
**Information technology —
Telecommunications and information
exchange between systems — Intermediate
system to Intermediate system intra-domain
routing information exchange protocol for
use in conjunction with the protocol for
providing the connectionless-mode
Network Service (ISO 8473)**

**AMENDMENT 2: Extensions for group
composition and related MST multicast
routing**

*Technologies de l'information — Communication de données et échange
d'informations entre systèmes — Protocole intra-domaine de routage d'un
système intermédiaire à un système intermédiaire à utiliser conjointement
avec le protocole fournissant le service de réseau en mode sans connexion
(ISO 8473)*

*AMENDEMENT 2: Extensions pour la composition de groupe et routage de
la diffusion sélective de MST liée*

IECNORM.COM : Click to view the full PDF of ISO/IEC 10589:1992/Amd 2:1999

© ISO/IEC 1999

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the publisher.

ISO/IEC Copyright Office • Case postale 56 • CH-1211 Genève 20 • Switzerland
Printed in Switzerland

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 3.

In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.

Amendment 2 to ISO/IEC 10589:1992 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 6, *Telecommunications and information exchange between systems*.

IECNORM.COM : Click to view the full PDF of ISO/IEC 10589:1992/Amd.2:1999

IECNORM.COM : Click to view the full PDF of ISO/IEC 10589:1992/Amd 2:1999

Information technology — Telecommunications and information exchange between systems — Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)

AMENDMENT 2: Extensions for group composition and related MST multicast routing

Add the following new annex G.

Annex G
(normative)

Extensions for group composition and related MST multicast routing

Introduction

The main intention of this annex is the specification of mechanisms for the distribution and handling of Minimum Spanning Tree (MST) multicast address information within the level 1 routing area as well as the specification of the computation of the MST used for routing of multicast NPDUs. The mechanisms specified herein enforce that each Intermediate system has the complete information about all existing address groups and their composition, that is the NSAP addresses of all members belonging to a particular address group. Since the distribution of this group information on larger networks would use a not negligible network bandwidth, the restriction on the level 1 area as mentioned above has been done. In a future step the MST multicasting mechanisms specified herein shall be expanded for the level 2 routing area. This kind of multicasting is referred to as MST multicasting within the document.

In ISO/IEC 10589:1992 there is no mean for using multicast addresses. This annex proposes the usage of multicast addresses, as defined in ISO 8348:1987 and the replication of NPDUs making use of a computed MST in those Intermediate systems, where the destination addresses divert to multiple links. This new definition can be used on all types of subnetworks and is downward compatible with the specification in the main body of this International Standard.

This Annex also refers to Annex D of ISO 9542, Addition of group composition information, which specifies a mechanism for the distribution of group composition information between Intermediate systems and End systems.

G.1 Scope

This annex describes procedures for the distribution and interpretation of MST multicast information between network entities residing in Intermediate Systems within a single routing domain.

G.2 Normative references

ISO 9542:1988/Annex D, Information technology - Telecommunications and information exchange between systems - End system to Intermediate system routing exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473) - Annex D: Addition of group composition information

C.-H. Chow; On Multicast Path Finding Algorithms; Proc. of Infocom. '91; pp 1274-1283

G.3 Definitions

For the purposes of this annex the following definition applies.

group: a set of systems, which can be identified by their network addresses.

G.4 Abbreviations

For the purposes of this annex, the following abbreviations apply in addition to those defined in the main body of the specification

ESGC PDU	End System Group Composition Protocol Data Unit (ISO 9542 Annex D)
GSP	Group State Protocol Data Unit
ISGC PDU	Intermediate System Group Composition Protocol Data Unit (ISO 9542 Annex D)
ID	Identification
L	Level
MO	Managed Object
MST	Minimum Spanning Tree
SPF	Shortest Path First

G.5 Overview of the protocol

G.5.1 Subnetwork independent functions

As a new subnetwork independent function the multicast routeing is introduced as follows:

- Group addresses are handled between Intermediate systems as a set of unicast addresses.
- The mechanism of handling NPDU's with a group address as the target address of this NPDU is available within all Intermediate systems of this routeing domain.
- For routeing of multicast NPDU's from the entry Intermediate system to the target Intermediate systems a minimum spanning tree (see annex C) is constructed. This routeing algorithm guarantees that in the case where non-broadcast subnetworks are used for connecting two adjacent Intermediate systems, NPDU's are only replicated at those Intermediate systems, where more than one subtree is used towards the target unicast NSAP addresses.
- MST multicast information is flooded within the intradomain routeing area using a new type of PDU's, called "Group State PDU's (GSPs)". These GSPs are handled like LSPs and are exchanged on all types of subnetwork connections, as proposed in Amendment 2 for LSPs (especially for exchanging LSPs and GSPs on DA circuits).
- Group addresses are either created, modified or deleted by systems operations at every Intermediate system or at an End system, exchanging group address information between End systems and their designated Intermediate system using additional PDU types "End System Group Composition PDU (ESGC)" and "Intermediate System Group Composition PDU (ISGC)" (extended operation). In the case the option is not used, End systems are not informed about the members of a group, only about the existence of a group. In this case End systems are only able to use group addresses (e.g. as target addresses for multicast NPDU's), to subscribe to an existing group address or to unsubscribe itself from a group address.

G.5.2 Design goals

This annex supports the following design goals:

- *Replication*: The replication function determines the number of different outgoing paths depending on the decomposition of the multicast addresses within a NPDU. Based on this number of different paths the NPDU is then replicated n times within an Intermediate System and these replicated NPDU's are sent once per determined outgoing path.
- *MST multicastness*: Exchange of group composition information in the optional or mandatory mode between End Systems and Intermediate Systems using ESGC PDU's and ISGC PDU's (see ISO 9542 Annex D), exchange of group composition information in the basic or extended mode among Intermediate Systems using GSPs (see mechanisms described in this amendment) and replication of NPDU's with multicast addresses in the basic or extended mode (see mechanisms described in this amendment).

G.5.3 Design non-goals

It is not a design goal of the procedures defined in this annex to guarantee delivery of all offered NPDU's to all defined destination addresses.

G.5.4 Enhancements to the Decision Process

If the decision process has recognised that a NPDU is to be routed to a group address, the following actions are needed:

- expand the group to all addresses contained in the group - (15) in figure G-1.
- create the routing decision according to the link state database for each of the group members and store it until replicated NPDU(s) are removed in the next step - (16) in figure G-1.

G.5.5 Enhancements to the Update Process

This process constructs, receives and propagates Link State PDUs and Group State PDUs. Each Link State PDU contains information about the identity and routing metric values of the adjacencies of the IS that originated the Link State PDU. Each Group State PDU contains a (new) group definition (or a part of it), that is just known at the adjacent IS.

The Update Process receives Link State, Group State and Sequence Numbers PDUs from the Receive Process - (4) in figure G-1. It places new routing information in the routing information base - (6) and propagates routing information to other Intermediate systems - (7) and (8).

General characteristics of the Update Process are:

- Link State PDUs are generated as a result of topological changes, and also periodically. They may also be generated indirectly as a result of System Management actions (such as changing one of the routing metrics for a circuit).
- Group State PDUs are generated as a result of changes in the Group State or based on Systems Management actions.
- Level 1 Link State PDUs are propagated to all Intermediate systems within an area, but not propagated out of an area.
- Level 2 Link State PDUs are propagated to all Level 2 Intermediate systems in the domain.
- Link State PDUs are not propagated outside of a domain.
- The update process, through a set of System Management parameters, enforces an upper bound on the amount of routing traffic overhead it generates

G.5.6 Enhancements to the Forwarding Process

After the decision process has resolved to which adjacencies the multicast NPDU is to be sent, the NPDU has to be replicated and sent to these adjacencies with the group address as the destination address - (18) in figure G-1.

G.5.7 Enhancements to the Receive Process

Additionally the Receive Process obtains its inputs from the group information handed to the group state database (both from ESs and ISs).

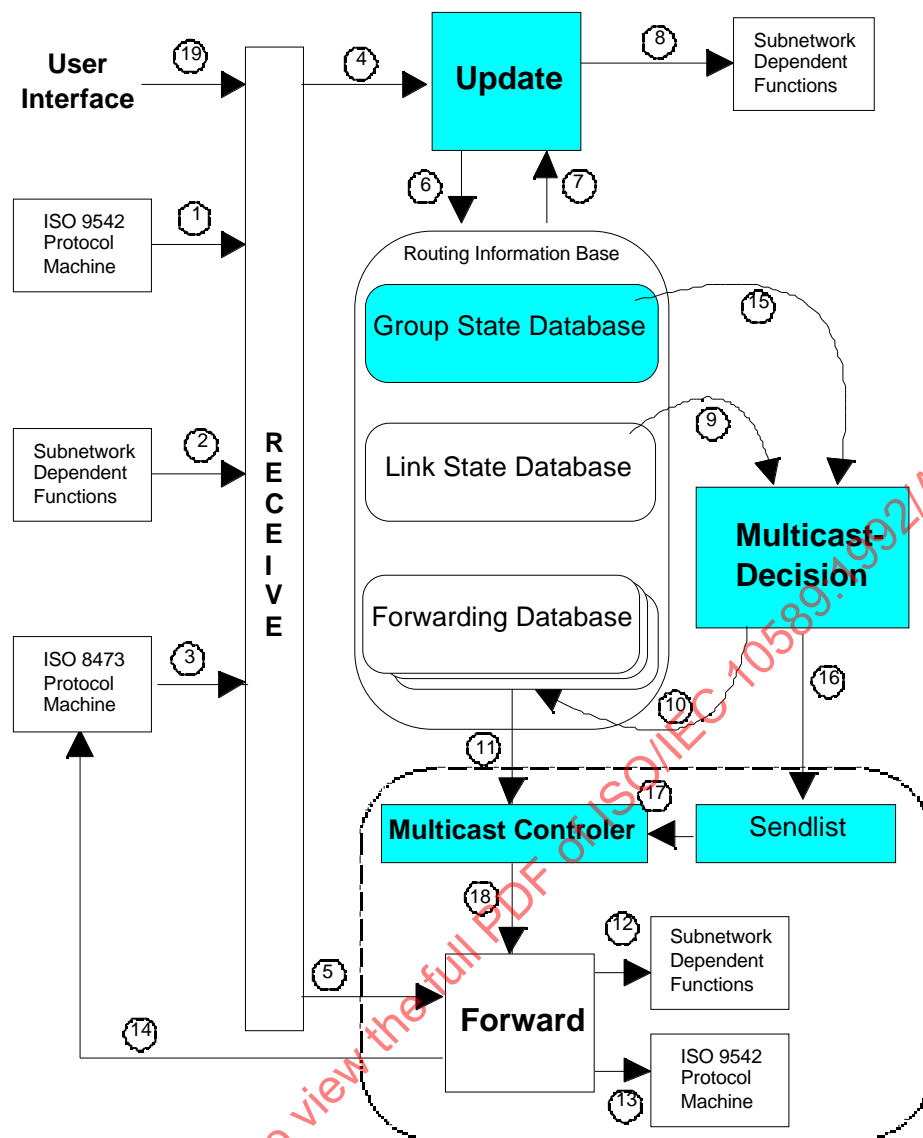


Figure G-1 — Decomposition of Subnetwork Independent Functions

G.6 Subnetwork independent functions

This Annex introduces a new process, the Sendlist Creation Process (L2, L1), and a new database, the Group State data base (L2, L1).

G.6.1 Decision process

This process uses the databases of Link State information and Group State information to calculate the forwarding database(s), from which the forwarding process can know the proper next hop(s) for each (replicated) NPDU. The Level 1 Link State Database is used for calculating the Level 1 Forwarding Database(s), and the Level 2 Link State Database is used for calculating the Level 2 Forwarding Database(s). The Group State Database is used to calculate in addition both the Level 1 and Level 2 Forwarding Database(s).

G.6.1.1 Input for the Decision Process

As a new input for the Decision Process this annex makes use of the Group State Database, which is a set of MST multicast information from the latest Group State PDUs from all known Intermediate Systems (within this

area, for Level 1, or within the level 2 subdomain, for Level 2), and from the latest Group Hello PDUs from all adjacent End Systems. This database is received from the Update Process.

G.6.1.2 Output for the Decision Process

This Annex generates the following new outputs for the Decision Process:

- Level 1 MST Multicast Forwarding Databases - one per routeing metric
- (Level 2 Intermediate systems only) Level 2 MST Multicast Forwarding Databases - one per routeing metric

G.6.1.3 MST multicast enhancements to the decision process

This clause contains the enhancements to the decision process necessary to support MST multicasting.

G.6.1.3.1 Exchange of GSPs

In addition to the exchange of LSPs, GSPs are sent to all Intermediate systems neighbours (also ISO 9542 ISGC PDUs are sent to all End Systems, if the extended mode of ISO/IEC 9542/Annex D is used) under two circumstances:

- a) The Intermediate System receives a GSP with a new multicast address or with a known multicast address but with a higher sequence number
- b) (optional): The MST multicast announcement timer triggers a periodical resent of the local group addresses.

The Update process is also capable of dividing a single logical GSP into a number of separate PDUs for the purpose of transmitting large group definitions.

G.6.1.3.2 Multicast routeing algorithm

The routeing algorithm used by the Decision Process is a Minimum Spanning Tree algorithm. Each Intermediate system executes the Minimum Spanning Tree algorithm autonomously to define a loopless tree of legal paths to all destination systems in a routeing domain. This routeing algorithm is specified in more detail in G.10.

G.6.1.3.3 Processing of multicast NPDUs

If the Decision Process recognises that a destination address is not a group address ((1) within figure G-2), no expanding is carried out. Otherwise (2) the first activity must be marking the source link, from which the NPDU is received. After the group list expansion a Sendlist is created. This list will contain all adjacencies to which the NPDU has to be forwarded. The following actions are carried out in a loop (3) for all members of a group:

- calculate the path to the members of the group list
- check whether a link exists in the Link State Database
- verify that the adjacency on this link is not the source link from which the NPDU is received (4)
- add the adjacency to the Sendlist, if it is not already in the list (5). This is to prevent sending duplicated PDUs to the same adjacency (6).
- create an Error Event to the System Management if the only available link is the source link (7). This event indicates a possible loop condition caused by backward routeing.

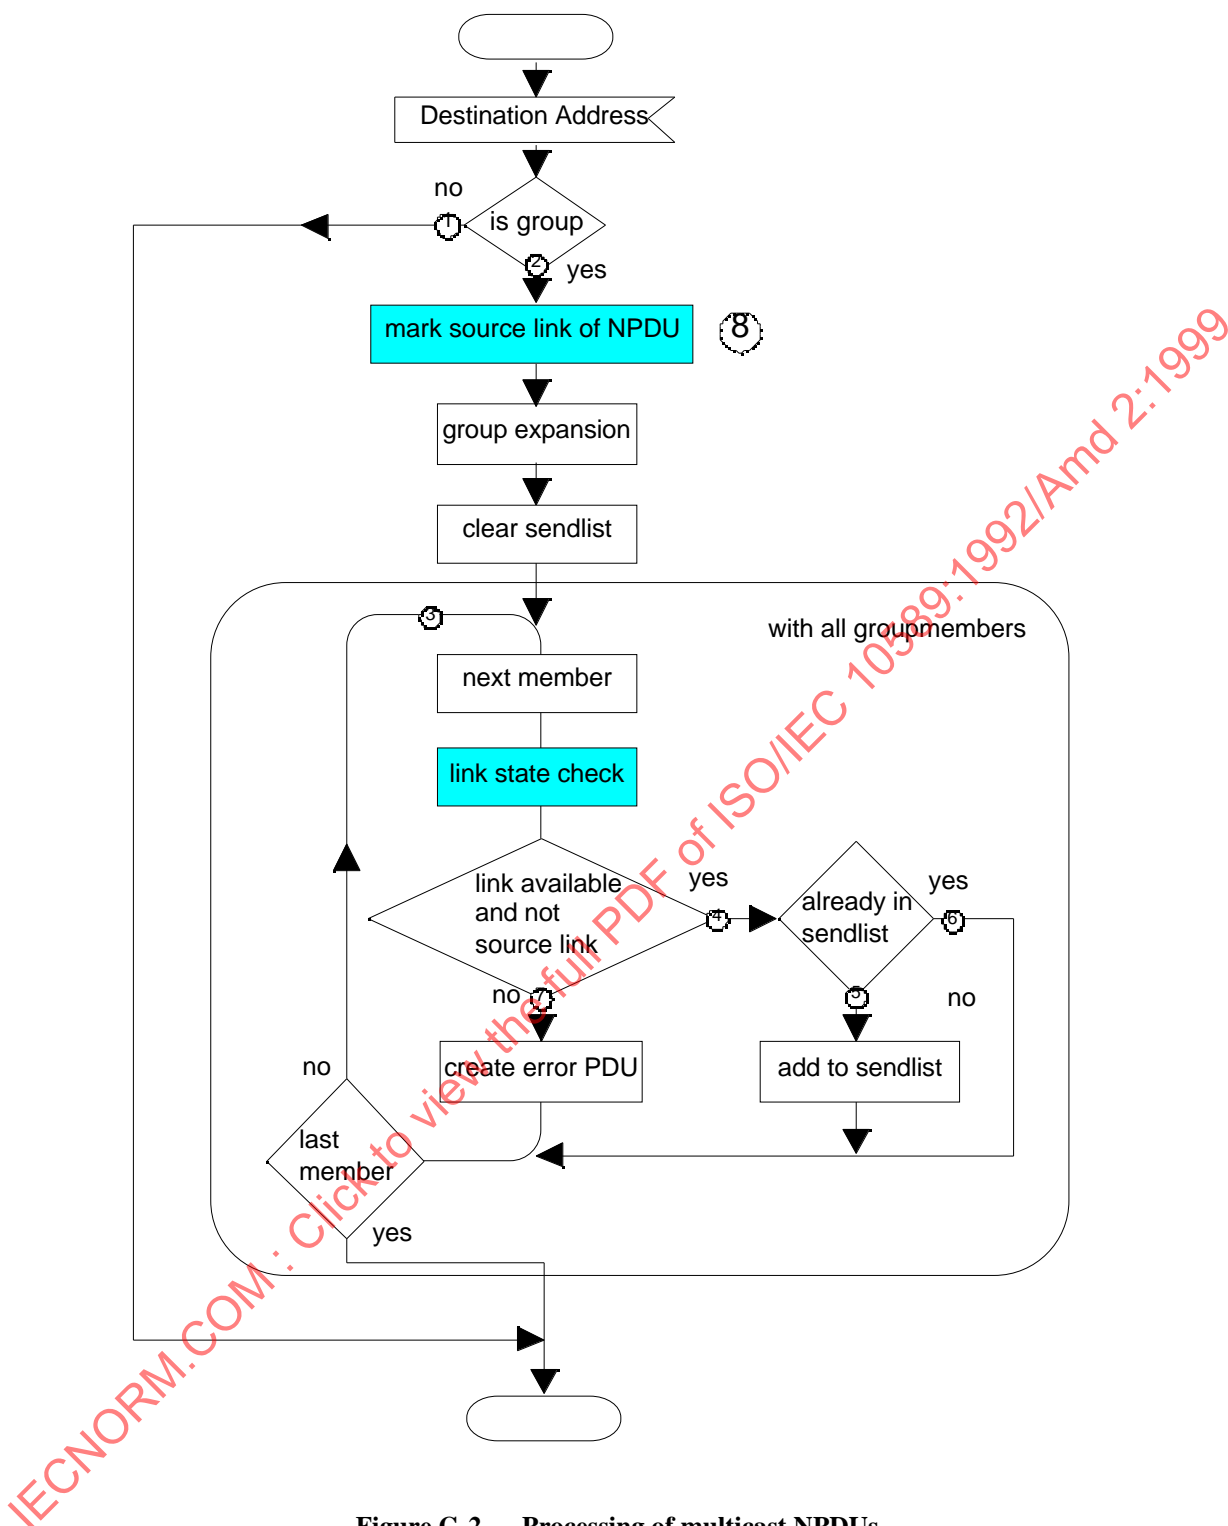


Figure G-2 — Processing of multicast NPDUs

G.6.2 Update process

G.6.2.1 Input for the Update Process

- As a new input for the Update Process this annex makes use of the Group State PDUs, which are passed by the Receive Process to the Update process along with an indication of which adjacency it was received from.

G.6.2.2 Output for the Decision Process

This annex generates the following new outputs for the Decision Process:

- Group State Database

Note: As a result of the Receive Process the Group State Database is maintained. A signal to the Decision Process is created as an event which is either the receipt of a Group State PDU with information different from the stored one, or the purging of groups from the database.

G.6.2.3 Parameters

This annex adds the following parameter to those defined in the main body of this International Standard.

maximumGSPGenerationInterval - This is the maximum amount of time allowed to elapse between generation of Group State PDUs by a source and will be used only for the optional periodic GSP generation. It shall be less than MaxAge.

A reasonable setting is 15 min.

G.6.2.4 Multiple GSPs

Because the GSP is limited in its size by ReceiveGSPBufferSize, it may not be possible to include all members of a big group in a single GSP. In such a case, a system may use multiple GSPs to convey this information. The recipient system recognises that they all pertain to a common originating system because they all use the same Sequence Number.

Because of interpreting the first area address in the variable length field as the group address, this entry must be the same in every following GSP.

G.6.2.5 Periodic GSP generation (optional)

The Update Process may optionally periodically re-generate and propagate on every circuit with an IS or ES adjacency the locally known group definitions. The Intermediate System may re-generate each GSP at intervals of at most maximumGSPGenerationLevel, with jitter applied as described in 10.1 of this specification.

G.6.2.6 Event driven GSP generation

An Intermediate System shall as normal operation generate a GSP when an event occurs which would cause the information content to change. The following events may cause such a change:

- receipt of a new group composition
- receipt of a modified group composition
- receipt of a deletion of an existing group
- modification of the Group State database by systems management actions

G.6.2.7 Propagation of GSPs

The Update Process is responsible for propagating Group State PDUs throughout the domain.

The basic mechanism is flooding, in which each Intermediate system propagates GSPs to all its neighbour Intermediate systems except the neighbour from which it received the PDU. Duplicates are detected and dropped.

Group State PDUs are received from the Receive Process.

Upon receipt of a Group State PDU the Update Process shall perform the following functions:

- a) An Intermediate System receiving a Group State PDU which is new shall
 - 1) store the Group State PDU into Group State database, and
 - 2) mark it as needing to be propagated upon all circuits except that upon which it was received.
- b) An Intermediate System receiving a Group State PDU with an multicast address still stored shall
 - 1) check whether the sequence number is higher
 - 2) overwrite the existing content if the sequence number is higher and flood the GSP on all circuits except the one it has been received on
 - 3) delete the new GSP if the sequence number is lower or equal the existing one

Note: There is a negligible small probability, that two different Intermediate systems will modify simultaneously a group composition. In this case each of them would use the same Sequence Number, therefore some Intermediate systems will get a GSP with this Sequence Number twice and delete it. This effect can be avoided by an Intermediate system, if it compares before the deletion of such GSPs the actual composition of an address group received within the GSP with the version stored in the local Group State database with the same Sequence Number. If those contents don't match, the Intermediate system can distribute the new group composition composed of both received modifications with an increased Sequence Number.

- 4) delete the existing content, if the sequence number is higher and the group members are zero
- c) An Intermediate System receiving a Group State PDU with an incorrect GSP Checksum or with an invalid PDU syntax shall
 - 1) generate a corrupted GSPReceived circuit event,
 - 2) overwrite the Checksum with zero.

G.6.2.8 Remaining lifetime field within GSPs:

If an Intermediate system generates a GSP, it shall specify in the Remaining Lifetime field the maximal number of Intermediate system hops a GSP can be forwarded until it is considered expired. Each Intermediate system, which receives a GSP, has to decrement this octet by 1 before it further forwards the GSP. If an Intermediate system receives a GSP with this field set to zero, it shall not further forward the GSP. This mechanism successfully prevents GSPs being endlessly forwarded.

G.6.3 Forwarding Process

G.6.3.1 Input for the Forwarding Process

This Annex adds the following new input for the Forwarding Process:

- Sendlist

G.6.3.2 Processing of multicast NPDUs

Based on the following figure G-3, the forwarding process is modified in the following way:

If the destination address is not a group address (1) according to the Forwarding Database, the above standard forwarding activities (3)-(4) are carried out unmodified. Otherwise (2) these actions are repeated for each entry in the Sendlist (5).

G.6.3.3 Basic operation

If the received PDU is a Group State PDU, the Receive Process shall pass it to the Update Process.

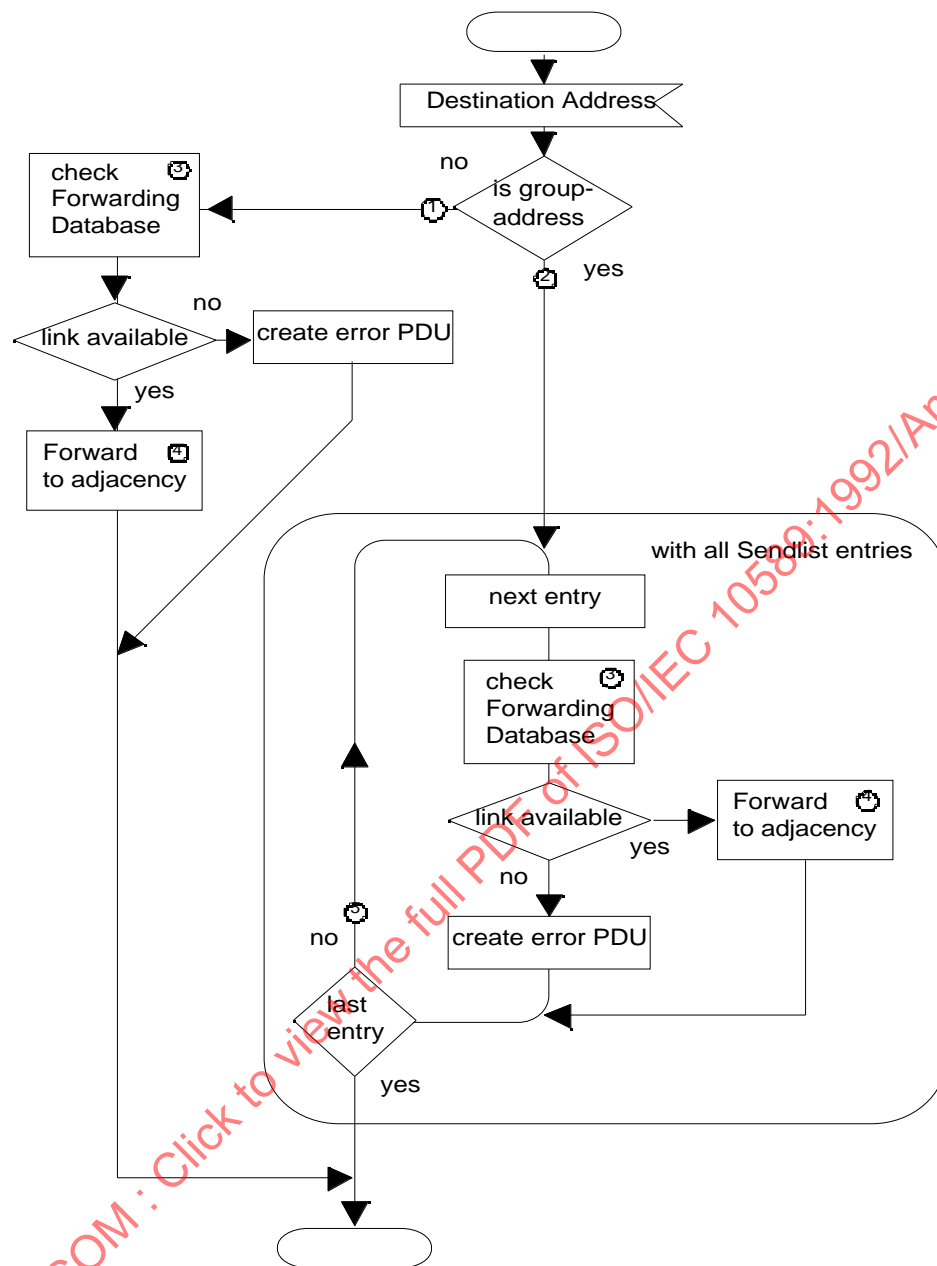


Figure G-3— MST multicast enhancements to the Forwarding Process

G.6.4 MST multicast functionality

G.6.4.1 Overview

The main reason of using a multicast address instead of single unicast addresses when transmitting a NPDU to a group of receivers is to save bandwidth on the network. Therefore it is important, not to use too much from this saved bandwidth for the management information containing the multicast address definitions. This paragraph specifies protocol modifications of ISO 10589:1992, which will achieve a reduction of the management information. While the creation of a group address requires the definition of all unicast addresses, which shall be members of a group during creation time, the modification of an existing group address can be done in two ways:

- modification of a group by replacing the existing definition totally by a new definition (extended definition)

- modification of a group by replacement of defined members (addition and/or deletion of specified members) (basic definition)

For this a new procedure of modification of an existing group is introduced, which only transmit the changes and not the complete new composition of a modified group.

To implement the possibility of transmitting only the changes of a modified group, an octet called Modification Flag is introduced in the VARIABLE LENGTH FIELD of a GSP-structure (see chap. 9). This octet determines the kind of modification, that is if the VARIABLE LENGTH FIELD contains a complete composition of a modified group, the systems which shall be added to an existing group or the systems which shall be deleted from an existing group.

Now there can be distinguished between two modes of protocol implementation. In one of them the basic functionality, in the other one the extended functionality of group management is available. It has to be guaranteed, that all Intermediate systems in a routing domain run their protocols in the same mode.

The following paragraphs describe the two functionalities of ISO 10589 concerning group management.

G.6.4.2 Basic Functionality

G.6.4.2.1 Functions for the composition and the administration of groups

- *create_group*: If an Intermediate system receives the information about the generation of a new group, it transmits the group composition using a GSP to the other Intermediate systems. In this GSP the Sequence Number is set to 1 and the VARIABLE LENGTH FIELD contains the complete composition of the group (if the group doesn't fit in a single GSP, it is transmitted in more GSPs). Group Length is set to the number of octets necessary for storing the Group Address, all member addresses and all address length indicators (each of them needs 1 octet) for the Group Address and the member addresses. The Modification Flag is set to 0. This symbolises the fact, that the VARIABLE LENGTH FIELD contains a complete group composition and not only the changes of a modified group.
- *delete_group*: If an Intermediate system receives the information, that an existing group is deleted, it transmits this information using a GSP to the other Intermediate Systems. In this GSP the Sequence Number Field contains the Sequence Number of the group stored in the local Group State Data Base increased by 1. The VARIABLE LENGTH FIELD contains only the Group Address to identify the deleted group, but no member of this group. Group Length is set to 0, since this group has been deleted. The Modification Flag is also set to 0.
- *add_member*: If an Intermediate system receives the information, that one or more members have been added to an existing group, it transmits this information using a GSP to the other Intermediate Systems. In this GSP the Sequence Number Field contains the Sequence Number of the group stored in the local Group State Data Base increased by 1. If not all added members fit in a single GSP, the Intermediate System transmits these changes in more GSPs with all of them using the same Sequence Number. The VARIABLE LENGTH FIELD consists of the Group Address to identify the appropriate group and only those members, which have been newly added to the existing group. Group Length is set to the length of the existing group increased by the number of octets necessary for storing the new members. The Modification Flag is set to 1, what symbolises, that the VARIABLE LENGTH FIELD doesn't contain a complete group composition, but the Group Address and a set of newly added members.
- *delete_member*: If an Intermediate system receives the information, that one or more members have been deleted from an existing group, it transmits this information using a GSP to the other Intermediate systems. In this GSP the Sequence Number Field contains the Sequence Number of the group stored in the local Group State Data Base increased by 1. If not all deleted members fit in a single GSP, the Intermediate System transmits these changes in more GSPs with all of them using the same Sequence Number. The VARIABLE LENGTH FIELD consists of the Group Address to identify the appropriate group and only those members, which have been deleted from the existing group. Group Length is set to the length of the existing group decreased by the number of octets necessary for storing the deleted members. The Modification Flag is set to 2, what symbolises, that the VARIABLE LENGTH FIELD doesn't contain a complete group composition, but the Group Address and a set of deleted members.

G.6.4.2.2 Information of End Systems

Executing the basic functionality an Intermediate system informs an End system neither about existing Group Addresses nor about the composition of a certain group, that is no group information received by the ISO 10589:1992 protocol will cause a protocol action in ISO 9542. ESGCs are only used for enabling End systems to join or leave an existing group.

G.6.4.3 Extended Functionality

G.6.4.3.1 Functions for the composition and the administration of groups

- *create_group*: This function remains unchanged. For a closer description see 7.6.2.1.
- *delete_group*: This function remains unchanged. For a closer description see 7.6.2.1.
- *modify_group*: If an Intermediate system receives the information, that the composition of an existing group has been modified, it transmits this information using a GSP to the other Intermediate systems. In this GSP the Sequence Number Field contains the Sequence Number of the group stored in the local Group State Data Base increased by 1. If not all members fit in a single GSP, the Intermediate System transmits these changes in more GSPs with all of them using the same Sequence Number as described above. The VARIABLE LENGTH FIELD consists of the complete composition of the modified group. Group Length is set to the number of octets necessary for storing this group. The Modification Flag is set to 0, what symbolises, that the VARIABLE LENGTH FIELD contains a complete group composition.

G.6.4.3.2 Information of End Systems

Executing the extended functionality an Intermediate system informs an End system about all existing Group Addresses and the composition of the respective groups. For this it uses the extended functionality of the ISO 9542 protocol. In this case group management is not restricted to Intermediate systems, but distributed on all systems in a routing domain. For this purpose ESGCs and ISGCs are used. With ESGCs an End system reports its locally known (e.g. changed group definitions) to its adjacent Intermediate system. With ISGCs an Intermediate system reports to a new operational End system the actual available group definitions (group name and members) and/or changed/deleted group definitions.

G.7 Structure and encoding of Group State PDUs

This section contains the additional PDU structures which is necessary to introduce MST multicast handling within ISO 10589:1992.

Group State PDUs are generated by Intermediate systems, and propagated throughout an area. The contents of the Group State PDU indicates the actual knowledge of an Intermediate system about group definitions.

				No. of Octets
Intradomain Routing Protocol Discriminator				1
Length Indicator				1
Version/Protocol ID Extension				1
ID Length				1
R	R	T	PDU Type	1
Version				1
Reserved				1
PDU Length				2
Remaining Lifetime				2
GSP ID				ID Length + 1
Sequence Number				2
Checksum				2
VARIABLE LENGTH FIELDS				VARIABLE

- Intradomain Routing Protocol Discriminator - architectural constant (see ISO 10589:1992: table 2)
- Length Indicator - Length of fixed header in octets
- Version/Protocol ID Extension - 1
- ID Length - Length of the ID field of group addresses used in this routing domain. This field shall take one of the following values:
 - An integer between 1 and 8, inclusive, indicating an ID field of the corresponding length
 - The value zero, which indicates a 6 octet ID field length
 - The value 255, which means a null ID field (i.e. zero length)
 All other values are illegal and shall not be used.
- PDU Type (bits 1 through 5) - 19
 - NOTE 64 Bits 7 and 8 are Reserved, which means they are transmitted as 0 and ignored on receipt, Bit 6 is used to differentiate between periodically exchanged GSPs (T=0) and event driven exchanged GSPs (T=1); this is drawn from ISO 10589:1992/Amd.2
- Version - 1
- Reserved - transmitted as zero, ignored on receipt
- PDU Length - Entire Length of this PDU, in octets, including header
- Remaining Lifetime - Number of hops before GSP considered expired
- GSP ID - the system ID of the source of the Group State PDU. It is constructed as follows:

		No. of Octets
Source ID		ID Length
GSP Number		1

NOTE The usage of a GSP ID is optional and may be used in future for acknowledgement mechanisms. For normal usage, the Source ID Length field is zero and the GSP Number field is ignored.

- Sequence Number - sequence number of GSP
- Checksum - Checksum of contents of GSP from Source ID to end. Checksum is computed as described in ISO 10589:1992: chap. 7.3.11.
- VARIABLE LENGTH FIELDS - fields of the form:

	No. of Octets
CODE	1
LENGTH	2
VALUE	LENGTH

Any codes in a received GSP that are not recognised are ignored and passed through unchanged.

Currently defined codes are:

Group Addresses - the name of a group and the set of NSAP addresses associated to this group of this Intermediate system.

- x CODE - 1
- x LENGTH - total length of the value field of this GSP
- x VALUE -

	No. of Octets
Group Length	2
Modification Flag	1
Length of Group Address	1
Group Address	Address Length
Address Length (1. member)	1
Address (1. member)	Address Length
Address Length (n. member)	1
Address (n. member)	Address Length

Group Length - Length of whole address field (may be distributed over more than 1 GSP) in octets (exclusive group length field).

Note If a group consists of $p > n$ members, where n is the last member in 1 GSP, then one or more GSPs are used, containing the remaining $(p-n)$ members; all these GSPs then contain the same elements at the beginning of the VARIABLE LENGTH FIELDS (Group Length, Length of Group Address, Group Address). Correlated GSPs are identified by the same Sequence Number. Modified or deleted GSPs must have a higher Sequence Number.

Address Length - Length of addresses in octets.

Address - Group or single NSAP Address.

Modification Flag: It indicates whether the subscriber handling for groups the full information of End systems about the construction of a group (extended) or not (basic). Mode and flag value can be drawn from the following table:

	Functionality	Modification Flag	VARIABLE LENGTH FIELD
create_group	Basic, Extended	0	Group Address and NSAP addresses of all members
delete_group	Basic, Extended	0	Group Address
add_member	Basic	1	Group Address and NSAP addresses of added members
delete_member	Basic	2	Group Address and NSAP addresses of deleted members
Modify_group	Extended	0	Group Address and NSAP addresses of all members

Authentication Information - information for performing authentication information of the originator of the PDU.

- x CODE - 10
- x LENGTH - variable from 1-254 octets
- x VALUE -

Authentication Type	No. of Octets 1
Authentication Value	VARIABLE

Authentication Type - a one octet identifier for the type of authentication to be carried out. The following values are defined:

0 - RESERVED

1 - Cleartext Password

2-254 - Routing Domain private authentication method

Authentication Value - determined by the value of the authentication type. If Cleartext Password as defined in this International Standard is used, then the authentication value is an octet string.

G.8 Requirements on the operation of ISO 9542

When the MST multicast and group composition functions are supported, the IS shall support

- the reception of ESGC PDUs on all types of circuits;
- the transmission of ISGC PDUs on all types of circuits.

G.9 System management enhancements

The following definitions for group managed objects are included:

Group address managed object

groupAddresses MANAGED OBJECT CLASS

DERIVED FROM "REC X.721 | ISO/IEC 10165-2: 1992": top;

CHARACTERISED BY groupAddressesPackage PACKAGE

BEHAVIOUR "ISO/IEC 10733":commonCreationDeletion-B, "ISO/IEC 10733":commonStateChange-B;

ATTRIBUTES

groupName GET-REPLACE,

operationalState GET,

groupStatus GET-REPLACE,

groupMembers GET-REPLACE;

NOTIFICATIONS

"Rec X.721 | ISO/IEC 10165-2 : 1992": objectCreation,

"Rec X.721 | ISO/IEC 10165-2 : 1992": objectDeletion;

REGISTERED AS {ISIS.moi groupAddress (9)};

Name binding:

groupAddresses -networkSubsystem-importedNAME BINDING

SUBORDINATE OBJECT CLASS groupAddresses AND SUBCLASSES;

NAMED BY SUPERIOR OBJECT CLASS "ISO/IEC 10733":networkSubsystem AND SUBCLASSES

WITH ATTRIBUTE groupName;

REGISTERED AS {ISIS.nboi groupAddresses-networkSubsystem-imported(9)};

groupName ::= OCTETSTRING

operationalState ::= ENUMERATED {disabled(0), enabled(1)}

groupStatus ::= ENUMERATED {create(0), complete(1),

```

                                uncomplete(2),
                                modify(3),
                                delete(4)}
groupMembers ::= SET OF OCTETSTRING

```

G.10 Minimum spanning tree algorithm

For MST multicast routing the "Minimum Spanning Tree (MST)" algorithm is used. This algorithm computes a minimum spanning tree for each group from the sink to all sources, based on the available LSP information.

The original algorithm [3] is applicable only for small networks (very small number of nodes), as the needed computational time for calculating the MST grows exponentially with the number of nodes. As explained in [3], a network consisting of 20 nodes needs 34 second for calculating the MST (for a given computing system)

Since using the unicast routing algorithms like the shortest path first SPF for multicast PDUs can result in infinite loops, it is necessary to develop a separate algorithm for this case. It is sure, that no loop will occur if all ISs in a routing area construct an identical tree containing all the vertices of the area and use this tree for creating their multicast forwarding data base. The tree used in this algorithm is the minimum spanning tree MST of the area. It would be optimal in saving bandwidth if we consider to send a PDU to all vertices in the network. If the PDU is only addressed to a group of all vertices there exists a better tree with regard to saving bandwidth during the transmission, but this tree, the MST which only contains the transmitter and all receivers of the PDU, has a too high computing complexity for bigger networks [3]. To ensure that all ISs in the area construct an identical MST, the distances of all edges are needed. If there are edges with equal distances a tiebreaker must be used to get an unequivocal order of the edges weights. In this algorithm we use the following tiebreaker:

- a) Intermediate systems are included into the Minimum Spanning Tree before End systems.
- b) The edge with the lower metric assigned has the shorter distance.
- c) If there are still two edges with equal distances compute the sum of the system IDs of the two vertices linked by the respective edge. The edge causing the lower sum has the shorter distance.
- d) If there are still two edges with equal distances examine the system IDs of all vertices attached to the edges separately. The edge linking the vertex with the lowest system ID has the shorter distance.
- e) If there are still two edges with equal distances we have the situation that two vertices are linked by more than one edge with each of them having the same metric. In this case we can select any edge for the one with the shortest distance, because this cannot result in the construction of different MSTs by the ISs of the area. This situation will probably never appear in real networks.

G.10.1 Overview of the modified MST algorithm

After making sure that all edges in the area having different distances it is possible for each IS to construct a single MST which is the same for all ISs and contains all vertices of the area. For this every IS needs two databases. One of them, called PATHS, is used after finishing the computation of the tree as MST multicast forwarding database. The entries of this database are couples of the $\langle N, \text{Adj}(N) \rangle$, with:

N: system ID of the vertex N

Adj(N): adjacency that the computing system X should use for forwarding to N.

The second database, called TENT, is used only during the computation of the tree. The entries of this database are triples of the form $\langle N, d(N), \text{Adj}(N) \rangle$ with:

N: system ID of the vertex N

$d(N)$: shortest possible distance of a single link, which would connect the vertex N to the tree which already exists at this time of the computation. If the connection of N to the tree over a single link is impossible the triple with the system ID of N is not present in TENT.

$Adj(N)$: adjacency that the computing system X should use for forwarding to N .

TENT can intuitively be thought of as a tentative placement of a vertex (system) N in PATHS, but the explicit placement in PATHS can first be done if it is sure, that no vertex has a shorter distance to the tree than N .

Additionally each vertex entry has to be marked as IS or ES and has to keep information about the previous vertex over which it was reached. Later in the algorithm this information will be used to make a correct decision.

Each computing IS in the area starts the algorithm by putting itself as first vertex in PATHS. TENT is then pre-loaded from the local adjacency database of this system. Afterwards it has to be calculated, which vertex has the shortest distance to the existing tree (at this point of time the tree only comprises the computing system itself). This vertex is added to PATHS. If it was an IS its neighbour systems are examined by reading all its LSPs.

This is done until all systems are placed in PATHS and no one remains in TENT.

G.10.2 Algorithm

There are three main steps of the algorithm separable:

- ✳ Initialisation (Step 0)
- ✳ Evaluation of the LSPs (Step 1)
- ✳ Selection of the shortest distance (Step 2).

Now the single steps are described in more detail.

Step 0:

- a) Initialise TENT and PATHS as empty.
- b) Add $\langle \text{SELF}, 0, W \rangle$ to PATHS, where SELF is the system ID of the computing IS and W is a special value indicating traffic to SELF is passed up to the Transport Layer.
- c) Now pre-load TENT with all neighbour systems N of SELF. This could be done by reading the local adjacency database. $d(N)$ is the metric stored with the respective adjacency and $Adj(N)$ is the adjacency number of the adjacency to N .
- d) If a neighbour system is already in TENT, compare the distance of the old and the new entry and keep only the entry with the shorter distance.
- e) If a neighbour system is not in TENT, then place it now.
- f) If all local adjacencies are examined, go to Step 2.

Step 1:

- a) Now examine all neighbour systems N listed in all LSPs of P of the system just placed on PATHS (P has been placed in PATHS during the last execution of Step 2). $d(N)$ is the metric of the link from P to N and the number of the adjacency to N , $Adj(N)$, is the same as the number of the adjacency to P , $Adj(P)$, because N could be reached from the computing system over P .
- b) If a neighbour system is already in PATHS, then do nothing.
- c) If a neighbour system is already in TENT, compare the distances of the old and the new entry and keep only the entry with the shorter distance.