FE LSM.
(R-115)	Link	Rx-UUS-IMA-FE	Rx Unusable seconds at FE: count of seconds with Rx Unusable indications from the FE LSM.
(R-116)	Link	Tx-FC	Count of NE Tx link failure alarm conditions.
(R-117)	Link	Rx-FC	Count of NE Rx link failure alarm conditions.
(O-13)	Link	Tx-FC-FE	Count of FE Tx link failure alarm conditions.
(O-14)	Link	Rx-FC-FE	Count of FE Rx link failure alarm conditions.
(O-15)	Link	OIF-IMA	Count of OIF anomalies during non-SES-IMA condition.
(O-16)	Link	Tx-Stuff-IMA	Count of Tx stuffing events during non-SES-IMA condition.
(O-17)	Link	Rx-Stuff-IMA	Count of Rx stuffing events during non-SES-IMA condition.
(R-118)	Group	GR-UAS-IMA	Count of one second intervals where the GTSM is Down.
(R-119)	Group	GR-FC	Count of NE group failure conditions.
(O-18)	Group	GR-FC-FE	Count of FE group failure conditions.
(R-120)	Group	GR-Timing-Mismatch	A possible configuration mismatch has occurred (CTC and CTC modes are not the same at both ends of the IMA virtual links.

Note 1. See Appendix F on page 126.

(O-19) The accumulation of IMA performance parameters over 15 minutes and 24 hours interval is optional.

(CR-7) If (O-19) is used, the current/previous and recent data shall be kept. For threshold crossing, the current data set shall be used.

12.2.3 IMA Failure Alarms

Table 20 on page 76 presents the IMA failure alarm definitions reported to the UM.

Table 20 IMA Failure Alarms

(R)/(O)	Link/ Group	Failure Alarm	Definition
(R-121)	Link	LIF	Persistence of LIF defect at NE (see note 1).
(R-122)	Link	LODS	Persistence of LODS defect at NE (see note 1).
(R-123)	Link	Tx-Mis-Connected	When the link is detected as mis-connected. This is reported mis-connected when the transmit end has determined the link is not connected to the same FE as the other links in the group. The detection is implementation specific.
(R-124)	Link	RFI-IMA	Persistence of RDI-IMA defect at NE (see note 1).
(O-20)	Link	Fault	Implementation specific FAULT declared at NE (see note 1).
(R-125)	Link	Tx-Unusable-FE	Tx link forced to Unusable due to FE behavior.
(R-126)	Link	Rx-Unusable-FE	Rx link forced to Unusable due to FE behavior.
(R-127)	Group	Start-up-FE	When the FE is starting-up (see note 2).
(R-128)	Group	Config-Abort	When the FE tries to use unacceptable configuration parameters.
(R-129)	Group	Config-Abort-FE	When the FE reports unacceptable configuration parameters.
(R-130)	Group	Insufficient-Links	When less than P_{Tx} transmit or P_{Rx} receive links are Active.
(R-131)	Group	Insufficient-Links-FE	When the FE reports that less than P_{Tx} transmit or P_{Rx} receive links are Active.
(R-132)	Group	Blocked-FE	When the FE reports that it is blocked.

Note 1. The persistence relevant for the failure detection and the failure clearance may be selectable by the UM separately and shall be adjusted to the IMA unit failure detection timing (implementation specific).

Note 2. The declaration of this failure alarm may be delayed to ensure the FE remains in Start-up.

Appendix B.2 on page 110 suggests a state machine to be used for processing error conditions.

(R-133) In the case of the LIF, LODS, RFI-IMA and Fault failure alarms, the IMA shall support 2.5 ± 0.5 seconds as a default persisting checking time to enter a failure alarm condition, and 10 ± 0.5 seconds as a default persisting clearing time to exit the same failure alarm condition.

(R-134) The clearing of the Fault Failure Alarm state is implementation dependent, but shall at a minimum require that the error which caused the fault no longer be present.

13. Test Pattern Procedure

The test pattern procedure is defined to provide extra supports to verify the connectivity of a link within an IMA group. It is based on the use of a test pattern sent over one link for which verification of the connectivity to the rest of the group is desired. The test pattern is expected to be looped over all the other links in the group at the FE. All the procedure is performed over the ICP cells exchanged between both ends of the IMA virtual links.

Figure 24 on page 77 shows an example where a test pattern is originated on one end to verify the connectivity of LID 2 link with the other links already in the group.

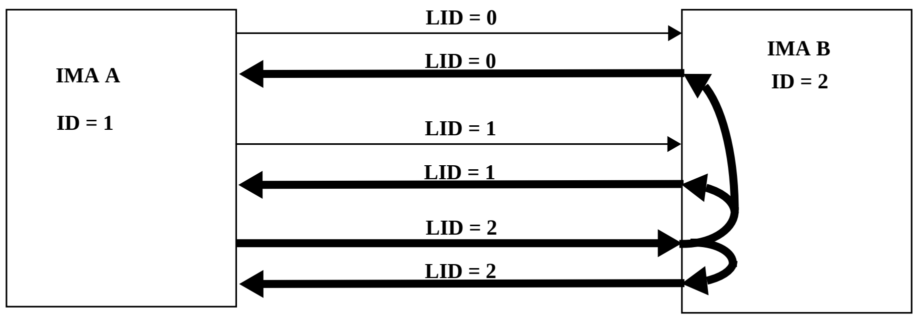

Links identified as LID = 2 are being added to existing groups using LID=0,1 links

Figure 24 Example of Looped Test Pattern

(O-21) The transmit IMA may optionally use the Test Pattern procedure.

The test pattern procedure is made available to ensure proper link connectivity. The following requirements shall be followed in order to exercise the Test Pattern procedure.

(CR-8) If (O-21) is used, the IMA interface shall use the Test Pattern command, Test Link and Tx Test Pattern fields in the ICP cell to activate the Test Pattern feature at the FE. It shall also send the same values over the next ICP cells sent over all the links (including test link). It shall specify the test link LID in the Test Link LID field.

(CR-9) If (O-21) is used, the IMA interface shall send the same value of the Test Pattern command, Test Link and Tx Test Pattern fields in at least 2 consecutive ICP cells.

(CR-10) If (O-21) is used, the IMA interface shall continue to send the new values of the Test Pattern command, Test Link and Tx Test Pattern fields as long as it wants the FE IMA to loop back the test pattern.

(R-135) The IMA interface shall monitor to incoming ICP cells on the links already recognized in the group to detect a change of the Test Pattern command. If the Test Pattern command is detected as active, the IMA interface shall monitor the Tx Test Link and Tx Test Pattern fields. It shall also verify that the Test Command bit is set to active on the test link. It shall then copy the Tx Test Pattern received on the test link, indicated by the FE IMA, into the Rx Test Pattern field on every subsequent ICP cell sent over all outgoing links in the group.

(R-136) The IMA interface shall keep sending the same Rx Test Pattern until it has been indicated to stop looping the pattern or to loop a new one from the same or another link.

(R-137) The IMA interface shall return the 0xff pattern over the Rx Test Pattern field when the incoming test command is inactive or the test link is not detected.

(R-138) The IMA interface shall only handle one test pattern per IMA group at any given time.

The following text suggests behaviors at the transmit and receive ends.

- Once sending the new value of the Tx Test Link and Tx Test Pattern, the IMA interface might have to wait sometime in order to let the pattern be detected by the receive end and to be looped back. The IMA interface should verify that the test pattern is returned over the Rx Test Pattern field over ICP cells received on any other links that have already being recognized part of the group on both ends.
- Loopback configuration on a giving link can be detected by verifying that the received Tx Test Pattern value read from the incoming ICP cells is the same as the one inserted in the outgoing ICP cells.
- The IMA interface should try several values randomly chosen to ensure that the FE has not been trying the same value at the same time.
- At start-up time, the IMA interface sends the 0xff pattern until receiving ICP cells indicating the first Test Pattern command.

Figure 25 on page 78 shows the Test Pattern procedure execution over time.

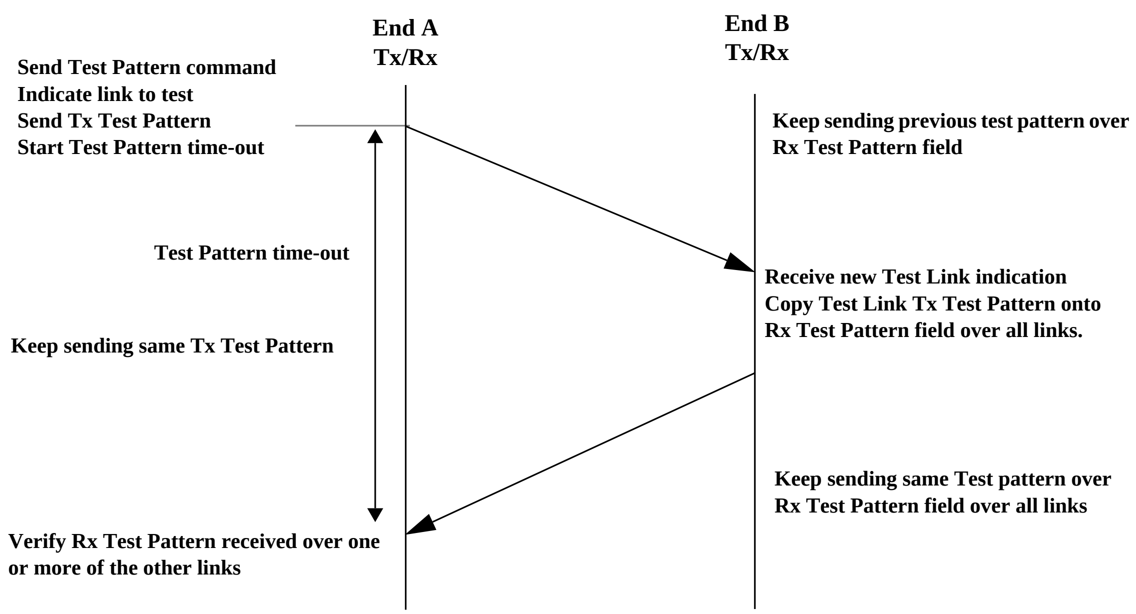

Figure 25 Test Pattern Procedure over Time

14. IMA Interaction with Plane Management

(R-139) The IMA shall process the following indications from/to Plane Management:

- *IMA group configuration*: this signal is received from Plane Management by the IMA entity and indicates which links are to be assigned to the IMA group.
- *Link addition/deletion*: this signal is received from Plane Management by the IMA entity and indicates that a link is to be added to or deleted from the IMA group.
- *IMA service operational status change*: this signal is sent by the IMA entity to the Plane Management to indicate a change of the operational status of the IMA (e.g., GTSM moving to Down state).
- *Tx/Rx cell rate change*: this signal is sent by the IMA entity to the Plane Management to indicate that the cell rate has changed in the transmit and/or receive directions (a link has been added to or deleted from the IMA group or a link has been inactivated due to a defect or a link recovered from a defect condition).

15. Management Information Base

(O-22) The Unit Management (UM) may be SNMP based.

(CR-11) If (O-22) is used, the MIBs defined in this section shall be implemented.

15.1 The Network Management Framework

IMA management is defined using SNMP MIBs. SNMP [15] is an Internet standard network management framework.

The Request For Comment (RFC) documents which are relevant here are

- RFC 1213, [17] which defines MIB-II, the core set of managed objects for the Internet suite of protocols
- RFC 1442 [19] which defines the Structure of Management Information (SMI), the mechanisms used for describing and naming objects for the purpose of management.
- RFCs 1443 through 1452, which define SNMPv2 [20][21][22][23][24][25][26][27][28][29].
- RFC 1573 [30], which defines the evolution of the Interfaces group of MIB-II.

15.2 Overview

This MIB provides the objects necessary for configuration, performance, and fault management of the IMA physical layer.

If IMA is used over DS1/E1 links, this MIB specification requires the implementation of the DS1 MIB defined in RFC 1406 [18] with a separate instance created for each DS1/E1 used by the IMA.

15.3 IMA Terminology

The errors/defects/failures listed in this document pertain to the IMA functionality only.

15.3.1 Error events

See Section 12.2.2.1 on page 74.

15.3.2 Defects

See Section 12.2.2.1 on page 74.

15.3.3 Performance Monitoring Parameters

See Section 12.2.2.2 on page 74.

15.3.4 Failure Alarm States

See Section 12.2.3 on page 76.

15.4 MIB-II and RFC 1573 Support

All SNMP agents which support IMA must implement MIB-II [19] and the mandatory groups of RFC 1573[30]. The goal of IMA is to present the illusion that a group of one or more links can be treated as a single physical interface of aggregate bandwidth. Real physical interfaces have entries in the MIB-II Interfaces

Table. To preserve a consistent management framework, it is highly desirable for each IMA Group to have entries in this table as well.

To identify an interface as belonging to an IMA Group, it must be tagged with an ifType constant: atmIma(107). This constant lets a network management application knows that additional information about the interface is available via the IMA MIB.

IMA Groups are created and destroyed by network management. So that agents can control IMA ifIndex allocation, IMA Group tables define an object imaGroupIfIndex that represents the ifIndex associated with the IMA interface. Both values are available through table entries:

- An imaGroupMappingTable entry converts the IMA ifIndex to an imaGroupIndex.
- An imaGroupTable entry contains an imaGroupIfIndex object that identifies the IMA ifIndex

The MIB-II interfaces associated with physical links exist independently of the IMA link definition. A physical link is designated to be an IMA link by the creation of a row in the IMA Link Table.

15.4.1 The ifStackTable

The relevant portion of the ifStackTable looks like this:

```
----------------------
ATM Layer
----------------------
IMA Group
----------------------
Physical Interface
----------------------
```

Note that:

- The ATM Layer may multiplex among multiple IMA groups.
- The IMA Group may perform inverse multiplexing over multiple physical interfaces.

15.4.2 Interpretations of Interface Tables for IMA Groups

The following items defined in the ifTable of RFC 1573 have IMA specific definitions:

ifIndex	IMA Interface index.
ifType	atmIma(107)
ifMtu	Transmission units are ATM cells. The value of ifMtu is 53 octets.
ifSpeed	The interface speed is represents by the following equation (where N_{on} represents, in the transmit direction, the number of links which currently transmit cells passed from the ATM layer, and in the receive direction, the number of links which currently receive ATM layer cells to be passed to the ATM layer; where TRLCR is the TRL cell rate providing by the TRL in the transmit and receive direction respectively; M is the IMA frame length in cell units):

$$\mathrm{ifSpeed} = N_{on} \times \mathrm{TRLCR} \times \left(\frac{M-1}{M}\right) \times \left(\frac{2048}{2049}\right)$$

ifPhysAddress	A zero-length octet string

ifOperStatus This reflects the state of the Group Traffic State Machine.

RFC 1573 provides the full descriptions of the status values.

For conformance to RFC 1573, the ifGeneralGroup and ifStackGroup are the only mandatory groups.

15.4.3 Interaction with ifAdminStatus

Setting the ifAdminStatus of the IMA Group interface to “down” or “testing” will cause the inhibiting condition which affects the Group State Machine.

Setting the ifAdminStatus of a link interface to ‘down’ or ‘testing’ will cause the Blocking condition which affects the link’s Transmit and Receive State Machines.

15.5 MIB Definition

```
IMA-MIB DEFINITIONS ::= BEGIN

IMPORTS
   ifIndex
      FROM RFC1213-MIB
   enterprises
      FROM RFC1155-SMI
   MODULE-IDENTITY, OBJECT-TYPE, Integer32, NOTIFICATION-TYPE
      FROM SNMPv2-SMI
   MODULE-COMPLIANCE, OBJECT-GROUP
      FROM SNMPv2-CONF
   TEXTUAL-CONVENTION, DateAndTime, RowStatus
      FROM SNMPv2-TC;

-----------------------------------------------------------------------

-- The implementation of this MIB is mandatory for all network devices
-- which have ATM Inverse Multiplexer network interface module.

atmfImaMib    MODULE-IDENTITY
      LAST-UPDATED "9701092245Z"
      ORGANIZATION "ATM Forum"
      CONTACT-INFO
      "ATM Forum
      World Headquarters
      2570 West El Camino Real
      Suite 304
      Mountain View, CA 94040-1313
      USA
      Phone: +1 415 949 6700
      Fax: +1 415 949 6705
      email: info@atmforum.com"
      DESCRIPTION
        "The MIB module to describe an ATM Inverse Multiplexer interface"
      ::= { atmForumNetworkManagement TBD }

--
   atmForum OBJECT IDENTIFIER ::= { enterprises 353 }
   atmForumNetworkManagement OBJECT IDENTIFIER ::= { atmForum 12 }
   atmfImaMibObjects OBJECT IDENTIFIER ::= { atmfImaMib 1}
```

```
-- This MIB consists of the IMA group subtree.
-- The IMA group subtree consists of the number of IMA groups and a
-- table of IMA groups. Each entry in the table of IMA groups contains
-- information (configuration and status) specific to each group.

-- Textual conventions

IfIndex ::=    TEXTUAL-CONVENTION
      STATUS      current
      DESCRIPTION
        "ifIndex"
      SYNTAX      Integer32

MilliSeconds ::= TEXTUAL-CONVENTION
      STATUS       current
      DESCRIPTION
        "Time in milliseconds"
      SYNTAX      Integer32

ImaGroupState ::= TEXTUAL-CONVENTION
      STATUS       current
      DESCRIPTION
        "State of the IMA group."
      SYNTAX      INTEGER {
        notConfigured(1),
        startUp(2),
        startUpAck(3),
        configAbortUnsupportedM(4),
        configAbortIncompatibleSymmetry(5),
        configAbortOther(6),
        insufficientLinks(7),
        blocked(8),
        operational(9)}

ImaGroupFailureStatus ::= TEXTUAL-CONVENTION
      STATUS       current
      DESCRIPTION
        "Failure reason of an IMA group."
      SYNTAX      INTEGER {
        noFailure(1), -- unit is up
        startUpNe(2),
        startUpFe(3),
        invalidMValueNe(4),
        invalidMValueFe(5),
        failedAssymetricNe(6),
        failedAssymetricFe(7),
        insufficientLinksNe(8),
        insufficientLinksFe(9),
        blockedNe(10),
        blockedFe(11),
        otherFailure(12)}

ImaAlarmStatus ::= TEXTUAL-CONVENTION
      STATUS     current
      DESCRIPTION
        "A qualification of the IMA trap which indicates if the
        condition causing the trap has been detected (declared)
```

```
        or is no longer present (cleared)."
      SYNTAX      INTEGER {
        cleared(1),
        declared(2)}

ImaAlarmType ::= TEXTUAL-CONVENTION
      STATUS     current
      DESCRIPTION
        "An identification of the event that caused the generation
        of the IMA trap."
      SYNTAX      INTEGER {
        imaAlarmLinkLif(1),
        imaAlarmLinkLods(2),
        imaAlarmLinkTxMisConnect(3),
        imaAlarmLinkRfi(4),
        imaAlarmLinkFault(5),
        imaAlarmLinkTxUnusableFe(6),
        imaAlarmLinkRxUnusableFe(7),
        imaAlarmGroupStartupFe(8),
        imaAlarmGroupCfgAbort(9),
        imaAlarmGroupCfgAbortFe(10),
        imaAlarmGroupInsuffLinks(11),
        imaAlarmGroupInsuffLinksFe(12),
        imaAlarmGroupBlockedFe(13),
        imaAlarmGroupTimingSynch(14)}

ImaGroupTxClkMode  ::= TEXTUAL-CONVENTION
      STATUS      current
      DESCRIPTION
        "Indicate the transmit clock mode of the IMA group.
        There are two possible modes: the Common Transmit
        Clock (CTC) and the Independent Transmit Clock (ITC).
        The CTC mode corresponds to the case when the transmit clock
        of all IMA links are derived from the same source. The ITC
        configuration corresponds to the case where there is at least
        one IMA link whose transmit clock is derived from a source
        different than at least another link transmit clock."
      SYNTAX      INTEGER {
        ctc(1),
        itc(2)}

ImaGroupSymmetry ::= TEXTUAL-CONVENTION
      STATUS      current
      DESCRIPTION
        "The symmetry mode adjusted during the group start-up."
      SYNTAX      INTEGER {
        symmetricOperation(1),
        asymmetricOperation(2),
        asymmetricConfiguration(3)}

ImaFrameLength ::= TEXTUAL-CONVENTION
      STATUS      current
      DESCRIPTION
        "Length of the IMA frames."
      SYNTAX      INTEGER {
        m32(32),
        m64(64),
```

```
        m128(128),
        m256(256)}

ImaLinkState ::= TEXTUAL-CONVENTION
      STATUS      current
      DESCRIPTION
        "State of a link belonging to an IMA group."
      SYNTAX      INTEGER {
        notInGroup(1),
        unusableNoGivenReason(2),
        unusableFault(3),
        unusableMisconnected(4),
        unusableBlocked(5),
        unusableFailed(6),
        usable(7),
        active(8)}

ImaLinkFailureStatus ::= TEXTUAL-CONVENTION
      STATUS      current
      DESCRIPTION
        "Local failure status of a link belonging to an IMA group."
      SYNTAX      INTEGER {
        noFailure(1),
        imaLinkFailure(2),
        lifFailure(3),
        lodsFailure(4),
        misConnected(5),
        blocked(6),
        fault(7),
        farEndTxLinkUnusable(8),
        farEndRxLinkUnusable(9)}

ImaTestProcStatus ::= TEXTUAL-CONVENTION
      STATUS      current
      DESCRIPTION
        "States of the Test Pattern Procedure."
      SYNTAX      INTEGER {
        disabled(1),
        operating(2),
        linkFail(3)}

-- The IMA Group subtree

imaGroupNumber    OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The number of IMA groups configured on this system."
      ::= { atmfImaMibObjects 1}

-- IMA Group Table contains configuration and status information

imaGroupTable     OBJECT-TYPE
      SYNTAX      SEQUENCE OF ImaGroupEntry
      MAX-ACCESS  not-accessible
      STATUS      current
```

```
      DESCRIPTION
        "The IMA Group Configuration table."
      ::= { atmfImaMibObjects 2 }

imaGroupEntry     OBJECT-TYPE
      SYNTAX      ImaGroupEntry
      MAX-ACCESS  not-accessible
      STATUS      current
      DESCRIPTION
        "An entry in the IMA Group table."
      INDEX       { imaGroupIndex }
      ::= { imaGroupTable 1 }

ImaGroupEntry ::= SEQUENCE {
      imaGroupIndex               Integer32,
      imaGroupRowStatus           RowStatus,
      imaGroupIfIndex             IfIndex,
      imaGroupNeState             ImaGroupState,
      imaGroupFeState             ImaGroupState,
      imaGroupFailureStatus       ImaGroupFailureStatus,
      imaGroupSymmetry            ImaGroupSymmetry,
      imaGroupMinNumTxLinks       INTEGER (1..32),
      imaGroupMinNumRxLinks       INTEGER (1..32),
      imaGroupNeTxClkMode         ImaGroupTxClkMode,
      imaGroupFeTxClkMode         ImaGroupTxClkMode,
      imaGroupTxTimingRefLink     IfIndex,
      imaGroupRxTimingRefLink     IfIndex,
      imaGroupLastChange          DateAndTime,
      imaGroupTxImaId             INTEGER (0..255),
      imaGroupRxImaId             INTEGER (0..255),
      imaGroupTxFrameLength       ImaFrameLength,
      imaGroupRxFrameLength       ImaFrameLength,
      imaGroupDiffDelayMax        MilliSeconds,
      imaGroupLeastDelayLink      IfIndex,
      imaGroupDiffDelayMaxObs     MilliSeconds,
      imaGroupAlphaValue          INTEGER (1..2),
      imaGroupBetaValue           INTEGER (1..5),
      imaGroupGammaValue          INTEGER (1..5),
      imaGroupRunningSecs         Integer32,
      imaGroupUnavailSecs         Integer32,
      imaGroupNeNumFailures       Integer32,
      imaGroupFeNumFailures       Integer32,
      imaGroupTxAvailCellRate     Integer32,
      imaGroupRxAvailCellRate     Integer32,
      imaGroupNumTxCfgLnks        Integer32,
      imaGroupNumRxCfgLnks        Integer32,
      imaGroupNumTxActLnks        Integer32,
      imaGroupNumRxActLnks        Integer32,
      imaGroupTestLinkIfIndex     IfIndex,
      imaGroupTestPattern         Integer32,
      imaGroupTestProcStatus      ImaTestProcStatus }

imaGroupIndex     OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
```

```
        "A unique value for the IMA Group."
      ::= { imaGroupEntry 1 }

imaGroupRowStatus OBJECT-TYPE
      SYNTAX      RowStatus
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "The imaGroupRowStatus object allows create and delete operations
        on imaGroupTable entries.

        To create a new conceptual row (or instance) of the
        imaGroupTable, imaGroupRowStatus must be set to 'createAndWait' or
        'createAndGo.'  A successful set of the following objects must be
        performed before the row can transition to 'Active' state:

          imaGroupTxImaId
          imaGroupTxFrameLength
          imaGroupMinNumTxLinks and
          imaGroupMinNumRxLinks

        To change (modify) an imaGroupTable entry, the imaGroupRowStatus
        object must be set to 'notInService'. Only then can the
        conceptual row be changed.

        To remove (delete) an imaGroupTable entry from this table, set
        imaGroupRowStatus to 'destroy.'

        Setting the imaGroupRowStatus to 'Active' has the effect of
        activating the Group Startup Procedure.  The Group Startup Procedure
        uses provisioned links that have imaLinkRowStatus set to 'Active' an
        imaLinkGroupIndex set the imaGroupIndex in this conceptual row.

        When the imaGroupRowStatus is not in 'Active' state, the Group State
        machine is in its 'Not Configured' state."
      ::= { imaGroupEntry 2 }

imaGroupIfIndex   OBJECT-TYPE
      SYNTAX      IfIndex
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "This object identifies the logical interface number
        ('ifIndex') assigned to this IMA group, and is used to identify
        corresponding rows in the Interfaces MIB.

        Note that re-initialization of the management agent may cause
        a client's 'imaGroupIfIndex' to change."
      ::= { imaGroupEntry 3 }

imaGroupNeState     OBJECT-TYPE
      SYNTAX      ImaGroupState
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The current operational state of the near-end IMA
        Group State Machine."
```

```
      ::= { imaGroupEntry 4 }

imaGroupFeState     OBJECT-TYPE
      SYNTAX      ImaGroupState
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The current operational state of the far-end IMA
        Group State Machine."
      ::= { imaGroupEntry 5 }

imaGroupFailureStatus       OBJECT-TYPE
      SYNTAX      ImaGroupFailureStatus
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
       "The current failure status of the IMA group (the reason why
        the GTSM is in the down state."
      ::= { imaGroupEntry 6 }

imaGroupSymmetry  OBJECT-TYPE
      SYNTAX      ImaGroupSymmetry
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "Symmetry of the IMA group."
      DEFVAL { symmetricOperation }
      ::= { imaGroupEntry 7 }

imaGroupMinNumTxLinks       OBJECT-TYPE
      SYNTAX      INTEGER (1..32)
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
       "Minimum number of transmit links required to be Active for
        the IMA group to be in the Up state."
      ::= { imaGroupEntry 8 }

imaGroupMinNumRxLinks       OBJECT-TYPE
      SYNTAX      INTEGER (1..32)
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "Minimum number of receive links required to be Active for
             the IMA group to be in the Up state."
      ::= { imaGroupEntry 9 }

imaGroupNeTxClkMode    OBJECT-TYPE
      SYNTAX      ImaGroupTxClkMode
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "Transmit clocking mode used by the near-end IMA group."
      DEFVAL { ctc }
      ::= { imaGroupEntry 10 }

imaGroupFeTxClkMode     OBJECT-TYPE
```

```
      SYNTAX      ImaGroupTxClkMode
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Transmit clocking mode used by the far-end IMA group."
      ::= { imaGroupEntry 11 }

imaGroupTxTimingRefLink     OBJECT-TYPE
      SYNTAX      IfIndex
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The ifIndex of the transmit timing reference link to
        be used by the near-end for IMA data cell clock recovery from
        the ATM layer."
      ::= { imaGroupEntry 12 }

imaGroupRxTimingRefLink    OBJECT-TYPE
      SYNTAX      IfIndex
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
         "The ifIndex of the receive timing reference link to be
         used by near-end for IMA data cell clock recovery toward
         the ATM layer."
      ::= { imaGroupEntry 13 }

imaGroupLastChange OBJECT-TYPE
      SYNTAX      DateAndTime
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The time-of-day the IMA group last changed operational
        state (i.e., value of imaNeGroupState changed)."
      ::= { imaGroupEntry 14 }

imaGroupTxImaId OBJECT-TYPE
      SYNTAX      INTEGER (0..255)
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "The IMA ID currently in use by the near-end IMA function."
      ::= { imaGroupEntry 15 }

imaGroupRxImaId OBJECT-TYPE
      SYNTAX      INTEGER (0..255)
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The IMA ID currently in use by the far-end IMA function."
      ::= { imaGroupEntry 16 }

imaGroupTxFrameLength    OBJECT-TYPE
      SYNTAX      ImaFrameLength
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
```

```
        "The frame length to be used by the IMA group in the transmit
        direction. Can only be set when the IMA group is startup."
      DEFVAL { m128 }
      ::= { imaGroupEntry 17 }

imaGroupRxFrameLength OBJECT-TYPE
      SYNTAX      ImaFrameLength
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Value of IMA frame length as received from remote IMA function."
      ::= { imaGroupEntry 18 }

imaGroupDiffDelayMax OBJECT-TYPE
      SYNTAX      MilliSeconds
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "The maximum number of milliseconds of delay differential among
        the links that will be tolerated on this interface."
      DEFVAL { 25 }
      ::= { imaGroupEntry 19 }

imaGroupLeastDelayLink OBJECT-TYPE
      SYNTAX      IfIndex
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The ifIndex of the link configured in the IMA
        group which has the smallest link propagation delay. This
        value has meaning only if at least 1 link has been configured
        in the IMA group."
      ::= { imaGroupEntry 20 }

imaGroupDiffDelayMaxObs OBJECT-TYPE
      SYNTAX      MilliSeconds
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The latest maximum differential delay observed (in
        milliseconds), i.e. between the links having the least and
        most link propagation delay, among the receive links that
        are currently configured in the IMA group."
      ::= { imaGroupEntry 21 }

imaGroupAlphaValue OBJECT-TYPE
      SYNTAX      INTEGER (1..2)
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "This indicates the 'alpha' value used to specify the
        number of consecutive invalid ICP cells to be detected
        before moving to the IMA HUNT state."
      DEFVAL { 2 }
      ::= { imaGroupEntry 22 }

imaGroupBetaValue OBJECT-TYPE
```

```
      SYNTAX      INTEGER (1..5)
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "This indicates the 'beta' value used to specify the
        number of consecutive errored ICP cells to be detected
        before moving to the IMA HUNT state."
      DEFVAL { 2 }
      ::= { imaGroupEntry 23 }

imaGroupGammaValue OBJECT-TYPE
      SYNTAX      INTEGER (1..5)
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "This indicates the 'gamma' value used to specify the
        number of consecutive valid ICP cells to be detected before
        moving to the IMA SYNC state from the PRESYNC state."
      DEFVAL { 1 }
      ::= { imaGroupEntry 24 }

imaGroupRunningSecs OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The amount of time (in seconds) since this IMA group has
        been in operation (up or down)."
      ::= { imaGroupEntry 25 }

imaGroupUnavailSecs OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Count of one second intervals where the IMA Group Traffic
        State Machine is Down."
      ::= { imaGroupEntry 26 }

imaGroupNeNumFailures OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The number of times a near-end group failure (Config-Abort,
        Insufficient-Links) has been reported since power-up or reboot."
      ::= { imaGroupEntry 27 }

imaGroupFeNumFailures OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The number of times a far-end group failure (Config-Abort-FE,
        Insufficient-Links-FE, Blocked-FE) has been reported since
        power-up or reboot. This is an optional attribute."
      ::= { imaGroupEntry 28 }
```

```
imaGroupTxAvailCellRate OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The current cell rate (truncated value in cells per second)
        provided by this IMA group in the transmit direction,
        considering all the transmit links in the Active state."
      ::= { imaGroupEntry 29 }

imaGroupRxAvailCellRate OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The current cell rate (truncated value in cells per second)
        provided by this IMA group in the receive direction,
        considering all the receive links in the Active state."
      ::= { imaGroupEntry 30 }

-- imaGroupNumTxCfgLnks is used by a network operator to tell how many
-- links are configured for transmit in the IMA group.

imaGroupNumTxCfgLnks  OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The number of links that are configured to transmit in this IMA
        group."
      ::= { imaGroupEntry 31 }

-- imaGroupNumRxCfgLnks is used by a network operator to tell how many
-- links are configured for receive in the IMA group.

imaGroupNumRxCfgLnks  OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The number of links that are configured to receive in this IMA
        group."
      ::= { imaGroupEntry 32 }

-- imaGroupNumTxActLnks is used by a network operator to tell how many
-- links which are configured for transmit are also Active.

imaGroupNumTxActLnks  OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The number of links which are configured to transmit and are
        currently Active in this IMA group."
      ::= { imaGroupEntry 33 }
```

```
-- imaGroupNumRxActLnks is used by a network operator to tell how many
-- links which are configured for receive are also Active.

imaGroupNumRxActLnks  OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The number of links which are configured to receive and are
        currently Active in this IMA group."
      ::= { imaGroupEntry 34 }

--
-- Test Pattern Procedure control objects. These objects are implemented
-- if the IMA implements the Test Pattern Procedure. In this case all
-- test pattern procedure related objects must be implemented. Specifically
-- these are
--
-- imaGroupTestLinkIfIndex
-- imaGroupTestPattern
-- imaGroupTestProcStatus
-- imaLinkTestProcStatus
-- imaLinkRxTestPattern
--

imaGroupTestLinkIfIndex  OBJECT-TYPE
      SYNTAX      IfIndex
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "This object is used to designate an interface as the Test Link.
        A value of -1 specifies that the implementation may choose the
        Test Link. In this case, the implementation may also choose the
        value of 'imaGroupTestPattern'."
      DEFVAL { -1 }
      ::= { imaGroupEntry 35 }

imaGroupTestPattern  OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "The value of this object is used to specify the test pattern
        in an IMA group loopback operation. A value in the range 0 to
        255 designates a specific pattern. A value of -1 specifies
        that the implementation may choose the value. In this case,
        the implementation may also choose the value of
        'imaGroupTestLinkIfIndex'."
      DEFVAL { -1 }
      ::= { imaGroupEntry 36 }

imaGroupTestProcStatus OBJECT-TYPE
      SYNTAX      ImaTestProcStatus
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "This object is used to enable or disable the
```

```
        Test Pattern Procedure, and to note whether at least one
        link failed the test.

        The test is started by setting operating(2) status. If any
        link should fail the test, the IMA will set the status to
        linkfail(3). The linkfail(3) state will persist until
        either the disabled(1) state is set or until no instance
        of imaLinkTestProcStatus has the value linkfail(3).

        Only the values disabled(1) and operating(2) may be written.
        Writing the operating(2) value will not cause clearing of
        the linkfail(3) state."
      DEFVAL { disabled }
      ::= { imaGroupEntry 37 }

--
-- The IMA Group - ifIndex Mapping Table
--
-- Note that the Group Index is different than the ifIndex.
--
-- The "ifIndex Mapping" table allows the Unit Management to
-- perform easy look-ups (no searches and sorts).

imaGroupMappingTable OBJECT-TYPE
      SYNTAX      SEQUENCE OF ImaGroupMappingEntry
      MAX-ACCESS  not-accessible
      STATUS      current
      DESCRIPTION
        "A table mapping the 'ifIndex' values of 'imaGroupIfIndex'
        to the 'imaGroupIndex' values of the corresponding
        IMA group."
      ::= { atmfImaMibObjects 3 }

imaGroupMappingEntry OBJECT-TYPE
      SYNTAX      ImaGroupMappingEntry
      MAX-ACCESS  not-accessible
      STATUS      current
      DESCRIPTION
        "Each row describes one ifIndex --> imaGroupIndex mapping."
      INDEX { ifIndex }
      ::= { imaGroupMappingTable 1 }

ImaGroupMappingEntry ::=
      SEQUENCE {
         imaGroupMappingIndex    Integer32
      }

imaGroupMappingIndex OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The imaGroupIndex of the IMA Group which implements the
        specified interface."
      ::= { imaGroupMappingEntry 1 }
```

```
-- The IMA link status and configuration table
-- Each entry in the table contains status information about a link
-- which is part of an IMA group.

imaLinkTable      OBJECT-TYPE
      SYNTAX      SEQUENCE OF ImaLinkEntry
      MAX-ACCESS  not-accessible
      STATUS      current
      DESCRIPTION
        "The IMA group Link Status and Configuration table."
      ::= { atmfImaMibObjects 4 }

imaLinkEntry      OBJECT-TYPE
      SYNTAX      ImaLinkEntry
      MAX-ACCESS  not-accessible
      STATUS      current
      DESCRIPTION
        "An entry in the IMA Group Link table."
      INDEX   { imaLinkIfIndex }
      ::= { imaLinkTable 1 }

ImaLinkEntry ::=      SEQUENCE {
      imaLinkIfIndex              Integer32,
      imaLinkRowStatus            RowStatus,
      imaLinkGroupIndex           Integer32,
      imaLinkNeTxState            ImaLinkState,
      imaLinkNeRxState            ImaLinkState,
      imaLinkFeTxState            ImaLinkState,
      imaLinkFeRxState            ImaLinkState,
      imaLinkNeRxFailureStatus    ImaLinkFailureStatus,
      imaLinkFeRxFailureStatus    ImaLinkFailureStatus,
      imaLinkTxLid                INTEGER (0..31),
      imaLinkRxLid                INTEGER (0..31),
      imaLinkRelDelay             MilliSeconds,
      imaLinkImaViolations        Integer32,
      imaLinkOifAnomalies         Integer32,
      imaLinkNeSevErroredSec      Integer32,
      imaLinkFeSevErroredSec      Integer32,
      imaLinkNeUnavailSec         Integer32,
      imaLinkFeUnavailSec         Integer32,
      imaLinkNeTxUnusableSec      Integer32,
      imaLinkNeRxUnusableSec      Integer32,
      imaLinkFeTxUnusableSec      Integer32,
      imaLinkFeRxUnusableSec      Integer32,
      imaLinkNeTxNumFailures      Integer32,
      imaLinkNeRxNumFailures      Integer32,
      imaLinkFeTxNumFailures      Integer32,
      imaLinkFeRxNumFailures      Integer32,
      imaLinkTxStuffs             Integer32,
      imaLinkRxStuffs             Integer32,
      imaLinkRxTestPattern        INTEGER (0..255),
      imaLinkTestProcStatus       ImaTestProcStatus }

imaLinkIfIndex    OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
```

```
      DESCRIPTION
        "This corresponds to the 'ifIndex' of the MIB-II interface
        on which this link is established. This object also
        corresponds to the logical number ('ifIndex') assigned to
        this IMA link."
      ::= { imaLinkEntry 1 }

imaLinkRowStatus OBJECT-TYPE
      SYNTAX      RowStatus
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "The imaLinkRowStatus object allows create, change and delete
        operations on imaLinkTable entries.

        To create a new conceptual row (or instance) of the imaLinkTable,
        imaLinkRowStatus must be set to 'createAndWait' or 'createAndGo.'
        A successful set of the object imaLinkGroupIndex must be
        performed before the RowStatus of this new row can transition
        to the 'Active' state. The imaLinkGroupIndex provides the
        association between a physical IMA link to an IMA group.

        To change (modify) an imaLinkTable entry, the imaLinkRowStatus
        object must be set to 'notInService'. Only then can the
        conceptual row be changed. Setting object imaLinkGroupIndex to
        a different value has the effect of changing the association
        between a physical IMA link and an IMA group. To place the
        link 'in group', the imaLinkRowStatus object is set to 'Active'.
        While the row is not in 'Active' state, both the Transmit and
        Receive IMA state machines are in the 'Not In Group' state.

        To remove (delete) an imaLinkTable entry from this table, set
        this object to 'destroy.'

        A network manager may elect to perform a delete operation
        followed immediately by a create operation instead of a
        modify operation. The net effect of the delete/create is to
        re-instantiate (reset) agent counters defined in this table."
      ::= { imaLinkEntry 2 }

imaLinkGroupIndex OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-create
      STATUS      current
      DESCRIPTION
        "The value which identifies the IMA group (imaGroupIndex)
        of which this link is a member."
      ::= { imaLinkEntry 3 }

imaLinkNeTxState  OBJECT-TYPE
      SYNTAX      ImaLinkState
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The current state of the near-end transmit link."
      ::= { imaLinkEntry 4 }
```

```
imaLinkNeRxState  OBJECT-TYPE
      SYNTAX      ImaLinkState
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The current state of the near-end receive link."
      ::= { imaLinkEntry 5 }

imaLinkFeTxState  OBJECT-TYPE
      SYNTAX      ImaLinkState
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The current state of the far-end transmit link as reported
        via ICP cells."
      ::= { imaLinkEntry 6 }

imaLinkFeRxState  OBJECT-TYPE
      SYNTAX      ImaLinkState
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The current state of the far-end receive link as reported
        via ICP cells."
      ::= { imaLinkEntry 7 }

imaLinkNeRxFailureStatus       OBJECT-TYPE
      SYNTAX      ImaLinkFailureStatus
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The current link failure status of the near-end receive link."
      ::= { imaLinkEntry 8 }

imaLinkFeRxFailureStatus       OBJECT-TYPE
      SYNTAX      ImaLinkFailureStatus
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The current link failure status of the far-end receive link
        as reported via ICP cells."
      ::= { imaLinkEntry 9 }

imaLinkTxLid      OBJECT-TYPE
      SYNTAX      INTEGER (0..31)
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The outgoing LID used currently on the link by the local end.
        This value has meaning only if the link belongs to an IMA group."
      ::= { imaLinkEntry 10 }

imaLinkRxLid      OBJECT-TYPE
      SYNTAX      INTEGER (0..31)
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
```

```
        "The incoming LID used currently on the link by the remote
        end as reported via ICP cells. This value has meaning only
        if the link belongs to an IMA group."
      ::= { imaLinkEntry 11 }

imaLinkRelDelay   OBJECT-TYPE
      SYNTAX      MilliSeconds
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The latest measured delay on this link relative to the link, in
        the same IMA group, with the least delay."
      ::= { imaLinkEntry 12 }

imaLinkImaViolations       OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "ICP violations: count of errored, invalid or missing ICP
        cells during non-SES-IMA condition."
      ::= { imaLinkEntry 13 }

imaLinkOifAnomalies        OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The number of OIF anomalies during non-SES-IMA condition at the
        near-end. This is an optional attribute."
      ::= { imaLinkEntry 14 }

imaLinkNeSevErroredSec       OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Count of one second intervals containing several IV-IMA, or
        one or more link defects (e.g., LOS, OOF/LOF, AIS, LCD), LIF, LODS
        defects during non-UAS-IMA condition."
      ::= { imaLinkEntry 15 }

imaLinkFeSevErroredSec       OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Count of one second intervals containing one or more
        RDI-IMA defects."
      ::= { imaLinkEntry 16 }

imaLinkNeUnavailSec       OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Count of unavailable seconds at near-end: unavailability begins
```

```
        at the onset of 10 contiguous SES-IMA and ends at the onset
        of 10 contiguous seconds with no SES-IMA."
      ::= { imaLinkEntry 17 }

imaLinkFeUnavailSec       OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Count of unavailable seconds at far-end: unavailability begins
        at the onset of 10 contiguous SES-IMA-FE and ends at the onset of
        10 contiguous seconds with no SES-IMA-FE."
      ::= { imaLinkEntry 18 }

imaLinkNeTxUnusableSec     OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Tx Unusable seconds: count of Unusable seconds at the near
        end Tx LSM."
      ::= { imaLinkEntry 19 }

imaLinkNeRxUnusableSec     OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Rx Unusable seconds: count of Unusable seconds at the near
        end Rx LSM."
      ::= { imaLinkEntry 20 }

imaLinkFeTxUnusableSec     OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Tx Unusable seconds at far-end: count of seconds with Unusable
        indications from the far-end Tx LSM."
      ::= { imaLinkEntry 21 }

imaLinkFeRxUnusableSec     OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Rx Unusable seconds at far-end: count of seconds with Unusable
        indications from the far-end Rx LSM."
      ::= { imaLinkEntry 22 }

imaLinkNeTxNumFailures       OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The number of times a near-end transmit failure alarm condition
        has been entered on this link."
```

```
      ::= { imaLinkEntry 23 }

imaLinkNeRxNumFailures       OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The number of times a near-end receive failure alarm condition
        has been entered on this link (i.e., LIF, LODS, Mis-Connected,
        RDI-IMA)."
      ::= { imaLinkEntry 24 }

imaLinkFeTxNumFailures       OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The number of times a far-end transmit failure alarm condition
        has been entered on this link. This is an optional attribute."
      ::= { imaLinkEntry 25 }

imaLinkFeRxNumFailures       OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The number of times a far-end receive failure alarm condition
        has been entered on this link (i.e., LIF, LODS, Mis-Connected, RDI-IMA).
        This is an optional attribute."
      ::= { imaLinkEntry 26 }

imaLinkTxStuffs  OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
       "Counts of stuff events inserted in the transmit direction.
       This is an optional attribute."
      ::= { imaLinkEntry 27 }

imaLinkRxStuffs  OBJECT-TYPE
      SYNTAX      Integer32
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Counts of stuff events detected in the receive direction.
        This is an optional attribute."
      ::= { imaLinkEntry 28 }


--
-- Test Pattern Procedure control objects. These objects are implemented
-- if the IMA implements the Test Pattern Procedure. In this case all
-- test pattern procedure related objects must be implemented. Specifically
-- these are
--
-- imaGroupTestLinkIfIndex
```

```
-- imaGroupTestPattern
-- imaGroupTestProcStatus
-- imaLinkTestProcStatus
-- imaLinkRxTestPattern
--

imaLinkRxTestPattern  OBJECT-TYPE
      SYNTAX      INTEGER (0..255)
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "This object identifies the test pattern received in the
        ICP Cell (octet 17) on the link during the IMA Test Pattern
        Procedure. This value may then be compared to the transmitted
        test pattern."
      ::= { imaLinkEntry 29 }

imaLinkTestProcStatus OBJECT-TYPE
      SYNTAX      ImaTestProcStatus
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "This value indicates the current state of the Test Pattern
        Procedure. If the value is disabled(1), the test is not
        running. A value of operating(2) means the test is running
        and no error has been found on this interface. A value of
        linkfail(3) means an error has been detected on this link
        during the test.

        Once an error occurs, the linkfail(3) value is latched until
        either this object is read or until the imaGroupTestProcStatus
        is moved to disabled(1). Once read, if the error no longer
        persists, a subsequent read will report the value operating(2)."
      ::= { imaLinkEntry 30 }


--
-- IMA Failure Alarms
--
-- Implementation of the imaFailureAlarm Trap is optional.
--

imaAlarmStatus OBJECT-TYPE
      SYNTAX      ImaAlarmStatus
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "Status of the IMA alarm."
       ::= { atmfImaMibObjects 5 }

imaAlarmType OBJECT-TYPE
      SYNTAX      ImaAlarmType
      MAX-ACCESS  read-only
      STATUS      current
      DESCRIPTION
        "The Type of IMA alarm declared or cleared. The value of
        ImaAlarmType identifies the type of alarm according to the
```

```
        definitions in the IMA specification."
       ::= { atmfImaMibObjects 6 }

imaFailureAlarm NOTIFICATION-TYPE
      OBJECTS    { ifIndex, imaAlarmStatus, imaAlarmType }
      STATUS     current
      DESCRIPTION
        "The imaFailureAlarm provides a method for an agent implementing IMA
        to notify an NMS of an alarm condition."
      ::= { atmfImaMibObjects 7 }

-- Conformance Information

   atmfImaMibConformance OBJECT IDENTIFIER ::= { atmfImaMib 2 }

   atmfImaMibGroups        OBJECT IDENTIFIER ::= { atmfImaMibConformance 1 }

   atmfImaMibCompliances   OBJECT IDENTIFIER ::= { atmfImaMibConformance 2 }


-- Compliance Statements

atmfImaMibCompliance       MODULE-COMPLIANCE
      STATUS      current
      DESCRIPTION
        "The compliance statement for network elements implementing
        Inverse Multiplexing for ATM."

      MODULE IMA-MIB

-- Mandatory Part

      MANDATORY-GROUPS  { atmfImaGroupGroup, atmfImaLinkGroup,
         atmfImaGroupMappingTableGroup }

      GROUP imaTestPatternGroup
      DESCRIPTION
        "This group is mandatory only for implementations that support
        the Test Pattern Procedure. This group contains objects which
        control and report on the Test Pattern Procedure."

-- Compliance Part

   OBJECT     imaGroupSymmetry
   MIN-ACCESS read-only
   DESCRIPTION
     "Write access is not required."

   OBJECT     imaGroupMinNumTxLinks
   MIN-ACCESS read-only
   DESCRIPTION
     "Write access is not required."

   OBJECT     imaGroupMinNumRxLinks
   MIN-ACCESS read-only
   DESCRIPTION
     "Write access is not required."
```

```
   OBJECT     imaGroupTxImaId
   MIN-ACCESS read-only
   DESCRIPTION
     "Write access is not required."

   OBJECT    imaGroupTxFrameLength
   MIN-ACCESS   read-only
   DESCRIPTION
     "Write access is not required."

   OBJECT     imaGroupDiffDelayMax
   MIN-ACCESS read-only
   DESCRIPTION
     "Write access is not required."

   OBJECT     imaGroupAlphaValue
   MIN-ACCESS read-only
   DESCRIPTION
     "Write access is not required."

   OBJECT     imaGroupBetaValue
   MIN-ACCESS read-only
   DESCRIPTION
     "Write access is not required."

   OBJECT     imaGroupGammaValue
   MIN-ACCESS read-only
   DESCRIPTION
     "Write access is not required."

   OBJECT     imaGroupTestLinkIfIndex
   MIN-ACCESS read-only
   DESCRIPTION
     "Write access is not required."

   OBJECT     imaGroupTestPattern
   MIN-ACCESS read-only
   DESCRIPTION
     "Write access is not required."

   ::= { atmfImaMibCompliances 1 }

-- Units of Conformance

   atmfImaGroupGroup    OBJECT-GROUP
   OBJECTS {
      imaGroupNumber, imaGroupRowStatus, imaGroupIfIndex,
      imaGroupIndex, imaGroupNeState, imaGroupFeState, imaGroupFailureStatus,
      imaGroupSymmetry, imaGroupMinNumTxLinks, imaGroupMinNumRxLinks,
      imaGroupNeTxClkMode, imaGroupFeTxClkMode,
      imaGroupTxTimingRefLink, imaGroupRxTimingRefLink, imaGroupLastChange,
      imaGroupTxImaId, imaGroupRxImaId, imaGroupTxFrameLength,
      imaGroupRxFrameLength, imaGroupDiffDelayMax, imaGroupLeastDelayLink,
      imaGroupDiffDelayMaxObs, imaGroupAlphaValue, imaGroupBetaValue,
      imaGroupGammaValue, imaGroupRunningSecs, imaGroupUnavailSecs,
      imaGroupNeNumFailures, imaGroupFeNumFailures, imaGroupTxAvailCellRate,
```

```
      imaGroupRxAvailCellRate, imaGroupNumTxCfgLnks, imaGroupNumRxCfgLnks,
      imaGroupNumTxActLnks, imaGroupNumRxActLnks }
   STATUS     current
   DESCRIPTION
     "A set of objects providing configuration and status information for
     an IMA group definition."
   ::= { atmfImaMibGroups 1 }

   atmfImaLinkGroup    OBJECT-GROUP
   OBJECTS {
      imaLinkIfIndex, imaLinkRowStatus, imaLinkGroupIndex,
      imaLinkNeTxState, imaLinkNeRxState, imaLinkFeTxState,
      imaLinkFeRxState, imaLinkNeRxFailureStatus, imaLinkFeRxFailureStatus,
      imaLinkTxLid, imaLinkRxLid, imaLinkRelDelay,
      imaLinkImaViolations, imaLinkOifAnomalies, imaLinkNeSevErroredSec,
      imaLinkFeSevErroredSec, imaLinkNeUnavailSec, imaLinkFeUnavailSec,
      imaLinkNeTxUnusableSec, imaLinkNeRxUnusableSec, imaLinkFeTxUnusableSec,
      imaLinkFeRxUnusableSec, imaLinkNeTxNumFailures, imaLinkNeRxNumFailures,
      imaLinkFeTxNumFailures, imaLinkFeRxNumFailures, imaLinkTxStuffs,
      imaLinkRxStuffs }
   STATUS     current
   DESCRIPTION
     "A set of objects providing status information for an IMA link."
   ::= { atmfImaMibGroups 2 }

   atmfImaGroupMappingTableGroup OBJECT-GROUP
   OBJECTS { imaGroupMappingIndex }
   STATUS     current
   DESCRIPTION
     "A table mapping the 'ifIndex' values of 'imaGroupIfIndex'
     to the 'imaGroupIndex' values of the corresponding
     IMA group."
   ::= { atmfImaMibGroups 3 }

   imaTestPatternGroup OBJECT-GROUP
    OBJECTS { imaGroupTestLinkIfIndex, imaGroupTestPattern,
              imaGroupTestProcStatus, imaLinkTestProcStatus,
              imaLinkRxTestPattern }
   STATUS     current
   DESCRIPTION
     "Objects in the imaGroupTable and imaLinkTable which control and
     report on the Test Pattern Procedure. These objects must be
     implemented if the IMA Test Pattern Procedure is supported."
   ::= { atmfImaMibGroups 4 }

   imaAlarmGroup OBJECT-GROUP
   OBJECTS { imaAlarmStatus, imaAlarmType }
   STATUS     current
   DESCRIPTION
     "Objects used in the imaFailureAlarm notification."
   ::= { atmfImaMibGroups 5 }

END
```

16. References

[1] ANSI T1.101-1987 “Synchronization Interface Standards for Digital Networks”.

[2] ANSI T1.102-1993 “Digital Hierarchy - Electrical Interfaces”.

[3] ANSI T1.231-1993 “Digital Hierarchy - Layer 1 In-Service Digital Transmission Performance Monitoring”.

[4] ANSI T1.403-1995 “Network-to-Customer Installation-DS1 Metallic Interface”.

[5] ANSI T1.408-1990 “Integrated Services Digital Network (ISDN) Primary Rate - Customer Installation Metallic Interfaces Layer 1 Specification.

[6] ANSI T1.646-1995, Broadband-ISDN - Physical Layer Specification for User-Network Interfaces Including DS1/ATM, 1995.

[7] ATM Forum AF-BICI-0013.003 “B-ISDN Inter Carrier Interface (B-ICI) Specification Version 2.0 (Integrated), December 1995.

[8] ATM Forum AF-PHY-0016.000 “DS1 Physical Layer Specification”, September, 1994.

[9] ATM Forum AF-PHY-0064.000 “E1 Physical Layer Interface Specification”, July 1996.

[10] ATM Forum AF-PNNI-0055.000 “Private Network-Network Interface Specification Version 1.0”, March 1996.

[11] ATM Forum User-Network Interface Specification, Version 3.1, May 1994.

[12] Harel, D., 1987, “State Charts: A visual formalism for complex systems”, Science of Computer Programming, Volume 8.

[13] IETF RFC 1155, “Structure and Identification of Management Information for TCP/IP-based Internets”, May 1990.

[14] IETF RFC 1156, “Management Information Base for Network Management of TCP/IP-based internets”, May 1990.

[15] IETF RFC 1157, “A Simple Network Management Protocol (SNMP)”, May 1990.

[16] IETF RFC 1212, “Concise MIB definitions”, March 1991.

[17] IETF RFC 1213, “Management Information Base for Network Management of TCP/IP-based internets: MIB-II”, March 1991.

[18] IETF RFC 1406, “Definitions of Managed Objects for the DS1 and E1 interface types”, January 1993.

[19] IEFT RFC 1442, “Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2)”, April 1993.

[20] IEFT RFC 1443, “Textual Conventions for version 2 of the Simple Network Management Protocol (SNMPv2)”, April 1993.

[21] IEFT RFC 1444, “Conformance Statements for version 2 of the Simple Network Management Protocol (SNMPv2)”, April 1993.

[22] IEFT RFC 1445, “Administrative Model for version 2 of the Simple Network Management Protocol (SNMPv2)”, April 1993.

[23] IEFT RFC 1446, “Security Protocols for version 2 of the Simple Network Management Protocol (SNMPv2)”, April 1993.

[24] IEFT RFC 1447, “Party MIB for version 2 of the Simple Network Management Protocol (SNMPv2)”, April 1993.

[25] IEFT RFC 1448, “Protocol Operations for version 2 of the Simple Network Management Protocol (SNMPv2)”, April 1993.

[26] IEFT RFC 1449, “Transport Mappings for version 2 of the Simple Network Management Protocol (SNMPv2)”, April 1993.

[27] IEFT RFC 1450, “Management Information Base for version 2 of the Simple Network Management Protocol (SNMPv2)”, April 1993.

[28] IEFT RFC 1451, “Manager-to-Manager Management Information Base”, April 1993.

[29] IEFT RFC 1452, “Coexistence between version 1 and version 2 of the Internet-standard Network Management Framework”, April 1993.

[30] IETF RFC 1573 “Evolution of the Interfaces Group of MIB-II”, January 1994.

[31] ISO/IEC 11573.1994(E) “Information technology - Telecommunications and information exchange between systems - Synchronization methods and technical requirements for Private Integrated Service Networks”.

[32] ITU-T Recommendation G.703 “Physical/Electrical Characteristics of Hierarchical Digital Interfaces”, 1991.

[33] ITU-T Recommendation G.810, “Considerations on Timing and Synchronization Issues”, 1988.

[34] ITU-T Recommendation G.811, “Timing Requirements at the Outputs of Primary Reference Clocks Suitable for Plesiochronous Operation of International Digital Links”, 1988.

[35] ITU-T Recommendation G.823, “The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy, 1988.

[36] ITU-T Recommendation G.824, “The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy, 1988.

[37] ITU-T Recommendation G.826, “Error performance parameters and objectives for international constant bit rate digital paths at or above the primary rate, 1993.

[38] ITU-T Recommendation I.321, “B-ISDN Protocol Reference Model and its applications”, March 1993.

[39] ITU-T Recommendation I.361, “B-ISDN ATM Layer Specification”, March 1993.

[40] ITU-T Recommendation I.431 “Primary Rate User-Network Interface - Layer 1 specification”, March 1993.

[41] ITU-T Recommendation I.432 Series, “B-ISDN User-Network Interface - Physical Layer Specification”, April 1996.

[42] ITU-T Recommendation I.610, “B-ISDN Operation and Maintenance Principles and Functions”, 1995.

Appendix A Example of Link Delay Synchronization

This section describes an example on how link delay synchronization may be achieved.

The delay compensation at the receive end may be done using the IMA framing number and ICP cell offset contained in the incoming ICP cells.

Figure 26 on page 107 shows an IMA group with three links called Link 0, Link 1 and Link 2. The link multiplexing order is 0-> 1-> 2-> 0, etc. The ICP cells have been identified according to the LID of the link on which they are sent. The cells other than the ICP cells have been numbered in the order they are sent over the virtual link.

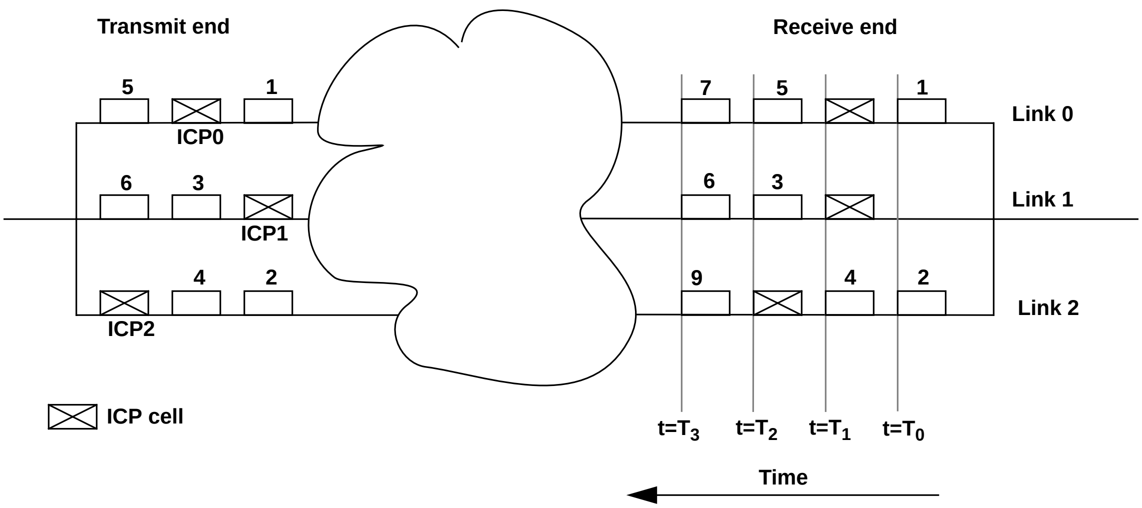

Figure 26 Link Delay Compensation example

The transmit end has staggered the ICP cells giving them ICP Cell Offset of 1, 0 and 2 for links 0, 1 and 2 respectively. Again, these values are used for illustration purposes only. It is recommended that ICP cell position from link-to-link be spaced out as evenly as possible in an IMA frame. At the receive end, link 0 and link 2 have the same amount of propagation delay but link 1 has a delay one cell time longer than link 0 or link 2. At the receive end, each link has its own circular buffer which should be deep enough to tolerate the maximum link delay variation.

Figure 27 on page 108 to Figure 30 on page 109 shows an example of the processing of the incoming cells at the receive end. One Delay Compensation Buffer (DCB) is defined per link. The shaded portions of DCBs on Figure 27 to Figure 30 represent unfilled portions of the buffer allocated for each link. Each section of the buffer represents a single cell buffer. The unshaded portion represents portions filled with cells.

Figure 28 to Figure 30 also shows the contents of the three Delay Compensation Buffers (DCB) on the receive end at time t = T_1, T_2, T_3 and T_6. The write pointer of each link shows the next location where a cell would be put when received from that link. At time T_0, no ICP cells have been received yet on any links so no cells are written to the buffers. Whenever a cell is written into a buffer, the write pointer is incremented. A read (playback) pointer increments as cells are read out of the buffer. Both pointers wrap around when they reach the end of the buffer. Until the first ICP cell is received, no cell is written into the buffer. When the first ICP cell is received, the write pointer is initialized to the value of the IMA frame number plus ICP cell offset contained in that ICP cell. Every cell received after that ICP cell is written into the buffer.

All of the delay compensation is done with the adjustment of the write pointers. The read pointers are always aligned and the cells are read out of the buffers in the assigned link multiplexing order. The differential delay between any two links can be simply calculated using the difference between the write pointer values for those two links. Extending this idea to the entire set of links, the maximum differential delay for the link group is simply the difference between the minimum and maximum write pointer values for that group. If we assume the use of DS1 links and given a maximum allowable delay delta of 25 ms, the maximum difference acceptable between the write pointers will be 91cells. The accuracy of this technique is ± 1 cell time (e.g., 276 microseconds at DS1 rate) which is satisfactory for this application.

The process of reading from the buffer is controlled by the IMA cell clock. In round-robin order and at the IMA clock tick, a cell is removed from the DCB (the pointer is advanced). If the cell is an ICP cell, the next link in order is immediately selected (and so on until a cell is found that is not an ICP cell). If the cell is a Filler cell, it is discarded and nothing is passed to the ATM layer. It the cell is not a Filler cell, it is passed to the ATM layer.

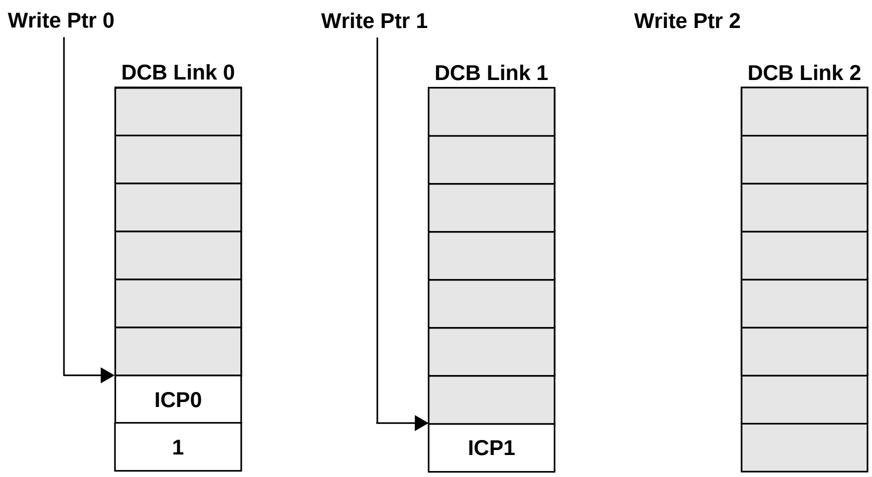

Figure 27 Snapshot of Buffers at T_1

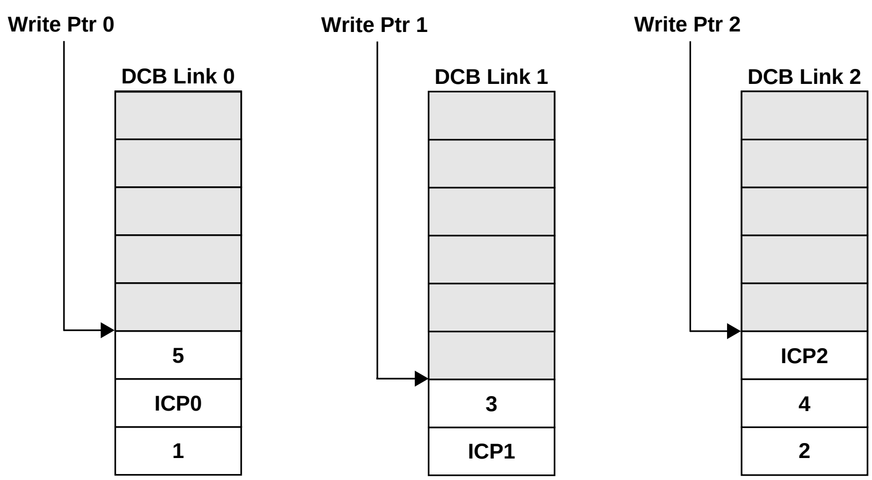

Figure 28 Snapshot of Buffers at T_2

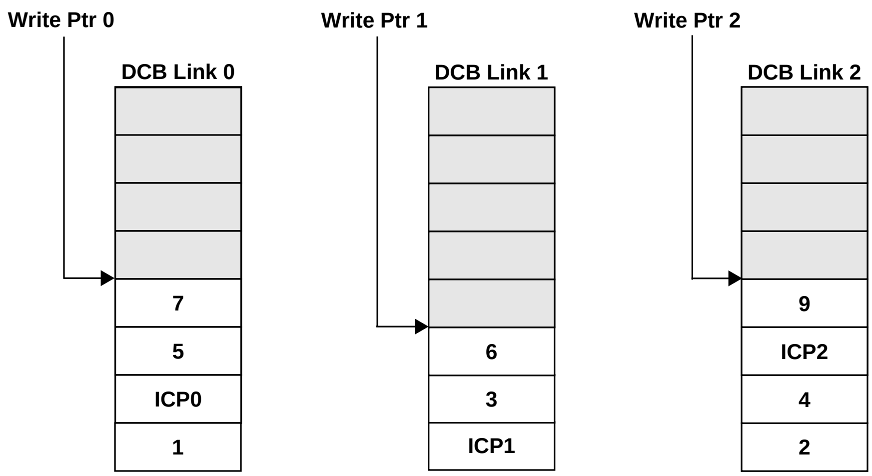

Figure 29 Snapshot of Buffers at T_3

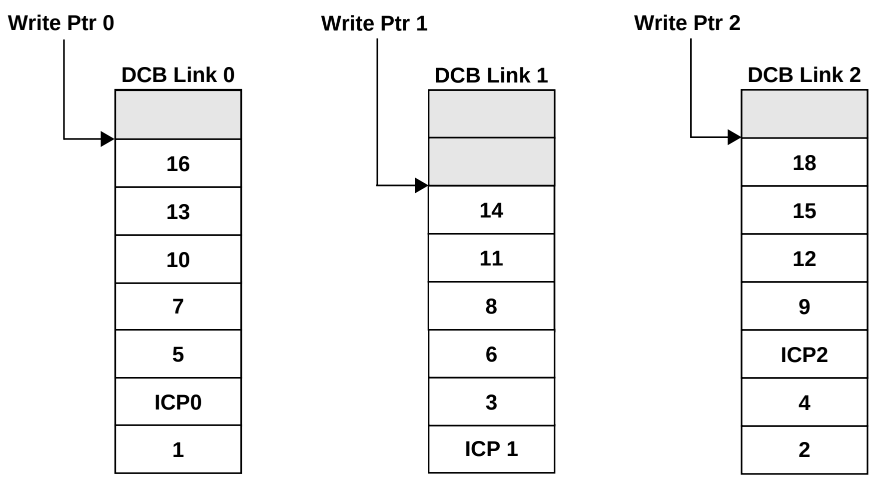

Figure 30 Snapshot of Buffers at T_6

Appendix B Implementation Consideration

This section may be used as guidelines for IMA implementers.

B.1 Dynamic Changes in ATM Layer Cell Rate

This section provides some guidelines to treat dynamic changes of IMA data cell rate (IDCR).

The ATM Layer needs to react to dynamic changes in the IDCR, particularly to decreases in the cell rate since decreases in the cell rate are problematic. Connections at the previous (larger) cell rate reserved some quantity of bandwidth (cell rate). If the interface cell rate drops (due to actions within the IMA), then the aggregate reserved cell rate of the connections may exceed the interface cell rate and result in higher Cell Loss Ratio (CLR) and delay than contracted.

Dynamic changes can occur in three types of circumstances: (1) a one or more links fail, (2) one or more link are removed administratively, and (3) one or more link are added administratively.

In case (1), the IMA interface may or may not experience an increase of CLR and delay on some or all connections due to a decrease of available bandwidth. In some cases, both ends of the virtual links may not release connections and let congestion control discard cells if severe congestion occurs. In other cases, each end may release connections to preserve QoS on the most important connections. The handling of connections upon dynamic bandwidth decrease has been under consideration by several ATM Forum Working Groups (NM, TM, PNNI, and SIG Working Groups).

In case (2), the IMA is instructed by the management control plane to drop a link. In this case, the administrative plane function should have taken steps to ensure that the connections on the interface would still pass Connection Admission Control (CAC) at the reduced interface cell rate before instructing the IMA to drop the link. It is conceivable that the administrative function may even preempt connections so those remaining will fall within CAC limits. In any case, close coordination between the CAC and the IMA control functions is necessary if dynamic additions/deletions are to occur.

In case (3), the IMA is instructed by the management control plane to add a link. In this case, the administrative plane function should have informed the CAC that additional bandwidth is present. Otherwise new connections will not be admitted, even though more bandwidth is available.

The treatment of CAC in the ATM Forum UNI and PNNI specifications [11][10] is implementation dependent. Thus the treatment of CAC under dynamic bandwidth changes is left to the discretion of the ATM interface implementer.

B.2 State Machine for De-Bouncing Error Condition Report

The Rx Failed condition, described in Section 10.1.4 on page 47, and the link failure alarm, such as defined in Section 12.2.3 on page 76, may be a direct mapping of the defect state machine presented in Figure 20 on page 70, but this may cause lengthening of short errors. Repetitious bursts of errors will also affect the data path if links are brought back in too rapidly. Since there were disparate views as to integration times, a generic state model can be employed, with events loosely defined. Figure 31 on page 111 shows a simplified Rx Failed integration state machine. The integration times can be left as vendor-defined parameters.

Another standard technique is the weighted integration into the error condition (e.g., Rx Failed or failure condition). For each sampling period, if a defect is present, a value is added to an integration counter. If no defect is present, a smaller value is subtracted. In this way, errored samples are weighted more heavily. A typical ratio is 10:1.

The state machine could also be used for de-bouncing any other error condition report such as “Fault” condition report.

The model described above does not apply well to the IMA-specific checks such as IMA framing or link delay synchronization. These take longer to detect than defects, and more latitude may be required. This state machine can be used as an example, leaving X and Y and specific integration techniques to the implementer.

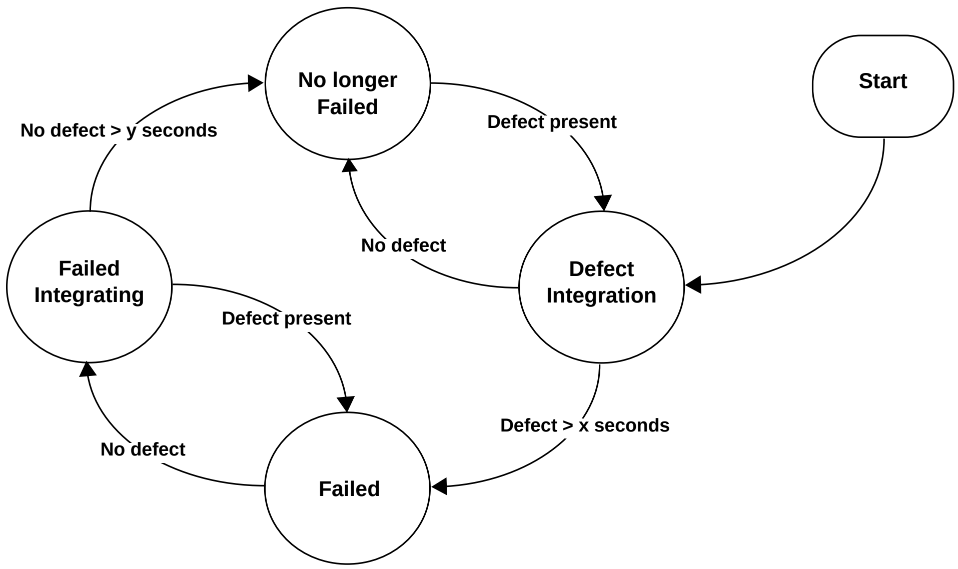

Figure 31 Rx Failed State Machine

B.3 Maximum Link Differential Delay Limitation

This section clarifies an aspect of the IMA baseline text regarding the maximum link differential delay that can be supported by the IMA interface due to the nature of the IMA Frame. Having the following characteristic of the IMA frame and a newly defined IMA super-frame:

- IMA frame size: M cells (32, 64, 128 or 256)
- IMA frame sequence number (IFSN) range: 0 to 255
- IMA Super-Frame (ISF) definition: starts with IMA frame with sequence number 0 and ends with IMA frame with sequence number 255.
- ISF length: 256*M

The ISF is defined to illustrate that the IFSN repeats every 256 IMA frames. The IFSN is used by the receiver to compensate for links with longer propagation delay. The receiver then always relies on the IFSN contained in the incoming ICP cells on each links. If the differential delay between the links with the shortest and longest propagation delay is larger than one half of the ISF transmission time, the receiver will not be able to determine which link has the shortest delay. The difference has to be modulo 256, which means that the receiver could think that the difference is either greater or smaller than one half of the ISF. The only way to avoid the confusion is to ensure that the differential link delay between the links with the shortest and the longest propagation delays is never greater than one half of the time to transmit the ISF.

Table on page 112 shows the maximum link differential delay that can be supported by the IMA interface in the case DS1/E1 links are used to form an IMA virtual link.

Table 21 Maximum Differential Delay for Different M Values over various Links

M (cells per IMA Frame)	Maximum Differential delay (in milliseconds)		
	DS1	E1	DS3 (Direct Mapping)
32	1130	905	39.3
64	2261	1809	78.6
128	4522	3618	175.1
256	9044	7236	314.3

Although the sizes of the delay listed above is extremely high, it is important to highlight this limitation of the IMA protocol. This would be particularly important if the IMA protocol is implemented over higher speed interfaces.

The information can also be expressed another way to give the maximum data rates that can be supported if the required maximum link differential delay is limited to 25 milliseconds. Table 22 on page 112 shows the maximum link rates that may be supported for different M. The data rate is the rate of the cell stream excluding any framing overhead.

Table 22 Maximum Data Rate for a Maximum Link differential Link Delay of 25 msec

M (cells per IMA Frame)	Maximum Data Rate for a maximum link differential delay of 25 msec	
	Bit Rate (Mbps)	Cell Rate (kcell/sec)
32	69.5	163.84
64	139	327.68
128	278	655.36
256	555	1310.72

B.4 Clock Variation Tolerance for IMA

There are three aspects to be considered for determining the required clock variation tolerance for proper IMA operation:

- required number of consecutive ICP cells required for performing stuffing,
- periodically stuffing on the timing reference link every 2048 cells, and
- maximum insertion rate of stuff event insertion rate

B.4.1 Tolerance Due to IMA Stuffing Procedure

The specification allows for indication of the stuff action over the previous 4 ICP cells (Section 7 on page 36). This excludes the two ICP cells that comprises the stuff event. This implies a maximum frequency of adjustment of once every 5 ICP cells, or every 5 IMA frames.

For an IMA frame with the default value of M=128, 5 IMA frames are 640 cells. This implies that links can have a total variance of 1 cell time out of 640 cell times, or 1/640. Converting this to ppm by multiplying by 10^6 gives 1562. The frequency variation is then 781 ppm. Table on page 113 gives these calculations for various values of M.

Table 23 M Implications on PPM

M (cells per IMA Frame)	5M (number of cells)	Total variance (ppm)	Frequency variation (± ppm)
32	160	6250	3125
64	320	3125	1562
128	640	1562	781
256	1280	781	390

The smallest limit on ppm variation, ± 390 ppm, is substantially more than that for the one required on DS1/E1 links (± 32 ppm and ± 50 ppm respectively). Accordingly the IMA stuff procedure does not limit link selection further, and in fact can adjust for differences in link frequencies substantially more than are allowable.

Since there is the SICP cell itself, the adjustment period is 5M+1 cells.

B.4.2 Tolerance Due to Periodically Stuffing on the Timing Reference Link

Equation (EQ. 1) on page 38 requires the insertion of one cell every 2048 ICP, Filler and ATM layer cells transmitted over the timing reference link (TRL). This value, 2048, has been chosen because it is the next binary number (2^{11}) that allows to use the optional advanced stuff indication procedure for M equals to 256 (1280 is the minimum required as described in the previous section).

In the case the TRL is the faster link in the group, link frequency differences up to 488 ppm can be tolerated between the TRL and the slowest link in the group. This is assuming that this is no need to insert stuff events on the slowest link.

B.4.3 Tolerance Due to Maximum Stuff Event Insertion Rate

The maximum stuff event insertion rate is defined to be one stuff event every 5M ICP, Filler and ATM layer cells for accommodating configurations using the optional four advanced indications. Enforcing such value also allows to avoid cases where some implementations would not constantly regulate the insertion of the stuff events over time.

In the case where the TRL is the slowest link in the group, the choice of 5M where M equals to 128 allows link frequency differences up to 293 ppm between the TRL and the faster link in the group. This value is derived from the fact that there is a frequency difference of 781 ppm between a link on which a stuff event is inserted every 1280 ICP, Filler and ATM layer cells and another link on which there is no stuff event being inserted, and a frequency difference of 488 ppm between the TRL and a link on which there is no stuff event being inserted. The frequency difference between the TRL and the faster link, 293 ppm, is still quite larger than the ppm requirement of most link interfaces. Table 24 on page 114 shows the maximum tolerated frequency difference between the TRL and the fastest link.

Table 24 Maximum Tolerated Frequency Difference between the TRL and the Fastest Link

M (cells per IMA Frame)	5M (number of cells)	Total variance (ppm)	Maximum tolerated frequency difference between TRL and faster link
32	160	6250	5762
64	320	3125	2637
128	640	1562	1074
256	1280	781	293

B.5 Example of IMA Data Cell Rate Recovery Implementation

This section describes an example of IMA Data Cell Rate (IDCR) recovery mechanism. It suggests that the IMA receive interface to the ATM layer limit itself to specifying the cell order, the IMA data cell rate and aggregate CDV. The actual method to clock an individual cell from the IMA sublayer to the ATM layer is implementation and application specific.

Figure 32 on page 115 shows a partial block diagram of the receive IMA. Portions shown are relevant to the IDCR recovery of the IMA group. It consists of the Delay Compensation Buffers (DCB), the IMA Data Cell Clock (IDCC), the Read Pointer Functional Block (RPFB) and the output interface to the ATM layer.

The shaded portions of the DCB, in Figure 32, signifies unfilled portions of the buffer allocated for each link. Each section of the buffer represents a single cell buffer. The unshaded portion represents portions filled with cells. The write pointer of each link shows the next location where a cell would be put when received from that link. The read pointer points to the next cell that would be queued to the ATM layer interface.

Figure 32 shows the DBC for an IMA group composed of three links 0, 1, 2 (therefore N=3) with M=8 cells. M equal to 8 is used to simplify the figure. Link 0 is the link with the greatest delay with 3 cells in its buffer. Links 0, 1, and 2 have a least 3 cells in their DCB’s.

When encountering an ICP cell, the RPFB will ignore the cell and will immediately advance to the next link in order without having to wait for the next IDCC tick. When encountering a stuff event, the RPFB will ignore both ICP cells part of the stuff event and will immediately advance to the next link in order. The RPFB is not responsible to examine the ICP cell content. It can be assumed that incoming ICP cells are passed to another functional block of the IMA interface for further processing.

When encountering a Filler cell, the RPFB will ignore the cell and wait until the next IDCC tick to advance to the next link in order.

Figure 33 on page 116 shows the insertion of stuff cells every “Y” IMA frames on the TRL, where Y is 2048/M, or 256 if this example.

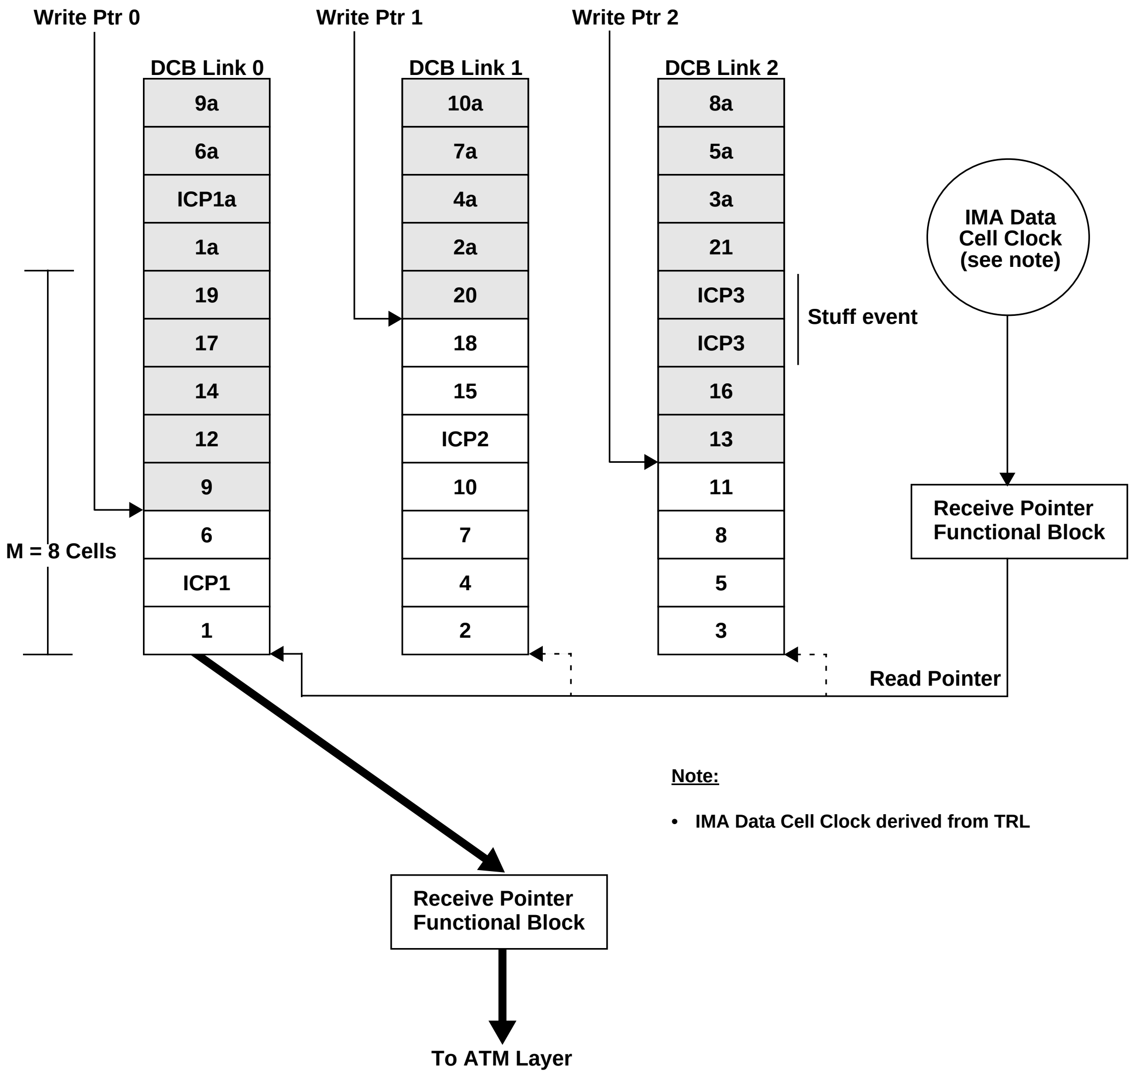

Figure 32 Serving Delay Compensation Buffers

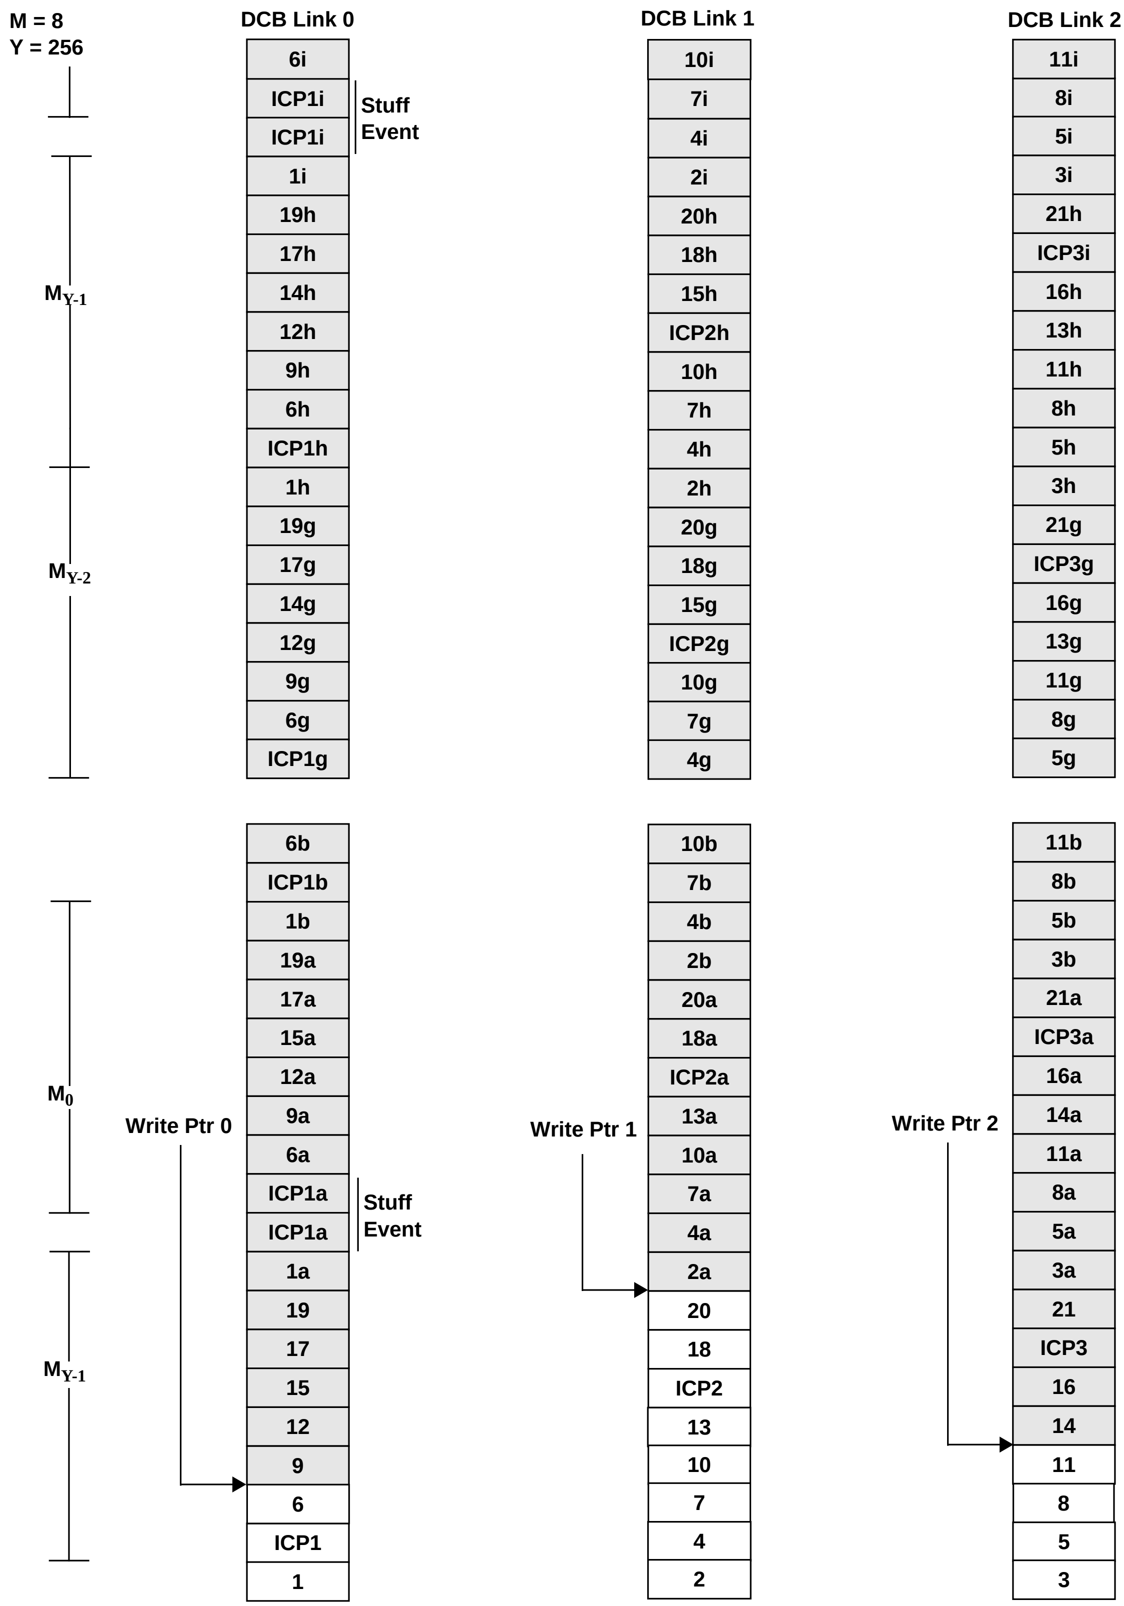

Figure 33 Insertion of Stuff Cell every Y*M Cells on the TRL

Appendix C Asymmetric Operation

C.1 Introduction

This appendix covers extensions of the IMA specification to the operations and configuration to cover two aspects of asymmetric links.

- Asymmetric configuration covers the configuration and static elements of IMA to allow the provision of asymmetric links, for example a group with 3 links in one direction and 2 in the other.
- Asymmetric operation treats the operation of the links independently separating the transmit and receive behavior and allowing the IMA to continue to use a link in one direction that has failed in the other.

C.2 Transmit and Receive State Transitions

The Transmit (Tx) and Receive (Rx) state transitions should operate independently, with independent setting of the Link Information fields in the ICP cells for the Tx and Rx ends. An IMA group set in the asymmetric operation but not supporting asymmetric configuration would always expect the Rx and Tx links to be either both in a group or both out of a group.

The NE Tx state transition does not consider the NE Rx state or the FE Tx state when deciding on the state of the NE Tx link (i.e., note 8 in Table 8 on page 49 applies).

The NE Rx state transition does not consider the NE Tx state or the FE Rx state when deciding on the state of the NE Rx link (i.e., note 8 in Table 9 on page 51 applies).

C.3 Link Start-up

There are two cases, one where the receive end is ready first, and a second where the transmit end is ready first. This is shown in Figure 34 on page 118 and Figure 35 on page 118.

C.4 Group Start-Up

For any group, two values will be provided, P_{Tx} and P_{Rx}. P_{Tx} and P_{Rx} as appearing in Figure 14 on page 60, may be different.

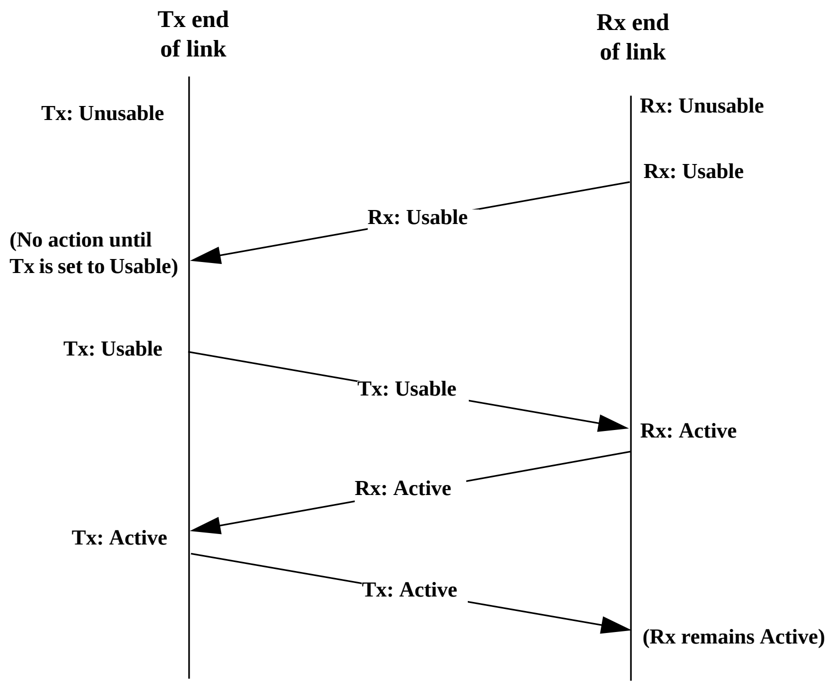

Note: Only the relevant NE state transitions and exchanges are shown.

Figure 34 Exchange of the NE Tx and Rx State during Link Start-up - Rx Ready First

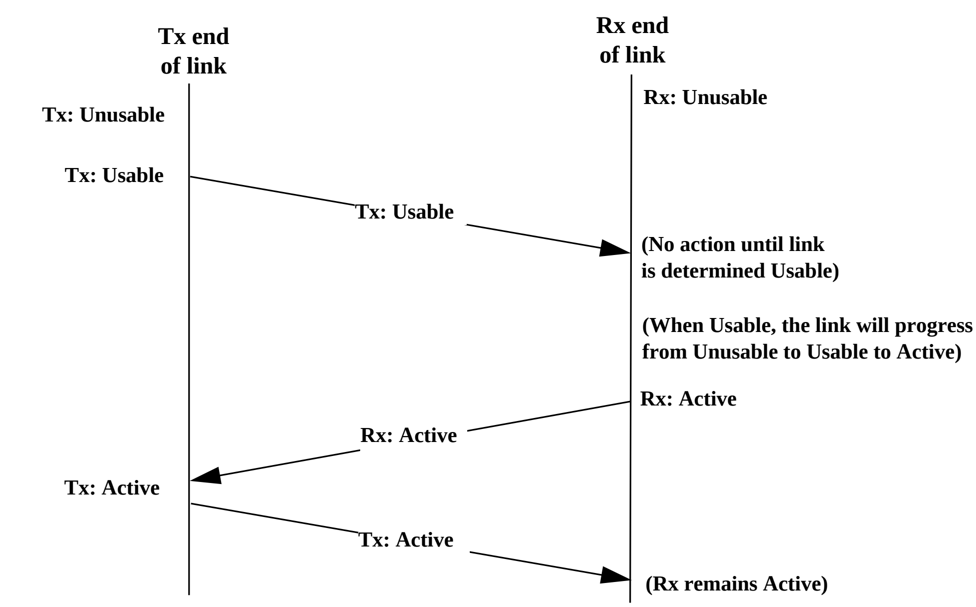

Note: Only the relevant NE state transitions and exchanges are shown.

Figure 35 Link Start-up - Tx Ready First

Appendix D Examples of Link State Exchanges

This section presents examples of state exchanges between the two ends of the IMA virtual link based on the LSM defined in Section 10.1 on page 43. The link addition, activation, deletion and de-activation cases are covered. Both symmetrical and asymmetrical configurations are also considered.

Figure 37 on page 120 shows the case where a link is added and activated simultaneously at both ends. Overlapping identical handshakes occur in both directions, controlling each direction independently.

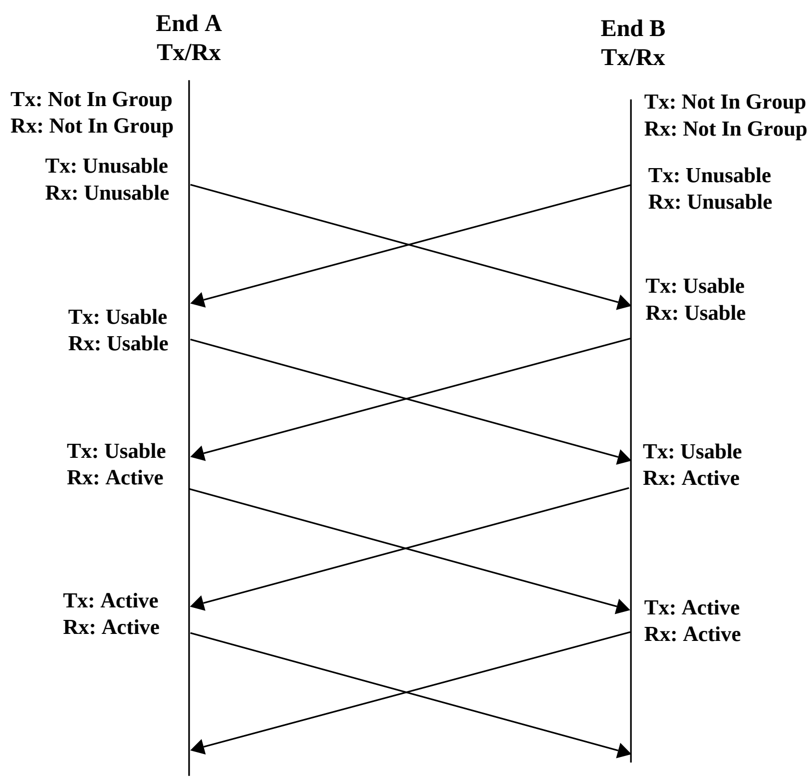

Note:
- **Only the NE states newly entered and transmitted over the ICP cells are shown.**
- **The transitions to Tx=Usable, Rx=Active and Tx=Active depend on the group wide synchronized addition/slow recovery procedure of links**

Figure 36 Link addition with Overlapping Sequence (Symmetrical Operation)

Figure 37 on page 120 shows the case when both ends are not overlapping. One end is moving to Usable before the other end. End B may safely skip to the receive Unusable state, and the response is immediate transition to receive Active state at the local end. Full operation is established in a single round-trip.

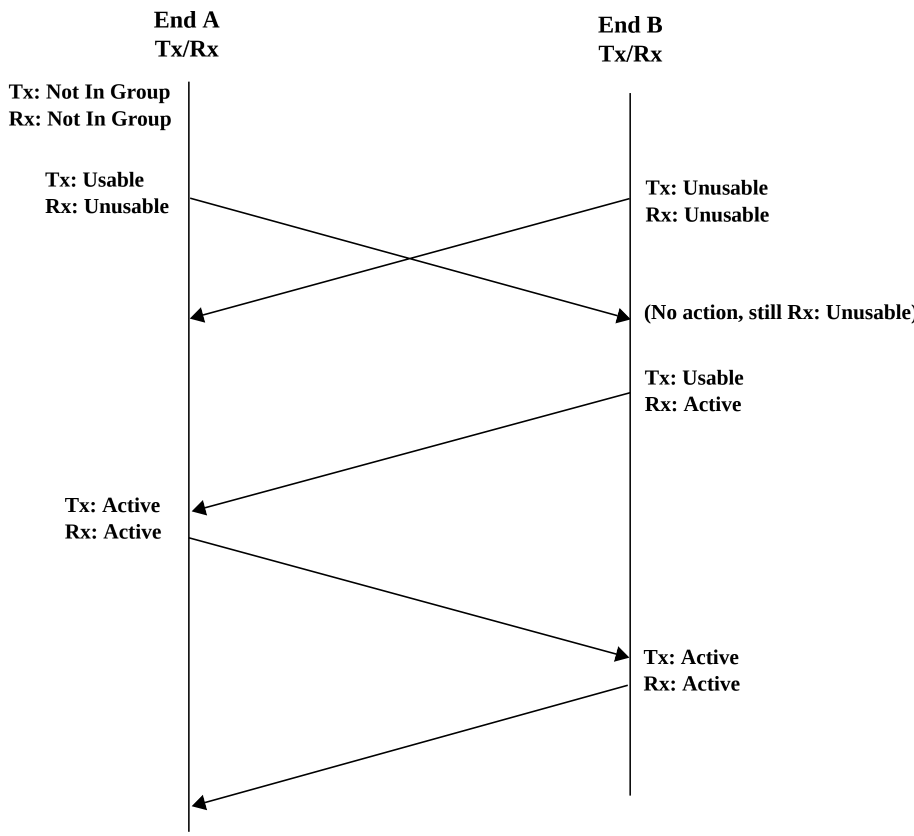

Note:

- **Only the NE states newly entered and transmitted over the ICP cells are shown.**
- **The transitions to Tx=Usable, Rx=Active and Tx=Active depend on the group wide synchronized addition/slow recovery procedure of links**

Figure 37 Link addition with Non-Overlapping Sequence (Symmetrical Operation)

Figure 38 on page 121 and Figure 39 on page 121 show the state exchange initiated by end A for deleting a link in the transmit direction for both the symmetrical and asymmetrical operations. Note that the receive link remains Active in the case the groups are asymmetrically operated.

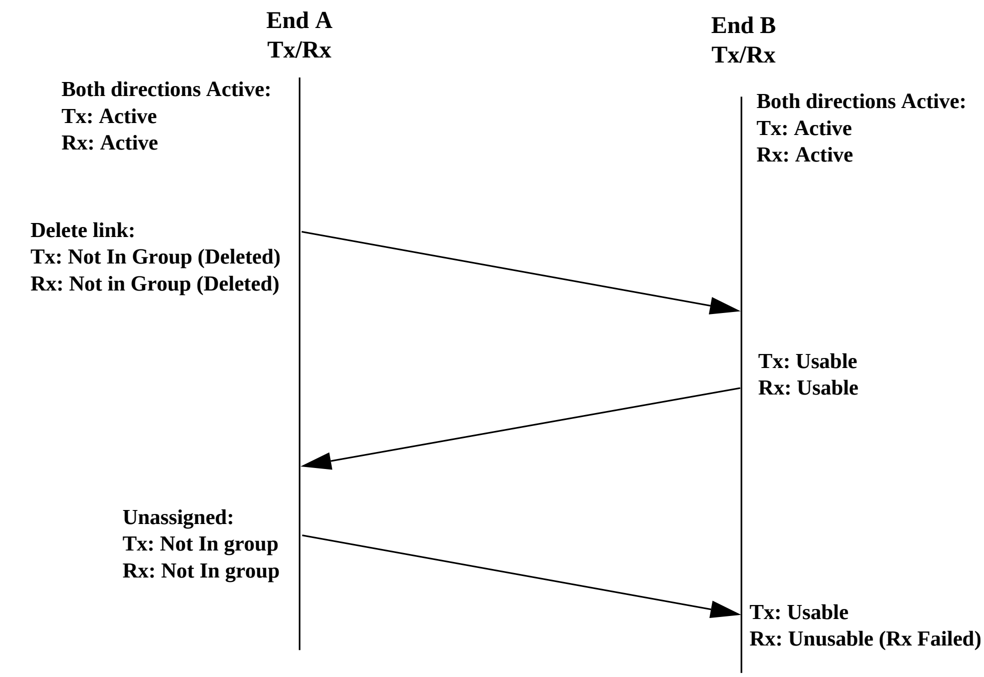

Note: Only the NE states newly entered and transmitted over the ICP cells are shown.

Figure 38 Transmit Link Deletion at One End (Symmetrical Operational)

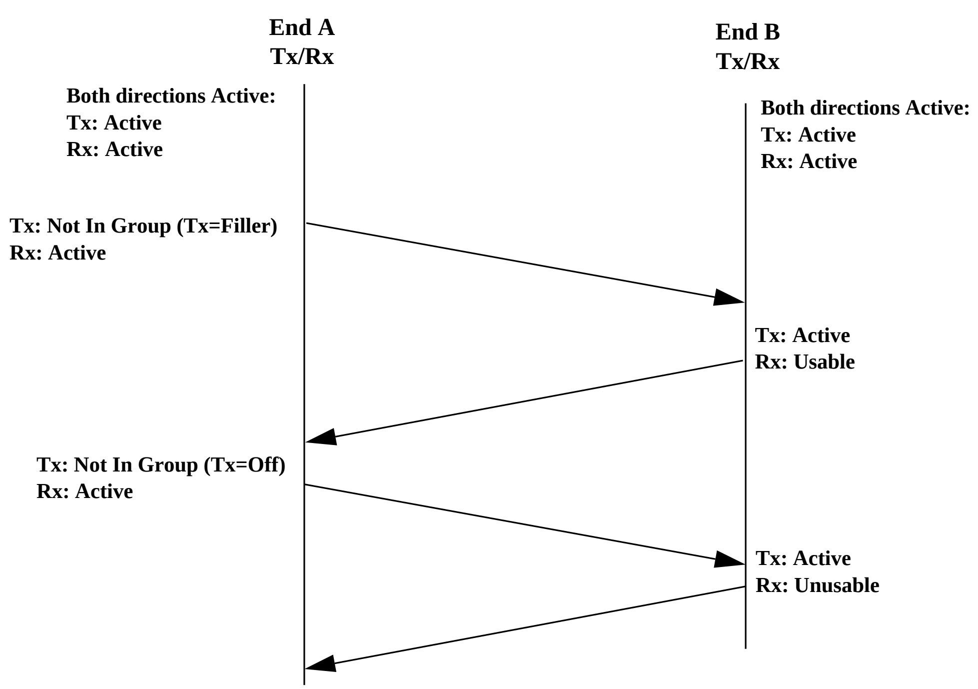

Note: Only the NE states newly entered and transmitted over the ICP cells are shown.

Figure 39 Transmit Link Deletion at One End (Asymmetrical Operational)

Figure 40 on page 122 shows the state exchange in the case of link de-activation in the case of a transmit failure in one direction under symmetrical configuration. Figure 41 on page 122 shows the scenario for the asymmetrical operation.

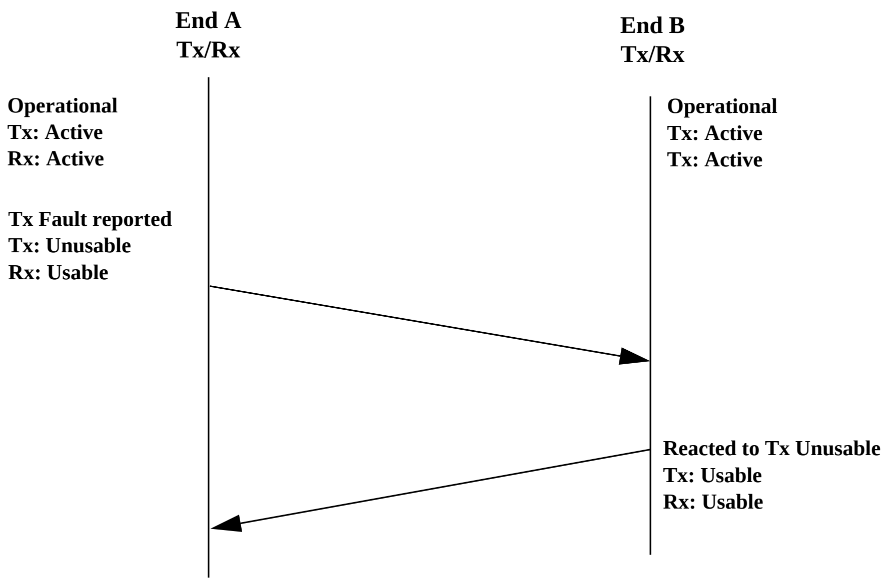

Note: Only the NE states newly entered and transmitted over the ICP cells are shown.

Figure 40 Link De-activation due to Tx Fault (Symmetrical Operation)

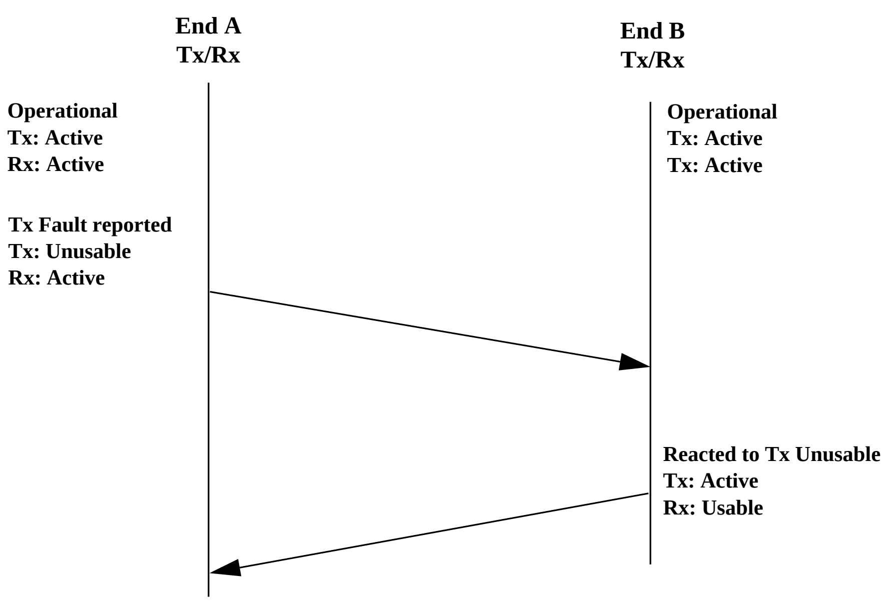

Note: Only the NE states newly entered and transmitted over the ICP cells are shown.

Figure 41 Link De-activation due to Tx Fault (Symmetrical Operation)

Figure 42 on page 123 and Figure 43 on page 124 shows a link de-activation when the Rx Failed condition is entered and the later link re-activation after Rx Failed condition is exited, under symmetrical and asymmetrical operation respectively.

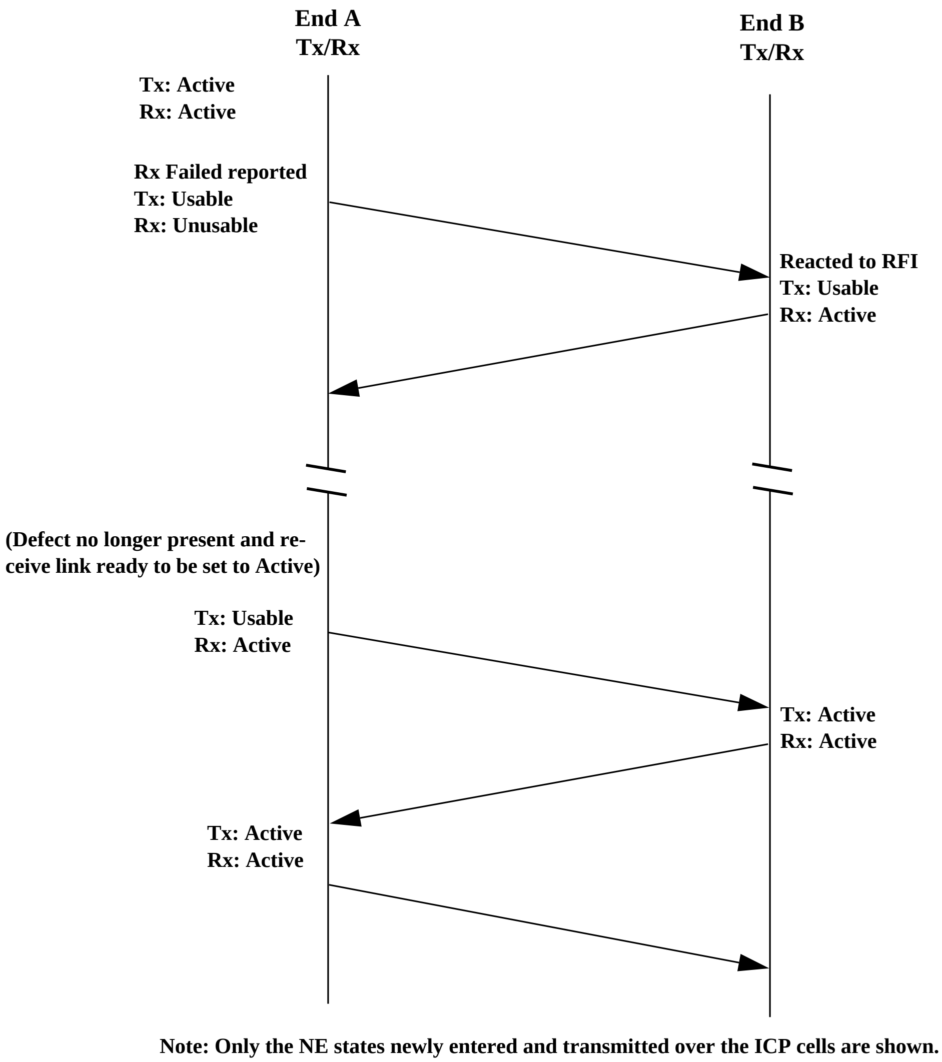

Figure 42 Link De-activation/Recovery under Link Failed Condition (Symmetrical Operation)

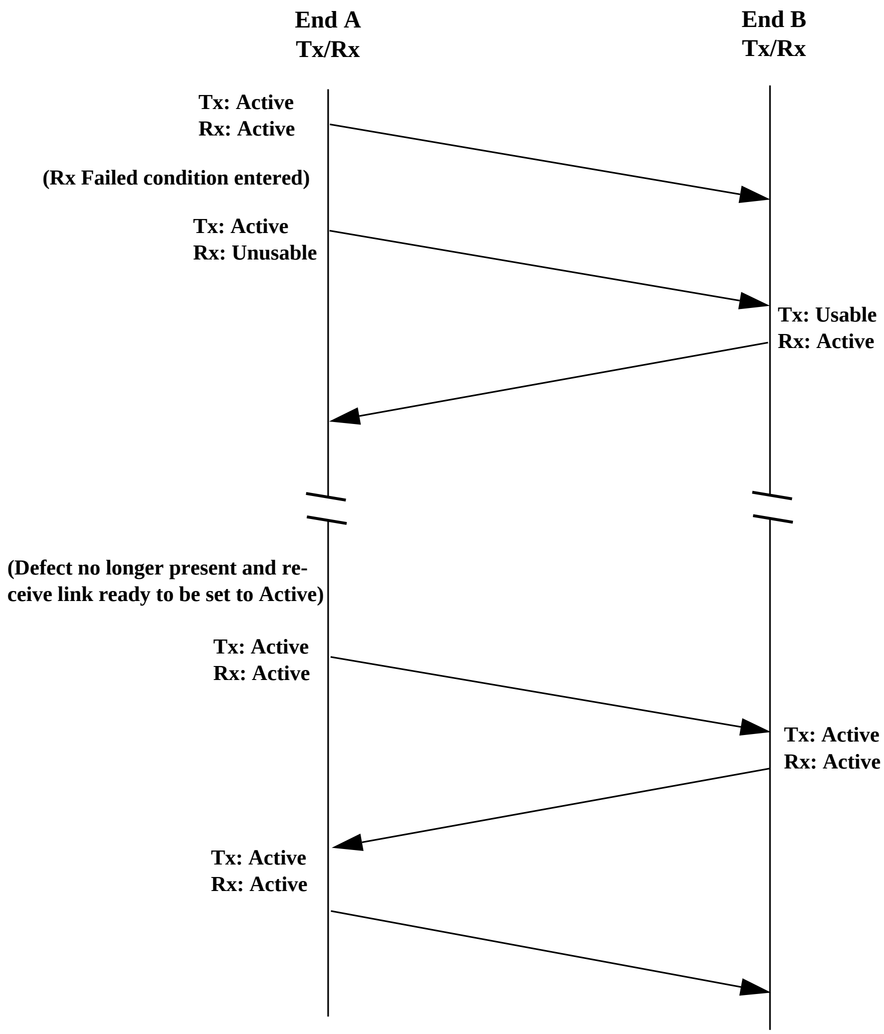

Note: Only the NE states newly entered and transmitted over the ICP cells are shown.

Figure 43 Link De-activation/Recovery under Link Failed Condition (Asymmetrical Operation)

Appendix E IMA OAM Behaviors

This appendix provides information about the OAM behaviors of the IMA interface. Figure 44 on page 125 shows the OAM behaviors of a single link and an IMA interface. The following can be noted:

- A defect is reported at a layer (F1, F2, F3 F4 or F5) once detecting a defect at the same layer or when AIS is reported by the lower layer
- Once a defect is detected at a layer, an RDI is reported in the opposite direction
- Once a defect is detected at a layer, AIS is reported to the next upper layer

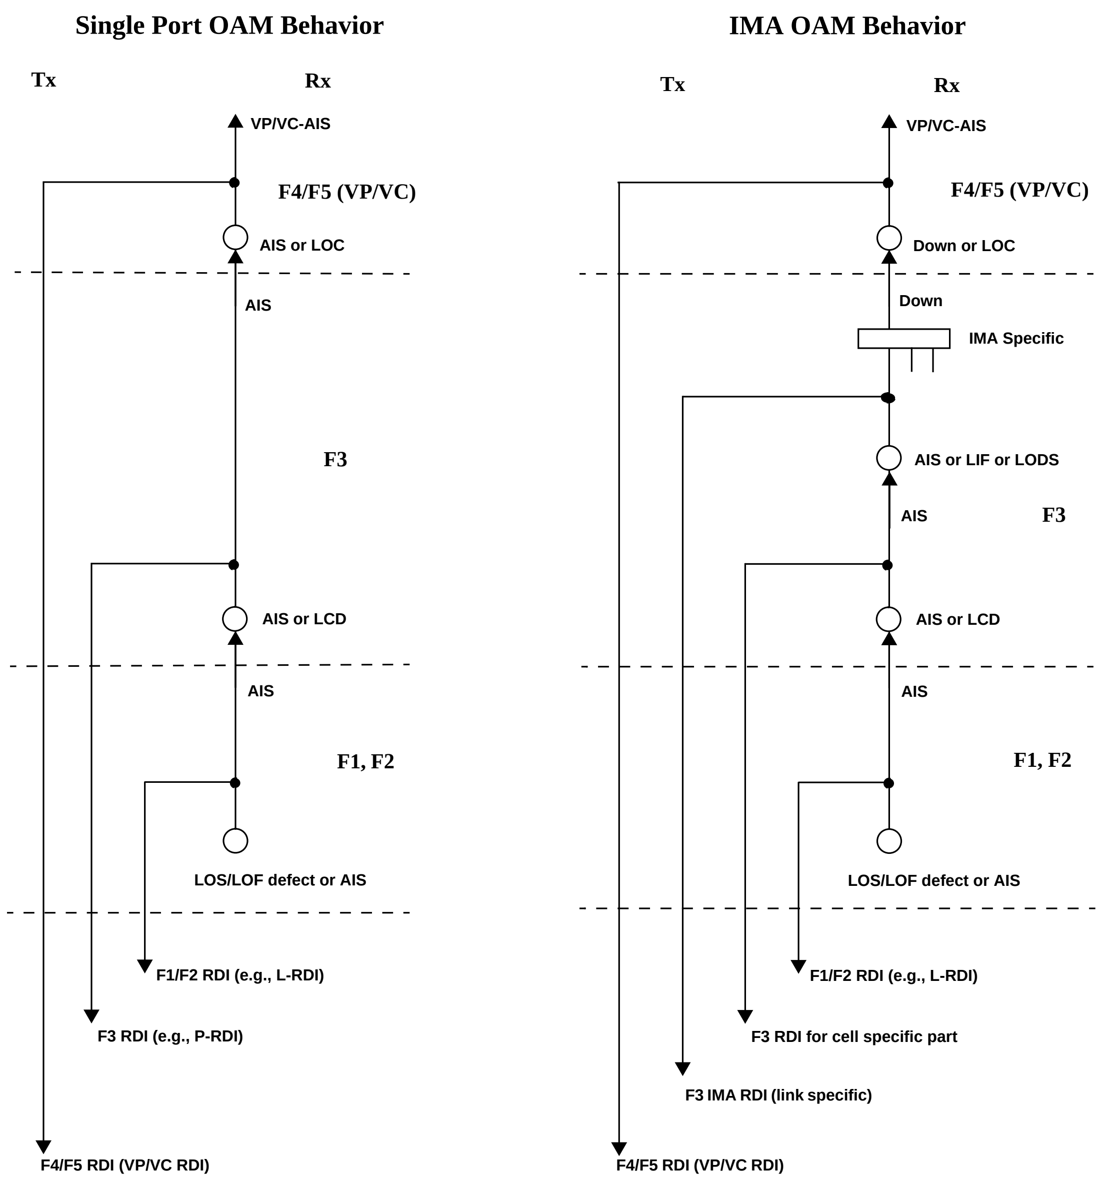

Figure 44 Single Link and IMA OAM Behaviors

Appendix F Numbers of IV-IMAs to Trigger SES-IMA

ITU-T Recommendation G.826 [37] specifies that SES shall be reported for second periods which contain ≥ 30% errored blocks. It also specifies that a lesser number may be used in the case the system has not the ability to detect and that value different from 30% may be used and these values may vary with transmission rate.

Table 25 on page 126 shows the number of errored ICP cells that could be used to derive SES if we require ≥ 30% errored blocks to report SES. This is excluding the presence of Filler and ATM layer cells. The table shows that the values of the SES threshold depends of M and the available LCR. The advantage is to have a threshold adapted to the available LCR.

Table 25 Proposed Numbers of IV-IMAs to Trigger SES

	DS1 (LCR = 3622.6 cells/second)		E1 (LCR = 4528.3 cells/second)	
M	**ICP cells/second**	**≥ 30%**	**ICP cells/second**	**≥ 30%**
32	113	34	142	43
64	56	17	71	22
128	28	9	36	11
256	14	5	18	6

Appendix G Detailed IMA Clocking Configurations

This appendix contains detailed information with respect to clocking options that are specific to various IMA link configurations, where the links are DS1 or E1 links. The information contained within this appendix builds upon the details as described within Section 4.3 on page 23 and addresses the scenarios where the transmit clocks of the links within the IMA group may or may not be derived from the same clock source (i.e., CTC or ITC mode), and where the IMA interface is used in public and private network configurations. The respective clocking configuration selected for an IMA application will be determined by the availability and provisioning of link services. It is suggested that, if possible, all the links within the IMA link group use a single source for transmit clock and, if possible, that source be traceable to a Primary Reference Source (PRS).

Throughout this appendix the terms Primary Reference Source (PRS)[5], Primary Reference Clock (PRC)[34], and Stratum 1 are synonymous.

G.1 IMA Group in CTC Mode

This section describes clocking options applicable to the configuration where the IMA interface is running in the Common Transmit Clock (CTC) mode. These clocking options can be categorized into two main sections, namely non-transparent and transparent, as shown in Figure 45 on page 127. The link reference model illustrates the different timing domains [31].

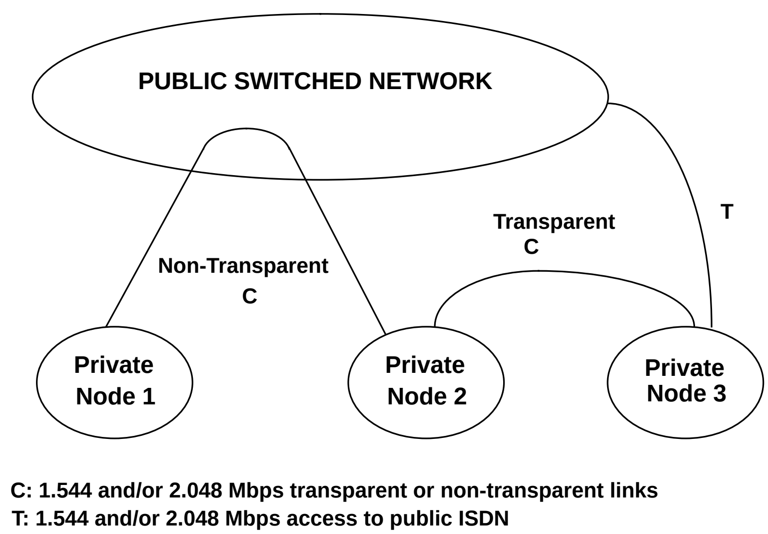

Figure 45 Link Timing Reference Model

Transparent is defined within reference [31] as: a link or group of links is transparent if the signal carried is not re-timed from a clock associated with the link(s). The timing of a signal passing across a transparent link may however be altered due to jitter, wander, filtering or error conditions.

G.1.1 Non-transparent Links

The following section will detail link configurations applicable to a non-transparent relationship. All these configurations are associated with the support of access to a synchronous public network.

G.1.1.1 Single Time Domain

The single-time domain clocking configuration is applicable to an application where all transmitters are within a single timing domain (e.g., all links from the same local area carrier). An example is given in Figure 46 on page 128.

The IMA units will derive their respective transmit clocks from one of their receive clock signals, which will normally be traceable to a Stratum 1.

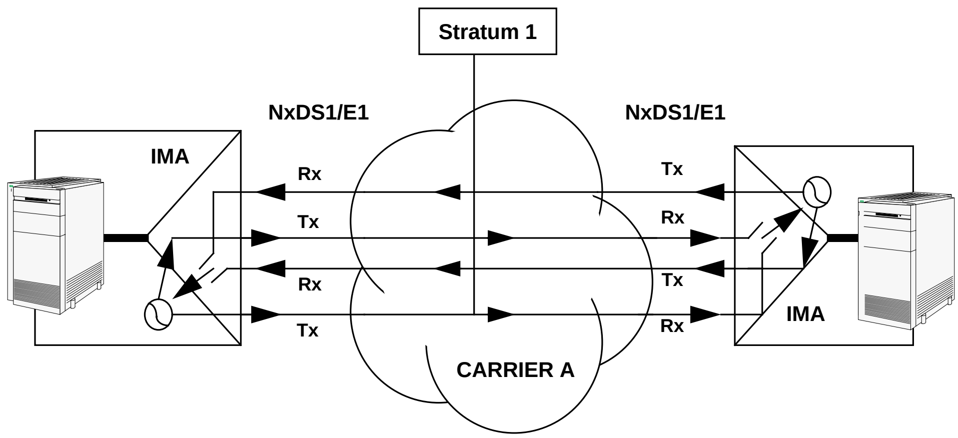

Figure 46 Single Time Domain

The DS1 interface will be expected to conform to the ANSI T1.408 [5] specification which imposes the following requirements:

- The IMA receiver and transmitted bit streams must be synchronized to the carrier's clock.
- In normal operation the bit stream is traceable to a Stratum 1 primary reference, which has a long term bit rate accuracy of $1x10^{-11}$. When synchronization by a Stratum 1 clock has been interrupted, the signal delivered by the network to the interface shall have a minimum accuracy of 4.6×10^{-6}, Stratum 3 accuracy, as specified in ANSI T1.408[5], Section 5.3.1.1.
- For multiple network interfaces (as with an IMA with multiple links), the transmission rate is determined by the signal received across only one link interface, or is derived externally from the carrier's Stratum 1, so that all interfaces are synchronized to the same master clock source.

An E1 interface has similar requirements that are specified in ITU-T Recommendations G.703 [32] and I.431 [40].

G.1.1.2 Two or More Timing Domains

The two or more timing domains clocking configuration addresses the scenario where the multiple IMA low speed links are obtained from a number of distinct and separate carriers, which have their own synchronization hierarchies (timing domains). This scenario may be quite common as it adds to the diversity and survivability of an IMA application. An example is shown in Figure 47 on page 129.

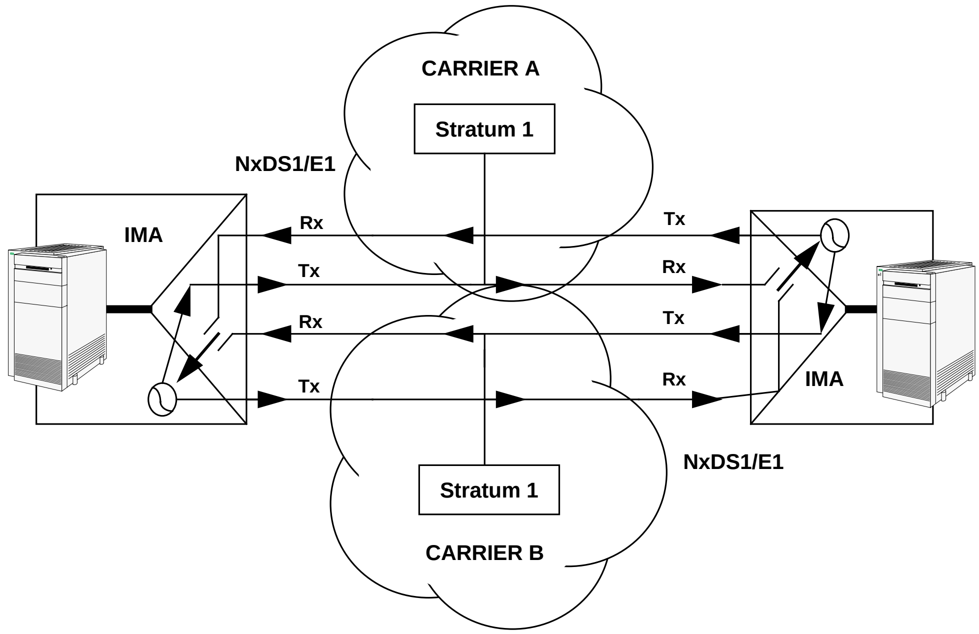

Figure 47 Two or More Timing Domains

The configuration will have each link or group of links synchronized to a separate master clock source, which again will normally be traceable to Stratum 1. In practice only one of the carrier master clocks will synchronize the IMA equipment for synchronization, e.g. Carrier A's Stratum 1 reference clock. This may result in virtually error-free traffic performance as the carriers are operating plesiochronously with each other, with average clock rate different by no more than 1 part per 10^{-11}. In this configuration, it is recommended to ensure all carriers use Stratum 1 traceable clocks when subscribing.

G.1.1.3 International or Inter Exchange Carrier (IEC) Timing Domains

The following configuration, shown in Figure 48 on page 130, can be referred to as the "Classic Plesiochronous" configuration with respect to a synchronous network. The scenario addressed is where the links traverse more than one timing domain. This will occur within international and inter-IEC topologies as well as within many larger networks. Each IMA's internal clock will synchronize to a DS1/E1 received from the host carrier and provide timing to the IMA transmitted signals, so that all the signals in the group of links are normally traceable to one of the carrier's Stratum 1. The inter-working and re-timing of the information stream across the different timing domains will be managed transparently to the application. If timing differences do accumulate then a controlled slip may be introduced by individual DS1/E1 receivers to compensate for the differences. Under fault free conditions the occurrence of a controlled slip will be rare, e.g. in a worst case scenario where the number of carriers (N) is 2, 1 controlled frame slip will occur per 72 days (per DS1/E1).

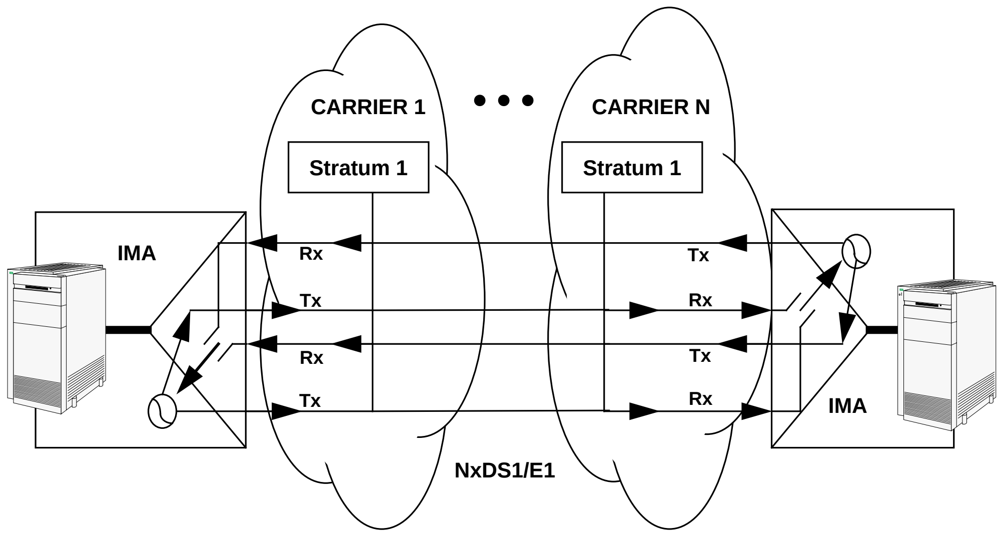

Figure 48 International or Inter IEC Timing Domains

G.1.2 Transparent Links

The following section details link configurations applicable to a transparent relationship, e.g. private leased lines. These configurations will require the customer premise to provide or act as the master clock source for the transmitted signals. All these configurations should be termed as being asynchronous, which indicates that they are not synchronized to a network wide clock source. Note that the respective local clock source may still be traceable to a network wide PRS but this is not a requirement.

G.1.2.1 Master/Master Configuration

The master/master configuration is applicable when addressing services which require the customer premise equipment to act as the clock source. In this configuration the transmit clock is not derived from the receive bit stream (see Figure 49 on page 130). The configuration may also be known as a “Transparent” or “Split timing” configuration [31]. The clocks are required to support an accuracy of 1.544Mbps ± 32 ppm for DS1 configuration [2] and 2.048 Mbps ± 50 ppm for E1 configuration [32].

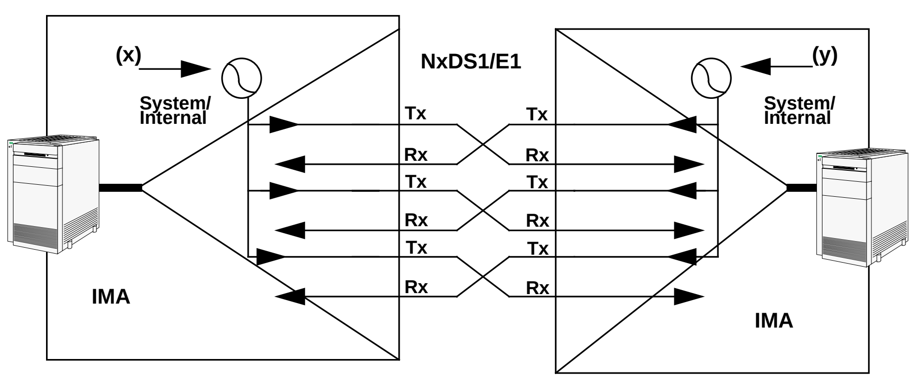

Figure 49 Master/Master Configuration

Practical configuration options with respect to the clock sources are detailed as follows, refer [31]. (x) and (y) are reference clocks. When (x) and (y) are providing a clock signal, two cases can occur, in addition to the free-running case:

1. (x) and (y) come from the same clock and are therefore synchronous and

2. (x) and (y) are from two different clock sources.

G.1.2.2 Master/Slave Configuration

The Master/Slave configuration, also referred as looped timing configuration, is the most common configuration with respect to the support of private networking scenarios. An example is represented in Figure 50 on page 131.

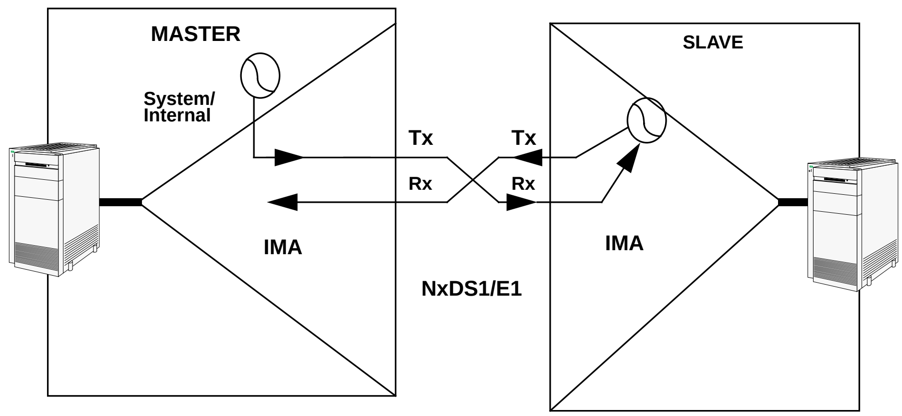

Figure 50 Master/Slave Configuration

As it can be determined from Figure 50, there is a single master clock source. The remote equipment (Slave) will synchronize to the master-provided transmitted signal. The accuracy of the master clock source is expected to be in the range of 1.544Mbps ± 32 ppm for DS1 configurations [2], and 2.048Mbps ± 50 ppm for E1 configurations [32].

G.2 Connections to Public ISDN

The access DS1/E1s may be either transparent or non-transparent. The ATM switch in the public network will provide timing. The overall effect is as if the links were non-transparent. Figure 46 on page 128 illustrates this case.

G.3 IMA Group in ITC Mode

Detailed within this section are clocking options applicable to the configuration where the IMA group is running in the Independent Transmit Clock (ITC) mode. This typically results in the scenario where some of the source clock signals are not traceable to a PRS.

G.3.1 Two or More Timing Domains - Un-synchronized

The clocking configuration, as detailed within Figure 51 on page 132, is similar to the configuration contained within Section G.1.1.2, where the multiple IMA low speed links are supported over a number of distinct and separate carriers in a non-transparent fashion. The main difference in this configuration is that each link within the IMA link group is synchronized via “looped” timing to each of the respective carriers’ reference clocks, whereas in the previous configuration (Section G.1.1.2) one of the carriers was selected as the reference. This configuration will result in virtually error-free performance as described in Section G.1.1.2.

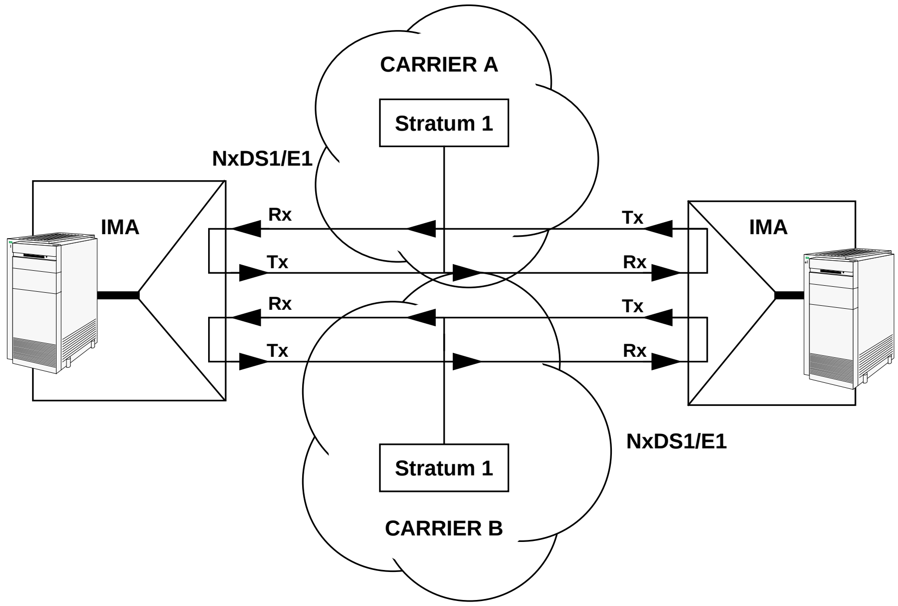

Figure 51 Two or More Timing Domains - Un-synchronized

G.3.2 Non-Transparent/Transparent Mixed Configuration

The non-transparent/transparent mixed configuration addresses the scenario where the links within the IMA link group are supported over a combination of transparent and non-transparent facilities. This may be the result of the following conditions where the majority of the links are supported as private leased line connections and additional links are managed/deployed via an on-demand basis by utilizing dialup and/or ISDN facilities. The additional links may be established either as a result of a demand for additional bandwidth or alternatively to address a restoration/survivability scenario.

Figure 52 on page 133 shows an example of mixed of transparent/non-transparent configurations. It displays the following configuration mix:

- the private leased line links will be supported as transparent links, in this case a master/slave clocking configuration has been selected.
- the dialup and/or ISDN lines will be supported as non-transparent links, in this case a master/slave.

The non-transparent/transparent mixed configuration, therefore, creates the situation where the links within the IMA link group are supported in a non-synchronous fashion. The private leased line links will be synchronized to a locally derived timing source, which maybe asynchronous to the carrier clock source; and the links supported over the dialup and/or ISDN facilities will be synchronized to the carrier’s master clock source.

If possible, the timing configuration represented in Figure 52 on page 133 should be avoided. In case this configuration is used, it is recommended to drive timing on the transparent links from one of the PRS traceable sources from a carrier(s) rather than from local sources. This changes the configuration to that covered in section Section G.1.1.1 on page 128 or Section G.1.1.2 on page 128 (depending on the number of independent Primary Reference Sources in the configuration).

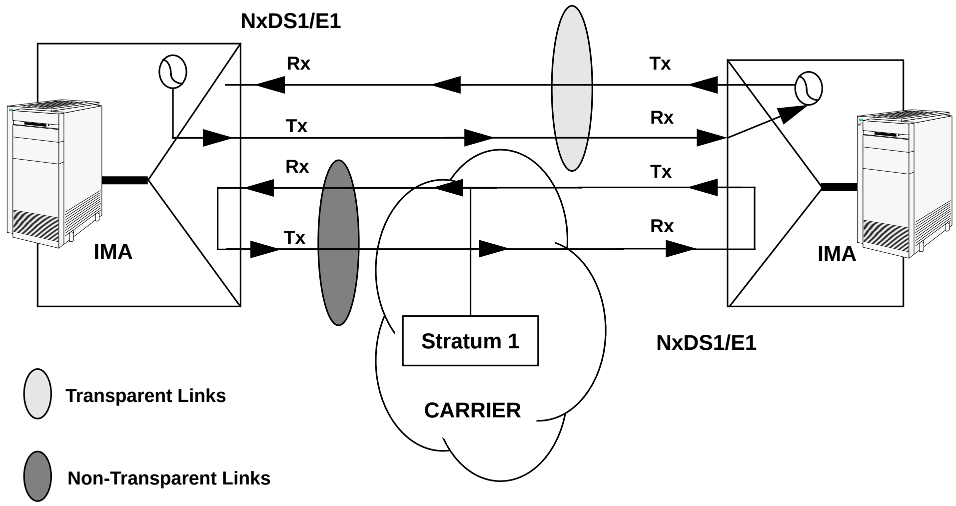

Figure 52 Non-Transparent/Transparent Mixed Configuration

End of the document.