Guidance Material for Sub-Volume 5 - ATN Internet SARPs

Issue 1.4

23rd September 1996

Foreword

This document provides Guidance Material for , "Sub-Volume 5" of the ATN SARPs, which itself contains the SARPs for the Internet (i.e. Network & Transport Layers) component of the ATN, as derived from the material in the ATN Manual (second edition) by the ATN Panel Working Group 2 (WG2).

Sub-Volume 1 contains introductory material to the ATN SARPs, and additionally "system level" provisions applicable to the Package as a whole. Sub-Volume 2 contains provisions for the initial set of Air/Ground applications i.e. Automatic Dependent Surveillance (ADS), Controller Pilot Data Link Communications (CPDLC), Flight Information Services (FIS) & Context Management (CM). Sub-Volume 3 contains provisions for the initial set of Ground/Ground applications i.e. Inter-Centre Communications (ICC) and Aeronautical Message Handling Service (AHMS). Sub-Volume 4 contains provisions for the Upper Layer Architecture (ULA) to be supported by ATN SARPs compliant End Systems.

Please note that the material in this document contains references to the documents of the International Organization for Standardization (ISO) and the International Telecommunication Union (ITU). In using these documents, due attention should be given to their publication dates as shown on the list of references.

List Of Contents

1. Introduction

1.1 Background

In January 1989, the Air Navigation Commission (ANC) expanded the terms of reference of the Secondary Surveillance Radar Improvements and Collision Avoidance Systems Panel (SICASP) to include the development of ICAO material as necessary to permit, to the maximum extent practicable, systems commonality and interoperability between ATS data links, including satellite data links.

The task emerged from the work of the Special Committee on Future Air Navigation Systems (FANS) which emphasised the need for the interchange of digital data over dissimilar aeronautical data links. The committee also recommended that the principles of the International Organisation for Standardisation (ISO) open systems interconnection (OSI) architecture be applied in developing aeronautical data links in order to provide for their interoperability.

Subsequent studies undertaken by the SICAS Panel resulted in the concept of the aeronautical telecommunication network (ATN) which is intended to support computer-to-computer communications operated by civil aviation authorities and aeronautical operating agencies. At its fourth meeting, the SICAS Panel developed a description of the ATN and recommended it be published as an ICAO manual. The first edition of the manual was published in 1991, and the second edition was subsequently developed by the SICAS Panel and recommended for publication at the fifth meeting of the panel, and is expected to be published by ICAO during 1995. The development of the ATN continues with the objective of recommending Standards and Recommended Practices (SARPs) and Guidance Material for the ATN during 1996, for inclusion in Annex 10 at that time.

Following the completion of the work on the ATN Manual (Second Edition) by the SICAS Panel, the Air Navigation Commission transferred the work of developing SARPs and Guidance Material to the ATN Panel (ATNP).

The concept developed by SICASP was of a multi-user multi-vendor internetwork, that would integrate the many different air-ground and ground-ground networking technologies that are currently in place and being planned. In addition, this internetwork would be designed to meet the exacting Quality of Service (QoS) requirements of aeronautical applications, and to respect the ITU and national regulations that applied to each air-ground data link.

1.2 Scope

This document provides guidance material for ATN Implementors, Service Providers and Users.

1.3 Purpose of Document

In line with normal ICAO practice, this document was developed as a companion document to Sub-volume V of the ATN Internet SARPs. It may be read alongside the SARPs, in order to provide a

greater understanding of the specification itself, or it may be read instead of the SARPs by readers that simply want to understand the ATN Concept rather than the detail of the specification.

1.4 Document Overview

This document is divided in seven chapters as follows:

Chapter 1 **Introduction**

This chapter

Chapter 2 **The ATN Concept**

An overview of the components and architecture of the ATN internet.

Chapter 3 **ATN Addressing**

This chapter discusses the ATN Internet Address and provides guidance for administrators on allocating ATN Addresses.

Chapter 4 **ATN Protocols and Functions**

This chapter discusses the protocols and procedures that are required to be implemented by ATN components, such as End Systems and Routers.

Chapter 5 **ATN Routing**

A major feature of the ATN is its ability to provide transparent support of mobile systems. This chapter presents the ATN Mobile Routing Strategy and discusses the ISO Inter-Domain Routing Protocol (IDRP) and how it is used to support this strategy.

Chapter 6 **Congestion Avoidance in the ATN Internetwork**

This chapter presents the congestion avoidance procedures implemented in the ATN. This is an important feature of the ATN essential to the support of safety related applications.

Chapter 7 **ATN Subnetworks**

The ATN is an internetwork that can, in principle, be used with many and indeed any networking technology. This chapters discusses the ICAO Mobile subnetworks that used specified for use with the ATN and how different types of ground subnetwork (e.g. X.25, Frame Relay, etc.) may by used by the ATN.

2.2 Technical Benefits

The ATN has been specified to meet the requirements of the Civil Aviation Community and gives the following technical benefits to its users:

- **Use of Existing Infrastructure** — The ATN is an internetwork built on top of existing networks through the use of routers as gateways between those networks. Investment in existing LANs, leased lines, CIDIN and X.25 networks is preserved. Furthermore, the ATN can make full use of emerging network technologies such as Frame Relay and Asynchronous Transfer Mode (ATM).

- **High Availability** — The ATN has been designed to provide a high availability network by ensuring that there is no single point of failure, and by permiting the availability of multiple alternative routes to the same destination with dynamic switching between alternatives. The same techniques apply to both fixed and mobile communications giving mobile communications an availability level that would have been unrealistic for older technologies based on directory lookups (e.g. ACARS).

- **Mobile Communications** — The ATN fully supports mobile communications over a wide variety of mobile communications networks including AMSS, VDL and Mode S. With the ATN, it is possible for a ground system to communicate with airborne avionics in any part of the world.

- **Prioritised end-to-end resource management** — All ATN user data is given a relative priority on the network in order to ensure that low priority data does not impede the flow of high priority data. Advanced Congestion Management techniques that “throttle back” low priority data when the network becomes near to saturation, ensure that high priority data always gets a low transit delay. The Congestion Management techniques used in the ATN are much superior to those used in the TCP/IP Internet. In the ATN, traffic load is balanced to the availability of communications resources, while in the TCP/IP Internet, a saturated network is regularly pushed into an overloaded state with consequential data loss.

- **Scaleability** — The ATN provides both a large Address Space and an approach to routing that ensures the scaleability of the network well beyond currently foreseen requirements.

- **Policy based Routing** — The ATN’s routing procedures support a wide range of Organisational and National policies, including the enforcing of restrictions on what types of traffic can pass over both ground and air/ground data links, and control over which air/ground data link types are used by which applications. Administrations and Organisations that interconnect the networks are free to enforce routing policies that control which types of data are exchanged and whose data is routed through their networks, and whose data is not.

- **Future Proofing** — The ATN is a way of using networking technologies can be readily extended to include new ground and air/ground data links technologies, with local rather than global impact of the

use of new networking technologies.

3. ATN Addressing

3.1 Introduction

The ATN naming and addressing scheme is based on the OSI Reference Model (ISO 7498-3) which supports the principles of unique and unambiguous identification of information objects and global address standardisation which are essential features for an international, mixed-user communications system as the ATN.

3.2 ATN Network Addressing Plan

A well defined network addressing plan has been established for the ATN which meets the needs of a variety of aeronautical data communication user groups, including ATS providers, airlines and international aeronautical communication service providers (IACSPs). Furthermore, it supports essential goals of ATN internal operation, such as efficient information reduction when exchanging address information (as part of the routing information) and unambiguous and complete address reconstruction from received address fragments (in the context management application).

In order to facilitate the address assignment and registration in the ATN, which is expected to comprise several thousands of objects, the ATN network addressing plan is hierarchically structured. This means that it is composed of a set of hierarchical address domains. Each address domain is a set of address formats and values which are administered by a single addressing authority. The ATN SARPs partition the overall ATN NSAP addressing domain (which is itself a sub-domain of the OSI addressing domain) into a number of addressing sub-domains, each with an identified addressing authority, in a recursive fashion. Each addressing authority is responsible for its own address sub-domain, and may further partition it into several subordinate sub-domains, and delegate authority for these sub-domains. This principle allows the establishment of sub-address spaces (i.e. the set of values within an addressing sub-domain) in a hierarchical fashion without the need to co-ordinate between sub-address spaces. This principle of hierarchically structured ATN sub-domains within the global OSI addressing domain is illustrated in Figure 1 for the example of the ATN NSAP addressing domain.

Figure 3-1 shows how the global OSI network addressing domain (which is itself a sub-domain of the global OSI addressing domain) is partitioned into several sub-domains, one of which is the ATN NSAP address sub-domain. This sub-domain is itself decomposed into a number of subordinate addressing sub-domains in a recursive fashion. Each such sub-domain is associated with an NSAP addressing authority which is responsible for this sub-domain, and may further delegate authority for those sub-domains into which it has partitioned its own addressing subdomain. This principle allows to construct ATN addresses as a sequence of individual address fields (see Figure 1), with each field corresponding to an addressing sub-domain. As these sub-domains are individually administered, the address field formats and values can be assigned without the need to co-ordinate between addressing authorities.

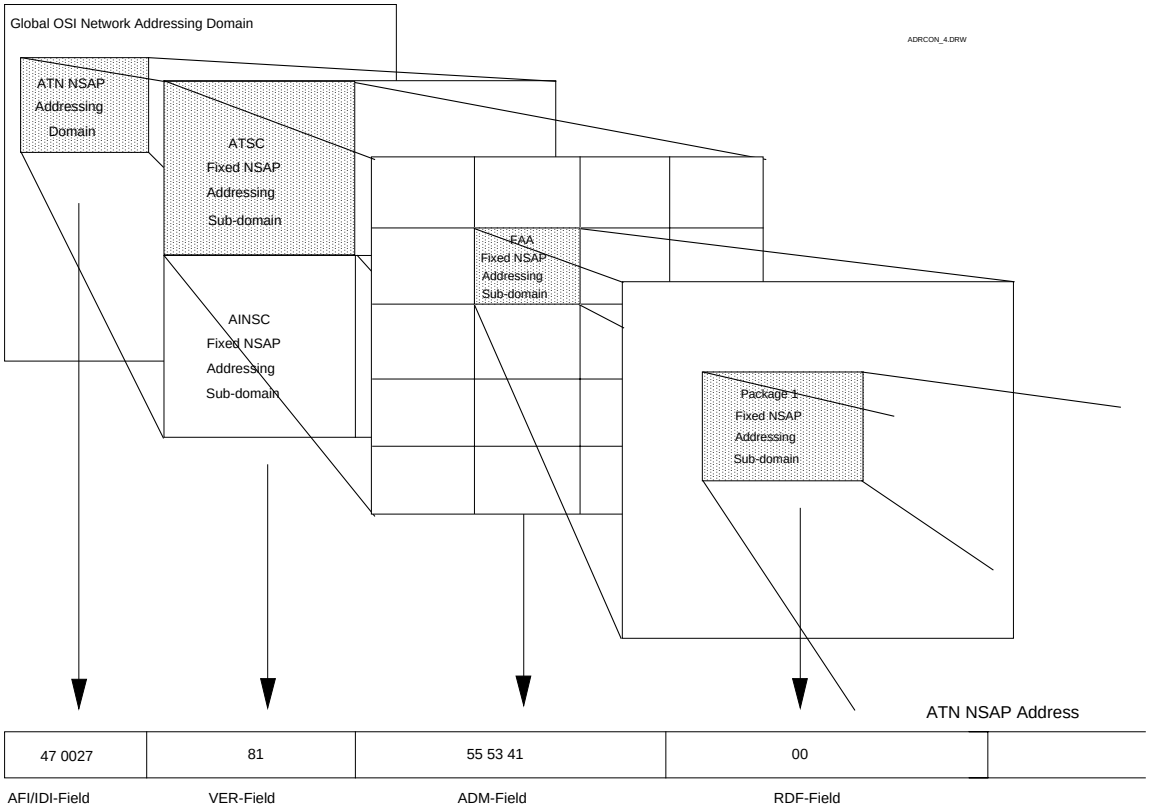

Figure 3-1 Hierarchical Decomposition of Addressing Domains Illustrated for ATN Network Addresses

3.3 Naming and Addressing Authorities

A naming/addressing authority defines the rules, including syntax (i.e. sizes and formats) and semantics (i.e. contents and interpretation), for specifying names/addresses within its naming/addressing domain and for the creation of further sub-domains. Furthermore, it allocates names/addresses within its domain according to specified rules, but does not perform the binding of the allocated names/addresses to the associated objects. This latter task is within the responsibility of the registration authority.

A naming/addressing authority may administer and allocate names/addresses itself, or, if it has partitioned its naming/addressing domain into naming/addressing sub-domains, may delegate the responsibility for naming/addressing within each such sub-domain to a sub-domain naming/addressing authority.

The overall naming/addressing authority for the ATN naming/addressing domains is ICAO which controls and manages these domains through the ATN SARPs.

Besides partitioning the ATN naming/addressing domains into appropriate sub-domains and specifying the syntax, semantics and encoding for these sub-domains, the ATN SARPs also directly allocate and register names/addresses within these sudomains, where appropriate or required. Furthermore, provisions have been made within the ATN SARPs which delegate full or partial responsibility for certain sub-domains (i.e. certain address fields) to organisations other than ICAO, such as IATA, regional ATS organisations and national ATS administrations. In order to comply with the ATN SARPs, these organisations have to implement the procedures for a naming/addressing authority and, where required, for a registration authority.

3.4 Name and Address Allocation

3.4.1 Address Allocation Principles

The general philosophy that is underlying the assignment of ATN network addresses is that the administration of the higher order address parts (i.e. the address domains which are close to the root of the hierarchical address structure) is performed by entities with a global scope, e.g. international organisations, such as ISO, ICAO and IATA. The further down in the hierarchical address structure one moves (i.e. the closer to tail of the address), the more the responsibility for address assignment and administration is delegated to entities with a more restricted scope, such as regional, national or local authorities. This means that the initial part of each ATN NSAP address and NET (i.e. the address prefix) is administered on an international level, the tail of the address is administered locally. Furthermore, this means that the process of allocating an address to an ATN object involves several authorities at different levels.

Figure 3-2 illustrates this distributed responsibility for address allocation using the example of an ATN TSAP address. This type of address is composed of 10 consecutive address fields comprising a total length of 21 or 22 bytes (depending on the length of the TSEL field which may be either one or two octets). According to the ATN addressing plan, address values within the first two fields (AFI and IDI) are assigned by ISO, within the next three fields (VER, ADM and RDF) by ICAO and IATA, or ICAO exclusively, and within the fields six to nine (ARS, LOC, SYS, NSEL) by ICAO Region authorities, State authorities and aeronautical organisations. Administration and address value assignment for the last field (TSEL) of an ATN TSAP address is done locally.

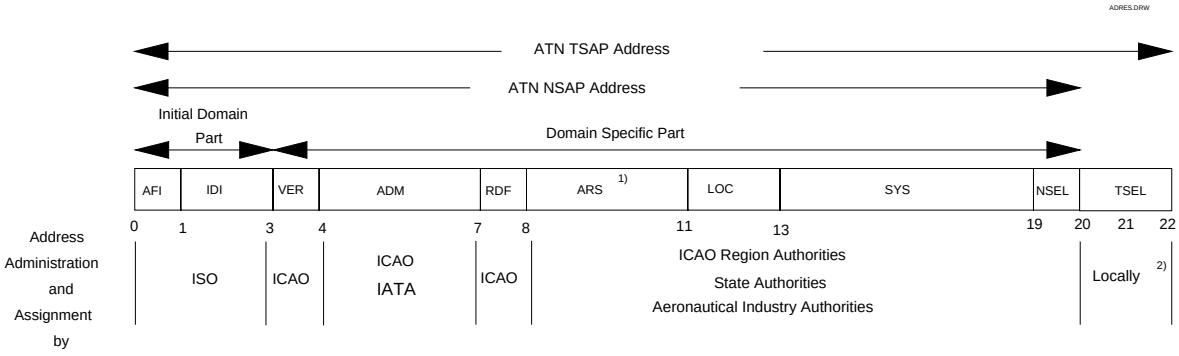

1) In mobile network addressing domains the value of the ARS field is the 24-bit ICAO Aircraft Identifier and is consequently assigned and administered by ICAO.

2) Allocation, assignment and administration activities may be transferred in parts or completely to the national State authority or aeronautical industry authority.

Figure 3-2 Responsible Addressing Authorities Illustrated for the Example of an ATN TSAP Address

It should be noted that, due to this hierarchical structure, each addressing authority manages an own addressing space. This addressing space is composed of those address values which establish the complete set of values to be constructed within the fields for which this authority is responsible. The registration authority for this address space is responsible to allocate and register addresses within this address space. The address space of this authority is linked into the overall ATN addressing space by appending the allocated values to the address prefix allocated by the superior address authority (i.e. the previous address fields).

The address registration function for the higher order fields of ATN addresses has already been partially performed in parallel to the development of the ATN SARPs. As a result of this, the values of the address prefix up to and including the RDF field (i.e. bytes 1 through 9 in Figure 2) of ATN addresses for ATSC systems are registered with ISO and ICAO and use either existing, already

established international registers or are published in Subvolume 5 of the ATN SARPs. In this sense, the ATN SARPs constitute in itself an international address register.

3.4.2 Responsibilities of Administrations

As illustrated in **Error! Reference source not found.**, a number of address fields of ATN addresses are expected to be registered and administered on a State level or by aeronautical industry authorities. Furthermore, provision are made in the ATN SARPs which foresee the delegation of the administrative responsibility for address assignment by ICAO to State authorities or aeronautical organisations.

Thus States are expected to assume full responsibility and administrative duties related to their own and/or delegated address space(s). The practical effect of this is that States, doing so, must establish the necessary administrative structure to carry out allocation, assignment and administration activities for ATN addresses, i.e. to place into operation an address (and naming) registration authority.

3.4.3 Registration Authority

A registration authority may be an organisation, a document, an automated facility, or any other body capable of name/address assignment that performs registration of one or more types of objects within its jurisdiction. The set of name/address values administered by a single registration authority constitutes the naming/addressing domain of this authority.

A registration authority is a specific instance of a naming/addressing authority (see section 3).

The role of a name/address registration authority is to :

1. assign and make available unambiguous names and addresses,

2. record definitions of the instances to which those names and addresses are assigned, and

3. propagate registered names and addresses to interested parties.

This includes activities such as:

- receiving proposals for register entries from applicants[1]

- processing received proposals for register entries, i.e. effect any necessary rationalisation or co-ordination of these proposals, such as check of responsibility of registration authority, verification of the qualification of the applicant, verification of the compliance of the proposal with the relevant provisions in the ATN SARPs

- recording names/addresses for each register entry that is accepted in accordance with the specified procedures for a register entry

- conveying the results of the registration, i.e. the decision taken on the proposal, to the appropriate applicant, and

- promulgating the register entries within its sphere of responsibility according to specified procedures and in a specified form.

---

[1] An applicant may be an ATS authority, a supra-national ATS organization, an airline, an aircraft operator, an aeronautical communication service provider or any other aeronautical industry organization.

A user of a registration authority may request an allocation of names/addresses from a registration authority, leaving the choice of names to the registration authority. Alternatively, a user of a registration authority may request an allocation of particular names/addresses. The registration authority may grant that request if it chooses, provided the names/addresses have not previously been issued.

The use of a name/address can be terminated by a registration authority and then the name/address re-used at a later time. The precise rules and constraints related to the re-use of a name/address to ensure unambiguous address allocation are within the responsibility of the registration authority.

Assigned and registered addresses have to be promulgated in order to enable communication partners to set up communication with the respective ATN system. The form how this address information is conveyed may reach from bilateral exchange of address registers on an "as-needed" basis to regular publication of official directories to a broad community, based on the individual security and reachability requirements of each participant in the ATN. (This promulgation of address information, which is a pure administrative matter and will most likely result in appropriate static configuration of ATN systems, should not be confused with the dynamic distribution of address information by routing protocols during ATN operation).

3.4.4 Delegation of Responsibilities

A State or organisation may chose to delegate its responsibility for its own addressing space to another State or organisation if it does not wish or if it is not practical to establish an own addressing authority and to carry out own administrative registration activities. In the case of delegation of addressing authority the respective State(s) or organisation(s) have to assume full administrative duties related to the delegated responsibilities. Appropriate interstate arrangements have to be established on a mutually agreed basis which cater for this transfer of authority.

3.4.5 Responsibilities of ICAO

Beside acting as addressing authority for a given portion of ATN addresses, the role of ICAO in the area of naming and addressing is one of international co-ordination, advice and consultation. Thus, ICAO may be expected to provide counsel to States and organisations having assumed such delegated responsibilities, in order to ensure that address administration is carried out in a manner that supports the orderly and efficient operation of the global ATN.

3.5 ATN Address Administration and Registration

In order to ensure unambiguous assignments, names and addresses must be registered by registration authorities within the context or environment in which they are to be used.

Name/address registration is the mechanism through which a name/address is assigned to an object in a way which makes the assignment available to interested parties. It is carried out by a naming/addressing registration authority (called registration authority in the following).

3.5.1 Subnetwork Address Administration and Registration

Registration of SNPAs is generally a consideration local to a subnetwork, but the SNPA addresses assigned to specific systems or services should be made available to all interested parties attached to a given subnetwork.

3.5.2 Network Address Administration and Registration

ICAO is the ultimate administrative authority of the ATN internetwork addressing plan and administers this plan through Section 5.4 of the ATN SARPs. Section 5.4 of the ATN SARPs

defines and administers the ATN NSAP address syntax (i.e. field boundaries, field sizes and field formats), the ATN NSAP address semantics (i.e. the field content and interpretation), and the ATN NSAP address encoding procedures (i.e. the representation of the abstract field syntax and semantics). Section 5.4 of the ATN SARPs delegates authority for the definition of the semantic content and the encoding of particular fields (or portions of address fields) as well as the assignment of address values for these fields to other authorities including ICAO Regions, States, and aeronautical or telecommunication standards organisations.

Four of the nine address fields of an ATN NSAP address have been assigned fixed values by Section 5.4 of the ATN SARPs and, consequently, do not need further registration. The values of the remaining five fields should be registered as follows:

ADM-Field

The ADM field shall be assigned a value representing either the AINSC or ATSC organisation (depending on the value of the VER field) which is responsible for the identified ATN Network Addressing Sub-Domain.

In the case of an ATSC Network Addressing Domain, the ADM field may identify either a State or an ICAO Region. When used for identifying a State, the ADM field contains the State's three-character alphanumeric ISO country designator defined in ISO 3166. When used to identify an ICAO Region, the first octet of the ADM field contains the ICAO Region Identifier defined in Section 5.4 of the ATN SARPs, while the value of the remaining two octets are assigned by the identified ICAO Region.

In the case of an AINSC Network Addressing Domain, the ADM field value should be an IATA 3-letter code for airlines (e.g. the "airline designator") and other stakeholders. IATA has already set up a registration procedure based on the use of current alphanumeric "airline designators" with extensions for "other stakeholders" compatible with the IATA Passenger Services.

One (or more) registration document(s) should be established and maintained, where values for the ADM field of the ATN NSAP address will be registered and published. This could be similar to the ISO 6523 standard document, where International Code Designators (ICD) are listed for addressing authorities registered with ISO (e.g. the value "0027" for ICAO can be found there). The registration document for ADM values may for instance contain tables with the layout illustrated in Table 3-1 and Table 3-2.

ICAO member / ATSC Administration	*ISO Country Code (3 letter code)*	*ADM value (hexadecimal)*
...		
Finland	FIN	46 49 4E
France	FRA	46 52 41
Germany	DEU	44 45 55
...		

Table 3-1 Example Allocation of ADM Field for ICAO Administrations

IATA member / AINSC Administration	*IATA Airline Code (3 letter code)*	*ADM value (hexadecimal)*

...		
Air France	AFR	41 46 52
British Airways	BAE	42 41 45
Lufthansa	DLH	44 4C 48
South African Airlines	SAL	53 41 4C
United Airlines	UAL	55 41 4C
...		

Table 3-2 Example Allocation of ADM Field for IATA Members

ARS-Field

In Fixed Network Addressing Domains, ARS values are assigned, administered and registered by the authority designated in the ADM field. A guideline for registration procedures and an outline for a registration document may be established by ICAO and/or IATA as the parent registration authorities.

Note - During the initial stage of Package-1, i.e. until appropriate registration authorities have been established by States, registration of ARS field values under the address space of ICAO will be within the responsibility of ATNP WG2.

Note - A common ICAO/IATA database comprising all registered NSAP address prefixes including the ARS value may be maintained and published through the appropriate ICAO and IATA channels. This database would allow to identify all registered ATN routing domains and the organisations responsible for each one. This database may be the basis of generalised directory services for post-Package 1 ATN deployments.

Note - In Mobile Network Addressing Domains, ARS values are defined a priori by the 24-bit ICAO Aircraft Identifier of the aircraft on which the addressed system is located.

LOC-, SYS-, N-SEL-Fields

Values for the LOC, SYS and N-SEL fields will be administered and registered either by the authority designated in the ADM field (i.e. an ICAO Region, State, or aeronautical industry organisation) or by appropriate sub-ordinate authorities to which registration authority has been devolved, if appropriate.

3.5.3 Transport Address Administration and Registration

According to Section 5.4 of the ATN SARPs, the TSAP selector of an ATN Transport address is administered on a local basis. This means that TSAP selector values are assigned by the organisation responsible for a given ATN End System within the constraints defined in Section 5.4 of the ATN SARPs.

The value for the TSAP selector field will be registered locally by a State's authority, or airline's or other aeronautical stakeholder's registration authority. In general, there is little need for ATN-wide coordination and publication of registered TSAP addresses. However, in order to support ATN directory services in the long term, the registered TSAP addresses of ATSC applications should be recorded and maintained in a possibly distributed ATN directory under the auspices of ICAO.

Note - The recording and global publication of the TSAP addresses associated with ground CM applications is an ATN requirement.

3.6 Address Allocation and Efficiency of ATN Operation

It is important to understand and to consider the impact of address assignment strategies on the quantity and frequency of exchange of routing information in the ATN, when performing the administrative role of an address authority. Routing information exchange efficiency may be adversely affected without due caution in the process of assigning network addresses.

In order to achieve the expected benefits concerning the reduction of address information in the routing traffic, in particular over severely bandwidth limited air/ground data links, co-ordination of address assignments within a region and/or a State will be required. This means that there may be the need for a centralised form of address allocation within a given region (at least related to air/ground communication objects). This might have the consequence that States have to delegate the responsibility for administration of a subset of their ATN systems to another State or an regional organisation.

3.7 Planning Aspects

It is important to consider that the establishment of an addressing and registration authority has to be planned and prepared well in advance to the operational phase of the concerned ATN systems. First, rules and methods for the registration process have to be defined and established and co-ordination with adjacent States or organisations may be required which have already reached a more advanced stage of their registration process. Secondly, addresses have to be assigned and these assignments have to be communicated even in the early phases of ATN deployment, when there are probably only some few ATN systems. Address "allocation" on a more or less random basis during these initial experimental and test phases should be avoided, as it is cumbersome to cancel existing assignments and re-allocate addresses to objects, if the overall ATN has reached a certain population. This is mainly due to the fact that the impact of such a re-allocation is not limited to the concerned object/system but impacts all its communication partners.

4. ATN Protocols and Functions

4.1 Introduction

There are two types of ATN System: the End Systems, which host the ATN Applications; and the Intermediate Systems that are the ATN Routers. Within these two basic types there are many variations. For example, there are some End Systems that are located on board aircraft and are part of the aircraft's avionics. There are also End Systems that are located in ATC Centres or are part of an airline's operational ground systems, and are the computers that host operational ATC and Airline applications. An End System is essentially any computer system that is connected to the ATN and implements the communications protocols necessary to access the ATN.

There are also many different types of ATN Router. In aircraft, airborne routers will also be part of an aircraft's avionics, and on the ground, ATN Routers will support both ground-ground and air-ground data communications. The various types of ATN Routers are classified in chapter two of the ATN Internet SARPs.

An ATN End System is required to support the ATN Transport Protocol, and the End System provisions for the Connectionless Network Protocol. In addition, the End System must implement the access protocol required for the subnetwork through which it accesses the ATN, and may also need to support the ISO 9542 ES-IS protocol. Support of ISO 9542 will be necessary if this is required by the ATN Router(s) through which the End System accesses the ATN, is a local matter as far as the ATN Internet SARPs are concerned.

An ATN Router is required to support the Intermediate System provisions for the Connectionless Network Protocol and most classes of ATN Router also require support of IDRP, although the support requirements for IDRP do differ depending on the role of the Router. Local considerations may also require support of the ISO 9542 EI-IS protocol and/or the ISO 10589 IS-IS protocol. ATN Routers must also implement the access protocol required for each subnetwork to which they are attached, and those attached to air-ground subnetworks are, additionally, required to implement the Route Initiation procedure specified in chapter 3 of the ATN Internet SARPs.

The remainder of this chapter is concerned with providing guidance for ATN Systems Implementors on the:

- Implementation of the Transport Protocol;
- Implementation of the Connectionless Network Protocol (CLNP);
- Implementation of the Inter-Domain Routing Protocol (IDRP); and
- implementation of the routing protocols that are outside of the scope of the ATN Internet SARPs, but which are nevertheless often required to meet local requirements.

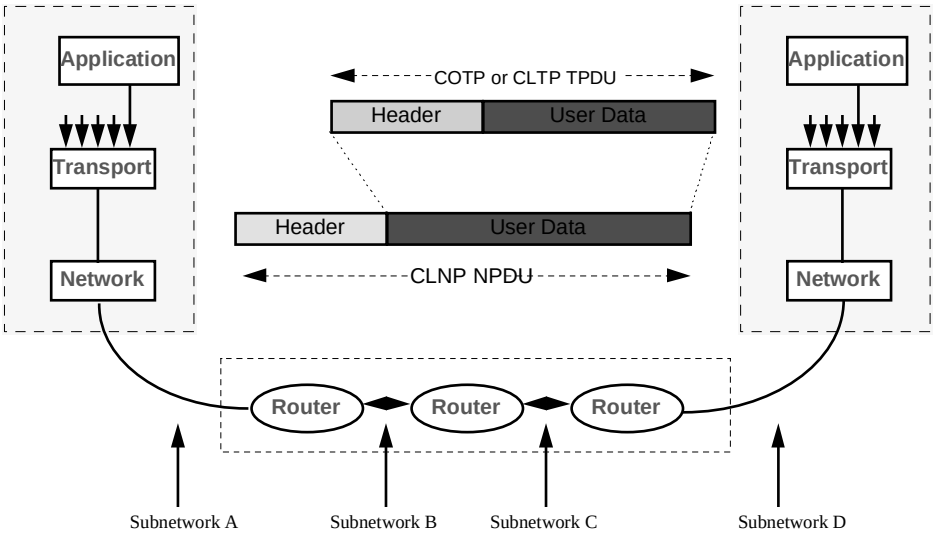

Figure 4-1 Scope of Transport Layer Interactions

4.2 Transport Layer Considerations

4.2.1 The ATN Transport Layer

4.2.1.1 Transport Layer Model

The OSI Transport Layer supports the end-to-end exchange of data between end systems, and serves as an interface between the application and the upper layers, which deal with the exchange of application messages, and the lower layers, which provide the necessary transmission and routing capabilities (see Figure 4-1). The applications and OSI upper layers that directly use transport layer services for the exchange of data are known as Transport Service users (TS-users).

TS-users receive a service which conceals the details in which reliable and cost effective transfer of data is achieved. This is achieved by the transport layer in an economical manner which is independent of the implementation specifics of the various subnetworks used, and of end system hardware and software implementation details. TS-users may choose to use either of two modes of transport service:

- The Connection Mode Transport Service (COTS), or
- The Connectionless Mode Transport Service (CLTS).

4.2.1.2 Transport Layer Protocols

The Transport Layer may consist of one or several transport entities, each implementing a different transport protocol. Two transport protocols are specified for use in ATN End Systems,: the Connection Oriented Transport Protocol (COTP) and the Connectionless Transport Protocol (CLTP). The COTS is support by the COTP and the CLTS by the CLTP.

A given ES may implement one or both of these, depending upon the requirements of the applications it contains. For example, if all of the applications in a given ES require only the COTS, then that ES does not need to implement the CLTP.

Both protocols support the exchange of application messages, henceforth referred to as Transport Service Data Units (TSDUs), while transferring each TSDU as one or more Transport Protocol Data Units (TPDUs), using the service provider by the ATN Network Layer.

- The **COTP** provides to its users an end-to-end connection mode service (i.e. the COTS), and is a conformant subset of the Class 4 Transport Protocol (TP4) specified in ISO 8073. This enables the reliable sequenced data transfer, where the user is guaranteed the byte order to data is preserved and that if a given TSDU is delivered then all previous messages will have been delivered. Both data integrity and data sequence integrity are supported by the COTP, together with end-to-end flow control.

- The **CLTP** provides a connectionless transport service (i.e. the CLTS), where no service guarantees are offered, other than preservation of the data integrity of each TSDU. Each CLTP TSDU is transferred as an event unrelated to the transfer of any other message and there is no guarantee either of delivery or that a TSDU may not overtake an earlier TSDU. The CLTP is conformant with ISO 8602.

4.2.1.3 Service Provided by the ATN Transport Layer

ATN Applications may choose to use either the COTS or the CLTS. The selection of the transport service used by an application is influenced by the communications characteristics and the quality of service requirements of that application. However, the choice of which mode to use cannot usually be left to the implementor. This must be specified by the application specification. It is the implementor's responsibility to implement the transport protocol necessary to support an End System's applications' requirements.

4.2.1.3.1 Service Provided by the COTS

As far as application designer's are concerned, the COTS is appropriate when users need to maintain an association, either because they need to transfer a lengthy data stream, or because the applications need to maintain a close binding (e.g. as a test of liveness). COTS is also appropriate for applications that place a higher importance on data sequence integrity than transit delay. The characteristics of the service provided by the connection mode protocol include the following:

- TS-users negotiate the establishment of a transport connection, prior to actual data transfer; this connection enables reliable data transfer between the two. An initial delay is associated with the establishment of a transport connection. During this phase, data cannot be exchanged.

- Maintenance of a transport connection will generally incur some additional costs associated with the transfer of TPDUs not associated with user data, such as acknowledgements. Acknowledgements are utilised for data acknowledgements, flow control purposes and keep-alive indicators.

- The order of submission of TSDUs is preserved on delivery.

- The underlying transport protocol provides facilities to detect and recover from end-to-end transmission errors within a TSDU.

- The underlying protocol is capable of segmenting TSDUs, allowing TSDU sizes larger than the maximum NSDU size. This has the potential for improving network performance, because network level (that is, the connectionless network protocol) segmentation is less efficient than transport segmentation.

- The underlying protocol has the capability to control the flow of TSDUs. This allows the receiver of information to adjust the rate of incoming TSDUs to meet local processing capabilities. In addition, this flow control can be exercised by a transport entity to react to varying network congestion problems, applying and relieving constraints to match resource limitations.

- Operation of the COTP requires system resources to maintain shared state and to monitor connection status.

4.2.1.3.2 Service Provided by the CLTS

The CLTS is appropriate when there is a requirement for time-critical data transfer, i.e. it is more desirable to discard data rather than apply flow control or retransmission techniques. The connectionless mode transport service is supported using the ISO 8602 protocol. The characteristics of the service provided by the CLTP include the following:

- No negotiation takes place before a TSDU is transmitted from one user to another. This mode does not have the delay associated with establishing a transport connection before data can be exchanged.

- There are no TPDUs transmitted other than those carrying user data.

- Each TSDU is transmitted independently from all others; TSDU delivery and TSDU delivery sequence are not guaranteed. There is no transport-layer recovery on detected errors.

- The transport protocol can employ facilities to detect end-to-end transmission errors within a TSDU. TSDUs containing detected errors are discarded.

- TSDU sizes are limited to the maximum NSDU size on each end system; no segmentation is performed by the connectionless mode transport protocol.

- Because there is no negotiated relationship between TS-users, the protocol does not have the capability to control the flow of TSDUs.

- The processing requirements for the connectionless transport protocol are minimal, since the transport protocol does not perform any TSDU sequencing or TSDU guarantee functions.

4.2.1.4 Transport Addresses

Users of the Transport Service are uniquely identified by their Transport Address (TSAP Address).

A TSAP address comprises two elements, an NSAP address and a TSAP-selector. The NSAP address provides the address of the transport protocol entity for a particular ES, such as the connection mode transport layer. The TSAP Selector then identifies one of the users of the transport protocol entity. Note that it is possible for the COTP and CLTP to share a common NSAP Address. However, if the End System supports other Transport Protocols (e.g. TCP), then these must use different NSAP Addresses.

4.2.1.5 Network Service Assumptions

The ATN Transport Layer operates using the connectionless network service provided by the ATN network layer. All TPDUs are transmitted and received as NSDUs using the N-UNITDATA service of the network layer. Each NSDU is considered independent of the others, and may arrive in a different order than was sent, in duplicate, or not at all. Although it is possible for NSDUs to be lost, the ATN is expected to have a low loss rate, based on the intrinsic reliability of the subnetworks

supporting communications. NSDU loss is only expected during times of network congestion, when NPDUs are discarded by congested routers.

4.2.1.6 ATN Security and Priority

The ATN SARPs specify the use of an ATN Security Label and the prioritisation of data. In the COTP an ATN Security Label applies to a transport connection rather than an individual TSDU, and all TSDU sent over a given transport connection must have the same ATN Security Label. On the other hand, in the CLTP, each TSDU may be assigned a separate ATN Security Label.

Similarly, in the COTP, a transport connection is given a priority, and all TSDUs sent over that transport connection have the same priority, while, in the CLTP, each TSDU may have a different priority.

The ATN Security Label and priority applicable to each TPDU are parameters of the N_UNITDATA service and are therefore encoded in the NDPU header, rather than each TPDU, and are referenced by the network layer forwarding function. For such reasons, TPDUs from transport connections with different ATN Security Labels, and/or priorities, cannot be concatentated.

4.2.2 Provision of the Connection Mode Transport Service

4.2.2.1 Overview

The operation of a Transport Connection (TC) is modelled as a pair of queues linking the two TSAPs to which the communicating TS-users are attached. For each TC, a pair of queues is considered to be available: one queue for the information flow from user A to user B, and one queue for the information flow from user B to user A. Each user of a TC is provided with the COTS.

The COTS may exist in four possible states: idle, connection establishment, data transfer, and connection release. In the idle state, there is no connection and data transfer cannot take place. In order to transfer data, a transport service user must request that a transport connection is established with the required remote transport service user, identified by its Transport Address. While an attempt is made to establish a transport connection, the COTS enters the connection establishment state.

During the connection establishment state, the transport entity attempts to establish contact with the remote transport service user. If it is successful, and the remote user agrees to the connection, then a transport connection is established, the data transfer state is entered, and data transfer may take place. If it is not successful then the COTS returns to the idle state.

Either user of a transport connection may, at any time, request that the transport connection is released. The COTS then enters the connection release state. This is only a transitory state as the connection is always released immediately the request is made with any in-transit data lost - it is the responsibility of the transport service users not to release the connection before all data has been transferred. The idle state is then re-entered.

The COTS is realised through the implementation of the COTP.

The ATN COTP uses the ISO 8073 class 4 procedures and is therefore able to operate over a CLNS, such as provided by the ATN network service. The Transport Protocol reacts to network status information and hides any problems from the TS-user.

For the transfer of TSDUs, the transport layer provides a known set of characteristics, as noted below.

- **TSDU Sequencing** — The ATN COTS guarantees that TSDUs will be delivered to the destination TS-user in the order they have been submitted by the source TS-user to the TS-provider. The only exception is expedited data which, being subject to a different flow control scheme, may overtake normal data.

- **TSDU Delivery Support** — The transport layer supports the delivery of a submitted TSDU to the destination TS-user. The only case where data may be lost is if the connection release phase has been entered by the local or remote TS-user and/or provider.

- **End-to-End Detection and Recovery of Error.** — Class 4 of the connection mode transport protocol provides mechanisms that support the detection and recovery of errors such as TPDU loss, duplication, or corruption. The error detection and recovery is done transparently to the user.

4.2.2.2 Connection Mode Transport Service Primitives

There are ten connection mode transport service primitives. In the connection establishment phase, the TS-user issues the T-CONNECT Request and the T-CONNECT Response; the TS-provider issues the T-CONNECT Indication and the T-CONNECT Confirmation. In the data transfer phase, the TS-user issues the T-DATA Request and the T-EXPEDITED DATA Request; the TS-provider issues the T-DATA Indication and the T-EXPEDITED DATA Indication. In the disconnect phase, the TS-user issues the T-DISCONNECT request; the TS-provider issues the T-DISCONNECT indication.

A TS primitive issued by one TS-user will, in general, result in receipt of an indication by the other TS-user. Figure 4-2 gives a summary of TS-primitive time-sequence diagrams for some typical scenarios.

Each of the Connection Mode TS primitives has one or more associated parameters. They will be discussed in detail in subsequent sections.

Note.— In Figure 4-2, *the flow of time is represented by the downward direction in individual figures. The sequential relation between two points of interaction is shown by a horizontal line which is discontinuous between the two vertical lines representing the flow of time (e.g. the T-CONNECT request primitive in (a) invoked by a TS-user at moment t1, is necessarily followed by a T-CONNECT indication primitive invoked by the remote TS-provider at moment t2). The absence of relationship is indicated by using a tilde (~).*

Figure 4-2 is derived from a state transition diagram which defines the allowed sequences of TS primitives at a TC endpoint. This state transition diagram pertains to the Transport Protocol Machine.

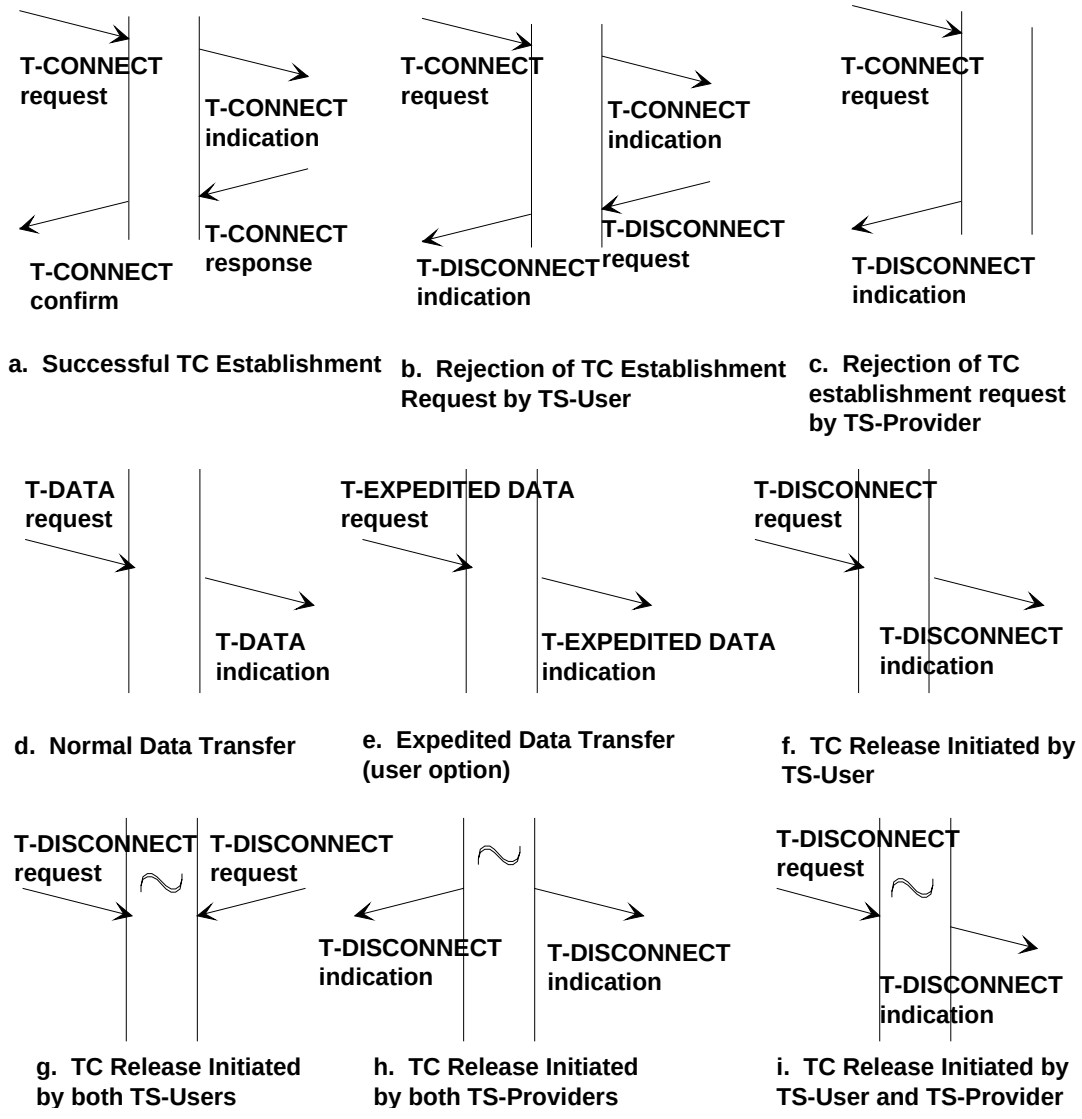

Figure 4-2 Transport Service Time Sequence Diagrams

4.2.2.3 The Connection Mode Transport Protocol (COTP)

4.2.2.3.1 Overview

COTP procedures support connection establishment, data transfer, and connection release. Although some type of connection management is handled by almost every layer, it is especially complex at the transport layer due to the unpredictability of network errors or delay.

There are two basic mechanisms used for transport connection management: the handshake-based mechanism and the timer-based mechanism. Handshake-based mechanisms use explicit exchanges in response to a given packet initiating an action, such as connection establishment. Timer-based mechanisms are, for example, used by the sender and receiver keeping track of the system state long enough to ensure that all PDUs from closed connections have left the system.

The handshake and timer-based mechanisms are combined to ensure that connection identifiers are

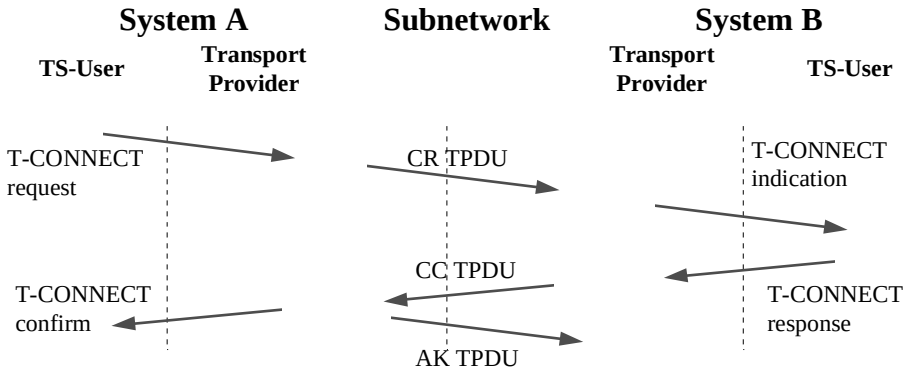

Figure 4-3 TPDU Exchanges for Connection Establishment

unique during the maximum time packets may remain in the system.

4.2.2.3.2 Connection Establishment

The COTP uses a three-way handshake mechanism in combination with a timer-based mechanism to ensure connection establishment in class 4. Figure 4-3 illustrates a typical transport connection establishment procedure. The service user, either the session layer or a specific application at system A, passes a T-CONNECT request primitive to its service provider (the transport layer) with appropriate parameters for setting up the connection. The transport layer entity of A then generates a connection request TPDU containing the parameter values and sends it to its peer transport layer entity at B. The transport entity at B generates a T-CONNECT indication primitive and passes it to its user.

If the user B accepts the connection establishment request, it generates a T-CONNECT response. The transport entity at B then transmits a connection confirm (CC) TPDU to the transport entity at A. Finally the transport entity at A informs its user that its connection establishment request has been accepted by invoking a T-CONNECT confirm primitive.

The transport entity at A also generates an acknowledgement (AK), or a data (DT), or expedited data (ED) TPDU (if there are data to be transferred), and sends it back to the transport entity at B. The connection is considered established only after the transport entity at B has received this acknowledgement or data TPDU.

If the connection request is initially refused by the TS-provider at A, a T-DISCONNECT indication is sent back to the TS-user at A as illustrated in Figure 4-4.

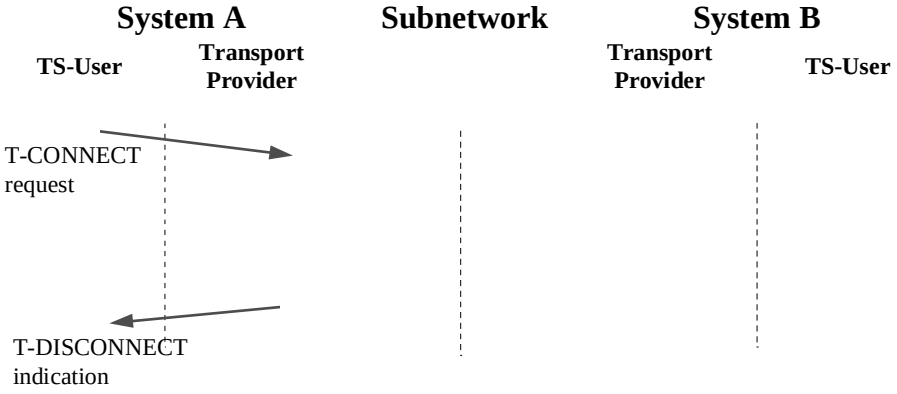

Figure 4-4 Connection Refusal by the TS Provider

To initiate communication with a peer, a TS-user invokes the T-CONNECT request primitive (see Figure 4-2). Upon arrival at the destination TSAP, a T-CONNECT indication is delivered to the destination ATN TS-user. The peer TS-user accepts the connection request by issuing a T-CONNECT response primitive. Finally, the calling TS-user receives a T-CONNECT confirm primitive and the connection is established. Simultaneous T-CONNECT requests typically result in a corresponding number of TCs. The parameters associated with the connection establishment primitives are listed in Table 4-1.

As part of the TC establishment phase, TS-users can negotiate the QoS parameters to be associated with a transport connection. Use of expedited data is also negotiated. QoS parameters are used to describe the desired characteristics of the data flow over the TC, rather than to provide mechanisms for the transport protocol to enforce specific characteristics. The use or non-use of expedited data is negotiated between TS-users, and will be selected based on TS-user requirements. Furthermore, some negotiations take place between TS-providers which are transparent to the TS-users. All the choices made during the connection establishment phase remain valid for the whole TC lifetime. The TC establishment procedure may fail due to:

- timeout procedures, such as when a TS-user does not respond to a connection request
- rejection by the TS-provider of an attempt to establish a TC (part c of Figure 4-2), for reasons such as invalid or unknown called TSAP address, lack of local or remote resources of the TS-provider etc., or,
- unwillingness of the called TS-user to accept the TC establishment request (part b of Figure 4-2).

The TC establishment may also fail due to either of the TS-users releasing the TC before the T-CONNECT confirm has been delivered to the calling TS-user.

4.2.2.3.2.1 Connection Request

A calling TS-user, when invoking a T-CONNECT request primitive, specifies the following parameters :

- **Called Transport Address:** The called transport address contains the addressing information necessary to reach the desired destination TS-user. An ATN called transport address comprises an ATN NSAP address and a TSAP Selector (also called TSAP-ID in ISO 8073).

- **Calling Transport Address:** The calling transport address contains the addressing information that identifies the TS-user invoking the T-CONNECT request. An ATN calling transport address comprises an ATN NSAP address and a TSAP selector.

- **Expedited data option:** By means of this parameter the communicating TS-users negotiate the use or non-use of the expedited data service for the TC in question. The calling TS-user initially specifies the use or non-use of expedited data. If non-use is initially proposed, the called TS-user cannot further negotiate its use. If its use is initially proposed, the called TS-user can either confirm use or can select non-use of the expedited data option.

- **Requested Quality of service:** QoS parameters are used to describe the desired characteristics of the data flow over the transport connection. The parameters which may be negotiated are transit delay, residual error rate, and priority.

- **TS-user-data:** A user can specify data from 1 to 32 octets in the connection establishment request. These data can be used by the TS-user in a manner agreed with the peer TS-user. For example, the information could be used to communicate authentication and access control information. It should be noted that the delivery of TS-user-data is not guaranteed. TS-user-data

Parameters	**Transport Service Primitive**			
	T-CONNECT Request	**T-CONNECT Indication**	**T-CONNECT Response**	**T-CONNECT Confirm**
Called Address	M	M(=)		
Calling Address	M	M(=)		
Responding Address			M	M(=)
Expedited Data Option	M	M(=)	M	M(=)
Quality of Service	M	M	M	M(=)
TS User Data	M	M(=)	M	M(=)
Security	O	O(=)	O	O(=)

Note: in the above table:

M *The parameter is mandatory*

(=) *The value of the parameter in the T-CONNECT Indication/Confirm is identical to the value of the corresponding parameter in the T-CONNECT Request/Response TS primitive*

O *Use of this parameter is a TS-user option*

Table 4-1 TC Establishment Primitives and Parameters

are not recommended for direct use by applications.

- **Security:** The security parameter may be used by the service user to indicate the value of the security label. The syntax and semantics of the ATN Security Label are specified in the ATN Internet SARPs..

Note 1.— Negotiation of options only proceeds in a "mandatory" direction. That is, the called TS-user can always negotiate to the mandatory aspect of any option.

Note 2.— In practice, not all of the parameters in a connection request must be explicitly specified, even though they exist in the service interface. For example, the invoking TS-user may only be required to specify the called transport address if the transport entity knows the calling address a priori. *Other parameters, if not specified, may take on default values. For example, most implementations today do not require explicit specification of QoS values. If not specified, one of two things may occur: QoS parameters may not be conveyed in the CR TPDU or the TE may select a standard set of parameters.*

4.2.2.3.2.2 Connection Indication

A T-CONNECT request issued by a TS-user results in a corresponding T-CONNECT indication to the destination ATN TS-user. The TS-provider, when issuing the T-CONNECT indication, specifies the following parameters:

- Calling and called address
- Expedited data option
- TS-user-data
- Indicated QoS

The values of the first three parameters are delivered unchanged by the TS-provider to the destination TS-user. The values of the indicated QoS parameters can be equal to or poorer than the requested QoS parameters selected by the calling user in the T-CONNECT request primitive. The value of a QoS parameter can be downgraded by either the transport entity serving the calling TS-user or the transport entity serving the called TS-user. This will happen if the transport entity has additional provisions implemented which monitor the ability to provide the requested QoS.

4.2.2.3.2.3 Connection Response

To accept the TC establishment, the called TS-user issues a T-CONNECT response primitive (otherwise, it invokes a T-DISCONNECT primitive and the connection is not established; see Figure 4-4). The associated parameters and their corresponding values are the same as in the T-CONNECT request.

4.2.2.3.2.4 Connection Confirm

A T-CONNECT response primitive at one TC endpoint starts the delivery of a T-CONNECT confirm primitive at the other TC endpoint. This primitive has exactly the same associated parameters as those of the T-CONNECT response primitive. The values of these parameters are also equal, that is, the TS-provider delivers these values unchanged to the calling TS-user. Once this primitive has been received by the calling TS-user, the connection is considered to be established.

4.2.2.3.3 Data Transfer

Once a connection has been successfully opened data transfe may take place. Normal data transfer is always full duplex with independent flow control in each direction. The Quality of Service is assumed to be the same in each direction.

TP4 implements a sliding window flow control mechanism enabling AKs to be returned while data are still being sent. An AK is returned when the acknowledgement timer set or reset after receipt of data expires. The acknowledgement timer mechanism enables multiple TPDUs to be acknowledged with the same AK TPDU. An example of normal data transfer is shown in Figure 4-5, which illustrates the transmission of a single transport service data unit via multiple TPDUs. After the establishment of the transport connection, the initial DT TPDU number is 0 (DT 0). An initial credit of 1 is assumed and transport entity A waits for an acknowledgement with more credit. Transport entity B returns AK 1, with a credit (CDT) of 2, allowing the transmission of two more TPDUs. When the EOT (end of TSDU) bit is set to 1 in the final DT TPDU, the sequence ends and the whole TSDU is delivered to user B. At the expiration of the acknowledgement timer, an AK is returned. This AK acknowledges up through the final TPDU.

The transport service provides for bidirectional exchange of TSDUs while preserving the integrity, sequence and boundaries of TSDUs. Two kinds of transfer service are offered by the ATN COTS provider: the normal data transfer service and the expedited data transfer service. Figure 4-2(d) describes the primitive sequences in a successful transfer of normal data.

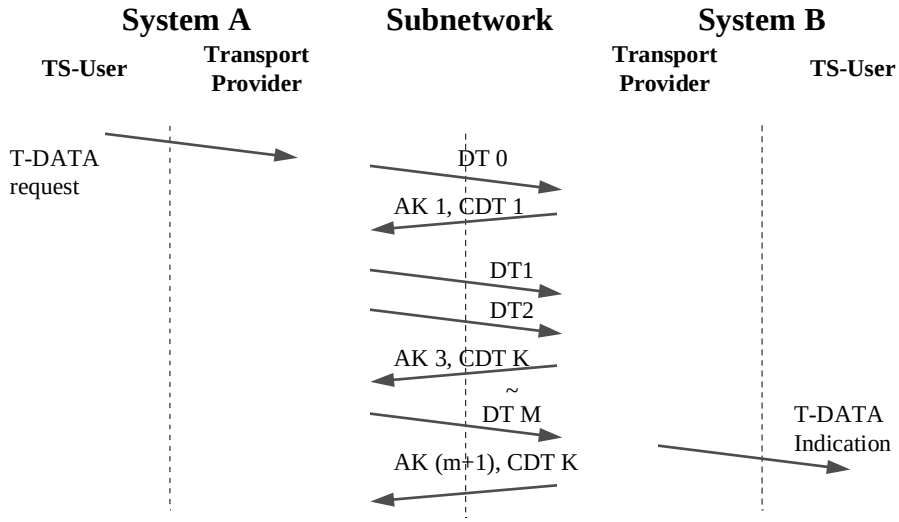

Figure 4-5 Normal Data Transfer

4.2.2.3.3.1 Data Request

A TS-user requests the transfer of a TSDU by invoking a T-DATA request primitive. This primitive has only one associated parameter: the TS-user-data parameter (i.e. the TSDU to be transmitted). A TSDU consists of an integral number of octets greater than zero; the length of a submitted TSDU is limited by implementation constraints only.

4.2.2.3.3.2 Data Indication

Upon arrival of the TSDU at the other TC endpoint, the TS-provider invokes a T-DATA indication primitive to the destination TS-user. The TS-user-data parameter of the T-DATA request primitive is delivered unchanged by the TS-provider to the destination TS-user.

4.2.2.3.4 Expedited Data Transfer

This service is available on a given TC only if its use has been requested by the calling TS-user and agreed to by the called TS-user during the TC establishment phase. The TS-provider guarantees that an expedited TSDU will not be delivered after any subsequently submitted normal TSDU or expedited TSDU on the same TC. The transfer of expedited TSDUs is subject to separate flow control from that applied to the data of the normal transfer service. Figure 4-2 (e) shows the sequence of primitives in a successful transfer of expedited data.

4.2.2.3.4.1 Expedited Data Request

A TS-user desiring to transmit an expedited TSDU invokes the T-EXPEDITED DATA request primitive. This primitive has only one associated parameter: the TS-user-data parameter (i.e. the TSDU to be transmitted).

An expedited TSDU consists of an integral number of octets between 1 and 16 inclusive .

4.2.2.3.4.2 Expedited Data Indication

Upon arrival at the destination, the TS-provider invokes a T-EXPEDITED DATA indication primitive which delivers the submitted TSDU (TS-user-data parameter) unchanged to the destination TS-user.

4.2.2.3.5 Connection Termination

Connection release can be performed at the initiative of either TS-user or TS-provider at any point in the lifetime of the transport connection. This is an abrupt release because the transport protocol does not have functions that support prior negotiation of termination and so data may be lost. Typical scenarios of connection release are demonstrated in Figure 4-2 (f) through (i).

The first scenario (f), is shown in more detail in Figure 4-6. User A sends a disconnect request (DR), the Transport entity at B sends a T-DISCONNECT indication to user B and the connection ends. A disconnect confirm (DC) TPDU is sent back from system B to system A.

In Figure 4-2 (g), the two users send a DR at the same time. In the third case (h), the transport layer itself (either the entity at B or at A) generates the DR. In the fourth case (i), user A sends a DR after the transport layer has initiated termination of the connection.

A TS-user may issue a connection termination primitive to refuse TC establishment or to release the established TC. The TS-provider never guarantees delivery of submitted data - it just guarantees order preservation - if it delivers a TSDU it guarantees to have delivered all previously submitted TSDUs. There is always an uncertainty over how much data has been lost once the release phase is entered and includes TSDUs submitted well before the release phase was entered. The degree of data loss is independent of the credit window, and depends on the length of the queue between TS-provider and TS-user. In particular, all data received after a transport entity has entered the release phase are discarded. The parameters associated with the connection termination primitives are summarised in Table 4-2.

Parameters	Transport Service Primitive	
	T-DISCONNECT Request	**T-DISCONNECT Indication**
Reason		M
TS User Data	M	M(=)

Note: in the above table:

M *The parameter is mandatory*

(=) *The value of the parameter is identical to the value of the corresponding parameter in the preceding TS primitive*

Table 4-2 TC Release Primitives and Parameters

4.2.2.3.5.1 Disconnect Request

A TS-user releases an established TC by invoking the T-DISCONNECT request primitive. This primitive has only one optional parameter: the TS-user-data parameter. The TS-user-data parameter is an integral number of octets in length between 1 and 64 inclusive. The content of this parameter may provide additional information on the reasons for the TC release request.

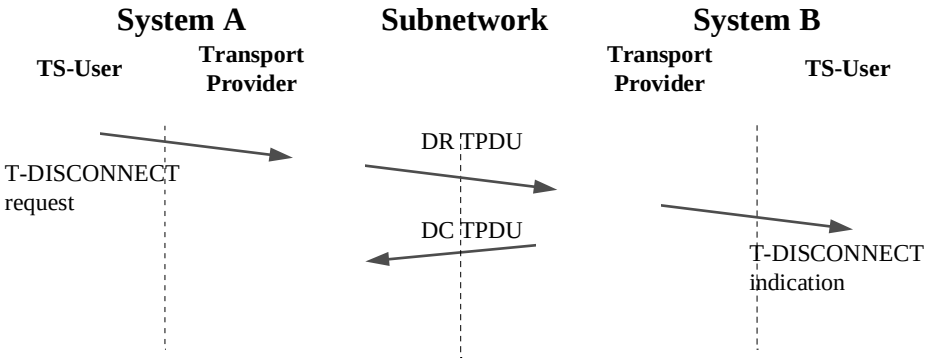

Figure 4-6 Transport Connection Termination

4.2.2.3.5.2 Disconnect Indication

The T-DISCONNECT indication primitive has different parameters, according to the originator of this primitive. If the T-DISCONNECT indication is invoked by the TS-provider as a result of a T-DISCONNECT request invoked by a TS-user at the other TC endpoint, this primitive has the following associated parameters:

- **TS-user-data:** This parameter is present only if it was also present in the T-DISCONNECT request primitive. These data are normally delivered unchanged by the TS-provider, except if the TS-provider initiates TC release before the T-DISCONNECT indication is delivered (see part (i) of Figure 4-2), or if TS-users initiate a T-DISCONNECT request simultaneously (see part (g) of Figure 4-2). In these cases these data may be lost.

- **Reason**: This parameter will take the value "remote TS-user invoked".

If the T-DISCONNECT indication is invoked by the TS-provider itself, the only associated parameter is the "Reason" parameter which takes the value "TS-provider-invoked" (in this case no TS-user-data parameter is present). Examples of reasons for a TS-provider-initiated release include: lack of local or remote resources of the TS-provider, misbehaviour of the TS-provider, called TS-user unknown, or called TS-user unavailable (if the release occurs during the connection establishment phase).

4.2.2.4 The ATN Security Label

ATN Security Functions are concerned with:

a) Protecting CNS/ATM applications from internal and external threats;

b) Ensuring that application Quality of Service and Routing Policy Requirements are maintained, including service availability; and,

c) Ensuring that air-ground subnetworks are used in accordance with ITU requirements.

The ATN Internet provides mechanisms to support items (b) and (c) above only. These mechanisms are defined to take place in a common domain of trust, and use a Security Label in the header of each CLNP Data PDU to convey information identifying the “traffic type” of the data and the application’s routing policy and/or strong QoS Requirements. Strong QoS Requirements may only be expressed by ATSC Applications, and they are expressed as an ATC Class identifier, encoded as part of the ATN Security Label.

Except when a transport connection is used to convey general communications data, each transport connection is associated with a single ATN Security Label. The value of this label is determined when the connection is initiated, and by the initiating TS-User. A responding TS-user may refuse to accept a transport connection associated with a given ATN Security, but cannot propose an alternative. It is also not possible to change an ATN Security Label during the lifetime of a transport connection.

The ATN Security Label is never actually encoded into a TPDU header. Instead, every NSDU passed to the Network Layer that contains a TPDU from a transport connection associated with an ATN Security Label is associated with the same ATN Security Label. This is passed as a parameter to the N-UNITDATA request, and then encoded into the NPDU header.

TPDUs from transport connections associated with different ATN Security Labels cannot be concatenated into the same NSDU.

Note. The mechanism by which the connection initiator specifies the appropriate ATN Security Label for a given transport connection is a local matter. For example, it may be identified by an extension to the transport service interface, be implicit in the choice of a given TSAP, or be identified using a Systems Management function. Similarly, the mechanism for determining the ATN Security Label associated with an incoming transport connection is a local matter.

4.2.2.5 ATN Transport Layer Quality of Service

QoS parameters are used to indicate the required characteristics of the underlying communications service supporting application information exchange. The transport layer may interpret the QoS parameters which are provided by a TS-user; these parameters may then affect the interactions of the transport layer with the network layer providing service.

QoS is of special importance to the aviation community because of the wide variation in service provided by the ATN network service. However, there are practical difficulties in a connectionless internet, as regards dynamic route selection based on differential QoS requirements. While dynamic

route selection is still a long term goal, in the near to medium term, application QoS requirements will be met through the following principles:

a) The capacity requirements of CNS/ATM-1 Applications will be met through a combination of network design and capacity planning, in order to ensure that network capacity both exists and is usable by CNS/ATM-1 Applications, and that their QoS Requirements will be met to the required availability.

b) The strong QoS Requirements of certain ATSC Applications will be met, without having to design the whole ATN to meet their QoS requirements, by reserving certain subnetwork paths for applications data of at least a given ATSC Class, as identified by the ATN Security Label associated with the data.

c) The strong QoS Requirements of certain AISC Applications will be met by respecting routing policy requirements, restricting their data to travel over only certain air/ground data links, expressed in the ATN Security Label associated with the data.

The only exception to this is *Residual Error Rate*. The ATN Internet provides an expected residual error rate of 1 in 10^8. This may be improved upon through use of the transport protocol checksum mechanism, and it is believed that with this additional mechanism, an undetected error rate of 1 in 10^{13} is achievable. Although checksum use is not explicitly indicated by a TS-user, its use can be defined either through configuration techniques or it can be inferred based on the QoS requirements of the TS-user.

Since checksums are contained in the TPDU header, implementation of checksums is a protocol performance issue. However, the checksum is essential for ensuring protection against undetected errors.

4.2.2.6 Priority

Although priority is defined by ISO 8072 to be part of QoS, it is important enough in the ATN to be treated separately.

The purpose of priority is to signal the relative importance and/or precedence of data, such that when a decision has to be made as to which data to action first, or when contention for access to shared resources has to be resolved, the decision or outcome can be determined unambiguously and in line with user requirements both within and between applications. In the ATN, priority is signalled separately by the application in the transport layer and network layer, and in ATN subnetworks. In each case, the semantics and use of priority may differ.

In the ATN Internet, priority has the essential role of ensuring that high priority safety related data is not delayed by low priority non-safety data, even when the network is overloaded with low priority data.

In the ATN Transport Layer, priority is concerned with the relationship between transport connections and determines the relative importance of a transport connection with respect to (a) the order in which TCs are to have their QoS degraded, if necessary, and (b) the order in which TCs are to be broken in order to recover resources. The transport connection priority is specified by the initiating TS-user either explicitly or implicitly, when the transport connection is established. As with the ATN Security Label, priority is not negotiable, and a responding TS-user must either accept the proposed priority or reject the connect request. TPDUs belonging to transport connections with different priorities cannot be concatenated.

When an ATN Transport Layer entity is unable to satisfy a request for a transport connection from either a local or remote TSAP, and which is due to insufficient local resources available to the transport layer entity, then it is required to terminate a lower priority transport connection, if any, in order to permit the establishment of a new higher priority transport connection.

Transaport Layer implementations may also use transport priority to arbitrate access to other resources (e.g. buffers). For example, this may be achieved by flow control applied to local users, by discarding received but unacknowledged TPDUs, by reducing credit windows, etc.

All TPDUs sent by an ATN Transport Layer Entity are transferred by the ATN Internet Layer, using the Network Priority that corresponds to the transport connection's priority according to Table 2-3 of the ATN Internet SARPs. The network priority is signalled by a parameter to the N-UNITDATA request, and the priority of an incoming NSDU is signalled by a parameter to the N-UNITDATA indication.

Transport Priority may be encoded into the CR TPDU. However, this is not essential and, if present must be equivalent to the network priority of the NSDU that conveys the CR TPDU. The priority of this NSDU determines the priority of the transport connection.

When specified, transport priority values should be integers in the range from 0 to 14, with priority level 0 as the highest priority. In such a case, there is a one-to-one correspondence with the CLNP priority values (see chapter 2 of the ATN Internet SARPs for further details on the mapping of Transport priority values to CLNP priority values).

4.2.2.7 Negotiation of Connection Parameters

The ISO transport layer allows areas of negotiation in the connection establishment phase. One of the negotiated features is the class of operation. Depending on the class selected, other features are also negotiated.

Negotiation in the transport layer is based on the following assumptions:

a. If a feature is not negotiated, the "default" option, or "mandatory" implementation of the option, is selected.

b. To suggest anything other than the default, the proposed value must be explicitly proposed in a connection request.

c. The responder has the choice of explicitly accepting the proposed value or possibly selecting a "lesser", or "mandatory" value. If the responder does not explicitly indicate the desired value, the default is in effect.

For example, one option for class four operation is the use of checksums. The default is use of checksums, and all implementations must be able to support use of checksums on a connection. To operate a connection without checksums, the requester must explicitly propose "non-use of checksums". If the responder does not explicitly reply with "non-use of checksums", then the checksum procedures are in effect for that connection.

Table 4-3 indicates the items that can be negotiated and their default, or mandatory, values in Class 4 operation.

Feature	Allowed Values	Default
Preferred TPDU Size, octets	Multiple of 128	128
Maximum TPDU Size, octets	128, 256, 512, 1024, 2048, 4096, 8192	128
TPDU Numbering Format	normal, extended	normal
Expedited Data	use, non-use	non-use
Checksum	use, non-use	use
Selective Acknowledgement	use, non-use	non-use
Request Acknowledgement	use, non-use	non-use

Table 4-3 Negotiable and Default Values for Class 4 Operation

4.2.2.7.1 Class Negotiation - Initiator.

The first ISO requirement for class negotiation states that "the preferred class in the CR TPDU may contain any of the classes supported by the implementation". This requirement is further constrained by connectionless network operation - for ATN implementations, the preferred class *must* be class 4.

In addition, a CR TPDU may contain an alternative class parameter. Since the only acceptable mode is class 4, there are no alternative classes allowed.

4.2.2.7.2 Class Negotiation - Responder.

There is only one appropriate class for operation in the connectionless network environment - class 4. An implementation of the ATN transport layer must respond with class 4 as the negotiated class.

4.2.2.7.3 TPDU Size Negotiation.

All transport entities must be able to support a TPDU size of 128 octets, the default required by ISO 8073. Larger sizes may also be supported, such as the recommended 1024-octet capability. 1024 octets is the minimum maximum-size value recommended for ATN usage. The actual TPDU size negotiated for a TC, however, may be smaller than the maximum size supported or the initial size proposed.

The larger TPDU size is recommended for application data exchanges involving large TSDUs. The optimum TPDU size may vary anywhere from 128 octets up to the maximum TSDU size required by a TS-user. The selection of a 1024-octet TPDU size ensures that no additional network segmentation will be performed on any TPDUs transmitted as NSDUs.

4.2.2.7.4 Use of Extended Format

The default format for TPDU numbering is the "normal" format, which involves the use of a seven-bit field. Extended format uses a 31-bit field. If there is no proposal in a connection request, the normal format is used. If the initiator proposes extended format, the responder may reply indicating use of normal format.

Generally, the extended format is used when an extremely large window of outstanding TSDUs is expected. This would occur, for example, on large data transfers with very little interaction between end users (e.g. reception of acknowledgements only after an extended interval). Large windows may also occur in the situation where a link has high capacity but long transit delays.

Thus, the use of normal formats is recommended for operation in the ATN because of the smaller resulting size of transport protocol headers. Note that as defined by ISO 8073, the ability to support normal formats is mandatory.

4.2.2.7.5 Expedited Data Transport Service

Support of the expedited data transport service is required by ISO 8073. Thus, all ATN implementations must have the capability to send and receive expedited data. Actual use of the feature is optional. Negotiation of the expedited data service is performed using the additional options selection parameter (bit 1) as shown in Table 4-3.

4.2.2.7.6 Non-use of Checksum

The default operation for a connection is to use checksums. If non-use is desired, the initiator must propose non-use of checksums and the responder must agree. Checksums are a valuable tool because they verify the end-to-end integrity of TPDUs, and thus all TSDUs.

Non-use of checksums may be selected, for example, to support transmission of low-fidelity graphical data. The initiator of a transport connection being used for this purpose may propose non-use of checksums if the cost of using checksums (both in terms of cost and transmission efficiency) is considered too high. It is recommended in such cases that the responding transport layer accept the non-use of checksums so that the efficiency gains can be realised.

There may be situations, however, when the responding transport entity would not agree to non-use of checksums. For example, if the responding entity has knowledge that the available QoS between the two end systems is not sufficient to support the needs of the TS-user, it may respond indicating that checksums are to be used.

Note.—The method of acquiring knowledge of available QoS is a local matter. For some applications, dynamic knowledge may be required. Other applications may have less stringent needs and will not require any dynamic information.

All ATN transport layer implementations must be able to propose either use or non-use of checksums in a CR TPDU. If non-use is proposed, all ATN transport layer implementations must be able to accept non-use. Mechanisms for determining when not to accept the non-use of checksums are not required.

4.2.2.7.7 Use of selective acknowledgement.

The default for selective acknowledgement is non-use. That is, selective acknowledgement must be explicitly proposed in a CR TPDU and accepted in the CC TPDU.

Because the selective acknowledgement feature reduces the need for retranslating TPDUs, it is recommended that transport layer implementations propose the use of selective acknowledgement in a CR TPDU. If a transport layer receives a CR TPDU proposing this option, it is recommended that the proposal be accepted in the CC TPDU.

Note.— Refer also to 4.2.2.10.4.3 for a description of the selective acknowledgement feature.

4.2.2.7.8 Use of Request of Acknowledgement.

The default for ROA is non-use, that is, ROA must be explicitly proposed in a CR TPDU and accepted in the CC TPDU. The ROA function allows a transport layer to request, on a per-TPDU basis, that the remote transport layer immediately acknowledge all TPDUs currently awaiting acknowledgement. This is especially useful in the case that a window is closing up, or if the sending transport layer is having buffer limitations, and needs to free up additional space. Thus, it is recommended that this option be proposed in a CR TPDU, and that it be accepted, if proposed, in the CC TPDU.

4.2.2.8 Error Handling

4.2.2.8.1 Action on Receipt of a Protocol Error

There are three possible actions of a transport implementation upon detection of a protocol error:

- The transport layer can issue an ER TPDU;
- The transport layer can terminate the transport connection (that is, issue a DR TPDU); or,
- The transport layer can discard the TPDU (that is, ignore the error).

Events which qualify as a protocol error are defined in ISO 8073. It is recommended that in event of a protocol error, that the transport layer issue an ER TPDU, and either discard the TPDU, or respond with a DR TPDU. This action ensures that the cause of a protocol error can be more readily identified.

4.2.2.8.2 Actions on Receipt of an Invalid or Undefined Parameter in a CR TPDU.

The actions upon receipt of an invalid parameter are defined as mandatory by ISO, and so must be performed by all ATN implementations of the transport layer.

ISO 8073 requires that, on receipt of an undefined parameter, that the parameter be ignored. This action, in combination with the general rules for negotiation allows compatibility between versions of the transport layer. For example, if a transport layer issues a CR proposing the selective acknowledgement option to a remote transport layer built to ISO 8073 (1988), the remote transport entity will not recognise the new option. Rather than declaring a protocol error, the remote entity would simply pass over the option and would continue to process the rest of the TPDU. A transport connection could then be established which operates without using selective acknowledgement.

If a recognised parameter has an invalid value, then an implementor may either ignore the error or declare a protocol error, at their own discretion. However, note that for class 4 over CLNS operation, if the parameter in question is the checksum, the transport layer is required to discard the TPDU.

4.2.2.8.3 Actions on Receipt of an Invalid or Undefined Parameter in a TPDU other than a CR TPDU.

For all other TPDUs, the decision s to whether to treat an undefined parameter as a protocol error or to ignore it is a local matter. In the case that a protocol error is defined, the implementation may either:

a) discard the TPDU silently;

b) issue an ER TPDU and either discard the TPDU or issue a DR TPDU; or,

c) immediately issue a DR TPDU.

4.2.2.9 Timers and Protocol Parameters

Although the implementation of most of the timers and protocol parameters is mandatory, there are no mandatory values for them, other than the maximum values which may be defined for each.

It is recommended for ground systems that timers be configurable on a per TC basis.

In general, the assignment of values for timers and parameters must be optimised based on operational testing of the applications. In such testing, incompatible timer values and optimum combinations can be identified. Implementations of the transport protocol should support configurable values for all timers and protocol parameters on a TC or TSAP basis, rather than having fixed values. This allows modification as operational experience is gained.

Note 1.— Refer to Table 4-4 *for the complete listing of timers and parameters.*

Note 2.— Refer also to 12.2.1.1 of ISO 8073 for more details on the timers.

Note 3.— In Table 4-4, the subscripts "R" and "L" refer to "remote" and "local", respectively. The variable E_{RL}, for example, refers to the maximum transit delay from the remote entity to the local entity. The variable E_{LR} is the maximum transit delay from the local entity to the remote entity. It is assumed that these values may be different.

Note 4- The example values in Table 4-4 have not been subject to validation and are for illustrative purposes only.

Several of the timers and variables listed in Table 4-4 are not directly configurable, but may be

Symbol	Name	Minimum	Example	Maximum
M_{RL}, M_{LR}	NSDU Lifetime, seconds	5	40	135
E_{RL} + E_{LR}	Maximum Round-trip Transit Delay, seconds	0	35	150
A_L, A_R	Acknowledgement Time, seconds	0	2	20
T1	Local Retransmission Time, seconds	12	37	300
R	Persistence Time, seconds	0	75	2710
N	Maximum Number of Transmissions	1	3	10
L	Time bound on reference and/or sequence numbers, seconds	160	160	3000
I	Inactivity Time, seconds	300	960	3000
W	Window Time, seconds	160	160	400

Table 4-4 Example Timer and Parameter Values and Ranges

determined based on the values of other timers and variables. That is:

- The NSDU lifetime variables, M_{RL} and M_{LR}, may have a general estimate, based on the lifetime values used for NPDUs. The NSDU lifetime value is the value used to delete aged packets from the ATN. It should be over three times the expected end-to-end time. The expected air-to-ground end-to-end time can be up to 30-40 seconds.

- The end-to-end delay variables, E_{RL} and E_{LR}, may be estimated only, or some mechanism may be available to determine these dynamically.

- The value for the local acknowledgement timer, A_L, may be determined based on application requirements. For example, applications supporting ATC may require immediate acknowledgement of TPDUs so that the uncertainty about delivery is minimised. The remote acknowledgement time variable A_R, for example, may not be known or it may be provided by the remote transport entity explicitly during the connection establishment phase. The value for A_L should be dynamically configurable.

- The local retransmission time, T1, is defined by ISO as:

  $$T1 = E_{LR} + E_{RL} + A_R + x,$$

  where x is the local processing time for a TPDU.

- The persistence time, R, is the maximum time a transport entity will attempt to retransmit a TPDU. The persistence time is larger, in general, than the maximum number of retransmissions, N-1, times the local retransmission time, T1.

- The maximum number of transmissions, N, is related to the expected transmission reliability of the end-to-end path, since exceeding N results in the termination of a transport connection. Too high a value, however, may result in wasted retransmissions if end-to-end communication is no longer possible.

- The maximum time to receive an acknowledgement of a given TPDU, L, is bounded by ISO as:

  $$L = M_{LR} + M_{RL} + R + A_R$$

- In general, a reference or sequence number should not be re-used for the time period L. The value of L, in combination with the expected traffic, may be used to determine if extended TPDU numbering is required.

- The inactivity timer, I, is set based on network delays and the expected QoS. Specification of this parameter is related to the use of the maximum number of transmissions parameter, N, since it is used to terminate transport connections.

- The window timer, W, determines when acknowledgements are sent in the case of no activity. Up-to-date window information is sent when W expires. It should be set smaller than the expected value of the remote value of I.

4.2.2.10 Transport Layer Protocol Conformance

This section provides background information and notes on the APRLs for the connection mode transport protocol and the encoding of TPDUs. The requirements for the connection mode transport protocol are defined using the APRL for the ISO 8073 protocol specified in the ATN Internet SARPs, which is derived from the PICS Proforma provided with ISO 8073. ATN specific extensions are also included in the APRL.

4.2.2.10.1 Base Standard

The base standard which applies to the ATN Transport Layer protocol is the 1992 version of ISO 8073.

During the development of the APRL, an important objective was to ensure backwards compatibility with ISO 8073: 1988, whilst permitting the use of the following features of the 1992 version which do not exist in the 1988 version:

1. A new parameter, "preferred maximum TPDU size", which was added to accommodate a larger set of sizes than was possible with the present parameter, "maximum TPDU size".

2. The Selective Acknowledgement option, which was added to allow a transport entity to acknowledge a non-contiguous set of TPDUs.

3. The Request Acknowledgement option, which was added to allow a transport entity to request that the remote entity acknowledge received TPDUs.

4. The inactivity time is now specified as two values, a "local" inactivity time and a "remote" inactivity time.

5. The values of the inactivity times can now be passed as parameters in the connection establishment phase.

4.2.2.10.2 Caveat to Conformance with Base Standard

The ISO 8073 PICS (D.6.2) identifies C4L as ISO:C2:0 reflecting that Class 4 over connectionless networks requires the implementation of class 2 for conformance purposes. However, ISO 8073 6.5.5.i indicates that Class 4 is the only valid class over the CLNS. There is no purpose for requiring Class 2 in the ATN environment as a connection mode network service is not provided. In respect of this item, ATN conformant implementations of ISO 8073 are therefore not necessarily in conformance with ISO 8073.

4.2.2.10.3 Initiator/Responder Capability for Protocol Classes 0-4

Predicates "IR1" and "IR2" are defined as an option set in the ISO PICS, which means that a conforming implementation of the transport protocol must be able to initiate a connection or respond to a connection request. The ATN Transport profile recommends that both capabilities be present. This capability will support the long-term utility of transport layer implementations in the ATN.

4.2.2.10.4 Notes on Required and Recommended Optional Functions

4.2.2.10.4.1 Extended TPDU Numbering

Support of extended TPDU numbering is recommended to allow support of ATS applications with high data rates or those operating over links with long delays. Normally, the transport protocol uses 7 bits for the TPDU number, resulting in a range of [0 - 127]. Extended TPDU numbering uses 31 bits for the TPDU number and expands this range to [0 - 2 147 483 647]. The extended numbering option is useful when there are a large number of TPDUs that may be unacknowledged at a time. This may occur, for example, when a large amount of data is transferred over a link which has long delays, or for the case when information transfer is primarily unidirectional. The other reason extended numbering is used is to support a high rate of TPDU transfer. TPDU numbers may not be re-used during the maximum period to receive an acknowledgement, L (see 4.2.2.9). If a large number of TPDUs (i.e. more than seven) is expected to be transmitted during the period L, and flow control is not acceptable, extended numbering is required to guarantee unique TPDU numbers. The cost of using extended TPDU numbering is an increased header on every TPDU that is transmitted

for a given connection. Thus, this option should not be exercised when the window sizes for normal TPDU numbering are sufficient.

4.2.2.10.4.2 Non-use of Checksum

Support of the non-use of checksum feature is required to allow applications that can tolerate some level of error to operate without the added cost of transmitting checksums with every TPDU. Checksums are used to verify the end-to-end integrity of data within a TPDU. By default, checksums are present in all TPDUs; non-use must be mutually agreed by both TS-users.

Note.— The transport layer provisions do not specify the conditions for an initiating transport layer entity to specify non-use of checksums. These are a local metter. The use or non-use of checksums is dependent on the characteristics of the TS-user-data flow.

4.2.2.10.4.3 Selective Acknowledgement

Support of the selective acknowledgement feature is recommended to improve the management of air-ground resources and to reduce unnecessary retransmissions of data. Selective acknowledgement allows the transport layer to acknowledge receipt of multiple TPDUs, even if there is one or more missing in a given sequence. For example, if the transport layer received TPDU numbers 4, 5, 6, 8, and 9, it can use the selective acknowledgement function to indicate receipt of all of these TPDUs, indicating that number 7 is not yet received. This provides the remote transport layer the information to retransmit only TPDU number seven, without having to retransmit 8 and 9.

4.2.2.10.4.4 Request of Acknowledgement

Support of the request of acknowledgement (ROA) function is recommended for ATN implementations. The ROA function allows a transport layer to request that the remote transport layer acknowledge all currently received TPDUs. This is especially useful in the case that either a transmit window is closing up, or the sending transport layer is having buffer limitations and needs to free up additional space.

4.2.2.10.4.5 Reduction of Credit Window

Support of the reduction of credit window feature is recommended to support congestion avoidance mechanisms in the transport layer.

4.2.2.10.4.6 Concatenation

Support of the concatenation function is recommended to improve use of air-ground resources. Concatenation of TPDUs may be performed when a number of TPDUs is to be sent to the same transport entity (for example, a DT TPDU and an AK TPDU). Multiple TPDUs may be concatenated and sent together in the same NSDU to the remote transport entity; the remote entity then separates the two TPDUs. Note, however, that concatenation of TPDUs may not be suitable with TS-users requiring minimal delays, since some TPDUs may be held until several are concatenated.

4.2.2.10.5 Notes on TPDU Support

4.2.2.10.5.1 Mandatory TPDUs.

All of the TPDUs defined by ISO for Class 4 operation over the connectionless network service are mandatory for the ATN transport layer.

4.2.2.10.5.2 Error TPDU Support.

The Error (ER) TPDU may be sent by a transport layer in response to an error condition, such as receiving a legal TPDU with illegal values. Transmission of the ER TPDU is not required by the transport protocol; the conditions which cause an entity to transmit one are left as a local matter. However, it can be very useful in providing diagnostic information, and has the added advantage that it makes clear which side of the transport connection detected the error and hence which implementation is the probable source of the error.

4.2.2.10.6 Notes on TPDU Parameter Support

4.2.2.10.6.1 Optional Parameters for the CR TPDU.

This section describes the ATN recommendations for support of the optional parameters which may be included with a CR TPDU. Note that no parameters are recommended that cannot be supported in both the 1992 and the 1988 versions of ISO 8073. The optional parameters for which ATN specific recommendation have been made are:

- **The called and calling TSAP-ID** parameters: Support is required in order to allow applications to be identified through the use of upper-layer selectors, rather than using *a priori* knowledge of the user based on the NSAP. The called TSAP-ID parameter contains the TSAP Selector portion of the called user's TSAP, and ensures unambiguous identification of the destination TS-user. The calling TSAP-ID allows the destination user to identify the calling TS-user, and initiate a call to the other user in the case that the transport connection is terminated.

- **TPDU size parameter**: The ability to use the TPDU size parameter is recommended. There are two different parameters which may be used to propose a TPDU size, the TPDU Size parameter (index I4CR9) and the Preferred Maximum TPDU Size parameter (index I4CR18). Either parameter may be used to negotiate a maximum TPDU size. The latter was added to the latest version of ISO 8073 to allow a larger range of TPDU sizes. Invocation of the Preferred Maximum TPDU Size parameter should only be done if the peer transport entity is known to implement the parameter. Otherwise, if the preferred maximum TPDU size parameter is not recognised, the maximum TPDU size will be the default value, 128 octets. Furthermore, indices TS1 and TS2 require that if a size for TPDUs is proposed, that the initiator must be capable of supporting all legal TPDU sizes smaller than the proposed size. For example, if the Preferred Maximum TPDU Size parameter was included in a CR to propose a TPDU size of 1,280 octets (128 octets times ten), the initiator must be prepared to use a negotiated TPDU size of (n*128) octets, where ($1 \leq n \leq 10$). If the Maximum TPDU size parameter is used, the negotiated size may be in the set [128, 256, 512, 1024, 2048, 4096, or 8192], as long as it is equal to or smaller than the proposed size. For operation over the ATN, it is recommended that a size of 1024 octets be negotiated for the maximum TPDU size. This value is derived from the requirements for the minimum SNSDU size. It eliminates the need for segmenting by the CLNP.

- **Preferred Maximum TPDU Size:** Support is recommended. The maximum preferred TPDU size that an initiator proposes may be any multiple of 128 octets. For operation over the ATN, it is recommended that a size of 1024 octets be negotiated for the preferred maximum TPDU size. This value is derived from the requirements for the minimum subnetwork service data unit (SNSDU) size

- **Version Number Parameter**: Support is not recommended. No specific use is seen for this parameter, and implementations should not expect that other ATN transport entities will use this optional parameter.

- **Protection Parameter:** Support is not currently recommended as no security mechanisms have been defined for the ATN besides use of the ATN Security Label, which is outside of the scope of this parameter. Use of this feature may be specified in later versions of the CNS/ATM SARPs, if a need for lower layer protection mechanisms had been identified.

- **Additional Option Selection Parameter:** The additional option selection parameter must be supported in a transport layer implementation, in order to allow negotiation of several transport layer optional functions.

- **Residual Error Rate And Transit Delay Parameters**: Support is not recommended for transport layer implementations, as these are design parameters of connectionless networks and cannot readily be selected dynamically.

- **Priority Parameter:** Support is recommended. In addition, the priority parameter should be present in a CR TPDU. Priority is an especially important feature in the ATN air-to-ground environment, as it is used to ensure that high priority (i.e. flight safety related data) is never impeded by lower priority, routine communications. Priority is non-negotiable in the ATN. TS-users should issue a DR TPDU if a different priority level is returned in the CC TPDU. There is a further recommendation in the ATN SARPs that the responding transport layer should respond with the same priority as was proposed. For transport implementations unable to specify priority, a default priority may be used. This default priority is the lowest transport priority (level 14), and is mapped to the lowest network priority level. Priority is used to separate classes of application traffic, and to ensure that in conditions of limited resources certain classes of traffic receive service in preference to others. Thus implementations unable to state priority will have their traffic discarded first in an ATN global congestion avoidance scheme. These priority mappings are also enforced by certain ATN Subnetwork Service Providers.

- **Acknowledgement Timer and Inactivity Time Parameters:** Support is recommended for both. These two parameters allow transport entities to better manage transport resources, and may be implicitly required in order to support applications (e.g. ADS) that demand well defined bounds on either data delivery, or an indication of transport connection less.

4.2.2.10.6.2 Optional Parameters for the CC TPDU.

Requirements and recommendations on the support of parameters for the CC TPDU follow those for the CR TPDU parameters. It is recommended that if both the preferred maximum TPDU size parameter and the Maximum TPDU size parameters are present in a CR TPDU, then the CC TPDU should respond using the Preferred Maximum TPDU size parameter only.

4.2.2.10.6.3 Optional Parameters for a Disconnect Request TPDU.

The Additional Information parameter (index I4DR4) in a DR TPDU is not recommended for ATN implementations of the transport layer.

4.2.2.10.6.4 Mandatory Parameter for a Data TPDU.

If the Request of Acknowledgement feature has been selected during the connection establishment phase, then the Request of Acknowledgement (ROA) parameter (index I4DT4) is mandatory in the DT TPDU.

4.2.2.10.6.5 Optional Parameters for an Acknowledgement TPDU.

The flow control confirmation parameter (index I4AK4) is recommended for ATN implementations of the transport layer.

4.2.2.10.6.6 Use of the Subsequence Number Parameter in the Acknowledgement TPDU

If the reduction of credit window capability is implemented, support of this parameter is required. Even if it is not implemented, support of the flow control confirmation parameter is recommended for use in congestion avoidance mechanisms.

4.2.2.10.6.7 Use of the Selective Acknowledgement Parameter in the AK TPDU

Support of this parameter is recommended for transport layer implementations. If selective acknowledgement has been selected for a given TC, then this parameter is optional in an AK TPDU.

4.2.2.10.6.8 Optional Parameters for an Error TPDU

The Invalid TPDU parameter (index I4ER3) in an ER TPDU is not recommended for ATN implementations of the transport layer.

4.2.2.10.6.9 User Data in Class 4 TPDUs

A TS-user may optionally include data in the CR, the CC, or the DR TPDUs. The ability to include data in the CR, CC, and DR TPDU is required for ATN implementations.

As defined by ISO, all transport layer implementations capable of initiating a CR must be able to receive user-data in the two possible responses: a CC TPDU or a DR TPDU. These data are passed on to the TS-user. Similarly, all transport layers capable of responding to a CR must be able to receive user-data within a CR TPDU.

4.2.2.11 Use of the Network Service

The transport layer uses the connectionless network service to exchange TPDUs with remote transport entities. This involves two network service primitives: the N-UNITDATA request, to send TPDUs, and the N-UNITDATA indication, to receive TPDUs.

4.2.2.11.1 Use of the N-UNITDATA Request

All TPDUs are transmitted using the N-UNITDATA request primitive. In general, the transmission of a single TPDU corresponds to an invocation of the N-UNITDATA request. If the transport layer performs TPDU concatenation, the combined set of TPDUs is sent via a single request.

The N-UNITDATA parameters are used as follows:

4.2.2.11.1.1 NS-user-data.

The transport layer sends a TPDU (or a concatenated set of TPDUs) as a single NSDU.

4.2.2.11.1.2 Network Service Access Point Addresses

Transport addresses are passed between the TS-user and the transport protocol entity. With the connection mode transport layer, transport addresses are passed during the connection establishment phase. The TS-user issuing a CR must provide the destination transport address and the source transport address. These addresses are interpreted by the transport layer when the user's connection request is translated into a CR TPDU and transmitted. The TSAP selectors of the source and destination transport addresses are transmitted within the CR TPDU. The NSAP addresses of the source and destination are used in the transport layer's invocation of the N-UNITDATA request that is used to transmit the CR TPDU.

4.2.2.11.1.3 Network Quality of Service.

4.2.2.11.1.3.1 Network Layer Protection

.The possible actions that can occur when the user specifies a protection parameter are:

a. the transport layer can use protection techniques peer-to-peer;

b. the transport layer can use network protection techniques by setting the network layer protection parameter;

c. the transport layer can use a combination of the above actions; or,

d. the transport layer can pass protection parameters but not interpret them.

The ATN effectively implements option (b) by passing the ATN Security Label to the network layer, as the protection parameter.

The value of the ATN Security Label specified by the connection initiator on the connect request, is used as the value of the NS protection parameter for the N-UNITDATA that contains the CR TPDU. The same value is then used for all subsequent N-UNITDATA requests used to convey TPDUs sent by both the connection initiator and the connection responder on that transport connection.

4.2.2.11.1.3.2 Network Layer Transit Delay, Cost, and Residual Error Probability

The ATN network layer QoS parameters include the relative ranking of cost, transit delay, and error. The TS-user interface supports the specification of transit delay and residual error rate. The cost parameter, however, is not one of the QoS parameters that are supported by the TS-user interface. The selection of the requested Network Layer QoS parameters can be done by configuration or dynamically.

However, general support of the network layer QoS parameters is not expected in the near to medium term. They may be specified by the sending transport layer, but are ignored by the network layer.

4.2.2.11.1.4 Network Layer Priority

When specified, the transport priority parameter has a one-to-one correspondence with network priority. Note that for the transport layer, priority level 0 is highest, while for the network layer, priority level 14 is highest. The relationship between transport priority and network priority is specified in chapter 2 of the ATN Internet SARPs.

The selection of the network priority may be done either on a dynamic basis or on a static configuration basis, depending on the application categories on the ES. If the transport layer supports levels of priority higher than 14, these should be assigned a network priority level of zero.

4.2.2.11.2 Use of the N-UNITDATA Indication

The transport layer receives all TPDUs via the N-UNITDATA indication. TPDUs are contained within the NS-user-data parameter of the N-UNITDATA indication. Note that if the remote transport layer is performing concatenation, there may be multiple TPDUs within a single NSDU.

The parameters of an incoming N-UNITDATA indication are interpreted as follows.

4.2.2.11.2.1 NS-user-data.

The transport layer assumes that the first TPDU begins at the first octet of the NS-user-data. If the length of the TPDU is less than the length of the NSDU, the transport layer assumes that there are one or more TPDUs following the first one.

4.2.2.11.2.2 Network Service Access Point Addresses.

The source and destination NSAP addresses are used to determine the source and destination transport addresses associated with a TPDU. In general, this is only required during the connection

establishment phase, before a TC identifier has been assigned. The transport addresses are determined by combining the NSAP addresses with the appropriate TSAP selectors. The selectors are contained in a CR or CC TPDU.

4.2.2.11.2.3 Network Quality of Service.

The connection mode transport layer does not need to interpret most of the indicated network layer QoS parameters associated with an N-UNITDATA indication, expect for the protection parameter conveying the ATN Security Label. The network layer priority is not interpreted, because, when its use has been specified by the TS-User, the transport priority is set explicitly. The network layer protection parameter is not used. The relative rankings of transit delay, residual error, and cost are after the fact, and do not have any effect on the transport protocol.

Another parameter passed in the indicated network layer QoS is the flag for the presence of congestion. For each NSDU delivered to the transport layer, the network layer must indicate whether the CE flag was set for any NPDUs associated with the NSDU. The presence of this condition indicates that there is congestion between the source and destination. This information is used in the ATN to implement Congestion Avoidance and is discussed in more detail in chapter six.

The value of the protection parameter received in an N-UNITDATA indication is interpreted as the ATN Security Label, and saved by the TS-provider and used with all subsequent N-UNITDATA requests on that transport connection.

4.2.3 The Connectionless Mode Transport Layer

The ATN CLTS is based on the ISO 8072/AD1 Standard Service Definition, and the ATN CLTS offers the necessary means for transferring TSDUs of limited size without prior transport connection establishment. The ATN CLTS offers transmission with no protection against losses, duplication or misordering of a TSDU. It is well suited to ATN applications requiring a one-time, one-way transfer of data, thus taking advantage of simpler mechanisms than those employed by the connection mode protocol.

4.2.3.1 Overview of the Connectionless Mode Transport Layer

The defining characteristic of CLTS transmission is the independent nature of each invocation of the Service. Each TSDU is independent in the sense that it bears no relationship to any other TSDU transmitted through the invocation of the connectionless mode service. It is also self-contained in that all of the information required to deliver the TSDU (destination address, quality of service selection options, etc.) is presented to the TS-provider, together with the user-data to be transmitted, in a single service access. Each unit of data transmitted is routed independently by the layer providing the connectionless mode service.

Certain elements of QoS associated with each instance of connectionless mode transmission, are requested from the TS-provider by the sending TS-user. The TS-provider does not guarantee any of the characteristics the user may set.

The connectionless mode transmission is the transmission of a single data unit from a source service access point to one or more destination access points without establishing a connection. By avoiding the overhead of transport connection establishment and connection management, it is possible to speed up the data exchanges and reduce transit delays of short TSDUs. The functions in the Transport Layer are those necessary to interface between the service available from the Network Layer and the service to be offered to the TS-users. The functions provided by the Transport Layer in connectionless mode are:

1 network service selection;

2 mapping of Transport address onto Network address;

3 TSDU delimiting (determine the beginning and end of a TSDU); and,

4 end-to-end error detection (implying the use of a specific mechanism) and the necessary monitoring of the QoS.

These functions will operate according to the type of subnetwork and the related network services. Only a pre-arranged association between the entities which determine the characteristics of the data to be transferred is required. No dynamic agreement is involved in an instance of the use of service.

4.2.3.1.1 Service Characteristics

The CLTP operates using the ATN connectionless mode network service. The procedure of data transfer is used for one-time, one-way transfer of a TSDU between TS-users. The protocol does not provide confirmation of receipt, TC establishment and release, or network connection establishment and release.

4.2.3.1.2 Data Transfer

The data transfer procedure is used for one-shot, one-way transfer of a TSDU between TS-users without confirmation of receipt, without transport connection establishment and release, and without network connection establishment and release.

The QoS parameter in the T-UNITDATA request is used to determine if a checksum mechanism should be used (including a checksum parameter). If a checksum is used, it is generated at the transmitter and verified at the receiver. TPDUs failing verification are discarded.

Receipt verification is unavailable, so any recovery is by a higher layer. Note that no segmenting of a TSDU into smaller TPDUs is permitted and large TSDUs (over 63,488 octets) are discarded.

As the ATN transport layer operates over a CLNS, only the following network service primitives are used : N-UNITDATA request and indication. There is no indication given to transport entities of the ability of the network entity (NE) to fulfil the service requirements given in the N-UNITDATA primitive. However, it can be a local matter to make TEs aware of the availability and characteristics (QoS) of the CLNS (e.g. through the use of the N-FACILITY management primitives set).

4.2.3.1.3 ATN Connectionless Mode Transport Service Model

The CLTS can be modelled in the abstract as a permanent association between the two TSAPs. Only one type of object, the unitdata object, can be passed to the TS-provider. The TS-provider may perform any or all of the following actions:

- discard objects,
- duplicate objects,
- change any order of independent service requests into a different order of service indications.

The existence of the association does not depend on the behaviour of the TS-users. The set of actions which are performed by the TS-provider on a particular association may depend on the TS-users' behaviour. However, these actions are taken by the TS-provider without notification to the TS-user. Awareness of the characteristics of an association is part of the TS-users' *a priori* knowledge of the ATN environment.

4.2.3.2 ATN Connectionless Mode Transport Layer Quality of Service

4.2.3.2.1 Use of Transport Layer QoS

The use of transport layer QoS parameters for the CLTS is similar to that of the connection-mode service. However, unlike the COTS, there is no concept of negotiation of requested transport layer QoS parameters. Each invocation of the T-UNITDATA service involves a set of requested transport layer QoS parameters by the source TS-user; the corresponding T-UNITDATA indication to the destination TS-user contains the indicated transport layer QoS parameters.

The TS-user can specify the requested transport layer QoS parameters, but there is no guarantee that the TSDU will have the requested level of service. Upon delivery of a TSDU, the transport layer provides the indicated transport layer QoS parameters. The indicated parameters are only an estimate of what may have been provided for that TSDU. The transport layer can determine the indicated transport layer QoS parameters by either *a priori* information or through a systems management interface which provides information on the expected QoS between two ESs.

4.2.3.2.2 Connectionless Mode Transport Layer QoS Parameters

Four QoS parameters are identified for the connectionless mode transport service: transit delay, residual error probability, priority and protection.

As with the connection mode, transit delay is not used, and, if specified, will be ignored. Two levels of residual error rate are provided, equivalent to use and non-use of the transport checksum. Both a priority and an ATN Security Label may be specified on a per TSDU basis.

4.2.3.2.3 Priority

This parameter enables the TS-user to specify the relative priority of a TSDU in relation to every other TSDU handled. A TSDU of higher priority is processed before a TSDU of lower priority by the TS-provider. This parameter specifies the order in which TSDUs should have their associated QoS downgraded, and the order in which they should be discarded in order to retrieve resources.

When specified, priority values should be integers in the range from 0 to 14, with priority level 0 as the highest priority. In such a case, there is a one-to-one correspondence with the CLNP priority values. Chapter 2 of the ATN Internet SARPs specifies the mapping of transport layer priority values to network layer priority values.

4.2.3.2.4 ATN Security Label

The security parameter is used by the service user to indicate the value of the security label to be associated with the TSDU. The syntax and semantics of the ATN Security Label are specified in the ATN Internet SARPs.

4.2.3.3 Connectionless Mode Transport Layer Service Primitives

Two TS primitives are used to provide the CLTS: the T-UNITDATA request primitive and the T-UNITDATA indication primitive. The sequence of primitives in a successful CLTS transmission is defined in Figure 4-7.

4.2.3.3.1 T-UNITDATA Request

An ATN TS user requests the transfer of a TSDU by invoking a T-UNITDATA request primitive. This primitive has the following associated parameters:

- **Source and Destination Address:** These are TSAP addresses and they are unique within the scope of TSAP addresses. The ATN transport addressing scheme is the same for COTS and CLTS providers i.e. each transport address is composed of an NSAP address and a TSAP Selector.

- **Quality of service:** The value of the QoS is a list of subparameters. The subparameters composing the CLTS QoS are presented in 4.2.3.2.1. The TS-provider does not guarantee that it can offer the requested QoS.

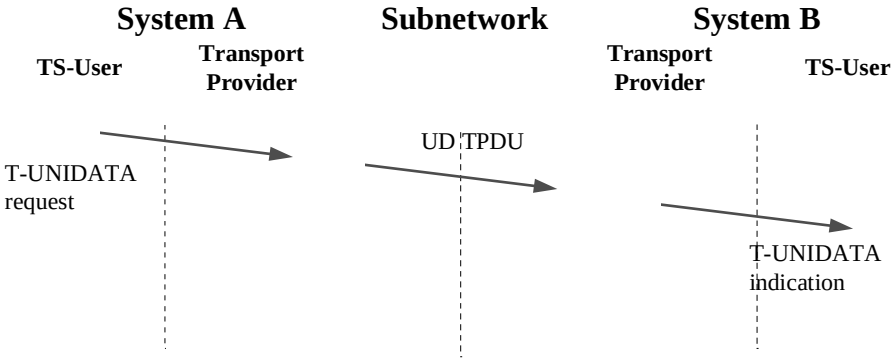

Figure 4-7 Sequence of Primitives and TPDU Exchange for Connectionless Data Transfer

- **TS-user-data:** These are the user-data (i.e. the TSDU) to be transmitted between TS-users. The ATN TS-user can transmit an integral number of octets greater than zero up to a limit of 63,488 octets (this amount is 1 K less than the maximum allowed ATN NSDU size). Using a TSDU size of more than 1024 octets may lead to CLNP segmentation and so, to more overhead on the mobile subnetworks.

- **Security:** The security parameter is used by the service user to indicate the value of the security label to be associated with the TSDU. The syntax and semantics of the ATN Security Label are specified in the ATN Internet SARPs.

With the connectionless mode transport layer, transport addresses are passed with each invocation of the T-UNITDATA primitive. The TS-user sending data must provide the destination transport address and the source transport address. The TSAP selectors of the source and destination transport addresses are transmitted within the header of the UD TPDU; the NSAPs of the source and destination are used in the transport layer's invocation of the N-UNITDATA request that is used to transmit the UD TPDU.

4.2.3.3.2 T-UNITDATA Indication

Upon arrival at the destination TSAP, a T-UNITDATA indication is delivered by the TS-provider to the destination TS-user. This primitive has exactly the same associated parameters as the T-UNITDATA request primitive. Their values are unchanged by the TS-provider, except for the QoS parameter which may have a different value from the value specified in the request primitive.

The QoS parameter value associated with the T-UNITDATA indication primitive, is based on the NS QoS indication and on the use of the checksum mechanism; it may be different from the value requested, if the TS- or NS-provider has the means to verify that the requested QoS has not been reached. Note that the TS-user-data parameter value is expected to be equal to the TSDU transmitted only if a checksum mechanism has been used for this TSDU.

4.2.3.4 Use of the Network Service

Note.— Refer to 4.2.2.11.1 for more background on selection of requested network layer QoS parameters.

4.2.3.4.1 Use of the N-UNITDATA Request

Each UD TPDU is transmitted using the N-UNITDATA request primitive. The transmission of a single TPDU corresponds to an invocation of the N-UNITDATA request. The N-UNITDATA parameters are used as follows:

4.2.3.4.1.1 NS-user-data.

The transport layer sends the UD TPDU as an NSDU.

4.2.3.4.1.2 Network Service Access Point Addresses.

Transport addresses are passed between the TS-user and the transport protocol entity. With the connectionless mode transport layer, transport addresses are allocated into two elements: the TSAP selector and the NSAP. The source and destination TSAPs are sent within the UD TPDU; the NSAPs of the source and destination TS-users are passed as the source and destination NSAPs within the invocation of the N-UNITDATA primitive.

4.2.3.4.1.3 Network Quality of Service.

QoS parameters are used to indicate the needed characteristics of the underlying communications service supporting application information exchange. The transport layer must interpret the QoS parameters which are provided by a TS-user; these parameters may then affect the interactions of the transport layer with the network layer providing service.

The determination of the network QoS parameters for transit delay, cost, and residual error probability can be done in a manner similar to that of the COTS. See 4.2.2.11.1.3.

The value of the ATN Security Label specified by the service user when invoking the T-UNITDATA service, is used as the value of the NS protection parameter for the N-UNITDATA that contains the UD TPDU.

4.2.3.4.1.4 Network Layer Priority.

There is no explicit priority parameter in a UD TPDU. To meet the ISO 8072 Service Specification, the CLTP entity translates the TS-user priority to network priority upon transmission of a TPDU and perform the inverse upon receipt. For example, to send a TSDU, the CLTP entity maps the TS-user Priority parameter to the network priority parameter, which is passed to the NE in the N-UNITDATA request. This passed parameter is used by the Network entity to set the Network NPDU

priority parameter. This mapping ensures that the TS-user requested priority is used for transmission of the TSDU.

Once the TSDU is received by the destination CLTS entity, the datagram transaction is complete. There are no requirements for the receiving TE to make any distinctions based on the received priority of a TPDU. The received priority value is not negotiated, so the receiving TS-user may or may not choose to modify its processing based on the indicated value of priority for a TSDU.

4.2.3.4.2 Use of the N-UNITDATA Indication

The transport layer receives all UD TPDUs via the N-UNITDATA indication. TPDUs are contained within the NS-user-data parameter of the N-UNITDATA indication. The N-UNITDATA parameters are interpreted as follows:

4.2.3.4.2.1 NS-user-data.

The transport layer assumes that the UD TPDU begins at the first octet of the NS-user-data.

4.2.3.4.2.2 Network Service Access Point Addresses.

The source and destination NSAPs are used to determine the source and destination transport addresses associated with a TPDU. With the CLTS, transport addresses are determined by combining the NSAPs with the appropriate TSAP selectors, which are contained in the header of the UD TPDU.

4.2.3.4.2.3 Network Quality of Service.

The connectionless mode transport layer does not need to interpret most of the indicated network QoS parameters associated with an N-UNITDATA indication, except for the network protection parameter. The relative rankings of transit delay, residual error, and cost are after the fact, and do not have any effect on the transport protocol.

Another parameter passed in the indicated network layer QoS is the flag for the presence of congestion. For each NSDU delivered to the transport layer, the network layer must indicate whether the CE flag was set for any NPDUs associated with the NSDU. The presence of this condition indicates that there is congestion between the source and destination ESs. Because the CLTP does not implement flow control mechanisms, there is little that can be done to treat the congestion. Some metering function could be implemented to reduce the rate of TSDUs submission by a local TS-user.

The value of the protection parameter received in an N-UNITDATA indication is provided to the TS-user with the received TSDU, as the ATN Security Label associated with the TSDU.

4.2.3.4.2.4 Priority

The CLTP must interpret the indicated network layer priority to determine the associated transport layer priority, since priority is not passed in the UD TPDU. See chapter 2 of the ATN Internet SARPs for the mapping between NL priority and TL priority.

4.3 CLNP Implementation Considerations

ISO 8473 describes a protocol for providing the connectionless mode network service. The ISO 8473 protocol is a SNICP capable of operating over many different sorts of subnetwork including X.25, ISDN and LANs. It is an internetworking protocol and may be used to create a connectionless internetwork integrating many different underlying subnetworks.

4.3.1 The Connectionless Mode Network Service

The OSI connectionless network service is the service provided to a network service user when the ISO 8473 connectionless network protocol is used as a SNICP. The operation of the connectionless network service is illustrated in Figure 4-8. It consists of a single end to end primitive - the N-UNITDATA service.

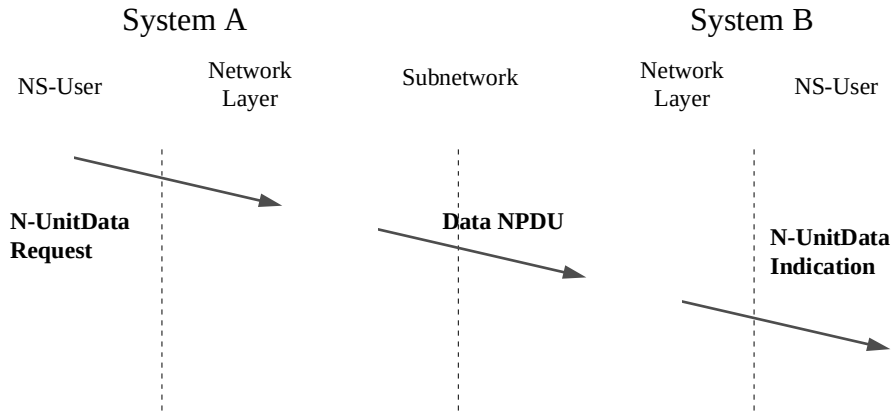

Figure 4-8 Connectionless Network Service

The service is requested by the sender who passes, as the service parameters, the user data (up to 64Kbytes), the network address of the destination, the sender's own source address, and an indication of the quality of service required, and, in the ATN, this includes the Network Priority of the data and the associated ATN Security Label.

The unitdata item is then passed through the network independently of any other data passed between the same source and destination and is finally delivered to the addressed destination. Delivery is not guaranteed and neither is the order of submission of successive unitdata items necessarily preserved. The network may discard a packet if the network is congested, and different packets may take different routes and hence have different transit times. It is the responsibility of a transport layer protocol to provide reliability, in the form of data and data sequence integrity management, if this is required.

4.3.2 The Connectionless Network Protocol

The protocol specified by ISO 8473 is the Connectionless Network Protocol (CLNP). The operation of the CLNP is straightforward and is as described below.

4.3.2.1 Satisfying the N-UNITDATA.Request

Once an NS-User has submitted an N-UNITDATA Request, the information passed with the request is formatted as a single packet, known as a Data Protocol Data Unit (Data PDU). As well as the user data, the PDU contains the source and destination addresses and the quality of service requests, priority and ATN Security Label.

Information controlling the maximum lifetime of the PDU in the network is also provided, in order to prevent PDUs existing forever in erroneous loops, and local management may also add information to specify all or part of the route that the PDU takes. As large Data PDUs may also need to be segmented en route to cope with subnetworks that support only a small packet size, there is thus information present to enable the unambiguous reassembly of segments when and if they arrive at the destination.

Once the PDU has been created, the ES or IS to receive the PDU is then chosen (the next hop) as well as the subnetwork over which the PDU to be sent. This is typically performed by consulting a local routing table. This may have been configured by a System Manager but, more likely, it is maintained by the ISO 9542 ES-IS Routing Information Exchange Protocol (see 4.4.1).

The NPDU is then sent to the chosen “next hop” ES or IS. Note that if the PDU is larger than the maximum packet size supported by the subnetwork then it is segmented prior to being sent.

The procedures for the transfer of an NPDU over a given subnetwork are specific to that subnetwork and are specified in a Subnetwork Dependent Convergence Function (SNDCF) appropriate to the subnetwork type. SNDCFs for the common subnetwork types are specified in ISO 8473. A special SNDCF has, however, been specified for ATN Mobile Subnetworks (see 4.4.4.1).

4.3.2.2 NPDU Forwarding

At each Intermediate System that receives the Data PDU, a similar decision to that made in the originating ES, is made as to which system is the next hop and over which subnetwork, out of those attached to the IS, the PDU will be sent. Segmentation may occur if necessary. Note that once a PDU has been segmented, its component parts are treated as if there were separate Data PDUs and may even be further fragmented.

An Intermediate System may discard a whole Data PDU or a segment. It may do this because of congestion, a security problem, because the PDU's lifetime has expired, or just because it cannot determine a suitable next hop for the PDU's destination.

The Routing Tables kept by an Intermediate System are typically much more complex than an End System’s, and are maintained by a dynamic routing information exchange protocol. These include ISO 9542 ES-IS (see 4.4.1), the ISO 10589 IS-IS Intra-Domain Routing Information Exchange Protocol (see 4.4.3.2), and the ISO 10747 Inter-Domain Routing Protocol (see 4.4.3.2.3).

4.3.2.3 At the Destination End System

When a Data PDU arrives at the End System that contains its destination, the PDU must first be re-assembled if it was previously segmented - assuming that all the constituent segments arrive within the PDU's lifetime - otherwise, the PDU will be discarded without being presented to the destination user.

Otherwise, once a whole PDU has arrived, it will be passed to the destination NS User, with the service primitive's parameters derived from the PDU contents, including the NPDU priority and Security Label. In the ATN, it is essential that these latter two parameters are made available to the NS-User, as they are required by the Transport Layer.

4.3.3 Addressing Consideration

The Source Address and Destination Address parameters used by the CLNP are OSI NSAP Addresses. These are variable length octet aligned addresses allocated from a global addressing plan that is ultimately administered by ISO, as specified in ISO 8348. The ATN Addressing Plan specified in the ATN SARPs is compliant with this addressing plan, and specifies a twenty octet NSAP Address syntax, together with the allocation procedures. As far as the CLNP is concerned, the actual syntax of the address is immaterial; the forwarding algorithm operates by comparing octet strings and through address prefix matching rules.

The encoding used by the ISO 8473 protocol to convey NSAP Addresses is the preferred binary encoding specified in ISO 8348.

4.3.3.1 Network Entity Titles

NSAP Addresses are used to identify NS-Users by way of the NSAP through which they access the Network Service. However, it is also sometimes necessary to address an NPDU to the Network Entity itself. This is necessary both for network management purposes and for certain routing techniques. Network Entities are identified and addressed by their Network Entity Title (NET).

A *NET* identifies a Network Entity in an end-system or intermediate-system. A NET has exactly the same format as an NSAP address, and is indistinguishable from an NSAP Address. NPDUs addressed to a Network Entity have its NET as their destination address.

NETs are also used widely by CLNP. For example, the entries in the *Source Routing* and *Recording of Route* parameters are NETs. The *Source Address* parameter in the Error Report (ER) NPDU is also a NET.

4.3.4 Other NS User Services

Although the service provided to the NS User is strictly speaking a unitdata service only, other information is typically available and useful for NS Users in making efficient use of the Network. Specifically, information on service characteristics may be accessed and indications on PDUs discarded while in transit.

The service characteristics information that may be made available includes:

- Quality of Service information i.e. an indication of the likely transit delay, protection from unauthorised access, cost and the residual error probability.
- Probability of sequence preservation
- Maximum PDU lifetime.

However, in the ATN, it is expected that such information will be known *a priori* by the Transport Layer and need not be available on a dynamic basis. Indeed, there is standard mechanism available to support the dynamic distribution of such Quality of Service Information.

4.3.5 Error Reports

Error reports may also be provided if PDUs are discarded while in transit. These are supported by a second PDU format - the Error PDU.

An Error PDU may be generated to report every Data PDU that is discarded. However, neither its generation nor its receipt are guaranteed.

In the ATN, Error PDUs received by an End System need to be made available to the NS-User as additional reports. This may be as an extension to the service interface or through a local management mechanism.

4.3.6 Quality of Service Maintenance

CLNP permits an NS-User to make specific QoS Requests in the form of relative preferences as to which QoS metrics to route a packet on. The use of such requests has been considered at length during the development of the ATN SARPs, and, due to practical difficulties in maintaining the necessary routing information, there are no near to medium term plans to make use of these facilities in the ATN.

4.3.7 Priority

Priority is an essential feature in the ATN Internet for ensuring that the performance targets for safety related data are met, whilst permitting the network also to be used by routine communications. Safety related data is always sent with a higher priority than routine data and is given preferential access to resources.

In the ATN Internet Layer itself, an NPDU of a higher priority is given preferred access to network resources. During periods of higher network utilisation, higher priority NPDUs may therefore be expected to be more likely to reach their destination (i.e. are less likely to be discarded by a congested router) and to have a lower transit delay (i.e. be more likely to be selected for transmission from an outgoing queue) than are lower priority packets.

ATN Internet Entities maintain their queues of outgoing NPDUs in strict priority order, such that a higher priority NPDU in an outgoing queue will always be selected for transmission in preference to a lower priority NPDU. Higher priority PDUs may thus overtake a lower priority PDU, and this effect will be especially noticeable during periods of network congestion; the network may appear congested to low priority data, whilst still appearing uncongested to higher priority data.

Furthermore, during periods of congestion, or when any other need arises to discard NPDUs currently held by an ATN Internet Entity, lower priority NPDUs are always discarded before higher priority NPDUs.

4.3.8 ATN Security

In the first phase of ATN deployment, security mechanisms are largely the responsibility of each application. However, in order to meet several Routing Control requirements, security related mechanisms are implemented in the ATN Internet. These take the form of routing decisions that are made with respect to a Security Label encoded in each NPDU header, and according to a set of routing policy rules specified in the ATN Internet SARPs.

The ATN Security Label conveys the following information:

I. A **Traffic Type**: this identifies the type or class of data and is used both to place other information in the security label in context, and as an input to access control rules. In the latter case, certain air-ground networks may limit their users to certain traffic types only. The Routing Control mechanisms will not route data of an unacceptable traffic type over such networks and will attempt to route such data around these subnetworks, if possible.

   A. ATN Operational Communications -ATSC
   B. ATN Operational Communications - AOC
   C. ATN Administrative Communications
   D. General Communications
   E. ATN Systems Management

II. An **ATSC Class**: this is valid for ATSC Traffic Types and identifies the class of subnetwork over which the data should be forwarded. Data may not be forwarded over a subnetwork with a higher ATSC Class than that indicated by its Security Label, and the ATN Internet aims to send data with a declared ATSC Class over a subnetwork supporting that ATSC Class. If no such subnetwork is available, then the next lowest available class subnetwork is chosen.

III. An **Air/Ground Subnetwork Preference**: this is valid for AINSC Traffic Types are identifies the Air/Ground subnetworks over the data may be forwarded and the relative preference of such subnetworks.

The ATN Internet Routing Control mechanisms are supported only by the Inter-Domain Routing Protocol. When routes are advertised between ATN Routers, they include a Security Information field that provides information on:

a) The Traffic Types permitted to use the route;

b) The Air/Ground Subnetwork(s) over which the route passes, if any; and

c) The ATSC Class of the route.

Using this information, ATN Routers are able to forward each NPDU in line with its Security Label and the routing control rules.

Within a Routing Domain, routers are expected to ignore the ATN Security Label. With many commercially available routers, it will be found that ignoring the security label is a configuration option.

4.3.9 ISO 8473 Mandatory Internetwork Protocol Functions

This section describes the functions which are performed as part of the ATN Internetwork Protocol within all Network entities conforming to ISO 8473. These are listed in Table 4-5, which also classifies these functions according their conformance requirement and by protocol subset. ATN Systems have to be able to support both the full and the non-segmenting subset.

The conformance requirements of each function are identified as a numeric type, as follows:

Type 1: These functions are supported in all implementations of the protocol.

Type 2: These functions are not required to be supported. If an implementation does not support a **Type 2** function and the function is selected within an NPDU, then the NPDU is discarded. If the ER flag is set within the NPDU header, then an error report is generated.

Type 3: These functions are not required to be supported. If an implementation does not support a **Type 3** function and the function is selected within an NPDU, then the NPDU is processed exactly as though the function had not been selected.

4.3.9.1 PDU Composition Function

The PDU COMPOSITION function is responsible for the construction of a Network protocol data unit according to the rules governing the encoding of NPDUs. PCI required for delivering the data unit to its destination is determined from current state and local information and from the parameters associated with the **N-UNITDATA** Request.

Network Protocol Address Information (NPAI) for the Source Address and Destination Address fields of the NPDU header is derived from the **NS-Source-Address** and **NS-Destination-Address** parameters. The **NS-Destination-Address** and **NS-Quality-of-Service** parameters, together with current state and local information, are used to determine which optional functions are to be selected. ATN **NS-Userdata** comprises the Data field of the protocol data unit.

During the composition of the protocol data unit, a Data Unit Identifier is assigned to distinguish this request to transmit **NS-Userdata** to a particular destination ATN NS user from other such requests. The originator of the NPDU chooses the Data Unit Identifier so that it remains unique (for this Source and Destination address pair) for the maximum lifetime of the Initial NPDU in the Network; this rule applies for any NPDUs derived from the Initial NPDU as a result of the

Protocol Function Name	Classification of Protocol Function		
	Full Protocol	Non-Segmenting Subset	Inactive Subset
PDU Composition	1	1	1
PDU Decomposition	1	1	1
Header Format Analysis	1	1	1
PDU Lifetime Control	1	1	N/A
Route PDU	1	1	N/A
Forward PDU	1	1	N/A
Segment PDU	1	N/A	N/A
Reassemble PDU	1	N/A	N/A
Discard PDU	1	1	N/A
Error Reporting	1	1	N/A
Header Error Correction	1	1	N/A
Security	2	2	N/A
Complete Source Routing	2	2	N/A
Complete Route Recording	2	2	N/A
Echo Request	2	2	N/A
Echo Response	2	2	N/A
Partial Source Routing	3	3	N/A
Partial Route Recording	3	3	N/A
Priority	3	3	N/A
QOS Maintenance	3	3	N/A
Congestion Notification	3	3	N/A
Padding	3	3	N/A

Table 4-5 ISO 8473 Protocol Functions

application of the SEGMENTATION function. Derived NPDUs correspond to the same Initial NPDU, and hence the same **N-UNITDATA** Request, if they have the same Source Address, Destination Address, and Data Unit Identifier. The total length of the NPDU in octets is determined by the originator and placed in the Total Length field of the NPDU header. This field is not changed in any Derived NPDU for the lifetime of the protocol data unit.

When the non-segmenting protocol subset is employed, neither the Total Length field nor the Data Unit Identifier field is present. The rules governing the NPDU composition function are modified in this case, and are as follows:

1. The total length of the NPDU in octets is determined by the originator and placed in the Segment Length field of the NPDU header.

2. The segmentation field is not changed for the lifetime of the NPDU.

3. No Data Unit Identification is provided.

The Data Unit Identifier is also used for functions such as error reporting.

4.3.9.2 PDU Decomposition Function

The PDU DECOMPOSITION function is responsible for removing the PCI from the NPDU, in preparation for processing of that information.

Information pertinent to the generation of the **N-UNITDATA** Indication is determined as follows:

1. The **NS-Source-Address** and **NS-Destination-Address** parameters of the **N-UNITDATA** Indication are recovered from the NPAI in the Source and Destination Address fields of the NPDU header.

2. The data field of a received NPDU is retained until all segments of the original service data unit have been received; collectively, these form the **NS-Userdata** parameter of the **N-UNITDATA** Indication.

3. Information relating to the QOS provided during the transmission of the NPDU is determined from the QOS and other information contained in the Options Part of the NPDU header. This information constitutes the **NS-Quality-of-Service** parameter of the **N-UNITDATA** Indication.

4.3.9.3 Header Format Analysis Function

The HEADER FORMAT ANALYSIS function determines whether the full protocol described in this chapter is employed, or one of the defined subsets thereof. If the Network protocol data unit has a Network Layer Protocol Identifier indicating that this is a standard version of the ATN CLNP, this function determines whether a received NPDU has reached its destination, using the Destination Address provided in the NPDU. If the Destination Address provided in the NPDU identifies an NSAP served by this Network entity, then the NPDU has reached its destination; if not, it must be forwarded.

If the Network protocol data unit has a Network Layer Protocol Identifier indicating that the Inactive Network Layer Protocol subset is in use, then no further analysis of the NPDU header is required and the NPDU is discarded.

4.3.9.4 PDU Lifetime Control Function

The PDU LIFETIME CONTROL function is used to enforce the maximum NPDU lifetime. This function is closely associated with the HEADER FORMAT ANALYSIS function. This function determines whether an NPDU received may be forwarded or whether its assigned lifetime has expired, in which case it is discarded.

The operation of the PDU LIFETIME CONTROL function evaluates and takes action based on the contents of the PDU Lifetime field in the NPDU header. This field contains, at any time, the remaining lifetime of the NPDU (represented in units of 500 milliseconds). The lifetime of the Initial NPDU is at least three (3) times the ATN Internet span or three (3) times the maximum expected transit delay (in units of 500 milliseconds), whichever is greater. This value is set by the originating Network entity, and placed in the PDU Lifetime field of the NPDU. When the SEGMENTATION function is applied to an NPDU, the value of the PDU Lifetime field of the Initial NPDU is copied into all of the Derived NPDUs.

The lifetime of the NPDU is decremented by every Network entity which processes the NPDU. When a Network entity processes an NPDU, it decrements the PDU Lifetime field by at least one count. The value of the PDU Lifetime field is decremented by more than one count if the sum of:

1. the transit delay in the underlying service from which the NPDU was received; and

2. the delay within the system processing the NPDU

exceeds or is estimated to exceed 500 milliseconds. In this case, the PDU Lifetime field is decremented by one for each additional 500 milliseconds of delay. The determination of delay is not required to be precise, but where a precise value cannot be ascertained, the value used is an overestimate, not an underestimate.

If the PDU Lifetime field reaches a value of zero before the NPDU is delivered to the destination, the NPDU is discarded. The ERROR REPORTING function is invoked, and results in the generation of any required ER NPDUs.

4.3.9.5 Route PDU Function

The ROUTE PDU function determines the Network entity to which a protocol data unit must be forwarded and the underlying service that must be used to reach that Network entity. The ROUTE PDU function is closely associated with the routing functions of the ES-IS and IS-IS routing information exchange protocols.

The ROUTE PDU function uses the Destination Address, the total length of the NPDU, and connectivity/topology information contained in the Routing Information Base in order to select a destination Network entity and underlying subnetwork service for forwarding an NPDU. Where segmentation is required, the ROUTE PDU function further determines over which underlying service the Derived NPDU segments must be sent in order to reach that Network entity. The results of the ROUTE PDU function are passed to the FORWARD PDU function (along with the NPDU itself) for further processing. Selection of the underlying service that must be used to reach the "next" system in the route is initially influenced by the **NS-Quality-of-Service** parameter of the **N-UNITDATA** Request, which specifies the QOS requested by the sending ATN NS user. The ROUTE PDU function determines whether this QOS is to be provided directly by the ATN CLNP (through the selection of the Quality of Service Maintenance parameter and other optional parameters) or through the QOS facilities offered by each of the underlying services, prior to invocation of the FORWARD PDU function. Route selection also takes into consideration the value of the Quality of Service Maintenance parameter, and other optional parameters provided in the NPDU.

4.3.9.6 Forward PDU Function

The FORWARD PDU function provides access to and control of local interfaces to supporting subnetworks and/or convergence protocols. The FORWARD PDU function issues an **subnetwork-UNITDATA** Request primitive, supplying the subnetwork or SNDCF identified by the ROUTE PDU function with the protocol data unit as user data to be transmitted, the address information required by that subnetwork or SNDCF to identify the adjacent system within the subnetwork-specific addressing domain (this may be an intermediate-system or the destination end-system), and QOS constraints (if any) to be considered in the processing of the user data. When the NPDU to be forwarded is longer than the maximum service data user size provided by the underlying service, the SEGMENTATION function is applied.

4.3.9.7 Segmentation Function

For an ATN Network Entity implementing the full protocol, segmentation is performed when the size of the PDU is greater than the maximum service data unit size supported by the underlying service to be used to transmit the NPDU. The underlying service may be provided indirectly by the Subnetwork Dependent Convergence Facility, or directly by the Subnetwork Access Protocol. Segmentation comprises the composing of two or more new NPDUs (Derived NPDUs) from the NPDU received. The NPDU received may be the Initial NPDU, or it may be a Derived NPDU.

All of the header information from the NPDU to be segmented, with the exception of the segment length and checksum fields of the fixed part, and the segment offset of the segmentation part, is duplicated in each Derived NPDU, including all of the address part, the data unit identifier and total length of the segmentation part, and the options part (if present). The rules for forwarding and

segmentation guarantee that the header length is the same for all segments (Derived NPDUs) of the Initial NPDU, and is the same as the header length of the Initial NPDU. The size of an NPDU header will not change due to operation of any protocol function. The user data encapsulated within the NPDU received is divided such that the Derived NPDUs satisfy the size requirements of the user data parameter field of the primitive used to access the underlying service.

Derived NPDUs are identified as being from the same Initial NPDU by means of

1. the source address,

2. the destination address, and

3. the data unit identifier.

The following fields of the NPDU header are used in conjunction with the Segmentation function:

Segment Offset:	Identifies the octet at which the segment begins, with respect to the start of the Initial NPDU.
Segment Length:	Specifies the number of octets in the Derived NPDU, including both header and data.
More Segments Flag:	Set to **[1]** if this Derived NPDU does not contain, as its final octet of user data, the final octet of the Initial NPDU.
Total Length	Specifies the entire length of the Initial NPDU, including both header and data.

Derived NPDUs may be further segmented without constraining the routing of the individual Derived NPDUs.

The Segmentation Permitted flag is set to **[1]** to indicate that segmentation is permitted. If the Initial NPDU is not to be segmented at any point during its lifetime in the Network, the flag is set to **[0]** by the source Network entity. The setting of the Segmentation Permitted flag cannot be changed by any other Network entity for the lifetime of the Initial NPDU and any Derived NPDUs.

4.3.9.8 Reassembly Function

The Reassembly function reconstructs the Initial NPDU from the Derived NPDUs generated by the operation of the Segmentation Function on the Initial NPDU (and, recursively, on subsequent Derived NPDUs).

A bound on the time during which segments (Derived NPDUs) of an Initial NPDU must be held at a reassembly point before being discarded is provided, so that reassembly resources may be released when it is no longer expected that any outstanding segments of the Initial NPDU will arrive at the reassembly point. Upon reception of a Derived NPDU, a reassembly timer is initiated with a value which indicates the amount of time which must elapse before any outstanding segments of the Initial NPDU are assumed to be lost. When this timer expires, all segments (Derived NPDUs) of the Initial NPDU held at the reassembly point are discarded, the resources allocated for those segments are freed, and if requested, an ER is generated. While the exact relationship between reassembly lifetime and NPDU lifetime is a local matter, the Reassembly Function should preserve the intent of the NPDU lifetime. Consequently, the reassembly function should discard NPDUs whose lifetime would otherwise have expired had they not been under the control of the reassembly function.

The Segmentation and Reassembly functions are intended to be used in such a way that the fewest possible segments are generated at each segmentation point and reassembly takes place at the final destination of an NPDU. However, other schemes which:

1. interact with the routing algorithm to favour paths on which fewer segments are generated;

2. generate more segments than absolutely required in order to avoid additional segmentation at some subsequent point; or

3. allow partial or full reassembly at some intermediate point along the route;

are not precluded. The information necessary to enable the use of one of these alternative strategies may be made available through the operation of a Network Layer Management function or by other means.

The originator of the Initial NPDU determines the value of the Segmentation Permitted flag in the Initial NPDU and all Derived NPDUs (if any). Partial or full reassembly in an ATN Intermediate-system cannot change this value in the Initial NPDU or any NPDU derived from it, and cannot therefore add or remove the segmentation part of the header.

4.3.9.9 Discard PDU Function

The DISCARD PDU function performs all of the actions necessary to free the resources reserved by the Network entity when an error condition prevents further processing of the NPDU. The DISCARD PDU function is executed when any of the following error conditions is encountered:

1. A violation of protocol procedure has occurred.

2. An NPDU is received whose checksum is inconsistent with its contents.

3. An NPDU is received, but due to local congestion, it cannot be processed.

4. An NPDU is received with a correct header checksum, but whose header contents are invalid.

5. An NPDU is received which cannot be segmented and cannot be forwarded because its length exceeds the maximum service data unit size supported by any underlying service available for transmission of the NPDU to the next Network entity on the chosen route.

6. An NPDU is received whose destination address is unreachable or unknown.

7. Incorrect or invalid source routing was specified. This may include a syntax error in the source routing field, an unknown or unreachable address in the source routing field, or a path which is not acceptable for other reasons.

8. An NPDU is received whose NPDU lifetime has expired or a segmented NPDU is received whose lifetime expires during reassembly.

9. An NPDU is received which contains an unsupported Type 2 option.

4.3.9.10 Error Reporting Function

The ERROR REPORTING function initiates the return of an ER NPDU to the source Network entity when a protocol data unit is discarded. The ER NPDU identifies a discarded NPDU, specifies the type of error detected, and identifies the discarding Network entity. Error Report procedures are not used to convey information regarding success or failure of delivery of an NPDU issued by a source Network entity.

The originator of a DT NPDU controls the generation of ER NPDUs. An ER flag in the original NPDU is set by the source Network entity to indicate that an ER NPDU is to be returned if the Initial NPDU or any NPDUs derived from it are discarded; if the flag is not set, Error Reports are

suppressed. The suppression of ER NPDUs is controlled by the originating Network entity and not by the ATN NS user.

The ERROR REPORTING function performs as follows:

1. An ER NPDU is not generated to report the discard of an ER NPDU.

2. An ER NPDU is not generated to report the discard of a DT NPDU unless that NPDU has the ER flag set to allow Error Reports.

3. The entire header of the Discarded NPDU is placed in the data field of the ER NPDU. The data field of the Discarded NPDU is not included in the data field of the ER NPDU.

4. If a DT NPDU is discarded for one of the reasons in Paragraph 4.3.9.9, and the ER flag has been set to allow Error Reports, an ER NPDU is generated.

If a DT NPDU with the E/R flag set to allow Error Reports is discarded for any other reason, an ER NPDU may be generated (as an implementation option).

4.3.9.10.1 Initiation of Error Reports

An ER NPDU is composed from information contained in the header of the discarded *Data* (DT) NPDU to which the Error Report refers. The content of the Source Address field of the discarded DT NPDU is used as the Destination Address of the ER NPDU. This value (which in the context of the DT NPDU was used as an NSAP Address) is used in the context of the ER NPDU as the NET of the Network entity that originated the DT NPDU. The NET of the originator of the ER NPDU is conveyed in the Source Address field of the header of the ER NPDU.

Segmentation of ER NPDUs is not permitted; hence, no Segmentation Part is present. The total length of the ER NPDU in octets is placed in the Segment Length field of the ER NPDU header. This field is not changed during the lifetime of the ER NPDU. If the originator of the ER NPDU determines that the size of the ER NPDU exceeds the maximum service data unit size of the underlying service, the ER NPDU is truncated to the maximum service data unit size and forwarded with no other change.

The requirement that the underlying service assumed by the CLNP must be capable of supporting a service data unit size of at least 512 octets guarantees that the entire header of the discarded DT NPDU can be conveyed in the data field of any ER NPDU.

4.3.9.10.2 Processing of Received Error Reports

When an ER NPDU is decomposed upon reaching its destination, information required to interpret and act upon the Error Report is obtained as follows:

1. The NET recovered from the NPAI in the Source Address field of the ER NPDU header is used to identify the Network entity which generated the Error Report.

2. The reason for generating the Error Report is extracted from the Options Part of the NPDU header.

3. The entire header of the discarded DT NPDU is extracted from the data field of the ER NPDU to assist in determining the nature of the error.

ER NPDUs are routed and forwarded by ATN Intermediate-system Network entities in the same way as DT NPDUs.

4.3.9.10.3 Relationship of Data NPDU Options to Error Report NPDUs

The generation of an Error Report is controlled by options that are present in the corresponding DT NPDU. The presence of options in the original DT NPDU that are not supported by the system which has discarded that NPDU may cause the suppression of an Error Report even if the original DT NPDU indicated that an Error Report should be generated in the event of a discard.

The processing of an Error Report is controlled by options which are present in the corresponding DT NPDU. In particular, options selected for the original DT NPDU affect which options are included in the corresponding ER NPDU.

The selection of options for an ER NPDU are specified as follows:

1. If the Priority Option or the QOS Maintenance Option is selected in the original DT NPDU, and the system generating the ER NPDU supports the option, then the ER NPDU specifies the option.

2. If the Security Option is selected in the DT NPDU, and the system generating the Error Report supports this option, then the ER NPDU specifies the option using the value that was specified in the original DT NPDU. If the system does not support the Security Option, an Error Report must not be generated for a DT NPDU that selects the Security Option.

3. The Record Route Option, if selected in the DT NPDU, is specified in the ER NPDU.

The values of the optional parameters above may be derived as a local matter, or they may be based upon the corresponding values in the original DT NPDU.

4.3.9.11 PDU Header Error Detection

The PDU HEADER ERROR DETECTION function protects against failure of ATN IS or ES entities due to the processing of erroneous information in the NPDU header. The PDU HEADER ERROR DETECTION function uses a checksum computed on the entire NPDU header. The checksum is verified at each point at which the NPDU header is processed. If the checksum calculation fails, the NPDU is discarded. If NPDU header fields are modified (e.g., due to operation of the PDU LIFETIME function), then the checksum is modified so that the checksum remains valid. The use of the Header Error Detection function is optional, and is selected by the originating Network entity. If the function is not used, the checksum field of the NPDU header is set to zero.

If the function is selected by the originating Network entity, the value of the checksum field causes the following conditions to be satisfied:

$$\sum a_{i\ (1 \le i \le L)} \pmod{255} = 0 \qquad \textbf{(9.1)}$$

$$\sum (L - i + 1)a_{i\ (1 \le i \le L)} \pmod{255} = 0 \qquad \textbf{(9.2)}$$

where L = the number of octets in the NPDU header, and a_i = the value of the octet at position i. The first octet in the NPDU header is considered to occupy position i = 1. When the function is in use, neither octet of the checksum field is set to zero.

An efficient algorithm for calculating and checking the checksum octets is provided in Annex D of ISO 8073 and ISO 8602. The checksum is easy to compute and does not impose a serious burden on implementations. However, it will not detect insertion or loss of leading or trailing zero octets, nor will it detect some forms of octet misordering.

4.3.10 ISO 8473 Optional Internetwork Protocol Functions

ISO 8473 internetwork protocol options are selected by the ATN ES Network entity which originates ISO 8473 NPDUs. As a part of the ISO 8473 header, options are conveyed between peer Network entities via ATN subnetworks, and are evaluated in turn by each receiving ATN intermediate-system. The information contained in options conveyed via the ISO 8473 CLNP header is delivered unchanged to each successive ATN entity along the end-to-end path between source and destination ES.

4.3.10.1 Padding Function

The PADDING Function allows extending the length of ISO 8473 NPDUs beyond the length required to convey the NSDU, in order to accommodate those ESs and ISs which place sizing constraints upon NPDUs to facilitate processing.

4.3.10.2 Security Function

The SECURITY function supports imposition of Network Layer security provisions by way of an options field conveyed within the ISO 8473 header. The information contained within this options field may be specified in a global context (i.e. by the international standard), or within the context of the addressing authority responsible for the assignment of the NPDU's source or destination NSAP Address. These contexts are known respectively as the Globally Unique, Source and Destination Unique Formats.

ATN Conformant systems are only required to recognise this options field when it is specified in the global context. Although a source or destination NSAP Address assigned using the ATN NSAP Addressing Plan could be used to identify ATN Security Information in the source or destination context, the SARPs does not mandate support of the source or destination specific formats for the ISO 8473 security parameter, and hence to avoid service irregularities, neither format should be used.

The Security options field is included in the ISO 8473 header by an ES when the NS User provides a Security Label with an NSDU. In the ATN, this is always encoded using the Globally Unique Format, and is the encoding of the ATN Security Label provided on the N-UNITDATA.Request. As discussed in 4.3.8, the security options parameter is referenced by the inter-domain forwarding function and used to determine the route that an NPDU follows. It is, however, never modified by an IS.

When the NPDU reaches its destination, the value of the Security options field is provided to the destination NS use as the Security Label associated with the NSDU.

4.3.10.3 Source Routing Function

The SOURCE ROUTING Function allows specification of a particular path (i.e., sequence of ISs) through which a particular NPDU either should pass or must pass. The former is described as Partial Source Routing, and the latter is described as Complete Source Routing. The path is defined by a supplied list of NETs, which is conveyed within the NPDU header.

4.3.10.4 Record Route Function

The RECORD ROUTE function records the path taken by an NPDU as it traverses a series of ATN ISs. A recorded route consists of a list of NETs held in a parameter within the options part of the NPDU header. The length of this parameter is determined by the originating Network entity, and does not change as the NPDU traverses the Network. The list is constructed as the NPDU is forwarded along a path towards its destination. Only the titles of ATN Intermediate-system Network entities are included in the recorded route; the NET of the originator of the NPDU is not recorded in the list.

When an ATN IS processes an NPDU containing the Record Route option, the IS adds its own NET at the end of the list of recorded NETs. An indicator is maintained to identify the next available octet to be used for recording of route. This indicator is updated as entries are added to the list using the following procedure:

1. The length of the entry to be added to the list is added to the value of the next available octet indicator, and this sum is compared with the length of the Record Route parameter.

2. If the addition of the entry to the list would exceed the size of the parameter, the next available octet indicator is set to indicate that route recording has been terminated. The NET is not added to the list.

3. If the addition of the entry would not exceed the size of the Record Route parameter, the next available octet indicator is updated with the new value, and the NET is added to the head of the list after the other entries have been moved.

Two forms of the RECORD ROUTE function are possible. The first form is referred to as Complete Route Recording. It requires that the list of NETs be a complete and accurate record of all ATN ISs visited by an NPDU (including Derived NPDUs), except when a shortage of space in the record route option field causes termination of recording of route, as described in Step 2 above. When Complete Route Recording is selected, NPDU reassembly at ATN ISs may be performed only when the Derived NPDUs that are reassembled all took the same route; otherwise, the NPDU is discarded, and if selected, an Error Report is generated. The second form is referred to as Partial Route Recording. It also requires a record of ATN ISs visited by an NPDU. When Partial Route Recording is selected, NPDU reassembly at ATN ISs is always permitted. When reassembly is performed at an ATN IS, the route recorded in any of the Derived NPDUs may be placed in the NPDU resulting from the reassembly.

When a shortage of space in the option field causes termination of the RECORD ROUTE function, the NPDU may still be forwarded to its final destination, without further addition of NETs.

The Record Route function is intended to be used in the diagnosis of subnetwork and/or routing problems.

4.3.10.5 Quality of Service Maintenance Function

The QUALITY OF SERVICE MAINTENANCE function allows the originating Network entity to indicate to ATN Intermediate-systems the relative importance of certain qualities of service for routing decisions made on an individual internetwork packet basis. This information is conveyed to ATN Intermediate-system Network entities in a parameter in the options part of the NPDU header. This option is used to resolve routing ties, where more than one path is available for routing of an NPDU toward its destination. Network entities make use of this information in selecting a route when more than one route satisfying other routing criteria is available.

The ISO 8473 CLNP QUALITY OF SERVICE MAINTENANCE function may be encoded in one of three ways, denoted Source Address Specific, Destination Address Specific and Globally Unique. The first two choices allow selection of an option coding scheme which is associated with the authority defining either source or destination NSAP addresses, while the latter choice uses an internationally agreed upon coding of the relative importance of three subnetwork QOS parameters. These qualities of service include Expense, Transit Delay and Residual Error Probability.

The **Globally Unique** format for the QUALITY OF SERVICE MAINTENANCE function indicates the relative importance of three subnetwork QOS parameters: Expense; Transit Delay; and Residual Error Probability. This option is expressed as a four bit mask within one octet in the protocol header; there is no specified default value for this mask. If no value for **Quality of Service Maintenance** is indicated within the CLNP packet, Network entities use local route selection rules, making their best effort to deliver the CLNP packet. The omission of the **Quality of Service**

Maintenance option is equivalent to requesting that ATN ISs optimise offered throughput. In those instances where the QOS requested cannot be maintained, ATN Network entities will attempt to deliver the NPDU at any available QOS.

4.3.10.6 Priority Function

The PRIORITY function provides a means whereby the resources of ATN ES and ATN IS Network entities, (i.e., outgoing transmission queues and buffers) can be used to process higher-priority NPDUs ahead of lower-priority NPDUs. The PRIORITY function influences the dynamic reordering of the CLNP packet queue within ATN ISs and ESs. This queue management technique allows the proper allocation of packets among available subnetworks, as well as the proper ordering of packets for transfer within a given subnetwork.

The PRIORITY function supports the use of a number between 0 and 14 to indicate the relative importance of each connectionless internetwork protocol packet. The highest Network layer priority is associated with CLNP Level 14, while the lowest priority is associated with CLNP Level 0; Level 15 is a reserved value. CLNP Priority 0 is the default priority, and is used where no priority value is explicitly indicated.

ATN use of the Priority Function is discussed in 4.3.7.

4.3.10.7 Congestion Notification Function

The CONGESTION NOTIFICATION FUNCTION allows originating ATN ESs to take appropriate action when congestion is experienced within the ATN internet.

An ATN IS is viewed as congested when inadequate buffer space is available to maintain and process output queues. ATN ISs detect and indicate congestion based upon the depth of the output queue selected for an NPDU (according to its destination address or other routing information).

ATN Intermediate-systems informs the originating Network entity of congestion between the source and destination NSAP through the use of a flag in the **QOS Maintenance Parameter** option header. When the depth of a particular output queue exceeds a certain proportion of the depth of that queue, an ATN Intermediate-system will start to discard NPDUs; at this time, the ATN Intermediate-system sets the *Congestion Experienced* flag in the next NPDU to be forwarded toward one or more source Network entities and continues to do so until the congestion condition is alleviated.

The value of the *Congestion Experienced* flag is initially set to zero **[0]** by the originator of the NPDU and is set to one **[1]** by any ATN Intermediate-system which processes the NPDU to indicate that that ATN Intermediate-system is experiencing congestion. The method of initiating Congestion Notification is discussed in chapter 6..

4.3.10.8 Echo Request and Response

The Echo Request function is invoked by Network Layer Management to obtain information about the dynamic state of the Network Layer with respect to (a) the reachability of specific Network entities, and (b) the characteristics of the path or paths that can be created between Network Entities through the operation of Network Layer routing functions. Together with the Echo Response function, it fulfils the same role as “Ping” and “Traceroute” in the Internet Protocol suite.

An Echo Request is generated as a result of a request made on a local management interface. Its destination is the NET of another Network Entity i.e. the Network Entity for which reachability is to be determined, or the route traced. When the Echo Request is received by that Network Entity, an Echo Response is returned to the sending Network Entity.

A returned Echo Response may then be analysed to determine information about the route between two network entities.

4.3.11 Notes on the CLNP APRLs

The following notes have been prepared to provide implementors with background information on conformance requirements which may differ from normal practice.

4.3.11.1 Security

Mandatory implementation of the security parameter is required to support ATN Routing Control functions. As a type 2 function, every ATN System must support this parameter is connectivity is to be maintained. However, within a Routing Domain, it is acceptable for the actual value of this parameter to be ignored.

4.3.11.2 Complete Route Recording

Complete Route Recording is not permitted on the ATN due to concerns over the packet sizes that could be required and the consequential impact on air-ground data links and the transfer of safety related data.

4.3.11.3 Source Routing

Neither Complete Source Routing nor Partial Source Routing are permitted on the ATN. This is because source routing could be used to overcome or otherwise interfere with ATN Routing Control.

4.3.11.4 Priority

Priority is a mandatory ATN requirement. All ATN Systems must not only recognise the priority parameter, but must also prioritise their output queues and implement priority based discard algorithms, if it is necessary to discard packets during periods of congestion. This feature is essential to ensure that safety related data is not impeded if the ATN is congested with routine data.

4.3.11.5 Padding

NPDU padding is not permitted on the ATN as it would interfere with the compression algorithm used by the Mobile SNDCF. The Local Reference Compression mechanism includes no facilities for compressing padding and such NPDUs are sent uncompressed, resulting in a significant increase in the overhead on air-ground data links.

4.4 The Implementation of the Routing Information Exchange Protocols

In support of the ISO 8473 connectionless network layer protocol, ISO has defined a family of three routing information exchange protocols, specified by ISO 9542, ISO/IEC 10589 and ISO/IEC 10747, respectively.

ISO 9542 specifies a protocol for use between ESs and ISs. This protocol enables ISs to identify the NSAP Addresses located on each adjacent ES, and for ESs to determine the location of each adjacent IS. ESs then have a simple routing decision in the absence of any precise knowledge about the location of a packet's destination: they choose an adjacent IS and send the packet to it. It is then the IS's responsibility to route the packet either to its destination, or to an IS nearer to it. When the packet is passed to an ES or IS that is also known to be adjacent to the originating ES, then ISO 9542 allows the IS to notify the ES of the direct path, so that it may be used for all further packets to that destination.

ISO/IEC 10589 is a routing information exchange protocol for use between ISs within the same RD. This protocol freely exchanges all routing information known by each IS to all other ISs. Each IS then has a complete routing map of the RD from which it can calculate optimal routes. This is a simple and robust approach that exploits the requirements for common routing procedures and trust. However, it is hence not suitable for inter-RD routing information exchange. ISO has thus defined a different routing information exchange protocol for communication between RDs. This is specified in ISO/IEC 10747, and is known as the Inter-Domain Routing Protocol (IDRP).

Reflecting the environment of limited trust and different route selection algorithms, rather than exchanging general topology data, IDRP exchanges processed data; IDRP advertises routes to destinations and enables an RD to advertise only the routes that it wants to. It is thus said to support policy based routing. Each RD implements its own routing policy which reflects its security policy and other technical considerations.

4.4.1 ES-IS Implementation Considerations

4.4.1.1 Overview

ISO 9542 specifies a very simple datagram protocol which is suitable for use on all sorts of networks, although it achieves its greatest potential on Broadcast subnetworks. The protocol supports two functions: Configuration Information and Redirection Information.

The Configuration Information function enables End Systems to discover the existence of Intermediate Systems and vice-versa. On broadcast subnetworks, such as an Ethernet, each End System regularly sends an "End System Hello" message reporting the network addresses it hosts to the multicast address *all intermediate systems*. Similarly, each Intermediate System regularly sends an "Intermediate System Hello" message reporting its own identity to the multicast address *all end systems*. End Systems and Intermediate Systems always listen to their respective multicast addresses and can hence "discover" the existence of Intermediate Systems and End Systems, respectively.

In OSI, End Systems have a very simple routing decision: if they do not know the location of the destination of a packet, they send it to any Intermediate System they have discovered through the Configuration Information function.

The Intermediate System should then relay the packet on to its destination. However, if the destination is on the same subnetwork as the source, or another Intermediate System would have been a better choice, then the Route Redirection Information function can be used to inform the End System of the better routing decision. A redirection message is sent to the End System by the Intermediate System, which identifies the subnetwork address that is more appropriate for the destination network address. The End System can then use this subnetwork address in future.

The protocol can also support routing in the absence of an End System. In such cases, instead of the End System sending the packet to any Intermediate System, it sends it to the multicast address *all end systems*. If the End System which is the packet's true destination receives the packet then it returns an End System Hello to the sender to report the correct subnetwork address, and communication can proceed.

The Configuration Information function may also be used on general topology subnetworks e.g. "X.25 Networks". In such cases, it can still be used to determine the addresses supported by each system, by passing Hello messages over a virtual circuit. However, dynamic discovery of the systems themselves is not really possible given that the DTE Addresses must be known before a virtual circuit can be established.

In the ATN, the Configuration Information function is also used with mobile networks. The air-ground data links specified by ICAO, all appear externally as X.25 data networks. However, the systems reachable over such networks may come and go depending on their geographic position. Their availability may be notified by a "Join Event", or it may be determined through a polling

strategy, a subnetwork connection established and communication take place. An exchange of ISO 9542 Configuration Information is required as part of this procedure.

4.4.1.2 ATN Use of ISO 9542

In the air-to-ground environment, the operation of the ISO 9542 protocol is mandatory, in order to allow adjacent ground and airborne routers connected via a mobile subnetwork to monitor connectivity changes.

The ISO 9542 routing protocol is the recommended protocol for performing these functions over ATN fixed subnetworks.

ISO 9542 is also required when ISO/IEC 10589 is implemented (see 0).

4.4.2 The ES-IS Protocol

4.4.2.1.1 PDU Formats and Use

ISO 9542 operates among the systems attached to a single subnetwork, independently from the routing organisation. It is used to allow systems on the subnetwork to discover each other (configuration), and if necessary to provide minimal routing information to ESs (route redirection).

ISO 9542 specifies three PDU types: the End System Hello (ESH) PDU, the Intermediate System Hello (ISH) PDU, and the Redirect (RD) PDU.

For each type of ISO 9542 PDU, Table 4-6, Table 4-7, Table 4-8 and Table 4-9 respectively indicate:

ISO 9542 PDUs	Generation of PDUs
ESH	By each ES: On timer expiry or on other events, such as the ES or a new local SNPA becoming operational, a distant ES or IS becoming operational, or after another ES has performed a Query Configuration function (Configuration Response)
ISH	By each IS: On timer expiry or on other events, such as the IS or a new local SNPA becoming operational or a distant ES or IS becoming operational (Configuration Notification)
RD	By any IS: After reception of a data PDU, when the IS detects that there is a better path to reach the destination NSAP, or that it cannot route to this destination NSAP (Request Redirect)

Table 4-7 Generation of ISO 9542 PDUs

1. the main contents of the PDU,
2. the type of systems which generates this PDU,
3. the event which triggers its generation,
4. the destination systems of this PDU,

ISO 9542 PDUs	**Main Contents**
ESH	**Source address parameter**: Address(es) of the NSAP(s) supported by the ES originating the ESH PDU (an ESH may convey any number of NSAPs supported by the ES in the limit of subnetwork data units size, but in the end, the ES must have reported information about all its NSAPs, via one or several ESHs)
ISH	**Source address parameter:** NET of the IS sending the ISH PDU (the protocol allows only one NET in each ISH)
RD	**Source address parameter:** NET of the IS sending the RD PDU (only one NET); **Destination address parameter:** Destination NSAP address of the PDUs affected by the redirection (and possibly a mask selecting a "class" of NSAPs); Subnetwork address of the new network entity (on the same subnetwork) to which the redirected PDUs will be sent for the first hop from the ES (better path to destination)

Table 4-6 ISO 9542 PDU Types

5. its functional role.

The basic transmission mechanism for ISO 9542 configuration information is broadcast. When the underlying subnetwork does not support broadcast or multi-cast the SNDCF may have to provide the required adaptation.

Two broadcast subnetwork destination addresses are possible:

I. "All ESs network entities", or

I. "All ISs network entities".

Consequently, in the "normal" use of the protocol, all the ISO 9542 PDUs generated by each ES are sent to all the ISs on the same subnetwork, and all the ISO 9542 PDUs generated by each IS are sent to all the ESs on the same subnetwork.

ISO 9542 PDUs	**Propagation of PDUs**
ESH	• Transmitted on each SNPA the ES is attached to (the transmitted PDUs may be different but they must provide the same information) • Transmitted from an ES in response to a query configuration • Transmitted to all the ISs on the subnetwork
ISH	• Transmitted on each SNPA the IS is attached to • to all the ESs on each subnetwork the IS is attached to
RD	• Transmitted by any IS • Transmitted to the ES originating the PDU when the IS knows a better path

Table 4-8 Propagation of ISO 9542 PDUs

ISO 9542 PDUs	Functional Role
ESH	**CONFIGURATION** • Allows all the ISs to discover the existence and reachability (SNPA) of an ES on the same subnetwork, along with the NSAPs this ES supports • Allows the ESs to discover the existence and reachability of another ES on the same subnetwork, along with the NSAPs this ES supports
ISH	**CONFIGURATION** • Allows all the ISs to discover the existence and reachability (SNPA) of an IS on the same subnetwork along with the NET of that IS (when ISO 9542 is used between ISs) • Allows all the ESs to discover the existence and reachability (SNPA) of an IS on the same subnetwork along with the NET of this IS
RD	**ROUTE REDIRECTION** • Allows an IS to inform the source ES (on the subnetwork) of a better path to reach a destination NSAP (by indicating another IS corresponding to a better first hop on the same subnetwork, or directly the destination ES if it is on the same subnetwork) • It may also relate to a "class" of NSAPs (using Address Masks)

Table 4-9 Role of ISO 9542 PDUs

4.4.2.1.2 Main protocol functions

ISO 9542 may be implemented by a simple state machine, and a single function is specified to respond to each incoming event. These functions are discussed below.

4.4.2.1.2.1 Report Configuration Function

This function is used by ESs and ISs to inform each other of their reachability and current subnetwork address(es). Additionally, the NET of ISs and the NSAP(s) of ESs are made available to other systems on the subnetwork. This function is invoked on timer expiry or on other event detection.

4.4.2.1.2.2 Record Configuration Function

The record configuration function is implemented in ESs and ISs. It is in charge of the receipt of ESH and ISH PDUs. This function extracts configuration information from the received packets and updates the local Network entity's RIB.

4.4.2.1.2.3 Flush Old Configuration Function

This function is executed to remove configuration entries in the RIB when this information becomes obsolete. The function is either executed on timer expiry or on other event detection (SNPA re-initialisation).

4.4.2.1.2.4 Query Configuration Function

This function is executed by an ES attached to a broadcast subnetwork when no IS is reachable on the subnetwork and when the ES Route PDU function is not able to determine the SNPA address associated with the current destination NSAP.

When the ES needs to route an NPDU to a destination NSAP whose SNPA is unknown, it performs a broadcast on the subnetwork by sending the NPDU to "All ES entities on the Subnetwork".

Either the destination ES is attached to the subnetwork and the originator ES receives an ESH from the destination system, or no ESH is received and the destination may be declared unreachable.

4.4.2.1.2.5 Configuration Response Function

This function is performed by an ES on receipt of a NPDU addressed to one of its NSAPs, with broadcast destination SNPA address. This is the result of another ES having performed the Query Configuration Function.

The receiving ES builds an ESH PDU and sends it back to the originator ES.

4.4.2.1.2.6 Configuration Notification Function

This function is performed by an ES or IS in order to quickly transmit configuration information (ESH or ISH) to a system which has newly become available and which has issued an ESH or ISH PDU. The Hello PDU is specifically addressed to the newly reachable system.

4.4.2.1.2.7 Request Redirect Function

This function is performed by an IS having received an NPDU from an ES on the subnetwork. It is used to inform the originator ES that this NPDU should directly have been sent to another system on the subnetwork.

The Redirect information contained in the *Redirect PDU* (RD PDU) issued by the IS informs the originator ES of a better path to the NPDU destination.

4.4.2.1.2.8 Record Redirect Function

This function is implemented in ESs and is in charge of recording the redirection information received from an IS. The local Network Entity RIB is updated by this function.

4.4.2.1.2.9 Refresh Redirect Function

The purpose of this function is to increase the longevity of a redirection without allowing an incorrect route to persist indefinitely. In an ES, on receipt of an NPDU the previous hop of which maps the next hop address stored with some redirection information, and the source of which maps the destination address stored with the redirection information, the corresponding redirection holding timer is reset.

4.4.2.1.2.10 Flush Old Redirect Function

This function is performed to remove redirection entries in the RIB when this information becomes obsolete. The function is either executed on timer expiry or on event detection (SNPA re-initialisation).

4.4.2.1.2.11 PDU Header Error Detection

This function is performed by ESs or ISs in order to protect themselves against failures due to the processing of erroneous information in the PDU header. This function performs computation and verification of a checksum and discards the PDU in case of inconsistency.

4.4.2.1.2.11.1 Protocol Error Processing Function

An ISO 9542 PDU which is not discarded by the PDU Header Error Detection Function is discarded by the Protocol Error Processing Function if its encoding does not comply with the provisions of the ISO 9542 protocol.

4.4.2.1.3 ISO 9542 Operation Among ESs

When ISO 9542 is used among the ESs of a single subnetwork, the ESH PDUs are transmitted with the same destination subnetwork address ("All ISs"), and the ESs that wish to receive information about the other ESs validate the reception of the ESHs by validating this address; thus they are aware of the existence and reachability of the other ESs.

This allows optimisation, namely by anticipating the information contained in the RD PDUs, when the destination NSAP is supported by an ES on the same subnetwork.

The operation of ISO 9542 among the ESs generates no additional information transmission (compared with the "standard operation").

4.4.2.1.4 ISO 9542 Operation as an Initiation Phase for the Routing Protocols

In the same way, when ISO 9542 operates among the ISs attached to a single subnetwork, the ISs validate the reception of the ISHs normally destined for the ESs, by validating the corresponding subnetwork address ("All ESs").

This allows the ISs to discover their neighbour ISs existence and reachability and may be used as an initialisation phase for the routing protocols.

4.4.2.2 ISO 9542 Operation over Fixed Ground Subnetworks

4.4.2.2.1 Overview

The use of ISO 9542 over ATN ground subnetworks is a recommended practice. However, either static routing information or other routing protocols could be used to provide the same type of functions as ISO 9542.

If ISO 9542 is not operated over ground subnetworks, a facility must fulfil the following requirements

1. each system must be able to discover the existence of neighbour systems attached to the same subnetwork,

2. the NSAP and SNPA addresses of neighbour ESs and the NET and SNPA addresses of neighbour ISs must be made available to each IS directly connected to the local subnetwork,

3. each IS must be able to dynamically monitor connectivity changes over the local subnetwork

4.4.2.2.2 General Topology Subnetworks

In the case ISO 9542 is operated over ground ATN subnetworks, it seems reasonable to advise against the support of configuration information over general topology subnetwork (non-broadcast

subnetwork). Furthermore, it can be very costly to simulate broadcast over non-broadcast subnetworks. However, in some cases (high-bandwidth subnetworks), this solution can be chosen.

On the other hand, the support of ISO 9542 redirection information on general topology subnetwork may be advised, since it is not costly and may prove useful to ascertain local topology.

4.4.2.2.3 Broadcast Subnetworks

As far as broadcast subnetworks are concerned, the full use of ISO 9542 is recommended, since this protocol was designed for operation over this kind of subnetwork. The use of ISO 9542 over broadcast subnetworks is not too costly and allows to dynamically ascertain local configuration changes.

4.4.2.2.4 Point to Point Subnetworks

As far as point to point subnetworks are concerned, the use of ISO 9542 is recommended, and especially the support of the configuration information. The use of ISO 9542 protocol over point-to-point subnetworks is not too costly.

4.4.2.3 ISO 9542 Operation over Air-ground Mobile Subnetworks

When a new aircraft enters the coverage of a ground router directly connected to a mobile subnetwork, an initialisation phase is triggered so that communication can be established between peer ground and airborne routers.

Once this initialisation phase has been performed, it is necessary for each router to forward its local NET information to the newly reachable routers on the subnetwork.

This action is performed via the exchange of an ISO 9542 ISH PDU, and is discussed in more detail in section two, which deals with the Route Initiation procedure.

4.4.2.4 Notes on the ISO 9542 APRLs

These notes provide background information for implementors on the ISO 9542 APRLs contained in the ATN Internet SARPs. It should also be noted that the APRLs are specific to the use of ISO 9542 to support Route Initiation over air-ground data links. There are no APRLs specified for other uses of ISO 9542 (e.g. to support ES to IS routing).

4.4.2.4.1 Route Redirection Information

Route Redirection Information has no role to play in Route Initiation and is hence excluded from the requirements.

4.4.2.4.2 Configuration Notification

Configuration Notification has no role to play in Route Initiation and is hence excluded from the requirements.

4.4.3 Intra-Domain Routing Implementation Considerations

Intra-Domain Routing operates internally and independently within each ATN Routing Domain. The protocol used to support Intra-Domain Routing within an ATN Routing Domain is a local issue, provided that the general ATN Routing requirements are met.

However, it is recommended that a Routing Domain operate ISO/IEC 10589 IS to IS Intra-Domain Routing Information Exchange Protocol (also called here "IS-IS") as its Intra-Domain Routing Protocol.

This part of the Guidance Material first describes general Intra-Domain Routing goals. The operation of ISO/IEC 10589 for intra-domain routing information propagation within the ATN RDs is then described. Note that the description of ISO/IEC 10589 operation essentially applies to the ATN fixed environment, i.e., to the ground ATN RDs, and in particular AINSC and ATSC RDs. If an alternative intra-domain routing protocol is used, then it must satisfy these goals.

4.4.3.1 ATN Intra-Domain Routing Goals

A. Intra-Domain Routing must be able to route CLNP packets within the local Routing Domain, in order to perform end-to-end routing in the ATN.

B. Intra-Domain Routing must be integrated within the general structure of ATN Routing. Particularly, it must operate within the ATN Network Layer of the ISs located within the Routing Domain.

C. Intra-Domain Routing must meet the following general routing goals:

1. ATN Intra-Domain Routing must be efficient (i.e. induce as little overhead as possible and fulfil the user needs),

2. ATN Intra-Domain Routing must cope with the differences between the interconnected subnetworks (e.g. bandwidth),

3. ATN Intra-Domain Routing must be resilient to failures and adaptable to configuration changes,

4. ATN Intra-Domain Routing must support error control and diagnosis.

4.4.3.1.1 General Requirements

A. The ATN Intra-Domain Routing may use any type of routing procedure, namely:

1. Static routing or quasi-static routing (allowing alternate paths), where pre-determined paths are loaded into the Routing Information database through System Management,

2. Centralised (dynamic) routing, where each system of the RD reports information about its local environment to a central facility, which in turn computes the routes and returns them to all the systems of the RD,

3. Distributed adaptive (dynamic) routing, where all the systems of the RD dynamically sense their local environment and directly exchange Routing Information among themselves, using an Intra-Domain Routing Information dissemination Protocol.

B. Routing Information should preferably be propagated by an Intra-Domain Routing Information Exchange Protocol. However, this is not mandatory, provided that the general Intra-Domain Routing requirements are met .

C. When used, the Intra-Domain Routing Information Exchange Protocol must provide mechanisms for the exchange of connectivity and topology information among ATN Routers within an RD. It must support dynamic configuration of ATN Internet Routing tables on a

domain-wide basis. (see Clause 6.2.3.2. of ISO/IEC 10589 Intra-Domain Routing Information Exchange Protocol).

D. Distributed adaptive routing should preferably be used for Intra-Domain routing in the ATN, for performance considerations. Indeed, these procedures are robust and they automatically and quickly adapt to configuration changes.

E. ISO/IEC 10589 IS to IS Intra-Domain Routing Information Exchange Protocol performs distributed adaptive routing, and more precisely link state routing, where each system independently computes its routes, using a path minimisation algorithm.

F. Intra-Domain Routing may be hierarchically organised to manage large RDs (like ISO/IEC 10589 IS to IS, that allows two intra-domain routing levels).

G. If ISO 9542 ES to IS Routing Protocol is used, it should cooperate with Intra-Domain Routing, so that the ISs of the local RD can dynamically determine their local environment.

H. A RD may use means other than a Routing Information Exchange Protocol to update the Routing Information database (e.g. for RDs with a very simple topology and a limited number of routers). However, the general requirements for ATN Routing must be met. Particularly, the performance should allow timely update of the RIB, for resilience and adaptability.

I. Routing Information dissemination throughout the RD, must allow each IS of the RD to build its local Routing Information database, so that this database can be used to route the CLNP packets within the local domain .

J. Intra-Domain Routing must operate within the Network Layer of each Router and End System of the local RD.

K. Intra-Domain Routing should preferably take into account the distinction made in ISO-OSI Routing between the ESs and the ISs roles, although this is not mandatory .

L. Intra-Domain Routing must be integrated within the ATN Routing Framework described in Chapter two. It must cooperate with the other elements contributing to the ATN Internetworking and Routing, namely the ATN NSAP Addressing Plan, the ATN Internetwork Protocol, and the other ATN Routing Protocols (ISO/IEC 10747 IDRP and ISO 9542 ES to IS Protocol), in order to meet the ATN Intra-Domain Routing Goals defined in 4.4.3.1.

4.4.3.1.2 Intra-Domain Requirements relevant to Inter-Domain Routing

1) Intra-Domain routing must be able to route CLNP packets issued by an ES belonging to the local RD or to an external RD and bound to a destination ES belonging to the local RD or to an external RD.

2) When the local RD acts as a Transit RD, routing of the CLNP packets by the local Intra-Domain Routing procedure may require the encapsulation of the CLNP packets within other CLNP packets conveying locally known NSAP addresses. The decision to encapsulate the CLNP packets and the encapsulation operations (including the locally known NSAP addresses determination) must be performed by Inter-Domain Routing, in the BIS where the packets enter the local RD. The reverse operation must be performed by Inter-Domain Routing, in the BIS where the packets leave the local RD.

Note.— It is important to note however, that when an CLNP packet crosses several RDs , the routing criteria within each RD may differ. Moreover, a RD may use routing metrics that are not

consistent with the QOS parameters conveyed within CLNP packets. Consequently, it may be impossible to optimise a given criterion all along the end-to-end path.

4.4.3.2 Overview of ISO 10589

ISO/IEC 10589. is for use within a single routing domain, and enables Intermediate Systems (ISs) to learn the topology of their local routing domain, and to identify the quality of service available over each potential path to a given destination.

ISs within a Routing Domain may discover each other dynamically using the ISO 9542 Intermediate System Hello message. They then use specific 10589 hello messages to determine each other's exact status.

The protocol supports a type of routing procedure known as a *link state routing*. In *link state routing*, Intermediate Systems broadcast information about their local environment to all other Intermediate Systems within the routing domain. Each system thereby builds up a complete "topological map" of the entire routing domain.

Under 10589, periodically, and whenever topology changes occur, each IS constructs a Link State Protocol Data Unit (LSP). This is then copied (flooded) to all other ISs within the same routing domain. Where possible, this is by direct transfer, but may involve ISs forwarding LSPs to other ISs, when ISs are not fully interconnected.

In general terms, an LSP identifies the generating IS's neighbour ISs (i.e. those which it has active communications links), the End Systems (ESs) to which the IS also has links, (discovered by ISO 9542) and the quality of service metrics pertinent to each link. Once an IS has available to it the current LSP from every active IS, it can construct the topological map of the routing domain, and then perform routing decisions using a suitable routing algorithm, such as "shortest path first".

Clearly, as the number of ISs and ESs increases, the overhead involved in LSP transfer will increase rapidly, and to ensure that the overhead does not become excessive the ISO standard structures a Routing Domain into one or more Routing Areas.

4.4.3.2.1 Routing Areas

A routing domain is made up of a set of routing areas, each characterised by a set of unique address prefixes known as the *area addresses*; all Network Addresses within the same routing area must be prefixed by one of these area address. When two ISs discover each other, they will determine whether or not they are in the same Routing Area.

Within a given routing area, each IS will generate an LSP specific to the routing area (Level 1 LSP), and flood it to all other ISs within the same routing area. This LSP identifies:

- the address prefixes of their local End Systems (and of the IS itself)
- the identity of adjacent ISs (i.e. those ISs in the local routing area with which the IS is in communication and can exchange ISO 8473 PDUs) and the associated quality of service parameters
- the identity of adjacent End Systems (i.e. those ESs in the local routing area with which the IS is in communication and can exchange NPDUs) and the associated quality of service parameters.

Through level 1 LSPs, each IS thus learns the current topology and connectivity of its local routing area. Note that Level 1 LSPs may be received from ISs in other routing areas, but these will be discarded when it is determined that there is no overlap in the area addresses covered.

Within each level 1 routing area some ISs also operate as level 2 routers, and identify themselves as such in their level 1 LSPs, and during the dynamic discovery phase.

Level 2 routers flood a second type of LSP (Level 2 LSP) to all other Level 2 routers in the routing domain (i.e. both within the local routing area and all other routing areas). A level 2 LSP identifies:

- the set of area addresses that characterise the local routing area
- the identity of adjacent level 2 ISs (i.e. the level 2 ISs in the routing domain with which the IS is in communication and can exchange NPDUs) and the associated quality of service parameters
- the address prefixes of any End Systems, or groups of End Systems, which are reachable through the level 2 IS, but are not included in the set of area addresses. These are typically address prefixes for destination in other routing domains and reachable through this IS.

Level 2 ISs are thus able to learn the current topology of the level 2 subdomain and hence the connectivity of level 1 routing areas. Access points to other routing domains are also identified. NPDUs destined for addresses outside of a local routing area, may be sent by a level 1 only IS to its nearest level 2 IS, and hence to a level 2 IS in the destination routing area, or to one from which the destination address is reachable. It may then be forwarded to the actual destination.

This two level hierarchy allows very large routing domains to be constructed. Most changes are typically limited to the local routing area, and only major changes affect level 2 routing, but without consequential level 1 LSP exchanges in other routing areas. The extent of routing information exchange is thus limited, with only a marginal effect on routing efficiency.

4.4.3.2.2 Partition Repair

Level 1 routing areas may become disjoint, either due to failures or mis-configuration, and 10589 has the ability to repair such failures by routing between level 1 routing area partitions through the level 2 subdomain. This is a necessary function since the level 1/level 2 structure is essentially an artificial one created to maintain efficiency, and it would be highly undesirable to prevent communication when a path exists and the only barrier to communication is a purely artificial constraint.

Partition repair is effected by the level 2 IS that is the partition designated intermediate system. All level 2 ISs within a non-disjoint level 1 routing area can identify each other through their level 1 LSPs, and rules exist to determine the partition designated intermediate system. Level 2 ISs report the current partition designated intermediate system for their local routing area in Level 2 LSPs.

If a partition designated intermediate system receives a level 2 LSP from an IS in the same routing area which reports a different partition designated intermediate system then a disjoint routing area is assumed. NPDUs to be transferred between the partitions are routed through each partition's partition designated intermediate system.

4.4.3.2.3 Support for Inter-Domain Routing

ISO 10589 also recognises that some ISs may also be Boundary ISs, that is they are at the periphery of Routing Domain and have links to other similar Boundary ISs in other Routing Domains. In order to support routing to such Boundary ISs, Level 2 LSPs may carry *Reachable Address Prefixes*. These are address prefixes that characterise the Routing Domains reachable through a given Boundary IS, and the intra-domain routing function is, using this information, able to route NPDUs addressed to systems in other Routing Domains, and via the appropriate Boundary IS.

4.4.4 IDRP Implementation Considerations

4.4.4.1 Overview

ISs within the same Routing Domain communicate with a high degree of mutual trust. They accept unquestioningly the routing information supplied to them, with the consequence that bad routing information will lead to routing problems. This is acceptable in this environment because all these systems will be under the same administrative authority. However, when "firewalls" are required between different parts of an Administrative Domain, or when communication between different Administrative Domains is necessary, then a different approach is required.

ISO/IEC 10747 specifies a routing information exchange protocol for use between Routing Domains i.e. when the environment is one of mutual distrust and/or when firewalls are required. The protocol does not operate between any IS, but only between specially designated *Boundary Intermediate Systems (BISs).* A BIS can be regarded as fulfilling the same role as the Internet's Exterior Gateway.

Multiple BISs within the same Routing Domain are permitted. Their behaviour is co-ordinated so that they operate as if they were the same BIS. A Routing Domain always provides consistent routing information regardless of how many BISs is supports.

The protocol - the Inter-domain Routing Protocol (IDRP) - is naturally connection mode and is specified to operate over ISO 8473. BISs *connect* to one another and exchange routing information over these BIS-BIS connections.

IDRP is a vector distant routing protocol. BISs advertise to another BIS, only the routes that they want to advertise to that BIS. The protocol is said to be policy driven, in that routes are only advertised when permitted by the effective Routing Policy, and contain only the information the Routing Policy allows to be advertised. IDRP is introduced in this section and presented in more detail in chapter five.

4.4.4.2 Routing Policy

Within an OSI RD, in general routing decisions are made on the basis of performance, taking into account the QOS available over a given subnetwork connection and the QOS required by the sender of an NPDU. However, routing between RDs is also subject to the imposition of Routing Policy, where a Routing Policy is a set of rules laid down by an Administrator responsible for a RD that primarily determine:

1. Whether the RD permits NPDUs for which neither the source nor the destination is in the RD to transit through the RD, and if so, the RDs to which transit facilities are offered;

2. The internal NSAP Addresses for which routes are advertised to adjacent RDs, and the scope of any further distribution.

Routing Policy is necessary because even when connectivity exits, when systems are owned by different organisations, those organisations will want to exercise control over the use made of connections so that only those users authorised to use a communications resource may do so, and that data only passes through physical systems and communications networks that are trusted to undertake the required task and provide the QOS demanded. For example, a CAA or Aeronautical Industry administrative domain may choose to restrict the outside ATN domains that may use its routing services based on security or other policy related requirements. In general, an ATN domain may receive Operational Communications, Administrative Communications and/or APC related traffic. Depending on its policies, a domain may choose to exclude the reception or transmission of these traffic types.

The overhead of routing policy is not always necessary, and that is why RDs exist. A RD is in general no more than a set of interconnected systems where routing may be performed on

performance considerations, and where a simple and robust intra-domain routing protocol may as a result be implemented.

4.4.4.3 BIS-BIS Communications

BISs exchange routing information in a pair-wise fashion. They use the services of ISO 8473 to communicate routing information; the BIS-BIS protocol includes procedures that ensure the reliable transport of routing information, including recovery from the loss of an ISO 8473 Data PDU. The BIS-BIS protocol is thus connection mode in operation and has similar features to the ISO 8073 Class 4 transport protocol.

BISs must establish BIS-BIS connections prior to the exchange of routing information. If more than one BIS is present in a RD then these BISs must form BIS-BIS connections with each other. The BISs within a RD form BIS-BIS connections with BISs in other RDs according to configuration information provided by a System Manager. When two RDs are linked by a BIS-BIS connection, then the RDs are said to be adjacent to each other. A BIS-BIS connection is established following the exchange of OPEN PDUs between two BISs.

Each BIS maintains two information bases per BIS-BIS connection. These are the adj-RIB-out and the adj-RIB-in. A BIS places the routes it wishes to advertise to another BIS in the adj-RIB-out. The BIS-BIS protocol then copies the contents of the adj-RIB-out to the corresponding adj-RIB-in in the remote BIS, and subsequently ensures that they remain identical. A BIS may then use the routes received into an adj-RIB-in as it wishes.

The BIS-BIS protocol uses the UPDATE PDU to copy routes from the adj-RIB-out. An UPDATE PDU may carry multiple routes and may advise on the removal or replacement of existing routes. When an UPDATE PDU is received, the BIS updates the appropriate Adj-RIBs-In. There is also a RIB REFRESH PDU for periodic re-synchronisation of the adj-RIB-out and adj-RIB-in.

The BIS-BIS protocol maintains the Adj-RIB-out and Adj-RIB-in synchronisation as long as the BIS-BIS connection exists. If the connection is lost then the associated information bases, and the routes are discarded.

The BIS-BIS protocol is full duplex and UPDATE PDUs are transferred in both directions. Contained in the UPDATE PDU is protocol control information to provide flow control and reliability through retransmission. When there are no routes to be exchanged, a separate KEEPALIVE PDU may be exchanged to keep the connection open. The BIS-BIS connection may be explicitly terminated through use of a CEASE PDU.

Routing Policy information is exchanged as part of a route to the extent of information limiting the scope of its onward distribution. However, the main impact of routing policy is on the manipulation of routes within a BIS.

4.5 SNDCF Implementation Considerations

The ATN specification is predicated on the use of the Connectionless Network Protocol (CLNP) specified in ISO 8473. CLNP provides the unifying end to end internetwork protocol. However, it is necessary to provide an adaption mechanism in order to use CLNP over each different type of subnetwork encountered. Such an adaptation mechanism is called a Subnetwork Dependent Convergence Function (SNDCF).

Such adaptations concen how CLNP packets are encapsulated for transmission over different types of subnetwork, and how ICAO specific requirements, such as priority are managed. On the receiving side, indications of subnetwork congestion may also be recorded by the CE-bit.

4.5.1 SNDCFs for Fixed Data Networks

ISO/IEC 8473 provides serveral standard SNDCFs for use with common subnetwork types including IEEE 802 compatible LANS and ISO 8208 WANs. These SNDCFs should be used whenever possible. Industry standard approaches have also been developed for other subnetwork types including Frame Relay and these should be used whenever possible. ICAO has also developed the specification for use of CLNP over the ICAO CIDIN subnetwork.

4.5.2 The Mobile SNDCF

The mobile networks are a key component of the ATN. Air Traffic Control (ATC) applications require a data link between an Air Traffic Control Centre and each aircraft under its control; this requirement is satisfied by the mobile networks. However, the usable bandwidth of each mobile network is low (of the order of 2400 bits/s or lower). ATC applications tend to consist of the regular exchange of short messages and, in such an environment, the size of the CLNP header becomes a serious overhead. Considering this, ICAO has developed a set of procedures, and supporting protocol, to provide compression of CLNP headers over low bandwidth data links.

Several different compression mechanisms are available for use over low bandwidth subnetworks and the Mobile SNDCF provides a common specification for the negotiation of an appropriate set of compression mechanisms for a given data link. The available compression mechanisms are:

a) The Local Reference (LREF) compression mechanism. When this specification is used, a CLNP header of the order of sixty octets can be compressed down to at most fourteen octets.

b) The ICAO Address Compression Algorithm (ACA). This is a stream based compression mechanism that identifies NSAP Addresses within a data stream and removes redundant data within each NSAP Address.

c) ITU-T recommendation V.42bis compression. This is an adaptation of the stream based LZW algorithm for compressing data streams by the replacement of commonly occuring strings with shorter symbols.

4.5.2.1 Overview of LREF Compression Algorithm

LREF compression is specified for use over a reliable virtual circuit. It is a dictionary based compression algorithm that replaces the NSAP Address pair and ATN Security Label in a CLNP header with a single integer (the local reference). Separate directories and hence local references may be used for each virtual circuit, or for groups of virtual circuit. The compression algorithm is described as follows.

Whenever a CLNP packet is queued for transmission over the virtual circuit, the local directory for that virtual circuit is queried to see if an entry exists for which:

a. the outward NSAP Address is identical to the packet's destination NSAP Address, and

b. the inward NSAP address is identical to the packet's source address, and

c. the protocol version number is the same as that contain in the packet header, and

d. either the security parameter is absent in both cases, or the security parameter in the directory is identical to that in the packet header.

If the above condition is satisfied, and the packet header does not contain the source routing or route recording optional parameters, or more than seven octets of padding, then the CLNP packet may be replaced by a compressed header.

The actual format of the compressed header is dependent on whether the segmentation part is present in the original packet header and, if so, whether the packet is a derived or initial PDU. In all these cases, the compressed header includes the priority (if present) and the QoS Maintenance bits (if present) in a packed form, and the local directory entry number, as the "local reference" field. The segmentation part, when present, is copied unchanged in to the compressed header.

When a packet with a compressed header is received, the local reference is extracted and the corresponding entry found in the local directory. The original PDU header is then reconstructed from the information contained in the local directory and the compressed header.

Note that the reconstruction of the packet header does not aim to restore the padding octets, if any, to their original values. For such reasons the algorithm is not applied to CLNP packets encapsulated by a security protocol such as NLSP, which generates an integrity check on the entire packet.

If, when a CLNP packet with a compressed header is received, the indicated local directory entry does not exist, then this is an error condition reported to the peer SNDCF by the local management protocol. An SNDCF Error PDU is specified for this purpose.

4.5.2.1.1 Creating Local Directory Entries

A local directory entry is created when a CLNP packet is queued for transfer over the virtual circuit and no suitable entry could be found in the local directory. An entry is then created using the source and destination NSAP Address (inward and outward NSAP Addresses), protocol identifier, and security parameter (if present) in the packet header. Each side of the connection has a range of entry numbers (local references) which it is permitted to allocate, and a suitable (unused) entry number is selected from that range, to correspond to the newly created directory entry.

The allocated directory entry number is then inserted into the packet header as a new optional parameter, and the packet header and segment lengths and header checksum adjusted to ensure that the header is syntactically correct. The packet is then transferred over the virtual circuit.

Whenever an uncompressed CLNP packet is received over a virtual circuit supporting the Mobile SNDCF, its header is inspected for the addition of such a local reference parameter. If found it is removed, the header and segment lengths and checksum adjusted appropriately, and a local directory entry created for that local reference using the source and destination NSAP Address (outward and inward NSAP Addresses), protocol identifier, and security parameter (if present) in the packet header. By such a mechanism the local directories are synchronised. As the definition of the inward and outward NSAP Addresses is asymmetric, a local reference may be used in either direction with the same, albeit reversed, semantics.

Once a local directory entry is created, it remains valid for the lifetime of the virtual circuit; the local directory is disposed of when the virtual circuit is cleared. Communication over mobile subnetworks is typically for a limited period, and directory sizes can generally be chosen such that there is sufficient capacity available for the lifetime of the virtual circuit. If the directory becomes full then packets between further NSAP pairs are simply sent uncompressed.

However, it is possible that in some circumstances, the communications path may be long lived and it will be necessary to re-use directory entries. To satisfy such requirements, the use of the local reference cancellation mechanism may be negotiated when the connection is established.

4.5.2.1.2 Re-use of Directory Entries

Two local management protocol packets are specified for this purpose. A local reference cancellation request PDU enables one side of the virtual circuit to identify a range of local references (under its control) that it wants to cancel, and hence make available for re-use. When such a PDU is received, the identified local references are cancelled, and a response PDU returned. Once a

response PDU has been received by the initiator of the cancellation request, then the local references can be re-used.

Certain error conditions may indicate that the local directories at each end of the virtual circuit have lost synchronisation. if this situation occurs then the virtual circuit is reset, and the local directories returned to their initial state.

4.5.2.2 Implementation Model

The current generation of ICAO Mobile Networks all provide a network access service compliant with ISO 8208 (ITU-T recommendation X.25). The CLNP specification already provides a set of procedures for passing CLNP packets over X.25 virtual circuits; ISO 8473 defines such procedures as a Subnetwork Dependent Convergence Function (SNDCF). The procedures for compression of CLNP headers over ICAO Mobile Subnetworks are based on the X.25 SNDCF, and indeed may be negotiated down to this SNDCF. The specification of these procedures is known as the Mobile SNDCF.

The implementation model for the Mobile SNDCF is illustrated in Figure 4-9 Implementation Model of the Mobile SNDCF. Note that the specification is not necessarily restricted to X.25. In

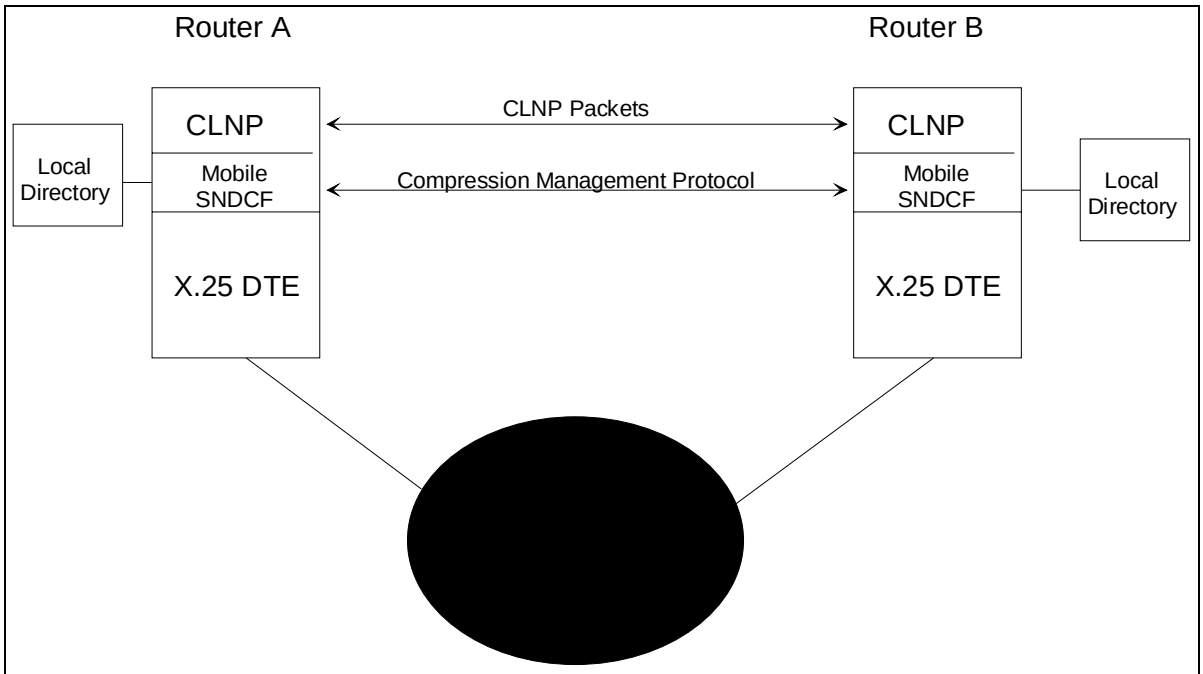

Figure 4-9 Implementation Model of the Mobile SNDCF

principle, this specification may be readily adapted to any connection mode data link.

The compression procedures are assumed to be implemented over a single data link between two routers, or a host and a router. In very simple topologies, they could be implemented between two hosts. The figure illustrates the typical case, which is between two routers, with the illustration of each router simplified such that only a single subnetwork stack is shown.

From an architectural perspective, the CLNP implementations in each router exchange CLNP Data and Error Packets over an X.25 virtual circuit using the procedures specified by the Mobile SNDCF. In addition, the implementations of the Mobile SNDCF also need to exchange information related to

the management of the compression algorithm. A local management protocol is specified for this; this protocol is passed over the same virtual circuit as are CLNP packets with compressed headers.

Note that the format of the compressed headers is such that they can be distinguished from normal CLNP packets, and well as IS-IS, ES-IS and NLSP packets, and the local management protocol.

In each router, the Mobile SNDCF maintains a local directory for use by the compression algorithm. A separate local directory is maintain for each virtual circuit over which CLNP header compression is in use. This is true even when more than one virtual circuit is concurrently available to the same router or host. The local directory contains the state information specific to the operation of the compression algorithm over a single virtual circuit, and the prime purpose of the local management protocol is to maintain synchronisation of the local directories at each end of a virtual circuit.

The local directory consists of entries numbered from zero to a maximum of 32767, each entry consisting of:

1. A pair of NSAP Addresses, known as the inward and outward NSAP Addresses respectively;

2. The ISO 8473 protocol version number;

3. The value of the security options parameter (see ISO 8473 Clause 7.5.3), which may be empty;

The directory is initially empty. The minimum directory size that may be supported is 128 entries.

Note that the algorithm is suitable only for uses of the security parameter that support "simple security", such as passwords or simple traffic class identifiers, which are likely to be constants for packets sent between the same NSAP pair. It is not suitable for "strong security" where the security parameter contains a checksum (encrypted or otherwise) binding the contents of the security parameter to the packet’s user data.

5. ATN Routing

5.1 Introduction

Within a Routing Domain, there are no special routing requirements for the ATN. Standard routing protocols, such as ISO/IEC 10589 may be used unmodified and the only problem that implementators are likely to encounter is the presence of the ATN Security Label. Some vendors products may not be able to handle this without modification to the product. However, many commercially available products can be configured to ignore a CLNP Security Parameter when present. Such a feature is essential for use with the ATN and routers within an ATN Routing will typically be configured to ignore the CLNP Security Parameter and hence the ATN Security Label.

However, routing between ATN RDs does need to consider ATN requirements and, generally, specially adapted ATN Routers will need to be used. In many cases, this adaptation is no more than the capability of using IDRP with the ATN Security Label. However, those routers that occupy key ATN roles, such as Air/Ground Routers and ATN Backbone Routers will also need to handle and apply ATN specific Routing Policies, in order to support routing to mobile systems.

This chapter is concerned with describing how IDRP works, how it is used in the ATN to support routing to both fixed and mobile systems, and the routing policies that have been adopted.

5.2 Background to IDRP

The OSI Routing Architecture described in ISO TR 9575 describes a routing architecture in which there are three different sets of requirements for routing protocols:

1. There is a need for the communication of routing information between End Systems and Intermediate Systems. This requirement is satisfied by ISO/IEC 9542.

2. There is a need for the communication of routing information between Intermediate Systems within the same Routing Domain. This requirement is satisfied by ISO/IEC 10589.

3. There is a need for the communication of routing information between Intermediate Systems in different Routing Domains. It is to satisfy this requirement that IDRP was developed.

In fulfilling the role of an inter-domain routing protocol, IDRP has to exchange routing information in what is described as a domain of limit trust. The information exchanged has to be limited to the minimum necessary to advertise the existence of a route without revealing any more about the internal topology of a Routing Domain, or its connectivity with other RDs. Furthermore, the information received by IDRP has to be interpreted according to local policy rather than accepted at face value, and the decision on whether to advertise a route is a matter of policy.

Scaleability is also a major consideration behind the development of IDRP. The inter-domain routing environment can potentially grow without limits, and IDRP must be able to cope with this without imposing limits on the growth of the internetwork.

In addition to meeting the requirements of ISO TR 9575, the ISO 10747 Inter-Domain Routing Protocol was also heavily influenced by the work done on policy based routing in the TCP/IP Internet and, as such is a direct descendant of the Border Gateway Protocol (BGP) family of routing protocols used between Internet Service Providers and large users.

5.3 Choice of IDRP for the ATN

IDRP was chosen for ATN use early on in the development of the ATN SARPs. At that time, it was still a draft standard and the aeronautical community was able to influence the development of IDRP in order to ensure that it fully met ICAO requirements.

IDRP was chosen because a need was identified for a routing protocol to support the routing of data to mobile systems wherever they may be. Such a protocol was required to:

a) Work in an environment comprising many different Service Providers, Administrations and other Organisations, both co-operating and competing to provide services to the aeronautical community.

b) Be reliable, with no single point of failure and permitting the concurrent availability of multiple alternative routes to a given mobile system.

c) Track the changes in connectivity and hence paths to mobile systems, in a timely manner, meeting the requirements of aeronautical applications.

d) Permit the operation of various organisational policies including control over the use of air/ground datalinks and controlled use of different ground data links by different classes of traffic.

Both Link State and Vector Distant models of routing information exchange protocols were studied. However, the Link State model was quickly rejected given the low bandwidth available on air/ground data links and the high amount of traffic expected with Link State Routing Protocols. On the other hand, the Vector Distant model appeared well suited to low bandwidth links as, in principle, only the minimum amount of routing information need be exchanged.

IDRP was specified as a vector distant protocol and had been designed to support multiple alternative routes and policy based routing. However, it lacked a mechanism to support choices of data links based on organisational policy. This required extra information to be carried in each route, and, following ICAO representations to ISO, a general purpose mechanism was added in the form of the Security Path Attribute. IDRP then fully met ICAO requirements for the ATN routing protocol and was adopted as such.

5.4 IDRP Overview

IDRP is a routing information exchange protocol that supports:

- The advertisement to routers in another Routing Domain of routes to local destinations;
- The re-advertisement of routes received from routers in other RDs to a router in another Routing Domain;
- The policy based interpretation of routing information received from other routers including a decision on a choice between alternative routes to the same destination;
- Policy based control over the advertisment and re-advertisement of routes;
- The realisation of large scaleable Internetworks.

IDRP is architecturally described by the protocol and process models shown respectively in Figure 5-1 and Figure 5-2, respectively.

As a routing information exchange protocol, IDRP is always implemented on an Intermediate System (IS). Further, such an IS is always at the boundaries of a Routing Domain and is therefore said to be a Boundary Intermediate System (BIS). The IDRP entity on a BIS may communicate with many other BISs simultaneously, both within its own Routing Domain, and in other RDs. This communication follows the connection mode i.e. the reliable exchange or routing information is support within the context of an agreed association supporting both flow control and error recovery, and is supported by a specially defined BIS-BIS protocol. The BIS-BIS protocol is a simplified version of the ISO Class 4 Transport Protocol, and uses the services of the Connectionless Network Protocol (CLNP) for data transfer between two Adjacent BISs.

Clearly, the BIS must have a way of routing CLNP PDUs to adjacent BISs that is not dependent upon IDRP routing information exchanges, and this imposes limitations on the interconnection scenarios for BISs. Within a Routing Domain, another routing information exchange protocol (such as ISO/IEC 10589) can be assumed to be available and hence the only requirement is that a path exists between two BISs; any number of subnetworks and routers may be traversed as long as the route is navigable using ISO/IEC 10589. However, between RDs, no such routing information exchange protocol is available. IDRP can therefore only be used to communicate between BISs in different RDs, when such BISs are directly interconnected by a real subnetwork (e.g. a leased line, X.25 virtual circuit, etc.), although a single IDRP adjacency may be supported by several subnetwork connections in parallel.

The CLNP forwarding information necessary for BIS-BIS communication is typically either configured into a router as a static route by a System Manager, or by using the “External Reachable Addresses” as defined in ISO/IEC 10589.

The messages exchanged by the BIS-BIS protocol are typically used to advertise one or more routes, where a route is said to comprise:

a) a set of destinations

b) information describing the path to such destinations.

These routes are then processed by IDRP both for use in local routing of data, and for re-advertisement to other BISs.

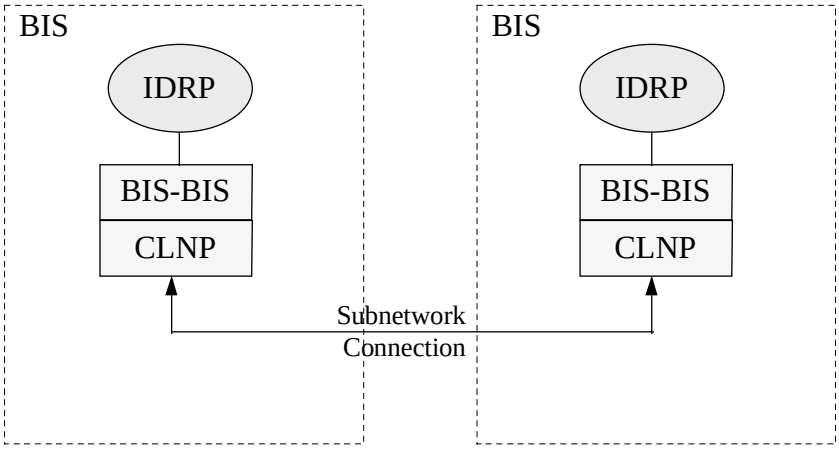

Figure 5-1 IDRP Protocol Model

IDRP process the routes it receives from adjacent BISs (and locally provided routes) according to the process model shown in Figure 5-2.

5.4.1 The Abj-RIB-in

All routes received from an adjacent BIS are first recorded in an input database known as the Adj-RIB-in, where there is a different Adj-RIB-in for each adjacent BIS with which the BIS is in communication. Indeed, there may even be multiple Adj-RIB-ins for a given adjacent BIS, when more than one set of "distinguishing path attributes" is supported (see 5.5). In such cases, there is a separate set of routes for each set of distinguishing path attributes.

The routes received from adjacent BISs are then processed by a Route Decision process. This acts upon all routes received so far. The Route Decision process firstly copies all routes received from BISs in other RDs to all the BISs in its local Routing Domain. This process is known as internal distribution and ensures that all BISs within a single Routing Domain share a common view of the outside world. Of course, it is possible that there may be two or more routes in different Adj-RIB-ins to the same destination. In such cases, the Route Decision process chooses the most preferable for internal distribution and ignores the rest.

The mechanism by which the most preferable route is computed is essentially a local matter, and is the first instance where we see the notion of *policy* appearing in IDRP. Local policy determines the order of preference of otherwise equal routes, and may even exclude certain routes because they are, perhaps, deemed unreliable, too costly, or there is no contractual agreement for their use.

5.4.2 The Loc-RIB

The Routing Decision process then selects routes for local use. This includes routes received from BISs in the local Routing Domain, external RDs and local routes provided either by a System Administrator or by intra-domain routing. This decision process is much like that described above where local policy is used to discriminate between routes to the same destination. The difference is in the scope of the routes acted upon and, in this case, the set of selected routes is placed in the Loc-RIB. The Loc-RIB is the Local Routing Information Base and there is one Loc-RIB for each set of distinguishing path attributes supported.

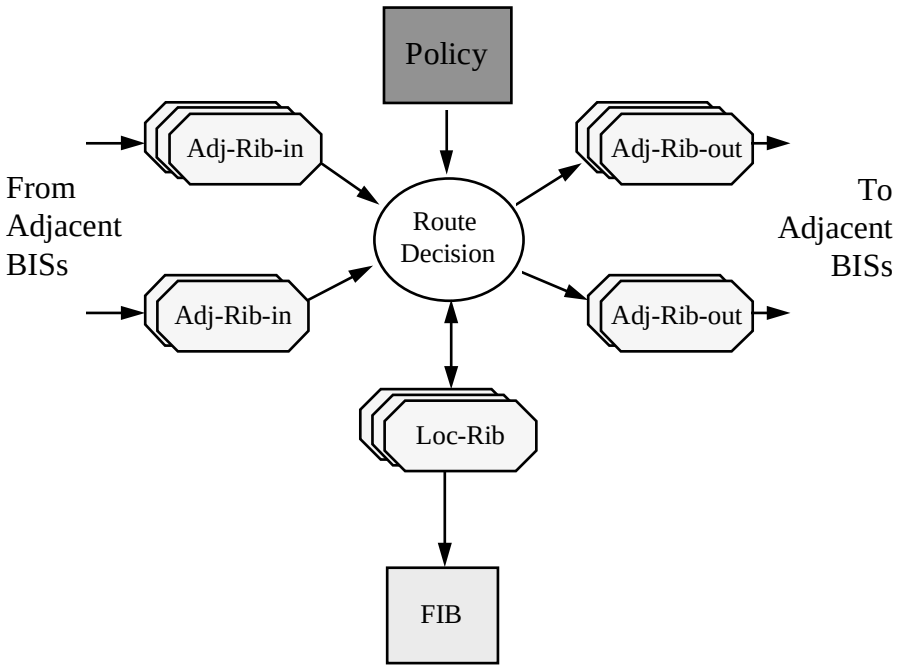

Figure 5-2 IDRP Process Model

The routes in the Loc-RIB are used to generate information for the BIS's Forwarding Information Base (FIB). This is the data structure used by CLNP for forwarding PDUs and it should be noted that IDRP is not the only source of information for the FIB. Intra-domain routing and the System Manager are other possible sources.

5.4.3 The Adj-RIB-out

The Loc-RIB is also the primary source of the routing information advertised to BISs in other RDs. For each known adjacent BIS a further set of routing policy rules has to be defined that determine which routes are selected from the Loc-RIB(s) for advertisement to each adjacent BIS. For each BIS, the selected routes are copied to another database - the Adj-RIB-out. From here, they may be advertised to the remote BIS. This process is known as external distribution and contrasts with the internal distribution mechanism used to copy received routes to BISs in the same RD.

As a minimum, the routes to local destinations are selected from the Loc-RIB(s) and copied to the Adj-RIB-out(s). A BIS's routing policy rules may also select routes received from BISs in other RDs and re-advertise them to a BIS in another Routing Domain. In the former case, the BIS does not, in consequence, offer any transit facilities for routing between other RDs, and the local Routing Domain is hence known as an End Routing Domain (ERD). In the latter case, transit facilities are offered and the local Routing Domain is known as a Transit Routing Domain (TRD).

It should be noted that the ATN explicitly prohibits the re-advertisement of routes where it is clear by examining the route's trace information that, to do so, would constitute a routing loop. This is very important as validation work has shown that if this is not done, false routes can be generated that persist for a lengthy period.

5.4.4 Route Aggregation

A further feature of IDRP is Route Aggregation. This is when routes in the same Adj-RIB-out are grouped together prior to their advertisement to another BIS, and merged or aggregated to form a single route. The routes that are to be aggregated are selected by Routing Policy, although the actual process itself is algorithmic and fully defined in the standard.

The benefit of this process is that it reduces the number of routes that need to be advertised to another BIS which, in turn, reduces the overhead of routing information exchange, and is an important contribution to ensuring scaleability. This is because, if an internet is to grow without bounds, then the amount of routing information that a sender needs to know about a given destination should decrease, the further away that destination is from the sender. Essentially, the granularity of routing information should get coarser as it is advertised from BIS to BIS, and Route Aggregation is the first stage in this process, reducing the number of routes advertised.

Route Aggregation is also automatically performed when more than one route is selected from a Loc-RIB that has identical NLRI (i.e. they have the same destination). For example, this will occur when two routes to the same destination have different security path information. In order to avoid the implementation of the full Route Aggregation procedures in routers that do not otherwise need them, the ATN SARPs have specified a simplified procedure known as *Route Merging*. This procedure is only appropriate when aggregating routes with identical NLRI and avoids have to implement the aggregation rules for the aggregation of the route trace information.

5.4.5 Route Information Reduction

The reduction of routing information is then completed by another process, known as Route Information Reduction. Route Information Reduction is again policy based, and is a mechanism by which the set of NSAP Address Prefixes that describe the destination of a routes is replaced by a set of shorter NSAP Address Prefixes. Typically, a whole set of prefixes is replaced by a single NSAP Address Prefix, and the policy rule that specifies such a replacement has been formulated taking

account of the known distribution of NSAP Addresses in a given part of an internet. Provided that NSAP Addresses have been allocated such that RDs that share common (shorter) NSAP Address Prefixes, are closer together in network topology terms, than RDs that are further apart, then Route Aggregation and Information Reduction rules can be formulated, that aggregate many routes together into a single route to a whole region of the internet, thus enabling the important objective of scaleability.

5.4.6 Routing Domain Confederations

The last important feature of IDRP worth describing is the Routing Domain Confederation (RDC). This is a generally useful concept that also helps in building scaleable internets. An RDC is simply a named set of RDs, and the formation of an RDC is done by mutual agreement. RDCs may contain RDs and RDCs and may be both nested and overlapping.

Routing Policy rules may reference RDCs as a convenient way of referring to groups of RDs in Routing Policies. However, their most important use is in providing well defined containment boundaries for Route Aggregation and Information Reduction, and in reducing the trace information that IDRP appends to every route. As containment boundaries, RDCs can readily identify the groups of RDs that share a common NSAP Address Prefix, and, ideally, an RDC boundary is positioned where Route Aggregation and Information Reduction is to be performed, enabling both a reduction in the number of routes, while ensuring minimal trace and addressing information.

5.5 The ATN Security Path Attribute

In IDRP, the information that describes a route, in addition to a route's destination, is known as the path information. In turn, the path information consists of a set of path attributes which provide information on, for example, where the route passes through (trace information), restrictions on to which RDs a route may be passed, and information about the Quality of Service available over the route, protection offered and access rights. The Quality of Service and Security path attributes are known as the distinguishing path attributes, as routes that have different combinations of such attributes but share the same destination, are still regarded as different routes.

The reason for this is to enable routers to make available routes to the same destination that may offer a different Quality of Service, or different Security. When NPDUs are forwarded, the sender's request for a distinct grade of service may then be matched with the routes available, and the most appropriate chosen. In IDRP terms, each distinct set of Distinguishing Path Attributes is known as a RIB attribute set or just *RIB-Att*. Each RIB-Att is regarded as describing a completely different domain of routes and a BIS will maintain a separate Loc-RIB for each RIB-Att it supports. Similarly, a distinct Adj-RIB-in and Adj-RIB-out is maintained for each RIB-Att in common with a given remote BIS.

The ATN does not make use of the Quality of Service path attributes. However, it does use the IDRP Security Path Attribute and uses this to label a route with information used to satisfy various user policies. ATN Routers therefore support two distinct RIB-Atts: the so called *empty RIB-Att* for routes that have no security information (and no other distinguishing path attributes), and a Security RIB-Att for those that do have security path information. The former routes are only used for General Communications, while the latter routes are used for both ATSC and AOC applications data.

The ATN Internet SARPs specify that the Security Information contained within an IDRP Security Path Attribute is used to convey information about the type of traffic that a route can carry and the Air/Ground Subnetworks that the route may pass over. This information is provided by two fields or *Security Tags*:

- The Air/Ground Subnetwork Security Tag, and
- The ATSC Class Security Tag.

5.5.1.1 The Air/Ground Subnetwork Security Tag

This tag is added to a route's security path information, whenever a route passes over an Air/Ground Data Link. The tag records the type of air/ground data link (e.g. Mode S, AMSS, etc.) and the traffic types of data that can pass over the data link (e.g. ATSC, AOC, etc.). If more than one type of Air/Ground Data Link concurrently supports access to the same aircraft, then a tag is added for each such data link.

This Security Tag is used:

a) to support the AOC user routing policy requests. These allow an application to specify which Air/Ground subnetwork type, out of those available, is used to convey the data, between air and ground. Such requests are also handled in a "strong" manner. That is, if the requested Air/Ground subnetwork type is not available, then the data is discarded.

b) to avoid data of a given traffic type and addressed to an airborne system, being routed to an Air/Ground Data Link that does not support the uplink of data of that type.

This Security Tag will only be found in routes to aircraft. It is never present in routes to ground destinations except in an Airborne Router. This includes routes that will be used by data that originated in an aircraft, has been downlinked to an Air/Ground Router, and is now in the ground portion of its journey. It cannot therefore be used as a general mechanism for determining the traffic types of data that may pass over a given route.

5.5.1.2 The ATSC Class Security Tag

This tag is added to a route when that route has been approved for ATSC data, and, additionally, identifies the ATSC Class supported. The tag is added when a route is created. It can be removed, or the ATSC Class reduced, but it can never be added to an existing route, nor can the ATSC Class be increased. The actual encoding of the ATSC Class is a bit-map, so that when routes to the same destination are aggregated, all supported ATSC Classes can be identified in the aggregated route.

This tag is used to support ATSC User specified routing policy requests. When data has a traffic type of ATSC, it can only be routed over an ATSC approved route, and this requirement is met by only forwarding such data over a route with an ATSC Class Security Tag present. Furthermore, when more than one possible route is available, the route is chosen that either:

a) Supports the same ATSC Class as indicated in the data's security label; or, if no such route can be found

b) Supports a higher ATSC Class; or, if no such route can be found

c) Supports a lower ATSC Class.

Editor's note: the following paragraph assumes the adoption of the change proposal in support of the WG3 requirement for routes that have an ATSC only semantic.

Two variants of the ATSC Class Security tag are specified, each providing a different semantic. The two semantics are:

- The route is available to both ATSC and non-ATSC data.
- The route is available to ATSC data only.

The value of the ATSC Class Security Tag may be modified en route in order to reflect local policies about the ATSC class support by a given data link and the type of traffic that may be carried over a data link. Such modifications always one way in that the class may be lowered and the data conveyed

made more restrictive, but the reverse is not permitted in order to avoid routing "black holes" developing.

5.6 The BIS-BIS Protocol

BISs communicate using a network layer protocol specified in ISO/IEC 10747. This is a connection mode protocol that uses ISO 8473 to communicate between BISs over both real and virtual (i.e. via one or more ISs) links.

The purpose of this protocol is to permit the reliable exchange of routes, between a pair of BISs. A route is passed between two BISs as the information content of an UPDATE BISPDU, which is itself transferred as the contents of a single ISO 8473 DT PDU. Routes once advertised may also later be withdrawn by another UPDATE BISPDU.

The BIS to BIS protocol itself is concerned with the reliable transfer of UPDATE BISPDUs.

5.6.1 BIS-BIS Connections

UPDATE BISPDUs may only be transferred when a connection is said to exist between a pair of BISs. A BIS-BIS connection may only be established when explicitly permitted by Systems Management action at both BISs, and once permission has been granted, an exchange of OPEN BISPDUs (again as the contents of a single ISO 8473 DT PDU) initialises the connection.

The OPEN BISPDUs enable the BISs to identify and authenticate each other; to identify the RDCs of which they are both members; and to identify the sets of distinguishing path attributes that they each support. Note that the exchange of OPEN BISPDUs is a symmetric process and only a single BIS-BIS connection results, even when two BISs simulteously issue an OPEN BISPDU.

Once a BIS-BIS connection is open, UPDATE BISPDUs may then be exchanged in order to enable one BIS to advertise routes to the other. Each UPDATE BISPDU carries sequencing and acknowledgement information in its header which enables each BIS to detect packet loss and bring about retransmission of lost UPDATE BISPDUs, and to support flow control between BISs.

As long as routes are being exchanged in both directions then all the protocol information necessary to maintain reliable communication is transferred in the header of the UPDATE BISPDU. However, if a BIS has no more routes to advertise, then the protocol provides what is known as the KEEPALIVE BISPDU. This permits protocol information to be exchanged in order to keep the connection open and permit data flow in one direction, when there is no data to send in the other. It is very similar to an UPDATE BISPDU, except that it consists purely of a protocol header and carries no data (i.e. a route).

The BIS-BIS protocol also includes an IDRP ERROR BISPDU to enable protocol errors to be reported from one BIS to the other, and a CEASE BISPDU in order to terminate a BIS-BIS connection.

5.6.2 RIB Refresh

Once routes are received by a BIS, as discussed above they are entered into the appropriate Adj-RIB-in. The Adj-RIB-in is constantly being updated as new routes are received and old ones are withdrawn. When BIS-BIS connections are long lived, there is the possibility that undetected errors may occur, and so that errors are not perpetuated, the BIS to BIS protocol permits what is known as a RIB Refresh.

A RIB Refresh consists of the transfer of a series of UPDATE BISPDUs corresponding to all the current routes advertised by the BIS providing the Refresh (i.e. the contents of the Adj-RIB-outs associated with the BIS-BIS connection), and delimited by the RIB REFRESH BISPDU, which is

part of the BIS-BIS protocol. During a refresh, the receiving BIS may compare the received routes against the RIB, and rectify any discrepancies.

A RIB Refresh may be performed automatically by the "refreshing" BIS, or solicited by the one receiving the refresh, again using the RIB Refresh BISPDU.

5.6.3 Route Combination

Route Combination is the combination of two or more routes into a single UPDATE BISPDU and is an optimisation intended to reduce the number of BISPDUs exchanged between two adjacent BISs. The principle is that when a BIS has two or more routes that need to be advertised to an adjacent BIS, and when these routes have the same NLRI, but different sets of distinguishing path attributes, then they may be combined into a single UPDATE BISPDU, which encodes common path attribute values once and once only for each combined route. By the same process, Route Withdrawals may also be included in the same UPDATE BISPDU as a newly advertised route.

When aggregated routes are modified such that the NLRI changes, the original aggregated route has to be formally withdrawn and its replacement advertised as a new route. To prevent discontinuities in the availability of the aggregated route, it is important that the withdrawal of the older route and its replacement take place simultaneously, otherwise the availability of the remainder of the aggregated route will be discontinuous with the risk of temporary loss of communications. Route Combination, in this case combining withdrawals and updates together, is thus essential to the proper operation of Route Aggregation.

5.6.4 Authentication and Security

Physical Security measures protecting ATN Routers subnetworks, and other components from attacks, including unauthorised access and physical attacks, will need to be employed by Administrations and other Organisations. Each will need to consider what measures are appropriate to local circumstances. Such mechanisms will be necessary to protect against Denial of Service attacks.

Encryption of data links may also be considered as a means of preventing unauthorised access, especially to prevent Denial of Service by preventing unauthorised access to routing information, and hence unauthorised modification of routing information. Such mechanisms may also be used to protect against the injection of unauthorised messages, although application specific mechanisms will probably be more appropriate for this.

However, when public data networks are used, or when mobile subnetworks using free radiating media, then protocol specific mechanisms are required in order to protect against unauthorised access. This includes authentication mechanisms used to protect against access by unauthorised users. In order to protect the routing information base, authentication of the provider of IDRP routes is viewed as extremely important.

The IDRP protocol supports a range of authentication mechanisms (referred to as authentication types 1, 2 and 3) implemented on a per BISPDU basis. Authentication type 1 provides an unencrypted checksum on each BISPDU, and so is not secure, although it gives protection against arbitrary errors. Type 2 provides protection against masquerade and modification by use of a checksum on each BISPDU which is encrypted using a mutually agreed encryption algorithm. Authentication type 3 uses a “validation field” in each routing protocol exchange to carry a Message Authentication Check (MAC), generated from an agreed password.

The ATN SARPs currently require type 1 authentication. However, it should be noted that this is not believed to be adequate to protect against threats to the routing information base, resulting from unauthorised access. Type 2 authentication is necessary for this, and may be mandated on a regional basis where it is believed that such a threat exists, together with an appropriate security mechanism, such the Digital Signature Standard specified in FIPS Pubs 186 and 180. No additional protocol

overhead is necessary to support type 2 authentication. The field used to convey the authentication information for type 2 authentication is also used for type 1 authentication.

Appropriate security mechanisms will also require the distribution and use of encryption keys. Key Management may be considered as a bilateral matter for ground-ground connections. For Air/Ground connections, a common approach will need to adopted in each region requiring type 2 authentication. For example, a single secret key may be used per region, and regularly changed (e.g. daily). However, in the future, it may be necessary to move to a key per aircraft, if the threat increases in significance.

5.7 The Route Decision Process

The IDRP Routing Decision Process is described as a three phase process, where each phase is, respectively, concerned with:

- The Selection of routes for Internal Distribution
- The Selection of Routes for Local Use
- The Selection and update of routes for External Distribution.

Each of these three phases is described below.

5.7.1 The Phase One Decision Process

The Phase One Decision Process acts on all newly received routes, and on all received indications of the withdrawal of an existing route. For each new route, it computes a degree of preference according to a local policy algorithm. If that route has the highest degree of preference out of all known routes to the same destination and same set of distinguishing path attributes, and it was received from a BIS in a different Routing Domain, then the route is copied to the Adj-RIB-out associated with each BIS in the local Routing Domain, for internal distribution to those BISs. By this means all BISs in the local Routing Domain are kept up-to-date about the availability of the preferred route to each destination. There is no need to similarly copy routes received BISs in the local Routing Domain, because all such BISs are assumed to be in direct communication and will receive such a route direct from the local BIS from which it came.

Similarly, if the withdrawal of a previously preferred route is received from a BIS in another Routing Domain, then that withdrawal is immediately copied to all other local BISs, so that they too may be made aware of the loss of such a route. An alternative but previously lower preference route may exist in another Adj-RIB-in and, if so, that route now becomes the preferred route and is copied, as above, to the Adj-RIB-out associated with each BIS in the local Routing Domain.

The Phase One Decision process also provides an opportunity for BISs in the same Routing Domain to check the consistent application of the local route selection policy. The computed degree of preference is passed with each route as part of the internal distribution procedure and is checked by phase one whenever it computes the degree of preference for a route received from a BIS in the local Routing Domain. Any lack of consistency is reported to Systems Management.

Note that there are also special rules for handling the security path attribute. Although there is only one Security RIB-Att, routes with different values if the Security Path Attribute satisfy different user policies and one cannot said to be preferable to the other. Because of this, when operating on routes under the Security RIB-Att, phase one will select the most preferable route for each destination and each value of the security path attribute for internal distribution.

5.7.2 The Phase Two Decision Process

The Phase Two Decision Process is responsible for choosing the routes to be made available for local use in the Loc-RIB. Essentially, the preferred route to each destination and for each RIB-Att, identified by phase one is copied into the corresponding Loc-RIB. Under the Security RIB-Att, the same special rules apply, and the Loc-RIB for the Security RIB-Att may include several routes to the same destination. In each case, these will be the preferred route for a given value of the security path attribute.

Indications of route withdrawal are also processed by the Phase Two Decision process. Withdrawn routes are removed from the appropriate Loc-RIB, and may be replaced by an alternative route to the same destination, if one is available.

5.7.3 The Phase Three Decision Process

The Phase Three Decision Process is responsible for selecting routes for External Distribution, and for the aggregation of certain groups of routes, and the application of Route Information Reduction. A process model for the IDRP Phase 3 Route Decision Process, including Route Information Reduction and Route Aggregation, is illustrated in Figure 5-3. This illustrates the data structures and processes needed to implement the Route Decision process.

Two PIB data structures are referenced: a list of “Route Selection Rules” and a list of “Reduction Rules”. The former is used for grouping routes together for the purposes of Route Aggregation, while the latter is for determining when Route Information Reduction of NLRI can be performed. In both cases, it will be necessary for the implementor to define a syntax to enable the text based definition of the rules, so that these data structures may then be created at system start up.

A “Route Selection” process is then specified to pass through the Loc_RIB applying first type 1 selection rules, and then applying type 2a and 2b selection rules to any routes in the Loc_RIB not selected by a type 1 rule. The rule types are defined as follows:

- A Type 1 rule is a rule that selects routes for aggregation i.e. all routes selected by a given type 1 are aggregated before being copied into the Adj-RIB-out.
- A Type 2a rule is an unconditional rule for which each route selected by such a rule is copied as an individual route into the Adj-RIB-Out, and
- A Type 2b rule is a conditional rule for which each route selected by such a rule is copied as an individual route into the Adj-RIB-Out, provided that the corresponding Adj-RIB-in also contains a specific route which is also present in the Loc-RIB (i.e. it has been selected for use by the BIS).

The routes selected by type 1 rules are grouped routes, i.e. the routes selected by each type 1 rule form a single group. Each group is then processed by a “Route Aggregation” process to create a single aggregated route for each such group. The aggregation process uses a library of aggregation functions to aggregate each type of path attribute.

Type 2b rules are defined in response to specific ATN requirements for supporting routes to mobile systems. In order to optimise route information distribution, it is necessary to formulate rules that advertise a route to a given BIS, only if that BIS is advertising the selected route to a particular destination. The type 2b rule is a class of rule that meets this requirement.

It should also be noted that some groups of routes cannot be aggregated, even if they have been selected by policy for aggregation. This is because the ISO standard specifically prohibits the aggregation of certain combinations of path attribute. The problem exists for routes that contain:

- DIST_LIST_INCL/EXCL path attributes
- different values of NEXT_HOP
- different values of MULTI_EXIT_DISC.

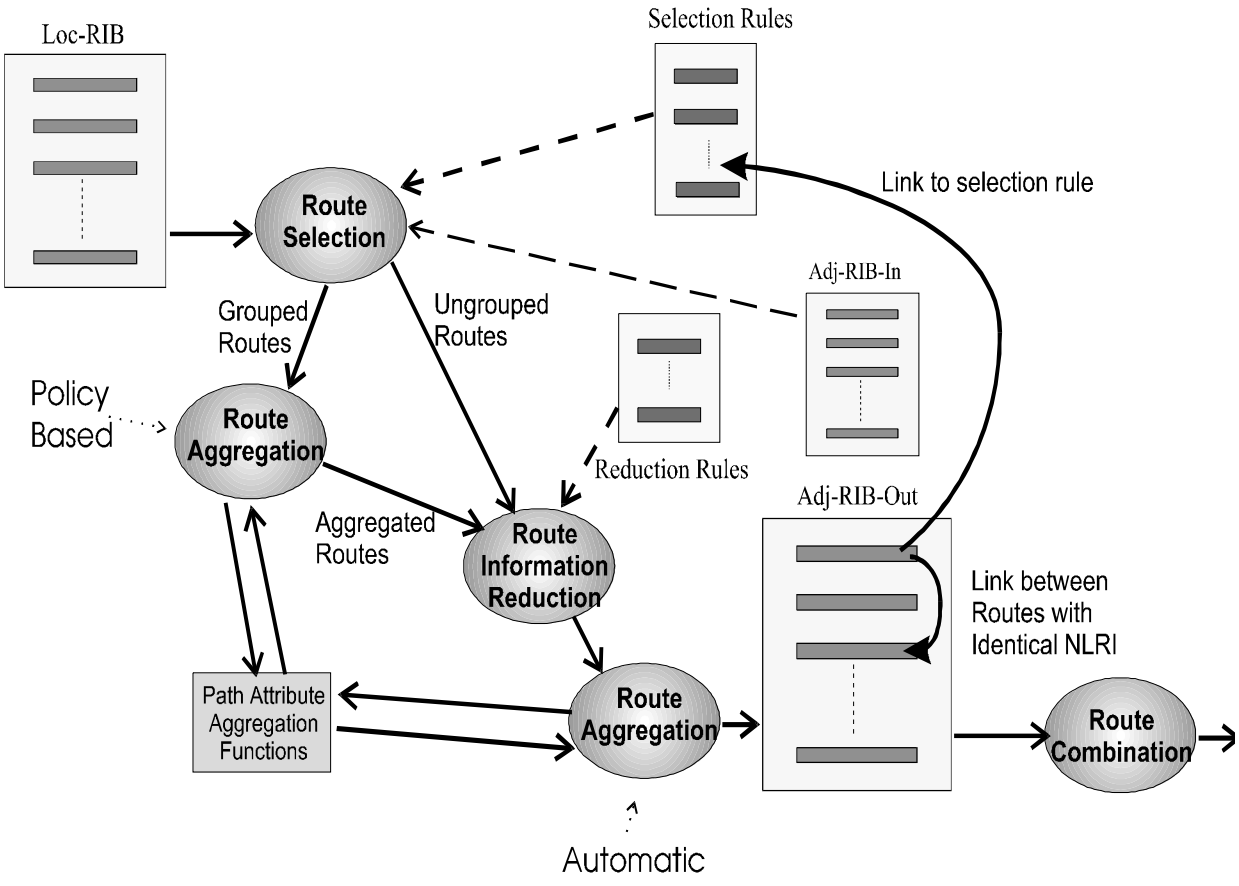

Figure 5-3 Generic Approach to Route Selection, Aggregation and Information Reduction

The outcome, in such cases, is a local matter. However, it is recommended that a deterministic outcome is always ensured.

The remaining routes selected by type 2 rules are ungrouped routes. Both ungrouped routes and the aggregated routes that result from the Route Aggregation process are then passed to a "Route Information Reduction" process. This process inspects the NLRI of each route presented to it and applies the reduction rules to it. The application of a reduction rule will, if the rule is satisfied, result in the replacement of one or more NSAP Address Prefixes in the route's NLRI, with a single shorter prefix. The rules are applied iteratively until no further reduction can take place.

Once the reduction rules have been applied, the routes are ready to be inserted into the Adj-RIB-out. However, it's at this point that a check must be made to see if some of these routes have identical NLRI. If they do then they must be aggregated prior to inserting them into the Adj-RIB-out. Note that the same problem may arise, that was discussed above concerning combinations path attributes that cannot be aggregated. In this case, the only solution may be to apply the Route Merging procedures that were specified in the ATN SARPs as a simplified Route Aggregation procedure.

When the routes are inserted into the Adj-RIB-out, they must be linked to the Selection Rule that originally selected it; this is necessary to support the latter processing of the route.

Prior to inserting the route, the inserting process must check the Adj-RIB-out to see if an existing route is present linked to the same Selection Rule. If this is a type 1 rule, the then new route is marked as replacing the route linked to that Selection Rule. If it is a type 2a or type 2b rule and there is an existing route in the Adj-RIB-out with the same NLRI as the new route, then again the new

route is marked as replacing the existing route. Note that in both cases, if the new route is identical to the existing route in both the path attributes it contains and their values then it does not replace the existing route. The existing route may be simply viewed as refreshed.

Indeed, once the phase 3 processes complete, any routes in an Adj-RIB_out that have been neither refreshed nor replaced, must be marked as withdrawn.

Finally, when a route is passed to the Update Send process for advertisement to an adjacent BIS, a "Route Combination" process is required. This will:

a) Ensure that a route withdrawal is always advertised in the same UPDATE BISPDU as the route, if any, that replaces it; and,

b) Ensure that when a route is advertised, it is combined with any routes with the same NLRI, and which are also queued for advertisement to the adjacent BIS.

A key feature of the above process model is that it enables routes to be selected for aggregation by any combination of selection filters, which do not necessarily make any reference to the routes' NLRI. However, it is believed that the process model can be simplified if it is always assumed that selection for Route Aggregation always includes a filter on the NLRI. Such a simplified model is illustrated in Figure 5-4.

The key simplification in this model is the removal of the second Route Aggregation process. This had had to be introduced to cope with the so called "Route Merging" requirement. This is when two or routes with identical NLRI are selected from the same loc_RIB for inclusion in an Adj-RIB-out. Such routes may have the same NLRI when they are contained in the loc-RIB provided that they differ in the security path attribute. However, this condition may also be a result of Route Information Reduction, and, as Route Information Reduction generally takes place after Route Aggregation, the need for a second Route Aggregation point arises.

However, if certain assumptions are made, it is possible to predict the need for routes to be aggregated because they will have identical NLRI after the Route Information Reduction phase. These assumptions are:

1. Route Information Reduction is only applied to aggregated routes (i.e. routes selected by type 1 rules).

2. Rules that select routes for aggregation and Route Information Reduction must always select routes that contain NLRI which would result from the application of the Route Information Reduction rule.

3. Routes selected by different type 1 rules cannot, as a result of Route Information Reduction, have identical NLRI.

With these assumptions in place, the process model illustrated in Figure 5-4 can be considered.

In this model, Route Selection is again shown separate from Route Aggregation. First, routes are selected from the Loc-RIB for advertisement to a given adjacent BIS, by applying the specified selection rules (type 1, type 2a and type 2b). From this set, routes selected by type 1 rules are queued for aggregation and Route Information Reduction before being entered into the Adj-RIB-out, as before; the remaining routes are copied directly to the adj-RIB-out.

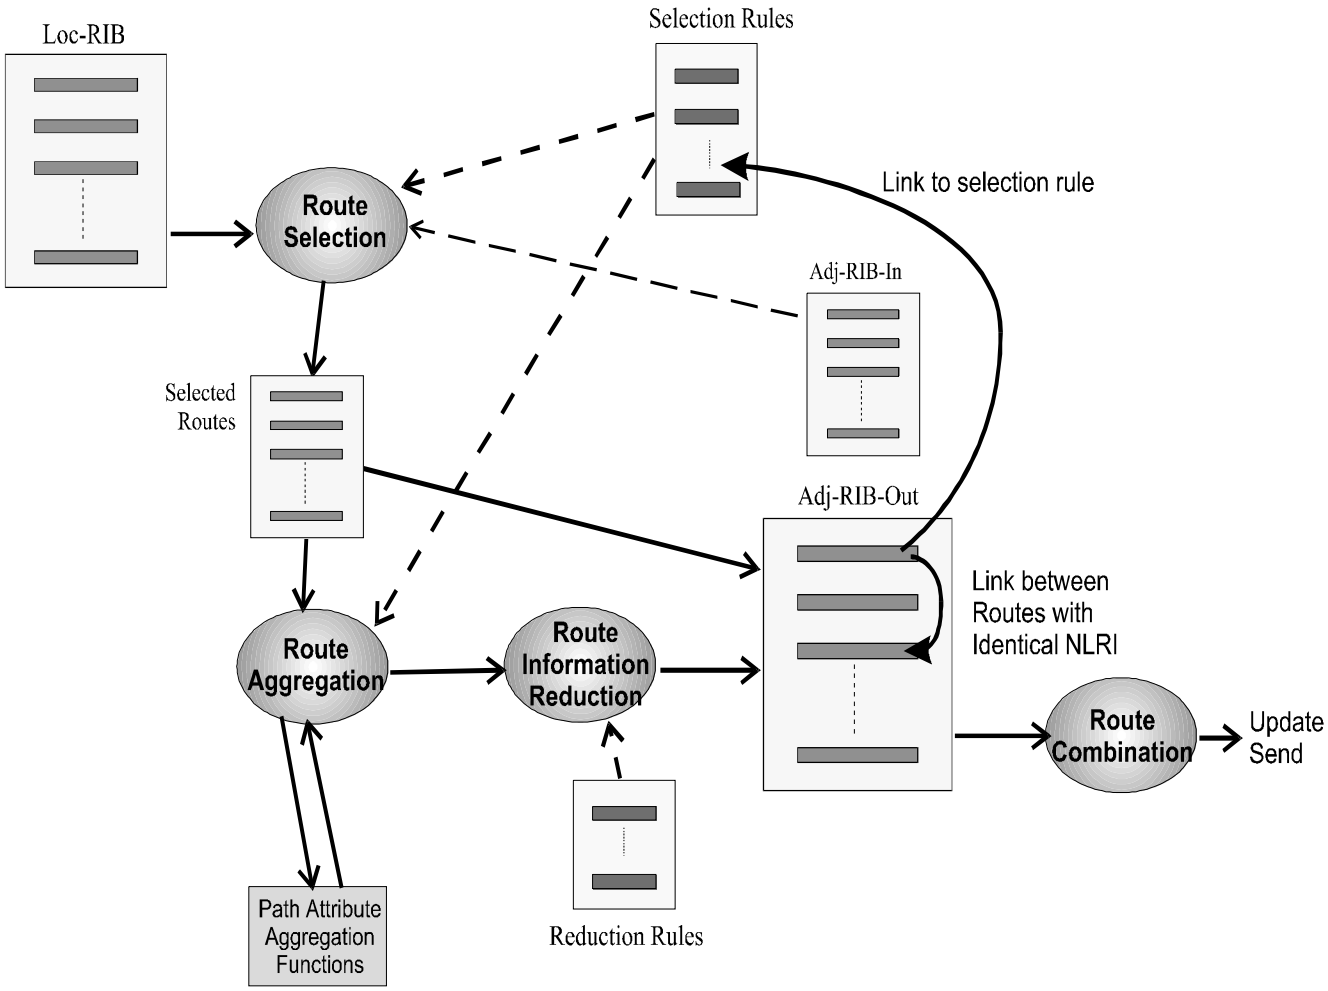

Figure 5-4 Simplified Model for Route Selection, Aggregation and Information Reduction

This procedure is perfectly satisfactory as long as there is no possibility of two routes with identical NLRI being placed in the adj-RIB-out. This can occur for two reasons. The first is that two routes with identical NLRI were selected from the same Loc-RIB. However, this situation can be readily handled by demanding that such routes are always selected for aggregation. However, the other case is more awkward to handle. This is when a route that was copied directly from the set of selected routes has the same NLRI as a route that was the result of Route Information Reduction.

This is where the above assumptions come in. The first is essentially aimed at ensuring that routes that are not aggregated do not end up with identical NLRI. This can only come about because of Route Information Reduction and prohibiting it in this case avoids the problem.

The second assumption ensures that a route copied directly to an adj-RIB-out cannot have the same NLRI as would result from Route Information Reduction being applied to a set of aggregated routes. The third assumption then ensures that this cannot happen as a result of two separate aggregations.

Each of these assumptions is a constraint that apply to the selection rules and which can be checked for when the rules are parsed by the phase 3 decision process.

5.8 Relationship to Intra-Domain Routing

A BIS is a gateway between the inter-domain environment and the intra-domain environment. It forwards NPDUs between the two environments and must also reflect routing information between the two environments.

All destinations within a single Routing Domain will be characterised by a limited set of NSAP Address Prefixes, and ideally such a set consists of a single NSAP Address Prefix. This is a static attribute of the Routing Domain and a BIS will advertise to other BIS a route to destinations within the local Routing Domain with this set of NSAP Address Prefixes as the destination of the route. Generally, there is no need for this route to be dynamically updated. The stability of routing information and the scaleability of the inter-domain environment depends on a certain amount of information hiding and, in particular, BISs will not reflect the actual availability of systems within their own RDs in the routes they advertise to other BISs. To put it simply, turning off a workstation or PC should not result in a change in routing information reported to other RDs.

However, when an Routing Domain has more than one BIS, there is a need to pass routing information from the inter-domain routing function to the intra-domain routing function, for onward advertisement in the level 2 domain as *Reachable Address Prefixes*. This is because intra-domain routers will need to know which BISs provide the best routes to external RDs. On the other hand, it will rarely be practicable or necessary to provide routing information on all known inter-domain destinations to the intra-domain routing function. The volume of information is likely to be far too much for this to be a realistic strategy.

Fortunately, a straightforward approach can be adopted for the intelligent passing of routing information to the intra-domain routing function. Furthermore, such an approach can be used to avoid the encapsulation of NPDUs passed between BISs in the same Routing Domain. The recommended procedure is as follows:

1) Initially, the inter-domain routing function makes available to the intra-domain routing function, as a Reachable Address Prefix, only the default route to all destination. This is a zero length NSAP Address Prefix.

2) Whenever the intra-domain routing function passes a PDU to the inter-domain routing function which is either

   a) decapsulated and then routed to another Routing Domain, or

   b) routed immediately to another Routing Domain, then

   the address prefix that characterises the route followed by the PDU is made available, as a Reachable Address Prefix, to the intra-domain routing function.

3) Whenever an inter-domain route is withdrawn then, if any of the address prefixes that characterise the destination of the route have been made available to the intra-domain routing function, then they must cease to be available for use as Reachable Address Prefixes.

4) Whenever a PDU is received by the inter-domain routing function from an adjacent routing domain, and needs to be routed to another BIS in the local Routing Domain, then the intra-domain routing function is queried to determine if a route other than the default route is available to the PDU's destination. If such a route is available, then the PDU is passed directly to the intra-domain routing function without encapsulation. Otherwise, the PDU is encapsulated, addressed to the NET of the BIS and passed to the intra-domain routing function.

5) Whenever a PDU is received by the inter-domain routing function from the intra-domain routing function and needs to be routed to another BIS in the local Routing Domain then the PDU must be encapsulated, addressed to the NET of the BIS and passed to the intra-domain routing function.

The consequence of the above approach is that BIS learn about the external destinations that systems inside the Routing Domain want to reach and, provide routes to such destinations exist, they are made available as *Reachable Address Prefixes*. The intra-domain routing function can then route such NPDUs direct to the appropriate BIS, rather than the nearest, which is a consequence of just advertising the default route. Furthermore, the same principles apply to NPDUs passing through the Routing Domain. The destinations for such NPDUs are similarly passed to the intra-domain routing function, and the encapsulation of such NPDUs thereby avoided. This is advantageous because encapsulation always carries the risk of unnecessary segmentation, with the overheads that that implies.

5.9 Route Selection, Aggregation and Information Reduction

The concepts of Route Selection, Aggregation and Information Reduction have already been introduced. However, while it has been stated that they have an important role to play in the scaleability of any internetwork, this role has not yet been fully explained. The purpose of this section is to illustrate how these mechanisms are used to implement a scaleable internetwork. The approach taken is deliberately informal, in order to present a complex subject in an accessible manner.

5.9.1 What is Route Aggregation?

Firstly, look at the signpost alongside in Figure 5-5, and imagine being confronted with it at a road junction. If you are going to one of the big cities indicated on it, then you're in luck. It points you in the right direction. But, if you are not, what do you do? Complain to the person that erected it?

Perhaps you do. You want to go to Berlin, and you're the kind of person that complains strongly if things aren't right. The person responsible for the signpost, reacts to customer demand and adds a sign for Berlin. Off you go, a satisfied customer.

The same then happens for people wanting to go to Rome, Toulouse, Sydney, Singapore, Peking, Cape Town, Rio de Janeiro, Seattle, Moscow, Dublin, Brisbane, Winchester, Prague, Bristol, Athens, Anchorage, Stornoway, Oslo, St Petersberg, and so on, until there is no further room on the signpost to hang another sign. What does our poor Signpost Manager do now?

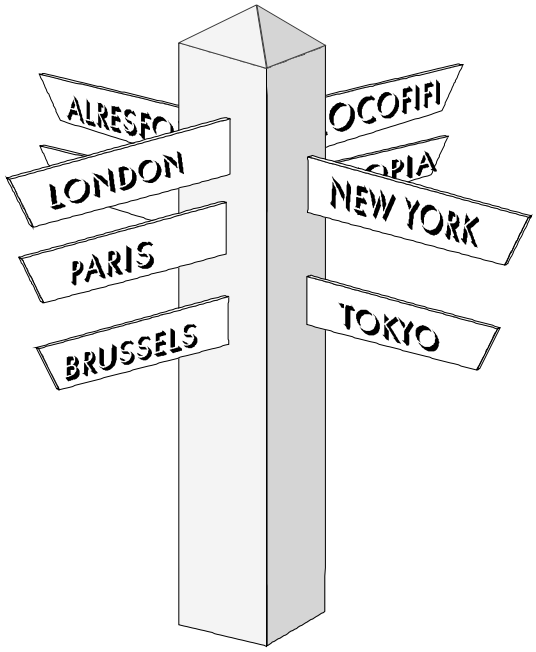

Figure 5-5 Signposting the Way

He could just erect a bigger signpost, but if he's bit cleverer, he may just realise that the problem is not one of insufficient signpost real estate, but really it's the granularity of information that is being provided. After all, London, Paris and Brussels are all in Europe, and hence could be replaced with a single sign indicating the direction to Europe, along with all the other cities and towns in Europe that are individually listed on the signpost.

In fact, this is a really bright idea, as it is not just the European cities that can be picked off in this way, but so can the Asian cities, the American ones, the African ones, and so on. Only those that really are local (i.e. on the same continent) need to be explicitly mentioned. What our bright signpost manager has realised is that his customers don't really need detailed information on the route for their individual destinations. There are only a few directions in which they can go anyway and, when he labels each direction with a suitable collective noun or group name, that properly and

unambiguously describes what is reachable in that direction, the signpost's users will get all the information they need. After this exercise in information reduction, our signpost ended up much like that in Figure 5-6.

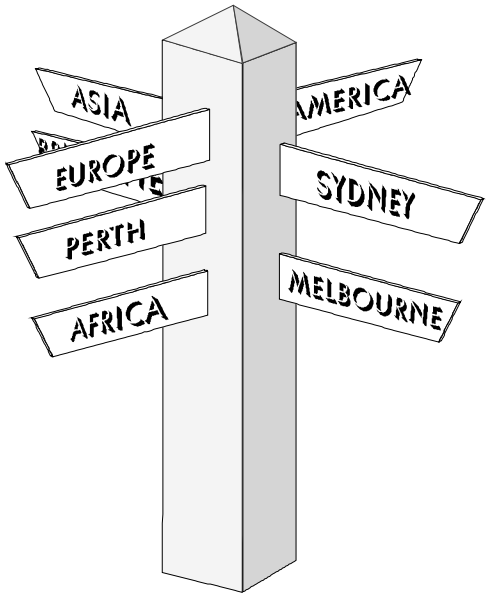

Figure 5-6 The Rationalised Signpost

This benefited the signpost's users, who didn't have to search through lots of different signs to find the one they wanted, and the signpost manager's company, as now, maintenance had been reduced to almost zero.

OK, so this is how road signs work, but is it really relevant to network routing?

Of course it is. Every router has an electronic signpost within it - its forwarding table. Each packet that it forwards, must find a sign telling it which direction to go in, otherwise it will be discarded. A Network Manager is akin to our Signpost Manager and must ensure that there is a suitable sign for every packet that needs to be routed.

By replacing whole groups of signs by a single sign, our Signpost Manager brought together the pointers to many different routes and merged them into a single pointer. In effect, he aggregated those routes - he performed *Route Aggregation*. In fact, he went one stage further. Not only did he bring the routes together, but he also replace the list of individual destinations, be a single common destination name. This procedure is properly known as *Route Information Reduction*.

5.9.2 Structured Addresses and Routing

From this you may conclude that routers adopt a principle similar to that illustrated in Figure 5-6, and minimise the amount of routing information by collecting routes together and signposting routes to appropriate group addresses. Unfortunate, you would not always be right in making such a conclusion.

For example, in the TCP/IP Internet, the routers implemented by the Internet Service Providers are much more like the signpost in Figure 5-5. There's a sign for every network in the world and, when they run out of space to add new signs, the only answer is to get a bigger signpost. In fact, even this isn't true, because for most Internet Service Providers, there aren't any bigger signposts anymore.

The reason why this is so is twofold. Firstly, the network addresses used in the TCP/IP Internet are rather on the small side at only 32-bits long. Secondly, such addresses have traditionally been allocated to networks without any regard to network topology. The first problem is due to the limited horizons of the early Internet developers. No one at that time thought the Internet would grow so big and a 32-bit address was chosen for engineering reasons (i.e. efficient processing) rather than with future growth in mind. The second problem is simply due to any recognition that there needed to be a way (in network address terms) of forming the structured addresses necessary to move away from the over-crowded signpost.

A Network Address is simply a binary number that uniquely identifies a single host computer on the Internet, However, network addresses are not simply names (like London or Paris) which, on their own tell you nothing about where the addressed location actually is. Network Addresses are first of all names of systems on a network, but they must also be parameters to a routing algorithm that is

implemented by every router in an internetwork, and their role as parameters constrains the scope for allocating network addresses.

In our signpost example, the address that we were trying to get to wasn't simply (e.g.) London, but in reality would be a structured address (e.g. 221b Baker Street, London, England, Europe). To find the addressed location, we would consult our first signpost:

- if the signpost is in London, then we start looking for a sign first to Baker Street;
- Otherwise, if the signpost is in England, we look for London;
- Otherwise, if the signpost is in Europe, we look for England;
- and finally, if the signpost is not even in Europe, we look for a sign to Europe.

This is the algorithm we employ to use signposts to help us find our destination. We employ it at every signpost we encounter on our journey and, if they are giving us the right information, we will eventually get to our destination.

In the TCP/IP Internet, a Network Address is similarly structured, but into only two parts. The first part is a unique network identifier and the second part uniquely identifies a Host Computer on the network identified by the first part.

Furthermore, the network identifiers were assigned on a "first come first served" basis. In the electronic signposts that exist in every Internet Router, there has to be a "sign" for every assigned network identifier, pointing along the route to that network. If network identifiers had been assigned (e.g.) that 1 to 100 were in North America, 101 to 200 were in Europe, and so on, then there would be opportunity for the "signposts" within each such router to be rationalised as in Figure 5-6. Within organisations, this is often done, with the Host Identifier split up into an internal (within the organisation) network identifier and a smaller Host Identifier However, at the level of the Internet Service Provider, there is a need to keep track of a route to each assigned network identifier, and this is a serious limitation on Internet growth.

If our electronic signposts are to be rationalised, then Network Addresses must be structured in a way that is much greater than simply Host on Network and so that we can address our systems as (e.g.) *Host* on *internal network*, in *organisation*, attached to *Internet Service Provider*, in *Country* or *Region*. Then, for example, the Routers in an Internet Service Provider (ISP) only need to have "signs" for their users, other ISPs in the same country or region, and an ISP in each other Country or region. The number of such "signs" is then unaffected by the attachment of a new organisation to another ISP i.e. the Internet can grow locally without global impact. This is a necessary condition for an Internet that is scaleable (can always grow bigger). Unfortunately, this is not a realistic proposition with addresses of only 32-bits.

5.9.3 The Allocation of Structured Addresses

By allocating network addresses arbitrarily (at least on a per network basis), the early developers of the TCP/IP Internet have compromised its later growth. Fortunately, for the ATN Internet, these problems were already known by the time that the ATN came to be developed and can thus be largely avoided.

The ATN specifies the use of the Connectionless Network Protocol (CLNP) instead of IP. This has the great advantage of large (variable length) addresses, and the ATN takes advantage of this to specify a 160 bit address format. Although it can be argued that such a long address is less efficient to process than a 32-bit address, 160 bits makes it much easier to ensure that similar network addresses are allocated to networks that are near each other in the ATN Internet, and can therefore be used to improve the overall routing efficiency.

This larger address space allows for a structured allocation of addresses to be made. The address may then be broken up into a number of fields (for the purpose of allocation), which then form a nested hierarchy. For example, in a left to right order, the fields may identify region, country, organisation, site, system. All Systems within a given organisation would than have addresses that share a common prefix and those on the same site also share a common (but longer) prefix. In the ATN, such addresses are known as NSAP Addresses and the prefixes are therefore called NSAP Address Prefixes.

With this approach, similar network addresses, as illustrated in Figure 5-7, imply that the addressed destinations are close together in the topology of the network. Indeed, how far down the address (seen as a bitstring) that the two addresses diverge, can be taken as a metric of closeness.

Indeed, in a scaleable Internetwork, such as the ATN, the Routers operate first by labelling routes with the address prefix(es) common to all destinations along the route, and perform routing simply by comparing destination network addresses against such address prefixes and forwarding each packet along the route labelled with the longest matching address prefix. This is very much like the use of a physical signpost described earlier.

Furthermore, as routing is done by such a simple prefix matching rule, the Routers do not themselves

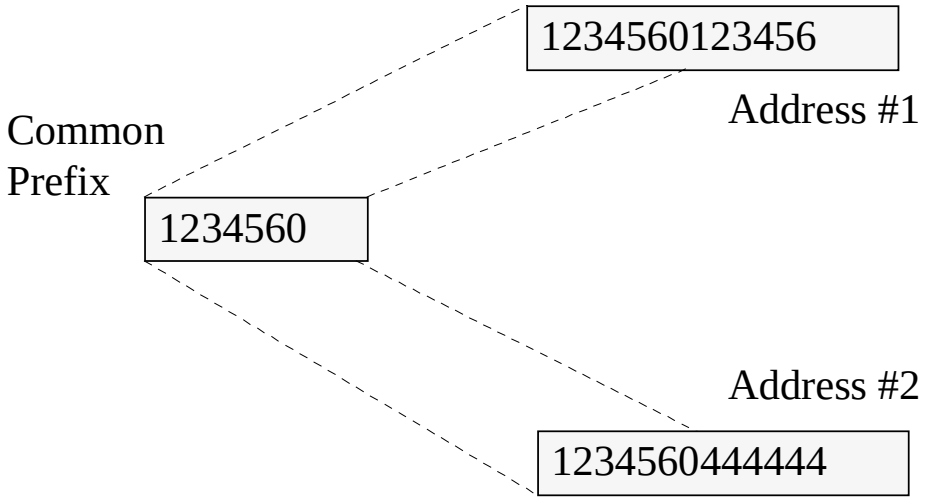

Figure 5-7 Similar Network Addresses

have any real need to know about the structure of the address. The structuring of a network address into a series of fields is therefore only for the purpose of address allocation and not for routing purposes. This is of course different to the way physical signposts are used and represents where our analogy and network routing diverge.

5.9.4 Towards a Scaleable Routing Concept

Our signpost analogy is really only one part of the routing concept. As illustrated in Figure 5-8, signposts are just waypoints along a route between a starting point and a journey's end and, formally, we define a route to be a combination of information that describes a path, and the NSAP Address that identifies the end point of the route. IDRP deals in such routes and allows BISs to keep each other informed about the routes that they offer.

Of course, IDRP's routes are not to actual destination systems. They are to the BISs at the edge of the Routing Domain that contains the destination system, and the NSAP Address of the route's end point is a Group Address - the common NSAP Address Prefix for all systems within that Routing Domain. Effectively, the BIS has brought together the individual routes to each system within Routing Domain into a single route, and replaced all the individual NSAP Addresses with the appropriate single NSAP Address Prefix. We already know these two processes to be called Route

Aggregation and Route Information Reduction, and these always occur implicitly, in a BIS, before a route to such internal destinations is advertised to the BISs of other Routing Domains.

The question now arises as to whether there is any merit in carrying out Route Aggregation and Route Information Reduction at any other points in route distribution. The answer is a definite yes.

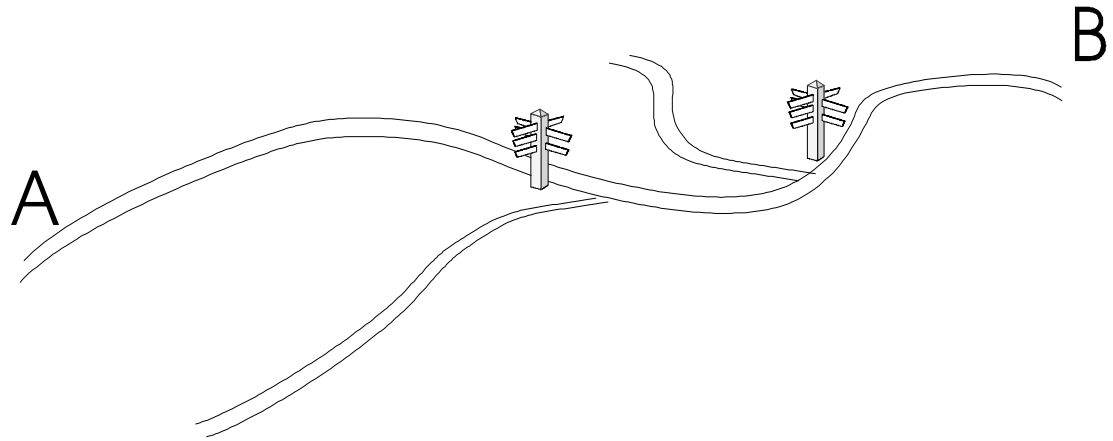

Figure 5-8 The Route - a path between A and B

Firstly, there is nothing magic about an 88-bit NSAP Address Prefix. That figure so happens to be a convenient breakpoint in the ATN Addressing Plan. In IDRP, NSAP Address Prefixes can be any number of bits in length. If routes to individual Routing Domains can be aggregated together, and their individual NSAP Address Prefixes replaced by a single shorter common prefix, then we have achieved a useful simplification not just for our local electronic signpost, but for all such signposts downstream of the point at which the routes were aggregated.

In fact, if we can achieve the general principle that the further away from a route's destination you are, the shorter the NSAP Address prefix is for the route's destination, then we have achieved the goal of a scaleable internetwork. This is because for an internetwork to be scaleable, that is to be able to grow without any serious limitation on its total size, we must never get into the situation that the TCP/IP Internet has got itself into, where there are routers which have to keep having bigger and bigger "signposts" as the internet grows. The internet then cannot grow any more, once these routers have the biggest signposts that can be purchased.

As long as the above principle is obeyed, growth can occur in the far away internet without affecting remote routers, and hence growth can continue in an almost unbounded fashion.

For example, consider the example in Figure 5-9. Here we have a service provider supporting several users, and it is assumed that the service provider has been allocated the NSAP Address Prefix "1234" for all NSAP Addresses that it allocates. It allocates the prefix "12340" to its own Routing Domain, and then allocates "12341", "12342". etc. to each of its users' Routing Domains. The systems with those Routing Domains are then allocated NSAP Addresses relative to the NSAP Address Prefixes assigned to each Routing Domain.

In each User's Routing Domain, a BIS forms a route to all systems within that Routing Domain. This is a route to all systems in the Routing Domain, and the route's destination is the NSAP Address Prefix assigned to the Routing Domain. This route is then advertised using IDRP to the Service Provider's BIS.

The Service Provider's BIS receives a so advertised route from each user's Routing Domain and can therefore build its own electronic signpost from each of these routes, "adding a sign" for each route advertised to it. This router could just re-advertise each such route on to a BIS operated by another service provider or its own users. However, because all these routes share a common NSAP Address

Prefix (“1234”) it is much more efficient to first aggregate the routes together, along with the route to the service provider’s own Routing Domain, and then apply the Route Information Reduction procedure to end up with a single route to “1234”. This is the route it then advertises on, instead of re-advertising the individual routes to each Routing Domain.

Not only is this efficient but, if for example, a new user’s Routing Domain is added (and given the next NSAP Address Prefix - “12344”), then this has no impact at all on the aggregated route or the number of routes maintained by the BIS in another Service Provider. The internetwork has grown locally without having a global impact, and this is what scaleability is all about.

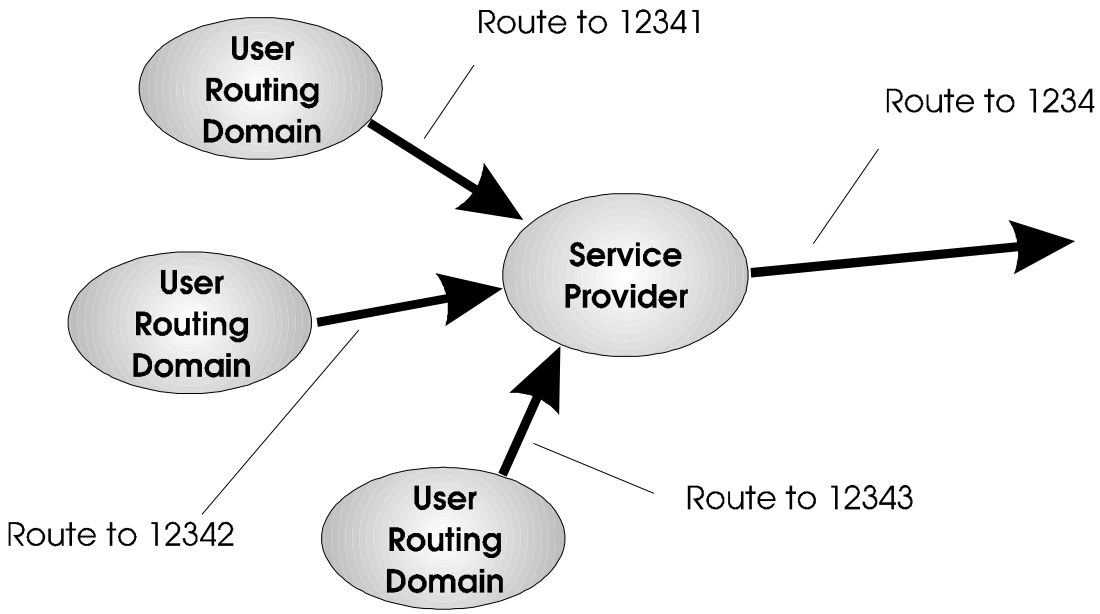

Figure 5-9 Aggregating Routes Together

This example can be readily extended. For example, if all of the Service Providers in a given country shared a common NSAP Address Prefix (e.g. “123”), then only a single route needs to be advertised internationally and which is common to all service providers. In fact, as long as the address allocation hierarchy reflects the way the network is organised, there will be many such opportunities for Route Aggregation and Route Information Reduction.

In the ATN, the addressing plan is so organised that each Administration has a single NSAP Address Prefix which will be common to all systems and Routing Domains that the maintain. Thus only a single route need be advertised between individual Administrations. Furthermore, provided that within a region, Administrations co-ordinate their addressing plans, it will be possible to form a single route to a given region keeping the overhead of inter-regional communications down to a minimum.

Looking ahead to 5.11.2, this principle is further exploited by the ATN Island concept. An ATN Island is essentially a regional grouping of Administrations with co-ordinated addressing plans. In such a situation, it is possible to form a single route to “the ATN Island”, and, indeed, it is recommended that this is done prior to route advertisement to aircraft, thus keeping down the routing overhead on low bandwidth air/ground data links to a bare minimum.

5.9.5 Containment Boundaries and Routing Domain Confederations

Route Aggregation and Route Information Reduction generally work very well by themselves. However, to help solve the problem of when to aggregate, we have already introduced the idea of a Containment Boundary (see 5.4.6). We need some way of defining the scope of a given NSAP

Address Prefix - that is to define a Containment Boundary that itself defines the limits of the domain of such an NSAP Address Prefix.

One obvious example of such a Containment Boundary is a Routing Domain. Each Routing Domain contains all systems identified by NSAP Addresses relative to the NSAP Address Prefix assigned to that Routing Domain. When routes exit a Routing Domain (i.e. at a BIS), the Containment Boundary is crossed, and the router knows *a priori* that it is appropriate to aggregate the individual routes together and form a single route with its destination being the common NSAP Address Prefix for the Routing Domain.

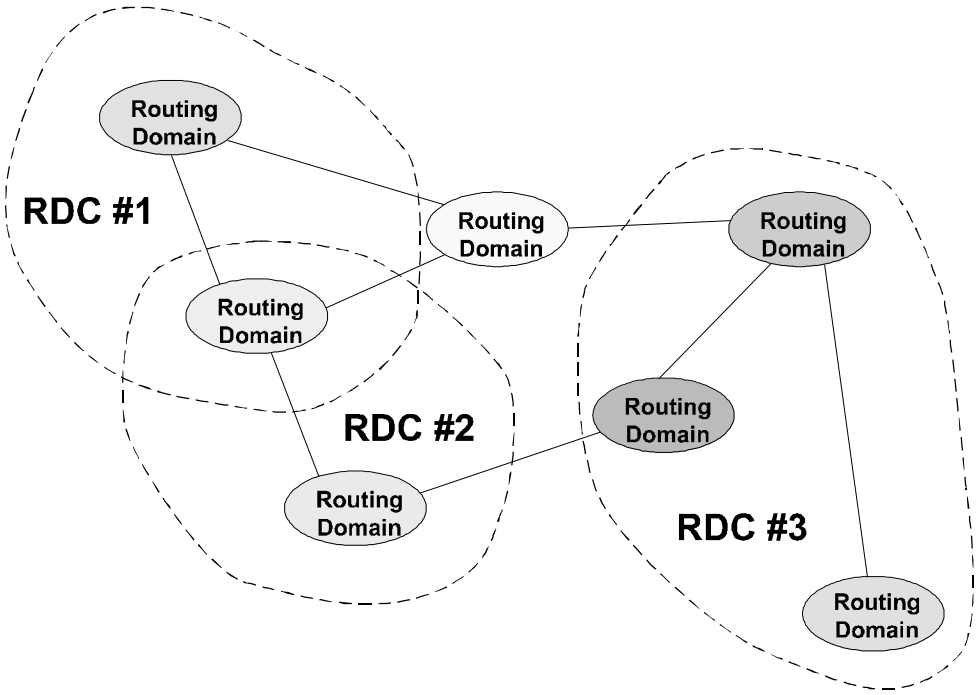

Figure 5-10 Routing Domain Confederations

In the example in 5.9.4 above, there is clearly some sort of Containment Boundary enclosing the Service Provider and its users. This can simply be a conventional boundary. However, IDRP does provide a means to make this more concrete in the shape of a Routing Domain Confederation (RDC).

An RDC is no more than a group of Routing Domains, as illustrated in Figure 5-10, and, at its simplest, is a means of collectively referring to a related group of Routing Domains. However, an RDC can usefully be defined to be a Containment Boundary for the domain of an NSAP Address Prefix. In the above example, we could have an RDC containing the Routing Domains of the Service Provider and its users.

With such an RDC, we can then implement a simple and effective rule for aggregating routes i.e. whenever a route that originates within the RDC is advertised across the RDC boundary, it is aggregated with all such routes to form a single route to a destination described by the common NSAP Address Prefix for all Routing Domains within the RDC. This is essentially what is happening in our example.

As happened in the example, more Routing Domains can be added to the RDC without affecting the route advertised external to the RDC. That is the internetwork has grown locally without global impact.

In the ATN, an ATN Island is an example of an RDC that contains all Routing Domains with a common NSAP Address Prefix i.e. common to all systems on the “Island”. Whenever a route is

advertised outside of the Island (e.g. to an aircraft) it becomes a candidate for aggregation with other such routes. As is described later in 5.11, RDCs, Address Allocation and Route Aggregation are used together to create a scaleable ATN supporting mobile routing.

5.10 Route Initiation

5.10.1 The Purpose of Route Initiation

ICAO has adopted the use of Policy Based Routing procedures for routing between ATN Routing Domains (RDs), including the support of routing to mobile systems. Dynamic Routing Information is exchanged using the procedures specified in ISO 10747 and used and disseminated according to local routing policies specified in accordance with the ATN SARPs. However, before routing information can be exchanged between any two Routing Domains, it is first necessary to establish a communications path between BISs in each of those RDs. The establishment of such a communications path is known as “Route Initiation”.

Route Initiation procedures are required whenever two ATN RDs need to be interconnected. Since the ATN SARPs specify that, on board an aircraft, the communications systems and the applications processors that they serve comprise a Routing Domain, Route Initiation procedures also apply to the establishment of air/ground communications.

Route Initiation commences when the decision is made to establish a communications path between two ATN RDs. Route Initiation finishes upon the initial exchange of routing information between the BISs, or the unsuccessful termination of the Route Initiation procedure.

Note: BISs within the same RD also exchange dynamic routing information using ISO 10747. The Route Initiation procedures are the same as for inter-domain connections except that both Routers will be under the control of the same administrator.

5.10.2 Ground-Ground Route Initiation

5.10.2.1 The Communications Environment

Ground-Ground communications typically use long lasting physical or logical communications paths. Route Initiation can normally be regarded as a rare event and will often be only semi-automated.

The communications networks in the ATN ground environment are outside the scope of the ATN SARPs, but can be assumed to include:

1. X.25 Public and Private Data Networks
2. Leased Lines
3. Integrated Services Digital Networks (ISDNs)
4. Frame Relay Services
5. The Public Switched Telephone Network (PSTN).

The actual choice of communications network is a matter for bilateral agreement between the organisations and states that wish to interconnect their RDs, and will depend on local availability, tariffs and policies. In many cases, high speed (e.g. V.32bis or V.34) Modems and the PSTN will be used as a backup for a dedicated data network.

The communications protocols used to provide the data link will also depend upon the communications network used and bilateral agreement. In the case of X.25 data networks, Frame Relay and communications services provided via the ISDN D-Channel, then the communications protocols are mandated by the data network provider. In the case of Leased Lines and the ISDN B-channel, then HDLC LAPB (ISO 7776) is the likely choice. For the PSTN, the asynchronous

communications provided by V.32bis and V.34 Modems makes the Point-to-Point Protocol (PPP) as specified in RFC 1548, the likely choice.

Note: Route Initiation is not necessarily synonymous with the establishment of an uninterrupted communications link between two BISs. For example, the speed at which an ISDN B-Channel is established is such that it may be practicable to break the communication circuit during idle periods and re-establish it when there is data to send, whilst still maintaining a logical communications path between the two BISs. Route Initiation is concerned with the establishment of the logical communications path.

5.10.2.2 Summary of Procedures

The sequence of procedures for a typical ground-ground Routing Initiation is illustrated in Figure 5-11, and summarised below. They are described in greater depth in the following sections. This illustrates the co-ordination of two Systems ("A" and "B") interconnecting over a common network. The procedures are:

1) Adjacent BIS MOs are established in both Systems. In each case, an MO is established to identify the other system and contains the parameters necessary to create and maintain a BIS-BIS connection with that system. Both systems will also have been configured with appropriate SNDCFs associated with each attached subnetwork.

2) A communications path is established over the subnetwork; typically one system is initiator and the other responder.

3) Establishment of the communications path is notified to the Systems Manager.

4) In response, the Systems Manager for each system adds a route to the local FIB and to the remote System, and

5) invokes the IDRP "Start Event" action, or re-run the decision process if a BIS-BIS connection already exists with the remote system.

6) On successful establishment of the BIS-BIS connection, Route Initiation completes.

Note: while the Systems Manager may be a real person explicitly issuing commands, the "Systems Manager" in the above description may alternatively be a procedural script carrying out an automatic action in response to a Systems Management Notification.

5.10.2.3 Initial Route Initiation

Route Initiation begins with the decision to establish a communications path between a pair of BISs, including the decision on which communications networks to use. The first procedure is to establish the underlying communications circuit between the BISs and hence to establish the logical communications path.

These procedures will be data network dependent and will require some sort of interaction between the respective Systems Managers. Typically, one BIS will need to be in a passive state awaiting an incoming event (e.g. an X.25 call indication or a PSTN Ring Indication), while the other takes an active role and initiates circuit establishment (e.g. by generating an X.25 call request, or "dialling" the telephone call).

When appropriate to the type of data network used, the QoS, Security and Priority requested on any such call request, should be satisfactory for the exchange of routing information.

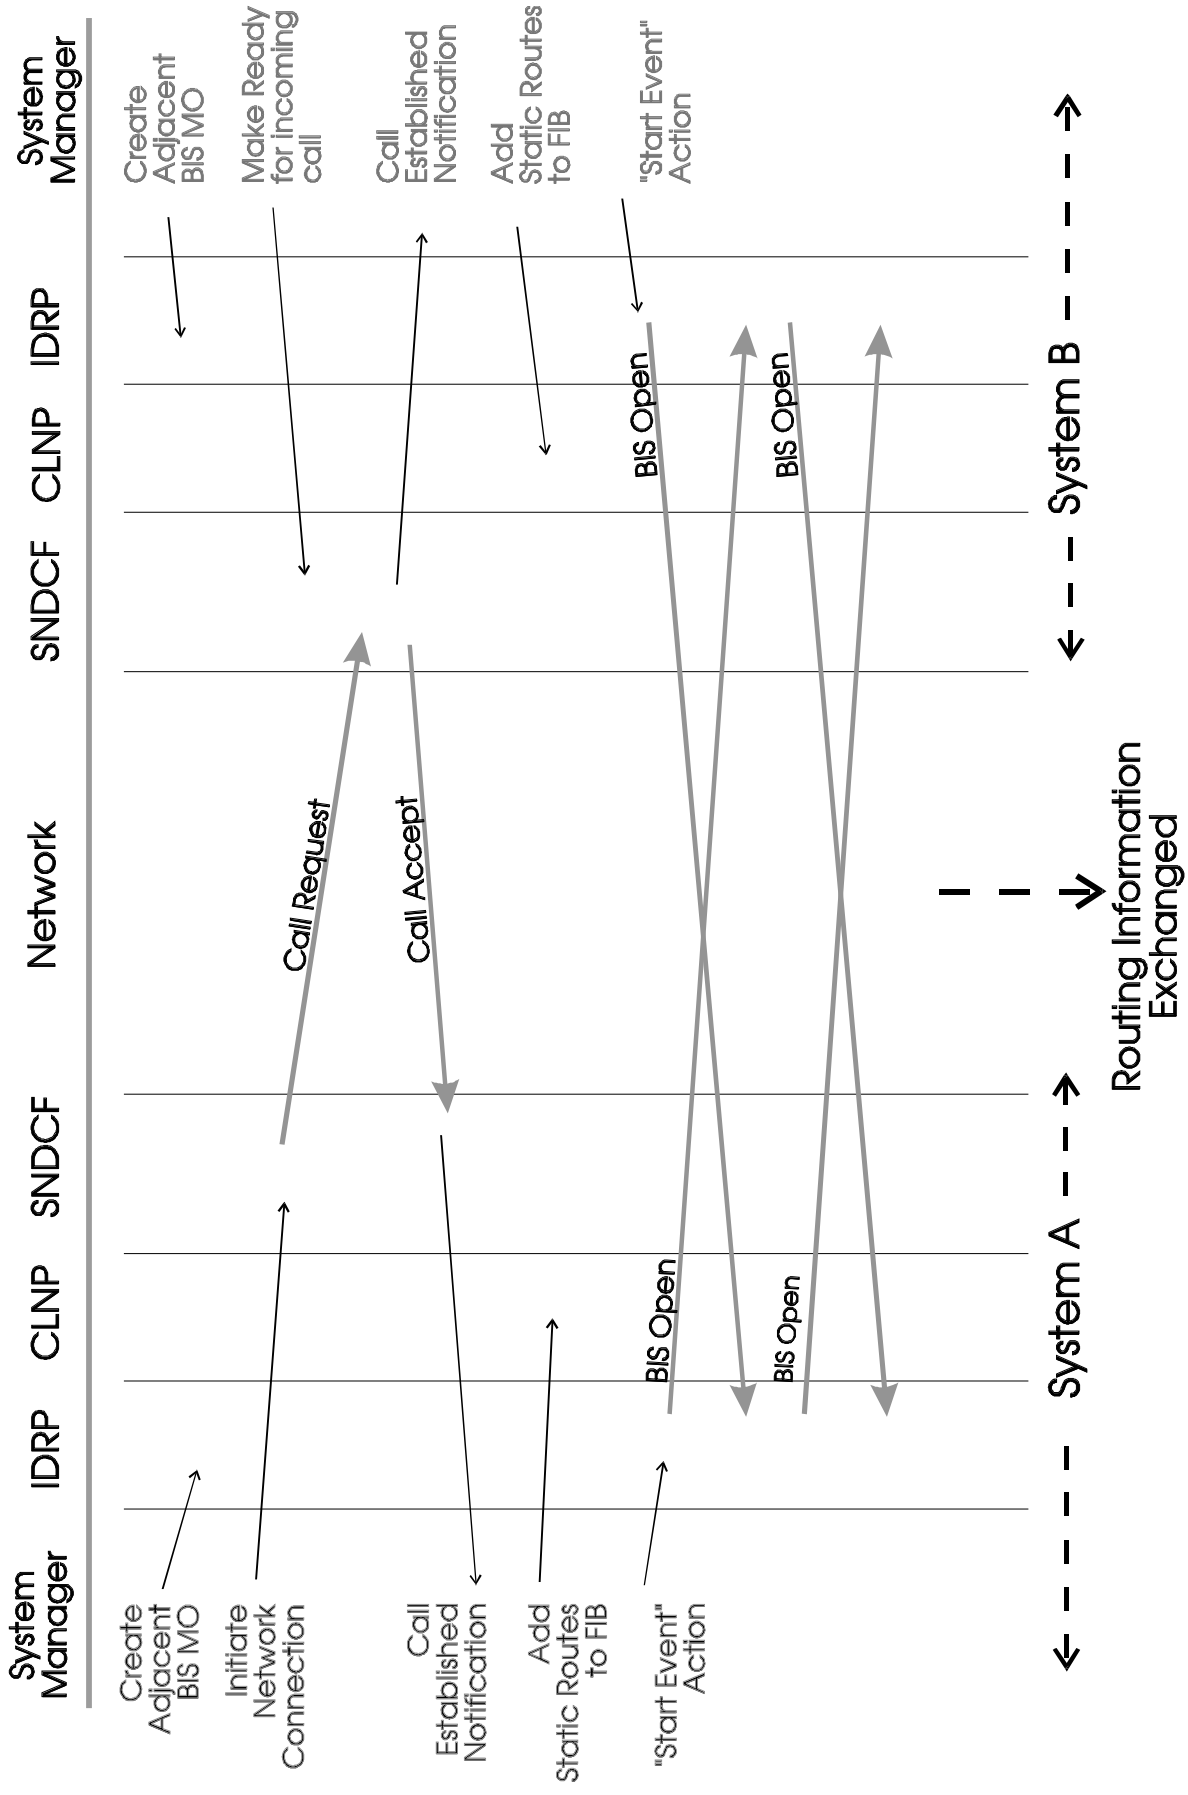

Figure 5-11 Ground-Ground Route Initiation Sequence

During this phase, there should normally be some validation to ensure that communications has been established with the correct remote system. This initial phase completes once the data link has been established.

5.10.2.4 Route Initiation in CLNP

The ATN SARPs specify the use of the Connectionless Network Protocol (CLNP) specified in ISO 8473 for ATN subnetwork independent communications. Establishing a data link (e.g. an X.25 virtual circuit) is a necessary condition for data to be exchanged between two BISs using CLNP, but not a sufficient condition. In order for the data link to be used by the CLNP Network Entity, and hence as a communications path for the forwarding of data packets, it is necessary to:

1. Assign an appropriate Subnetwork Dependent Convergence Function (SNDCF) to interface the data link to the Network Entity;

2. Update the Forwarding Information Base (FIB) to record statically known routes available over the data link and via the remote BIS.

The former is necessary in order to match the characteristics of the actual network and communications protocol used over that network to the characteristics assumed by the CLNP Network Entity. The second is necessary in order to permit the exchange of dynamic routing information.

The SNDCF is typically specified for a network type and associated at system configuration time with a physical communication port. In most cases, the assignment of the SNDCF is implicit in the network over which communications is established, and no explicit action will need to be carried out to assign the SNDCF. Indeed, most implementations will require assignment of the SNDCF prior to establishment of the data link. However, for some network types there may be alternatives chosen at connection establishment time.

The FIB may be updated with any statically known routes that are known *a priori* to exist via the newly established data link, where a route consists of an NSAP Address prefix paired with an identifier for a data link. When forwarding data packets, the CLNP network entity locates the longest matching NSAP Address Prefix in the FIB, when matched against the packet's destination NSAP Address, and then queues the packet for transmission over the associated data link. Multiple FIBs may also exist, matching different QoS and security requirements. So that Routing Information may be exchanged, the FIB associated with the QoS level used for the exchange of Routing Information, must be updated to include, as a minimum, a route to the network entity located on each BIS to which a data link has been established.

Therefore, once a data link has been established to a remote BIS, the System Manager must either directly, or via an automated procedure, insert into the FIB associated with the Security and QoS level used for the exchange of Routing Information, a route associating:

a) an NSAP Address prefix that is a prefix for the NET of the remote BIS at the other end of the newly established data link. As a minimum, this prefix may be the complete NET; and,

b) the data link to that remote BIS.

Note 1: the reverse must also take place when the data link is terminated i.e. the above route must be removed from the FIB.

Note 2: alternatively, such routes may be entered into the FIB at system initialisation. However, this strategy gives satisfactory results only if there is a single possible data path to the remote BIS.

5.10.2.5 Route Initiation in IDRP

Once a communications path has been established between two BISs and sufficient static routing information has been entered into the local FIB in order to enable the forwarding of data packets to the remote BIS itself, IDRP may be used to exchange dynamic routing information.

IDRP may only exchange dynamic routing information when a BIS-BIS connection has been established. This is a logical connection established by using the IDRP protocol, which in turn uses CLNP to transfer the protocol data units (BISPDUs) to the remote IDRP entity. A BIS-BIS connection supports the reliable transfer of dynamic routing information between BISs.

Prior to establishing a BIS-BIS connection it is necessary to create an "Adjacent BIS Managed Object" to provide the information necessary to establish and maintain a BIS-BIS connection with an explicitly identified remote BIS. The information held includes the NET of the remote BIS, authentication data, the specific IDRP procedures used to establish the BIS-BIS connection and timer values. One such MO exists for each remote BIS with which IDRP may exchange routes. Typically, this MO is setup in advance of the underlying communications path, and will usually be created once agreement to interconnect has been reached.

Once the FIB has been updated with a route to the remote BIS, the "start event" action is requested of the Adjacent BIS MO associated with that Remote BIS. This initiates the procedures for creating the BIS-BIS connection and is followed by the exchange of dynamic routing information. It is the final action of the Route Initiation procedure.

During establishment of the BIS-BIS connection either or both IDRP entities will take an active role in connection establishment, or one will be active and the other passive. The role, active or passive, is determined by information configured into the Adjacent BIS MO. If one IDRP entity is to be passive, then Systems Managers must ensure that the other is configured in the active role. If both IDRP entities are configured in the active role, then the BIS-BIS connection establishment procedures are less efficient, than if one is in the passive role. However, given that the loss of efficiency is small and typically of no consequence given that ground-ground BIS-BIS connections are usually long lived, Organisations and States are recommended by the SARPs to always configure the Adjacent BIS MOs for BIS-BIS connections between ground ATN BISs for BIS-BIS connection establishment in the active role. This is to avoid to risk of both being configured in the passive role by mistake.

However, there is one exception to the above. That is when the newly established communications path is to a remote BIS with which a BIS-BIS connection already exists. This is possible when multiple networks are available between the same pair of BISs. Multiple concurrent connections may be desirable in order to give high availability through redundancy and to provide additional data transfer capacity.

IDRP permits only a single BIS-BIS connection between a given pair of BISs, irrespective of the number of underlying connections and networks that may join them. Therefore, the Systems Manager should check to see if a BIS-BIS connection already exists to the remote BIS and only invoke the Start Event Action if one does not already exist. This action will in any case, be ignored if issued when a connection does already exist.

However, other action may be appropriate if there is a need to recognise the different QoS that may be available when a new communications path is opened up (or lost), or a change occurs in the Security Types that may be supported by alternative communications paths to the same remote BIS. In such cases, the SARPs require that the IDRP Decision Process be aware of the aggregate QoS and Security Restrictions over the communications paths to a given remote BIS (Adjacent BIS). The SARPs require the Decision Process to update the QoS on received routes (when processing the adj-RIB-in) to reflect the QoS of the communications path and to use this updated QoS when determining the degree of preference of the route and when re-advertising it.

The SARPs also require that the Decision Process does not place in the IDRP adj-RIB-out, any routes with Security Types incompatible with any restrictions that exist on the aggregate communications path. For example, if none of the available communications paths to a given remote BIS permits the transfer of "Administrative" data, then a route with a Security Type reflecting administrative data may not be placed in the Adj-Rib-out for that Router (and hence advertised to it).

Therefore, whenever an additional communications path to a given remote BIS becomes available (or is lost), the Systems Manager must cause the IDRP Decision Process to be re-run, instead of invoking the Start Event.

5.10.3 Air-Ground Route Initiation

Air-Ground Route Initiation is similar to ground-ground Route Initiation, but differs for the following reasons:

1. ICAO specified subnetworks are used for air-ground communications with their procedures for use mandated by SARPs rather than subject to bilateral negotiation.

2. Route Initiation typically starts as soon as communication is possible e.g. an aircraft coming into range of a Mode S Interrogator, and, in consequence Route Initiation starts as soon as the Systems Manager is notified of the possibility of communications (e.g. capture by a Mode S Interrogator).

3. It is not realistic to pre-configure Adjacent BIS MOs for every aircraft that may come into contact with a given ground ATN Router; these MOs must be set up as part of the Route Initiation Procedure.

4. Special procedures are necessary to identify the NET of a remote ground or airborne Router during the Route Initiation procedure as, in general, it is not possible to know this in advance.

5. Due to avionics limitations, not all aircraft will be able to implement IDRP and interim procedures inferring route availability over air-ground links must be accommodated.

5.10.3.1 Communications Environment

The following ICAO Air-Ground data networks are expected to be used to support the ATN:

1. The Aeronautical Mobile Satellite Service (AMSS)

2. The VHF Data Link (VDL)

3. The Mode S Data Network

In each case, ITU recommendation X.25 provides the data network access procedures, and the responsible ICAO Panel's have required that:

a) AMSS communications are "air initiated", that is the aircraft is responsible for initiating communication with the ground

b) VDL communications are similarly air initiated.

c) Mode S communications are "ground initiated" that is a ground ATN Router attached to a Mode S data network is responsible for initiating communications with an aircraft.

5.10.3.2 Summary of Procedures

The Air-Ground Route Initiation procedures are illustrated in

Figure 5-12, and summarised below. They are described in greater depth in the following sections. This figure illustrates the case where a Join Event is generated by the air-ground subnetwork. If the subnetwork cannot generate a Join Event then the procedures start with the Call Request, as part of a polling procedure. System “A” is the initiator and System “B” is the responder. If the air-ground subnetwork is air-initiated then System “A” represents the Airborne Router, and System “B” the Ground Router. If the air-ground subnetwork is ground-initiated, then System “A” represents the Ground Router, and System “B” the Airborne Router.

The Route Initiation Procedures are:

1) When an aircraft attaches to an air-ground subnetwork, a Join Event is generated, potentially to both Airborne and Ground Routers. If received by System “B”, the Join Event is ignored; System “B is ready to receive incoming calls as soon as it attaches to the Mobile Subnetwork.

2) System “A” acts on a Join Event by initiating the establishment of a virtual circuit to the address given by the Join Event, provided such a connection is permitted by local policy, or

3) if polling, System “A” issues a Call Request to the next address on its poll list.

4) When an incoming call is received by System “B”, it accepts the call if permitted to do so by local policy, and generates and sends an ISH PDU to System “A” over the newly established virtual circuit. This ISH PDU includes the NET of the System “B” Network Entity.

5) When System “A” receives a Call Accept, it too generates an ISH PDU, and sends it to System “B” over the newly established virtual circuit. This ISH PDU includes the NET of the System “A” Network Entity.

6) On receipt of the ISH PDU, both systems update their local FIB to include the routing information received on the PDU, and

7) if one does not already exist, the local IS-SME creates an Adjacent BIS MO for the remote system identified by the ISH PDU, and issues a “Start Event” action to that MO. The Adjacent BIS MO created in System “A” identifies the system as being in the passive role, while the System “B” MO identifies the system as being in the active role. Hence on receiving the start event, System “A” simply listens for an incoming BIS OPEN PDU, while System “B” generates one and sends it to System “A”. System “A” responds to the OPEN PDU, with its own OPEN PDU.

8) Alternatively, if a BIS-BIS connection already exists with the remote system, then the IDRP Decision Process is re-run.

9) Once the BIS Open PDUs have been exchanged, the Route Initiation procedures have been completed.

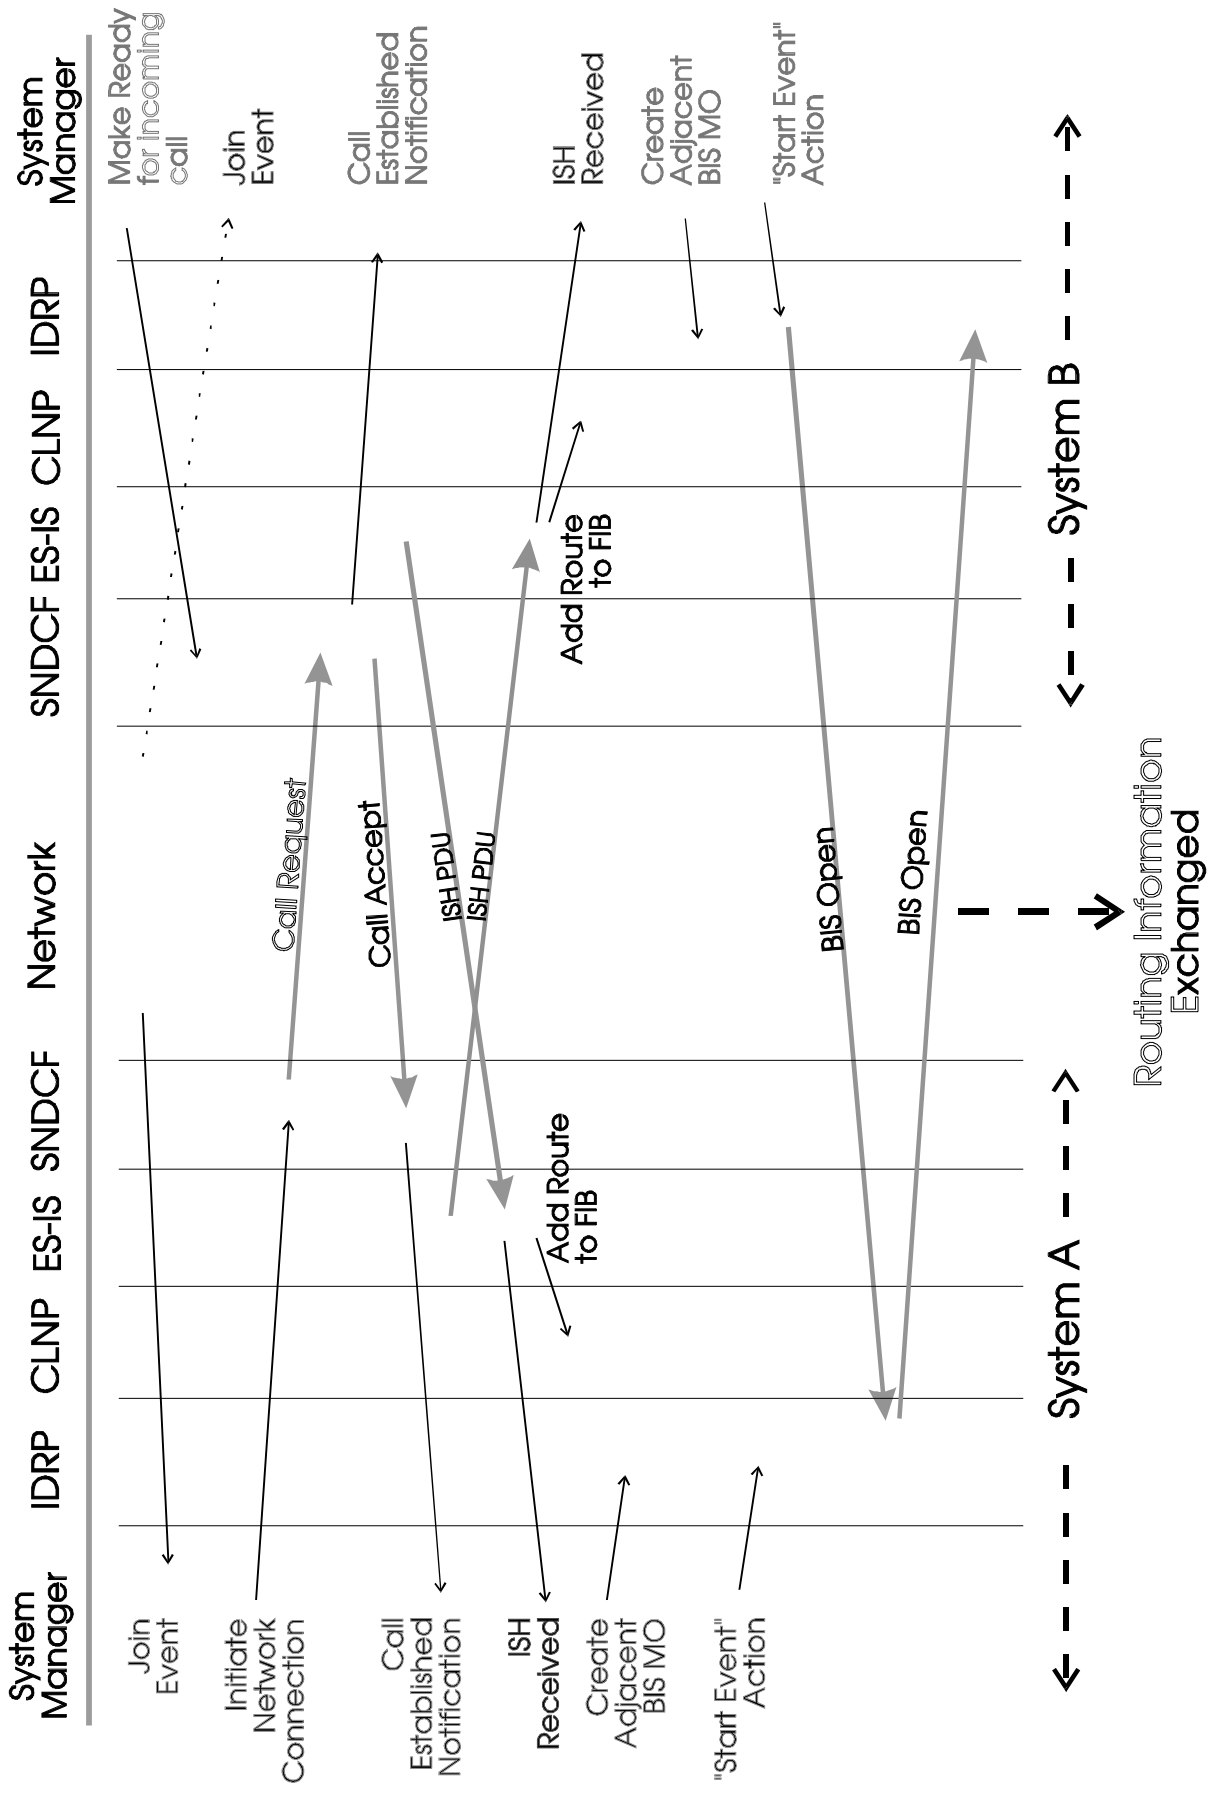

Figure 5-12 Air-Ground Route Initiation Procedures

5.10.3.3 Initial Route Initiation

In the air-ground environment, Route Initiation starts with the notification that an aircraft has come into contact with an air-ground subnetwork, and that a BIS-BIS connection should be established, so that dynamic routing information may be exchanged. In order to ensure the automatic and timely execution of these procedures, a management entity is required by the ATN SARPs to be implemented in each airborne Router and each ground Router with air-ground connectivity. This known as the "Intermediate System - Systems Management Entity" (IS-SME).

Note: The IS-SME is part of the Systems Management Agent for that Router and may also implement other functions outside of the scope of Routing Initiation.

The IS-SME may have to handle two different classes of air-ground subnetwork:

1) Air-Ground subnetworks that can recognise when an aircraft has come into contact with the subnetwork (e.g. logged on to a satellite, or captured by a Mode S Interrogator) and hence that a communications path may be established with that aircraft, and which report this event.

2) Air-Ground Subnetworks which have no mechanism for recognising the above event and/or reporting it.

In the former case, Route Initiation procedures commence when the air-ground subnetwork reports this event - known as the "join" event. In the latter case, Route Initiation additionally includes procedures to allow support for Route Initiation in the absence of such an indication.

Note: Only when air-ground communications are air-initiated is it possible to establish communications without a join event.

5.10.3.3.1 The Join Event

Ideally, the Join Event should be a Systems Management Notification sent to the IS-SME from a Management Entity in the subnetwork itself. This notification should provide the following information:

1) A subnetwork identifier allowing the BIS to associate the event with an air-ground subnetwork to which the Router is connected.

2) The address on that subnetwork of the remote airborne or ground Router.

3) The expected lifetime of the adjacency i.e. how long a communications path is expected to be available.

A Ground Router will typically receive a join event for each aircraft that joins each air-ground subnetwork to which the ground Router is attached. The receipt of such join events will therefore be a regular activity. An airborne Router will typically receive a join event for each ground Router on an air-ground network at the time it comes into contact with that air-ground subnetwork.

On receipt of a Join Event, an ATN Ground Router will, if communication is ground initiated, issue a call request to the subnetwork Address reported by the Join Event and thence establish a virtual circuit with the corresponding Airborne Router. An ATN Ground Router will ignore any Join Events received from air-initiated Air-Ground subnetworks.

Likewise, on receipt of a Join Event, an ATN Airborne Router will, if communication is air initiated, issue a call request to the subnetwork Address reported by the Join Event and thence establish a virtual circuit with the corresponding Ground Router. An ATN Airborne Router will ignore any Join Events received from ground-initiated Air-Ground subnetworks.

In each case, the QoS, Security and Priority requested on the call request should be satisfactory for the exchange of routing information. A local policy decision may also be taken to ignore a Join Event from certain sources.

5.10.3.3.2 The Join Event for Subnetworks that do not support ATN Systems Management

It is anticipated that not all ICAO air-ground subnetworks will support the ATN Systems Management protocols. In order to provide the equivalent of the join event, this Guidance Material provides the following guidance describing an alternative procedure for passing a join event to an air-ground Router. Future ICAO SARPs for air-ground subnetworks which do not specify support of ATN Systems Management should specify the following procedures or an equivalent procedure.

1) A communications path (e.g. a virtual circuit) is established between the ATN Router and a subnetwork processor (e.g. Mode S GDLP) by a Systems Manager and kept open as long as both Router and subnetwork are active.

2) Join events are passed from subnetwork processor to Router over this subnetwork connection and as discrete items of data (e.g. as a single packet), and passed to the IS-SME.

3) The Join Event packet is formatted as a sequence of fields according to 4th.

5.10.3.3.3 Procedures for Air-Ground Subnetworks that do not Provide a Join Event

With this class of subnetwork, it is necessary to adopt a polling strategy in order to establish air/ground communications, and an Airborne Router must "poll" a list of Ground Routers that has been configured by the System Manager.

A suitable "poll" is a periodically repeated Call Request packet addressed to the DTE Address of a Ground Router. Such call requests are regularly repeated until they are answered with a Call Accept from the addressed Ground Router, and an Airborne Router may cycle through a list of Ground Router DTE Addresses until a connection is established. The QoS, Security and Priority requested on this Call Request should be satisfactory for the exchange of routing information.

Once a virtual circuit has been established, the Router may cease to cycle through its poll list, until the connection terminates (e.g, because the aircraft goes out of range of the mobile subnetwork), when it must resume polling for another connection. However, this may lead to unnecessary gaps in communications availability. Furthermore, not all ground Routers will support all security types required by the aircraft. The airborne Router is thus recommended to continue to cycle through its poll list, even when subnetwork connections exist, and to poll the remaining DTE Addresses on the poll list. Polling need only stop when the Router has made sufficient air/ground connections to satisfy its requirements for each supported traffic type, QoS and availability. Polling may resume when these requirements cease to be met

Note: Typically, there will be many more Airborne Routers on a mobile subnetwork than there are Ground Routers, regardless of the subnetwork's coverage area. Hence, while an Airborne Router can be expected to be configured with a complete list of Ground Router DTE Addresses, it is unlikely to be practicable for a Ground Router to be configured with a complete list of Airborne Router DTE Addresses. This is why subnetworks which do not provide information to DTEs on the connectivity status of other DTEs are only considered suitable for air-initiated BIS-BIS connections.

Field	Size, octets	Format	Status	Contents
Message ID	1	binary	required	'1'
Length	1	binary	required	Total message length, in octets
Version	1	binary	required	'1'
Lifetime	2	binary	required	Lifetime of link, in seconds
SNPA	var	type/len/value	optional	Remote ATN Router DTE address(es) now available

Notes:

1. The length field defines the length of the entire message, including the message identifier field

2. The value of the lifetime field is determined by the subnetwork processor. This value should be set to the expected time (in seconds) that connectivity over the mobile subnetwork is expected. A typical value would be on the order of 600 - 1200 seconds (10 - 20 minutes). Note that if air/ground connectivity is still possible shortly before expiration of the lifetime, the SP should re-issue the routing initiation event.

3. The SNPA field contains the subnetwork address of the remote Router. For example, the routing initiation event delivered to the aircraft Router contains the SNPA of the ground Router(s). The actual SNPA may have a different format or length for each subnetwork (for an 8208 subnetwork, the SNPA is the equivalent to the DTE address). The three subfields, type, length, and value are set as follows:

   a) a one-octet type field is set to '1', indicating the field as type "SNPA"

   b) a one-octet length is set to the length of the remote Router SNPA address

4. the variable-length value contains the actual DTE address of the remote Router

5. Multiple SNPA fields may be included within a single routing initiation event to report the reachability of several Routers simultaneously .

6. The VER field should be set to '1'.

7. The value of the type field identifying the following data to be of type "SNPA" should be set to '1'

Table 5-1 Join Event Format

5.10.3.4 Route Initiation in CLNP

As a result of the handling of the Join Event or the "polling" procedure described above, a virtual circuit will have been established between Airborne and Ground Routers. The Mobile SNDCF specified in the ATN SARPs should also have been assigned to support the use of this virtual circuit by CLNP. As with ground-ground Route Initiation, it is now necessary for the IS-SME to add to each Router's FIB, a route to the NET of the remote Router's Network Entity, using the newly established virtual Circuit.

However, all each Router knows at this point is the DTE Address of the other Router. In order to avoid the maintenance problem inherent in managing lookup tables that would enable a correspondence to be made between a DTE Address and a NET, a dynamic procedure has been specified by the ATN SARPs.

An ISO 9542 IS Hello (ISH) PDU is used for this purpose. This is sent either as data, once the connection has been established, or as part of the Call Request/Call Confirm dialogue when "Fast Select" is supported by the air-ground subnetwork. Both Airborne and Ground Routers generate an ISH PDU that reports their NET to the other Router. On receipt of an ISH PDU, each Router updates its FIB with a route to the remote Router, using the NET supplied by the ISH PDU and associating this NET with the subnetwork connection over which the ISH was received, as the forwarding path.

Note: this procedure is also used to negotiate the interim procedures used when IDRP is not supported by the Airborne Router.

5.10.3.5 Route Initiation in IDRP

Route Initiation in IDRP in the air-ground case is then almost identical to the ground-ground case, except that the SARPs require that one Router is in the passive mode and the other in the active mode. This is because the efficiency improvement gained by this approach is worthwhile in the air-ground environment, and the active and passive roles can be unambiguously identified when ICAO air-ground data networks are used.

The SARPs specify that for air-initiated air-ground subnetworks (i.e. AMSS and VDL), that the Ground Router takes on the active role and the Airborne Router takes on the passive role. For ground-initiated air-ground subnetworks (i.e. Mode S), the SARPs specify that the Airborne Router takes on the active role and that the Ground Router takes on the passive role. This approach will permit the exchange of route initiation data to take place in the shortest timeframe.

The Adjacent BIS MO, if it does not already exist, must be created in response to a notification that an ISH PDU has been received over a new subnetwork connection. It is necessary to create this MO in response to receipt of the ISH PDU, because it is not realistic to pre-configure an Adjacent BIS MO for every Airborne or Ground Router to which it could be connected.

An IDRP "Start Event" is then invoked by the IS-SME, provided that a BIS-BIS connection does not already exists with the remote system. If a BIS-BIS connection does already exist then, as in the ground-ground case, and for the same reasons, the IS-SME must cause the IDRP Decision Process to be re-run.

5.10.4 Air-Ground Route Initiation without IDRP

Due to avionics limitations, the ATN SARPs permit, as an interim measure, the existence of ATN Airborne Routers which do not support IDRP. Modified Route Initiation procedures are specified to identify such Airborne Routers and thence to infer the routes that would have been distributed had IDRP been implemented.

Note 1: The identification of routes by inference is only possible because aircraft are required by the ATN SARPs to be End Routing Domains. That is they do not relay data between ground stations or to other aircraft, and hence only provide routes to their local Routing Domain.

Note 2: The consequence of this procedure is that aircraft cannot be dynamically informed about ground route availability. Therefore, until this interim measure has been withdrawn, the ground ATN environment must be constructed to ensure a higher level of availability than would have been necessary had dynamic information been available to all aircraft. This is because, when aircraft make assumptions about ground route availability, those ground routes must exist within the margins of tolerance necessary for air safety.

5.10.4.1 Summary of Procedures

The procedures for Air-Ground Route Initiation without IDRP are illustrated in Figure 5-13, and summarised below. They are described in greater depth in the following sections. The figure illustrates the case where Air-Ground Routing is ground-initiated. The Route Initiation Procedures are:

1) When an aircraft attaches to an air-ground subnetwork, a Join Event is generated, potentially to both Airborne and Ground Routers. If received by System “B” (the Airborne Router), the Join Event is ignored. System “B is ready to receive incoming calls as soon as it attaches to the Mobile Subnetwork.

2) System “A” (the Ground Router) acts on a Join Event by initiating the establishment of a virtual circuit to the address given by the Join Event, provided such a connection is permitted by local policy, or

3) if polling, System “A” issues a Call Request to the next address on its poll list.

4) When an incoming call is received by System “B”, it accepts the call if permitted to do so by local policy, and generates and sends an ISH PDU to System “A” over the newly established virtual circuit. This ISH PDU includes the NET of the System “B” Network Entity, with the NSEL set to the conventional value of hexadecimal **FE**.

5) When System “A” receives a Call Accept, it too generates an ISH PDU, and sends it to System “B” over the newly established virtual circuit. This ISH PDU includes the NET of the System “A” Network Entity.

6) On receipt of the ISH PDU, both systems update their local FIB to include the routing information received on the PDU, and

7) System “A” generates the derived routes using the NET of System “B”, inserts them into the IDRP RIB, and invokes the IDRP Decision Process.

8) System “B”, generates the derived routes from its local “look up” table and inserts them into its local FIB. If for any derived route, an alternative route exists via a different Ground Router to the same destination then only that with the highest degree of preference as indicated by the look up table is inserted in the FIB.

5.10.4.2 Initial Route Initiation

There is no difference in the initial Route Initiation procedures when IDRP is not used over the air-ground data link.

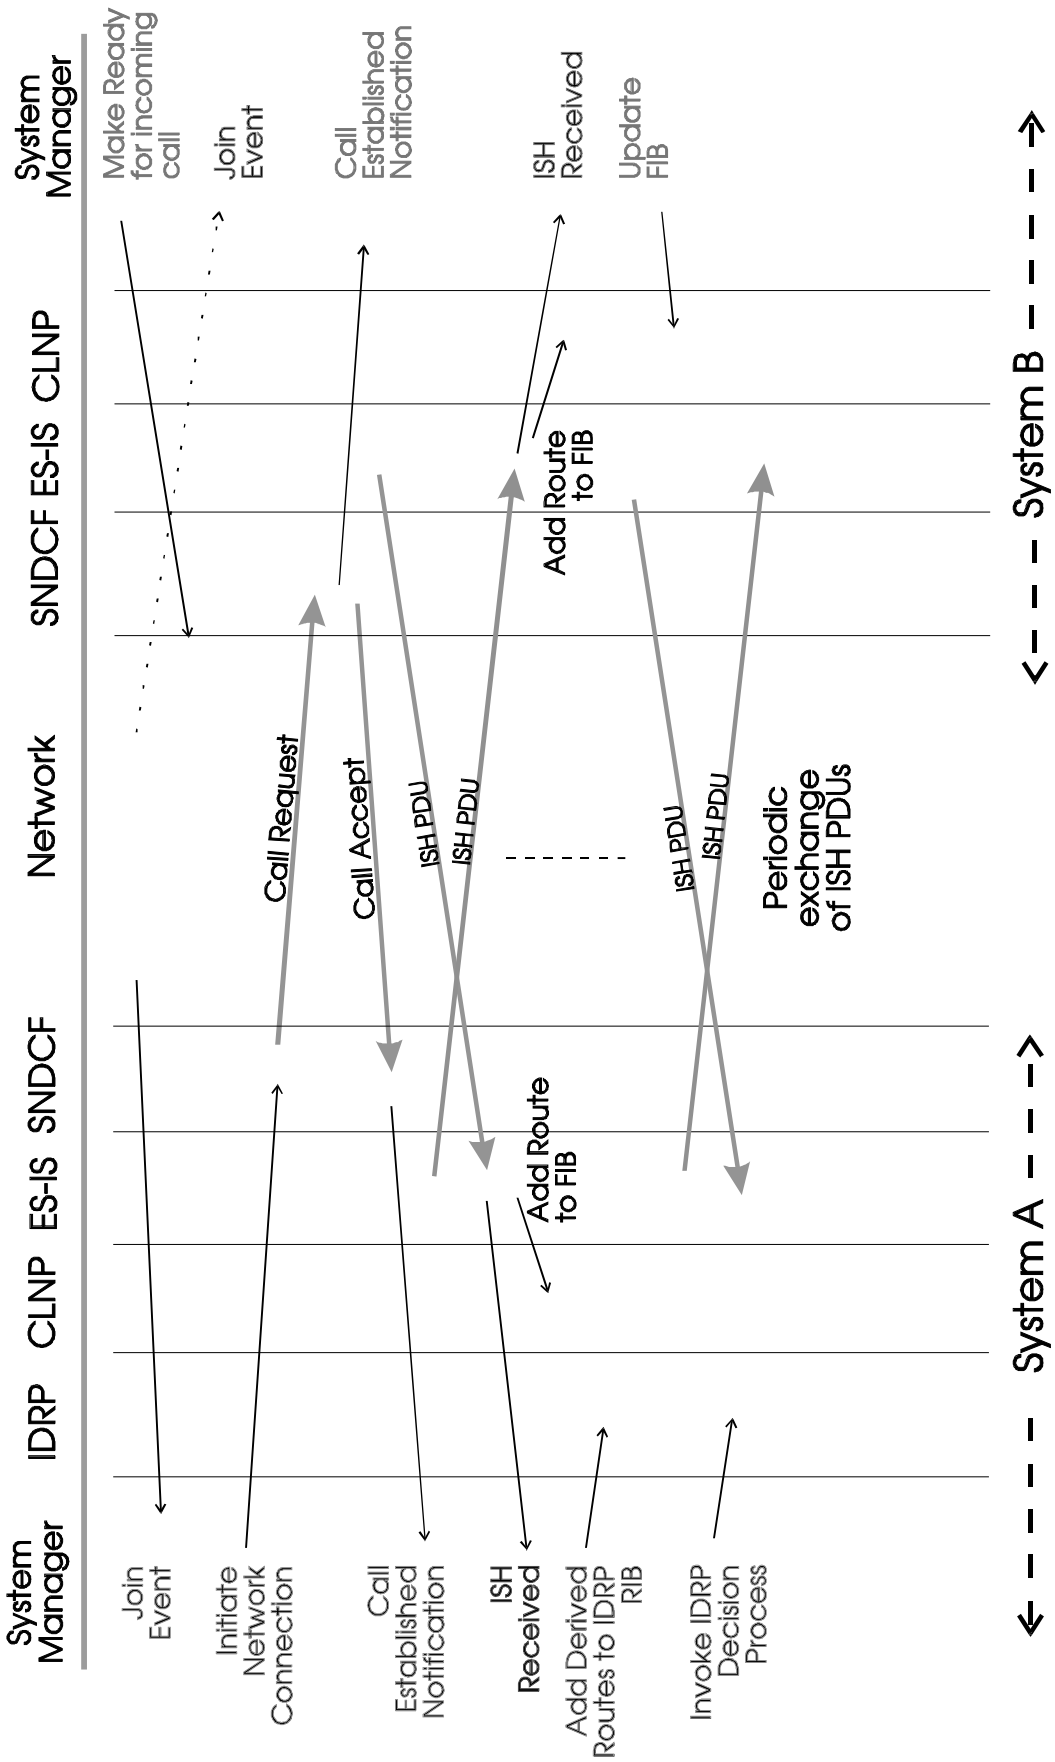

Figure 5-13 Air-Ground Route Initiation without IDRP

5.10.4.3 Route Initiation in CLNP

The ATN SARPs require that the NET of an ATN Router's Network Entity has a Network Selector (NSEL) of zero. This is in accordance with ISO 10589. The SARPs further specify that Airborne Router's that do not support IDRP over the air-ground data link, have an alias NET with an NSEL value of hexadecimal '**FE**', and that this NET is used in the ISH PDU passed over the air-ground data link.

Note: that support of a NET with an NSEL of zero is necessary in such Airborne Routers when, for example, they also support ISO 10589 within the aircraft.

Receipt of an ISH PDU with a NET that has an NSEL of hexadecimal '**FE**' indicates to the receiving Ground Router that the sending Airborne Router does not support IDRP. The IS-SME must then apply the special procedures detailed in the following section.

5.10.4.4 IS-SME Procedures without the use of IDRP

5.10.4.4.1 In the Ground Router

When the IS-SME receives a notification that an ISH PDU has been received from an Airborne Router that does not support IDRP, it must derive the routes that are available via the Airborne Router and add these routes to the local IDRP Routing Information Base (RIB). IDRP may then update the FIB and distribute these routes in the normal fashion.

The derivation of routes is possible because the aircraft is known to comprise an End Routing Domain, and from knowledge of the ATN Addressing Plan it is possible to determine an NSAP Address Prefix common to all systems in the aircraft from the NET of the Airborne Router. Further, from *a priori* knowledge of ITU restrictions that may apply to each air-ground data network and the Quality of Service offered by each such data network, the distinguishing path attributes appropriate to the routes may also be determined.

The number of routes derived by the Ground Router in respect of a specific Airborne Router will be determined by the number of different Application Security Types permitted by ITU restrictions to pass over the air-ground subnetwork multiplied by the number of QoS metrics appropriate to the network. Each such route will have as its Network Layer Reachability Information (NLRI), an NSAP Address Prefix constructed from the first eleven octets of the received NET. That is because the ATN Addressing Plan results in a common eleven octet prefix for all NSAP Addresses and NETs in one aircraft's Routing Domain, which may therefore be determined by inspection of any NSAP Address or NET from any system in that Routing Domain.

The IS-SME must then add those routes to the IDRP RIB and run the IDRP Decision Process, which then disseminates those routes and adds them to the FIB in line with the existing Routing Policy, and provided that they are a preferred route to the Airborne Router.

The actual strategy for doing this is implementation specific. However, a likely strategy is for the IDRP implementation to allocate special "adj-RIB-ins" (one per RIB-ATT) for holding routes received by mechanisms outside of the scope of IDRP. The Decision Process will then consider such routes along with those in "normal" adj-RIB-ins. The only distinguishing aspect of such routes is that they will include the "EXT_INFO" path attribute. This is a flag that enables Routing Policy to differentiate between routes that have been advertised by IDRP throughout and those which have been learned through some other mechanism, perhaps less reliable. As in the general case, the Decision Process must be able to associate this special Adj-RIB-in with the connections to the Airborne Router, and the QoS provided by these connections . This is so that when computing the degree of preference for each such route, or when copying them to the loc-RIB, the Decision Process can update their QoS to reflect the current communications paths that exist to the Airborne Router.

If additional subnetwork connections are opened up (or lost) to an Airborne Router then, instead of generating the derived routes, as before, the IS-SME must cause the IDRP Decision Process to be re-run.

Finally, in this interim role, the IS-SME must also determine when the assumed routes are no longer valid. This event occurs when either the air-ground subnetwork connection is lost or when the periodic exchange of ISH PDUs ceases. On the occurrence of either such event, the routes generated above must be withdrawn.

Note: that in contrast with the use of IDRP over an air-ground data link, when the ATN SARPs recommend that for reasons of efficient bandwidth utilisation, ISH PDUs are not periodically transmitted, in this case they must be periodically transmitted in order to maintain the "liveness" of the routes.

5.10.4.4.2 In the Airborne Router

The IS-SME procedures are in this case, similar to the ground case, except that:

a) the NLRI of the generated routes cannot be simply derived from the Ground Router's NET. This is because the Ground Router is typically part of a Transit Routing Domain, and the destinations of the onward routes that it offers will not have any known relationship to its NET.

b) The generated routes must be directly added to the FIB as IDRP is not present to do this on behalf of the IS-SME, or

c) if ISO 10589 is implemented, the generated Routes are used to generate Reachable Address MOs and the ISO 10589 entity is used to update the FIB.

In order to determine the NSAP Address Prefixes for the generated routes, lookup tables will have to be provided so that given the NET of a Ground Router, the Airborne Router can identify the NSAP Address Prefixes for destinations reachable via that Ground Router. Furthermore, such look up tables will have to provide:

i) restrictions on Security Types for such destinations that are additional to ITU restrictions imposed by the Air-Ground Subnetwork;

ii) The Capacity, Hop Count and QoS information for such destinations in a manner sufficient to enable alternative routes to be discriminated between. i.e. an indication of relative preference for each supported metric.

Operationally, there will be a need to ensure that such tables are up-to-date with information appropriate to the Flight Region(s) through which the aircraft will fly, prior to each flight. The actual implementation of this procedure is dependent on the systems involved.

The IS-SME will have to keep dynamic information on which routes are available via each Ground Router with which it is in contact. This information is derived from the look up table and *a priori* information for each Air-Ground Subnetwork supported. When multiple subnetwork connections exists to a given Ground Router then the routing information will be determined taking into account the characteristics of each such subnetwork.

When routes to the same destination are available via different Ground Routers, then the IS-SME will have to choose between them based on the degree of preference given by the look up tables.

The IS-SME is also responsible for maintaining the FIB with an up-to-date set of available preferred routes determined as above. It must add such routes to the FIB when they become available, and remove them when the reverse is true, Alternatively, if ISO 10589 is implemented, then the IS-SME may make such routes available to 10589 by creating a Reachable Address MO for each such route,

and removing the MO when the route ceases to be available. The ISO 10589 implementation may be relied upon to maintain the FIB with this routing information.

5.10.4.5 Management of the ISH PDU Holding Time

An ISH PDU exchange is a common feature for data link use, whether or not IDRP is also being used. However, in either case, it is important to set the ISH PDU Holding Time parameter with due care to avoid sending unnecessary ISH PDUs. In doing so, it is necessary to understand the main purposes of the ISH PDU exchange:

1. The ISH PDU exchange is first used to negotiate the use or non-use of IDRP.

2. The initial ISH PDU exchange is also used to avoid any pre-defined relationship between NETs and DTE Addresses. This is believed essential if ATN Airborne and Air/Ground Routers are to operate over many different types of air/ground data links with differing addressing plans, including future networks whose characteristics may not even be known for some time.

3. The ISH PDU can also be used to provide a check on the "liveness" of the data link, if the data link does not provide this as a built-in feature i.e. if the data link service does not provide timely information on the loss of a communications path. Note that ISH PDUs are sent on a per data link basis and not on a per adjacency basis and such liveness tests are specific to an individual data link.

The Holding Time is a parameter to an ISH PDU that specifies the maximum time for which the receiving network entity can retain the configuration routing information contained in the PDU. When an ISH PDU is received, the receiving network entity should start a timer which expires after the indicated Holding Time has elapsed. That timer is then restarted whenever a further ISH PDU is received from the same sender. If the timer does expire, then the receiving Network Entity will purge routing information about the NET contained in the ISH PDU, from its routing tables. The route to the indicated NET will therefore cease to be available. ISH PDUs must thus be retransmitted at a rate that is typically half that of the Holding Time, in order to ensure that the receiving Network Entity's routing information is up-to-date, and that routes are not lost through loss of a single ISH PDU.

When the procedures for the optional non-use of IDRP are employed, non-receipt of an ISH PDU within the expected time will additionally cause the downstream IDRP route to be withdrawn. When IDRP is being used, the same event will cause loss of communications between the adjacent BISs and, in consequence, the withdrawal of any routes advertised over the adjacency.

There are two factors involved in setting the ISH PDU Holding Time. The first is whether the underlying data link needs a "liveness" check. The second is the application requirement for notifying the using application, in a timely manner, of the loss of a communication path. Note that if a supported application requires a particularly rapid notification of the loss of a communications path then it may be necessary to have a regular exchange of ISH PDUs even when the data link also incorporates its own liveness check. That is if the data link's liveness check is not frequent enough for such an application.

In most cases, Airborne and Air/Ground Routers will set the ISH PDU Holding Time to the largest possible value (i.e. 65534). This will avoid unnecessary ISH PDU exchanges and hence costs. Only when *a priori* it is known that a data link does not have a suitably frequent check on liveness for the supported applications, should a shorter time be used. In such cases, the actual value for the Holding Time must necessarily depend upon application requirements.

Airborne Router implementors should note that Air/Ground Routers are generally in a better position to know *a priori* whether a short Holding Time is required. Airborne Routers implementors may therefore consider a pragmatic strategy whereby the first ISH PDU sent over a newly established data link always has a large Holding Time value set and then, if an ISH PDU is subsequently received from an Air/Ground Router with a short Holding Time, that Holding Time is also adopted by the

Airborne Router. That is, should an Airborne Router see an incoming ISH PDU with a short Holding Time, it should respond with an ISH PDU with the same Holding Time, and continue to use that short Holding Time on the same data link.

Implementors should also note that existing implementations of ISO 9542 were probably developed for the LAN environment and assume a low transmission cost and unreliable delivery. Such implementations will probably respond to an incoming ISH PDU from a previously unknown system with their own ISH PDU. Such behaviour is totally unnecessary on a reliable point-to-point data link and should be suppressed, if possible, in order to avoid the cost of transmission.

5.11 Support for Mobile Systems

5.11.1 Mobility and Routing Domains

The scaleability of an Internet is enhanced when Routing Domains near to each other are characterised by similar address prefixes. However, this is not an absolute requirement. Routing Domains can be adjacent, have totally dissimilar address prefixes and still interconnect successfully. Furthermore, with a dynamic routing protocol, such as IDRP, two Routing Domains need only to interconnect when they need to, and can both be active on the same network. The onward re-advertisement of routes can inform the rest of the ATN Internet about such a temporary connectivity while it exists, and the loss of connectivity when it occurs. A Routing Domain can thus temporarily join an Internet at one point of attachment, then disconnect and join the Internet at some other point, the only impact being in the efficiency of routing information distribution, and eventually on scaleability.

This property of the routing architecture and of IDRP, is exploited by the ATN to support Mobile Routing.

In the ATN, the systems onboard an aircraft form a Routing Domain unique to that aircraft and characterised by one address prefix for ATSC systems, and another for AISC systems. As an aircraft proceeds on its route, it interconnects with ground based Routing Domains over the various air/ground networks; the actual network used and Routing Domain interconnected with are dependent on the aircraft's actual position, and the airline's routing policy. Routing Information is then exchanged between ground Routing Domains, using IDRP, so that all ground Routing Domains are aware of the current route to that aircraft. This is illustrated in Figure 5-14.

In this example, there are four ground based Routing Domains RD1 through to RD4. RD1, RD2 and RD3 all support air/ground datalinks, while RD4 depends on the other three for air/ground communications. The aircraft currently has communications over air/ground datalinks with both RD2 and RD3.

Using IDRP, both RD2 and RD3 advertise a route to the aircraft's systems, to RD4. RD4 chooses between these two available routes using its own Routing Policy, which might, for example, favour the route through RD3. Similarly, the aircraft's router must choose between the routes to RD4 offered by RD2 and RD3. It need not make the same choice as RD4.

As the aircraft continues on its journey, it may lose communication with RD3. For example, it goes out of range of the VHF datalink it was using to communicate with RD3. RD3 informs RD4 of this situation by issuing the appropriate IDRP protocol action to withdraw the route, and RD4 now changes to using the route offered by RD2, as it is now the only route to the aircraft. The aircraft's router also recognises the loss of communication with RD3 and must now route all traffic via RD2.

Further on the journey, the aircraft comes into contact with an air/ground datalink offering communication with RD1. A datalink is established and routing information exchanged. RD1 now advertises the new route to the aircraft, to RD4. RD4 now once again has two routes to the aircraft and must make a choice between them using its local routing policy rules. It might, for example,

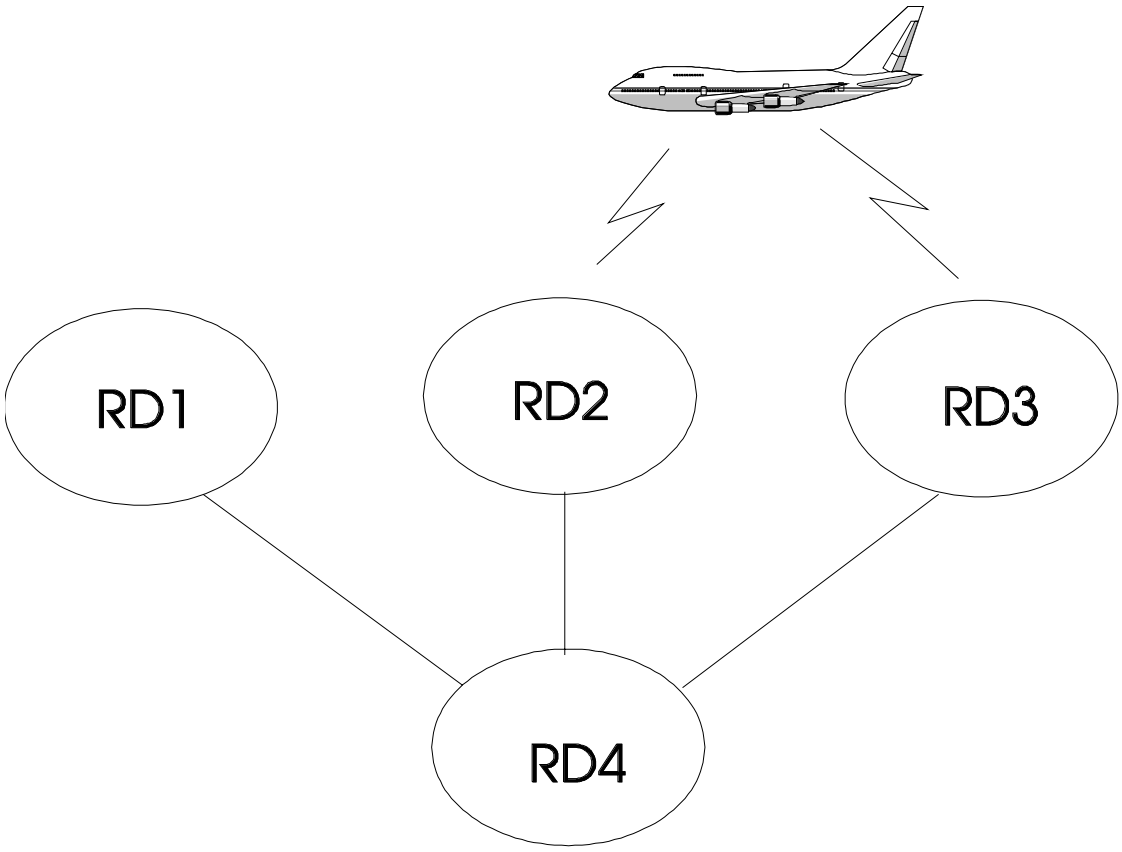

Figure 5-14 Mobile Routing Example

now prefer the route through RD1, in which case all data to the aircraft is now routed via RD1. The router in the aircraft also goes through a similar decision process.

While the topology of the ATN ground environment is much more complex than the above example, this is essentially how mobile communications is implemented by the ATN.

5.11.2 Containing the Impact of Mobility

While the principles of mobile routing outlined in the previous section are straightforward they are not scaleable using the existing IDRP mechanisms associated with Route Aggregation and RDCs. The problem is that even if an aircraft is given an address prefix similar to the address prefixes that characterise the ground Routing Domains at the start of its journey, such a similarity is unlikely to be maintained for the duration of the flight. Route Aggregation possibilities are thus very limited.

Instead, an alternative mechanism has been developed to permit mobility within a scaleable Internet architecture, building on two concepts: the ATN Island, and the “Home” domain (see 5.11.4 below). In addition, the ATN Addressing Plan specifies a common address prefix for all aircraft and, subordinate to that address prefix, specifies a unique address prefix for the aircraft belonging to each airline, and the General Aviation Aircraft of each country.

5.11.3 Routing to Mobiles within an ATN Island

The ATN island exists for the exclusive purpose of supporting routing to mobiles. An ATN Island is simply an ATN region comprising a number of Routing Domains, some of which support air/ground datalinks. These Routing Domains form an RDC, as illustrated in Figure 5-15, and an ATN Island is essentially an RDC in which certain Routing Policy rules are followed. All ATN Routing Domains that have air/ground datalink are members of an ATN Island and, although most ATN Routing Domains which do not have air/ground datalink capability will also be members of ATN Islands,

they do not have to be and can still have access to routes to aircraft if they are not a member of an ATN Island RDC. Routes to destinations in ground based Routing Domains will be exchanged by ATN Routing Domains, both within an Island and between Islands. However, this is outside of the context of the ATN Island.

Within each ATN Island, at least one Routing Domain forms the Island's *backbone*. This may be only one RD or may actually be an RDC comprising all backbone Routing Domains in the same ATN Island.

Within the ATN Island, the Backbone RDC provides a default route to *all aircraft*, as illustrated in Figure 5-14, this is advertised to all other Routing Domains within the Island as a route to the common address prefix for all aircraft.

Routing Domains with routes to aircraft then have a simple routing policy rule to determine to which adjacent Routing Domain they must advertise such a route[1]. This is the Routing Domain currently advertising the preferred route to *all aircraft*. This will be a backbone Routing Domain (or a Routing Domain that provides a route to the backbone). Either way the impact of such a policy rule is that the Backbone RDC is always informed about routes to all aircraft currently reachable via datalinks available to the Island's Routing Domains, and can thus act as default route providers for packets addressed to airborne systems.

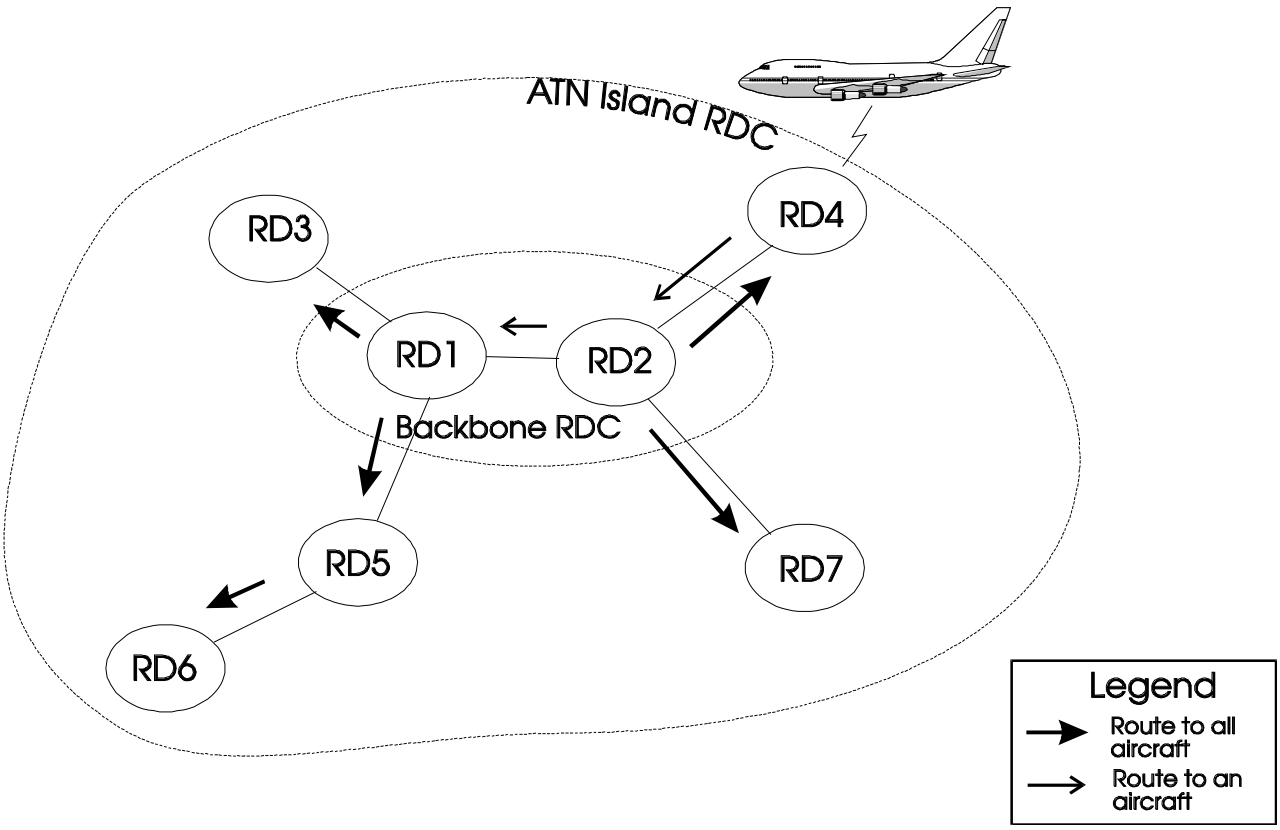

Figure 5-15 Mobile Routing Within an ATN Island

Routing Domains off the backbone also have a simple routing decision to make when they need to route a packet to a given aircraft. It is routed along the explicit route to the aircraft if it is known by them, or on the default route to all aircraft via the backbone. Routing with IDRP always prefers routes with the longest matching address prefix. Since the default route to all aircraft is always a

[1] A route to an aircraft is readily identifiable from the destination address prefix, as all address prefixes that characterise an aircraft Routing Domain descend from a unique address prefix.

shorter prefix of that for an explicit route to an aircraft, the explicit route to an aircraft will be preferred (since it will always have a longer matching address prefix). This routing strategy happens automatically without any special provisions.

The example above is not the only policy rule that can apply to routes to aircraft. Routes to aircraft can be advertised to any other Routing Domain within the Island, provided that a policy rule is set up to allow this. This may be because there is a known communication requirement which makes bypassing the backbone desirable, or because it is desirable to provide a second (hot standby) route to aircraft from the backbone. The architecture accommodates these requirements. The only limitation on this is that imposed by the overhead of supporting routes to mobiles (see 5.11.6 below).

Within the Backbone RDC, all Routing Domains must exchange all routes to aircraft, which are advertised to them, they are then able to act as default routers to any aircraft currently in communication with the ATN Island. However, because the backbone routers need to know routes to all such aircraft, their capacity places a limit on the number of aircraft that can be handled by an ATN Island and hence on the effective size of the Island.

The ATN Island is only the first part of achieving a scaleable routing architecture for mobile routing. Its true benefit is to focus the overhead of handling the potentially large number of routes to aircraft on a few specialised routers in the backbone. Off the backbone, a Routing Domain with an air/ground datalink needs only the capacity to handle the aircraft supported by its datalink, and there is a similar impact on Routing Domains that are Transit Routing Domains providing a route between the backbone and an air/ground datalink equipped Routing Domain. For all other Routing Domains on the Island, there is no impact on routing overhead due to aircraft.

In the absence of a backbone, all routers within the Island would need to be explicitly informed with a separate route to each aircraft, if they were to be able to route to any aircraft currently in contact with the Island. This is because there is very little probability of route aggregation with routes to aircraft.

5.11.4 Routing to Mobiles between ATN Islands

ATN Islands can be set up such that their geographical spread matches Air Traffic Control communication requirements and, for ATC purposes, there may not be a requirement to provide inter-Island communications in respect of aircraft. However, airline operational requirements are perceived to require this, and hence the mobile routing concept is developed to provide a greater level of scaleability.

The mechanism used to achieve this derives from the concept of the “Home” domain.

Aircraft for which inter-Island communications are required must have a “Home” domain, which is a Routing Domain in an ATN Island’s backbone. This “home” need not be in either the ATN Island through which the aircraft is currently reachable, or in the ATN Island with which communication is required. The role of the “Home” domain is to advertise a default route to all the aircraft belonging to an airline, or the General Aviation aircraft of a given country of registration. This default route is advertised to all other ATN Island’s backbone routers.

The operation of the “Home” domain is illustrated in Figure 5-16. In this example, ATN1 is the ATN Island acting as the “Home” for all aircraft belonging to the same airline as the aircraft illustrated as currently reachable via ATN4. ATN1 advertises the default route to all such aircraft to all Islands in which it is in contact and, depending on local policy this route may be re-advertised to other Islands. In the figure, ATN3 re-advertises the default route on to ATN4.

The backbone routers of an ATN Island have a simple policy rule to implement for each explicit route to an aircraft that they have available. If a default route to all the aircraft in the aircraft’s

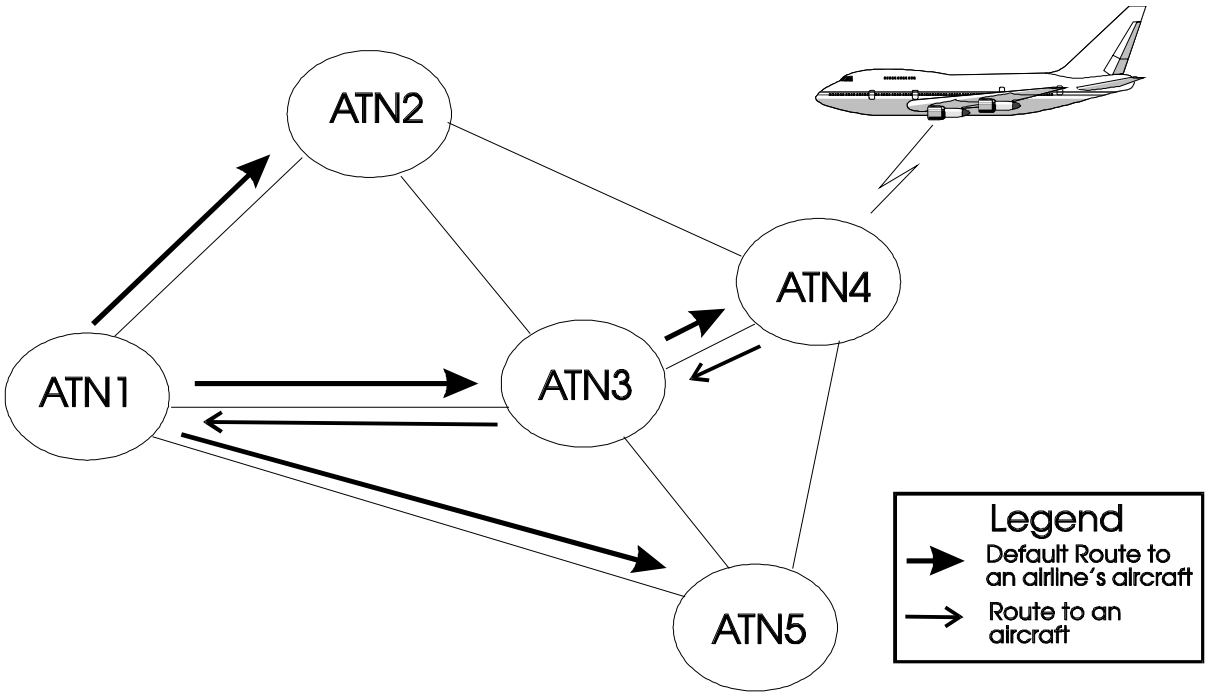

Figure 5-16 Inter-Island Routing

airline or country of registration exists[2] then the actual route to the aircraft is advertised to the Routing Domain advertising that default route. Otherwise, the explicit route is not advertised outside of the Island. In Figure 5-16, the route to the aircraft is first advertised by ATN4 to ATN3 and then re-advertised to ATN1. In each case, the same policy rule is applied.

The impact of this rule is that the "Home" is always kept aware of routes to all of "its" aircraft. As it is also providing the default route to such aircraft, routers on other ATN Islands (e.g. ATN2) that have packets to route to one of that "Home's" aircraft will by default send those packets to the "Home" Routing Domain (ATN1), where the actual route to the aircraft is known, and thus the packet can be successfully routed to the destination aircraft (via ATN3 and ATN4).

In the above example, this is clearly non-optimal as ATN4 can be reached directly from ATN2. However, the loss of optimal routing is acceptable as, otherwise a scaleable architecture could not have been developed.

The impact of this strategy on routing overhead, is that an ATN Island backbone has to be capable of handling routes to all aircraft currently in contact with the Island, and all aircraft for which it is the "Home".

However, this capacity handling requirement is independent of the total number of ATN Islands or the total number of aircraft. It is thus possible to add more ATN Islands, or aircraft belonging to airlines whose "Homes" are on other Islands, without affecting the capacity of an ATN island backbone (relating to the number of routes to aircraft). The routing architecture thus allows for a much larger number of mobile systems than that permitted by a single ATN Island.

[2] Such a route is generated by the "Home" Domain , and is readily identifiable from the destination address prefix, as all address prefixes that characterise an aircraft belonging to the same airline descend from a unique address prefix.

5.11.5 Impact on Air/Ground Datalinks

A final limiting factor on the ATN is the capacity of the air/ground datalinks. At present, these are low bandwidth communications channels and only the minimum routing information can be transferred over them.

IDRP is potentially an ideal protocol for this environment. Techniques such as RDCs and Route Aggregation can be used to minimise the information contained in each route. Furthermore, two or more routes to the same destination that differ only in security parameters, or service quality metrics, can be combined together into a single message keeping the actual information exchanged to a bare minimum.

In addition, IDRP is a connection mode protocol and, as such, once a route has been advertised between a pair of Boundary Intermediate Systems it does not have to be retransmitted during the lifetime of the connection. A BIS-BIS connection is kept alive by the regular exchange of small "keepalive" packets, and once routing information has been exchanged it remains valid for the lifetime of the connection without having to be retransmitted.

The ATN uses these properties of IDRP to keep the transfer of routing information over an air/ground datalink to a minimum. When the datalink is first established, the airborne router will advertise a route to internal destinations for each combination of traffic (security) type and QoS metric supported. These routes will be combined into a single protocol message and downlinked for onward distribution through the ground ATN.

The ground router will also uplink routes to the aircraft and to keep the information down to a minimum, a further RDC is defined, comprising all ground ATN Routing Domains. This RDC, the "ATN Fixed RDC" ensures that for each uplinked route, the path information is collapsed to a single identifier, that for the ATN Fixed RDC.

The actual routes uplinked are subject to the policy of the ground router's Routing Domain. However, it is anticipated that routes will be provided to at least:

- the local Routing Domain (typically that providing Air Traffic Services), and
- the ATN as a whole,

in addition to other routes as determined by local policy.

The airborne router will then be able to choose between the alternative routes (via different) ground routers to these destinations.

5.11.6 The Impact of Routing Updates

As indicated in the previous section, a scaleable routing architecture can be developed in support of mobile routing. It is now necessary to consider the factors that limit the number of routes to aircraft that an ATN Router can handle.

Each route known to a router occupies a certain amount of data storage and, while data store can be a limiting factor on the total number of routes handled, it is unlikely to be so in this case. The number of route updates that a router can handle is more than likely to be the limiting factor.

In the ground environment, route updates will usually only occur when changes occur in the local region of the Internet (changes further away are hidden by route aggregation). Typically the introduction of a new Routing Domain or interconnection, or the removal or loss of one of these will cause a change. However, the frequency of update is unlikely to be high.

However, with mobiles, such as aircraft, the situation is very different. Aircraft are constantly on the move, changing their point of attachment to the ATN, and hence generating routing updates. The

impact of these updates needs to be minimised if the number of aircraft that can be handled by an ATN Island is to be maximised, and an important and useful feature of IDRP can be exploited in order to help meet this objective.

5.11.6.1 "Hold Down" Timer Use

Vector distant routing protocols, such as IDRP, typically implement a "hold down" timer, which introduces a minimum delay between the receipt of a route and its re-advertisement. This timer is used to avoid instability due to frequent route changes, and the actual value of the timer is then usually a trade-off between a short timeout to give rapid response and a long timer to keep down routing overhead and minimise instability.

However, under IDRP, routing events that indicate a major change (i.e. new route or loss of a route) are not subject to a hold down timer, only those that report a minor change to an existing route are subject to a hold down timer. This means that IDRP is very responsive to connectivity changes while avoiding instability due to minor changes. For example, consider a simple extension to the previous example, illustrated in Figure 5-17.

In this example, RD4 provides a route to the aircraft, to RD5. When the aircraft loses contact with RD3, RD4 is immediately informed, as there is an effective zero length hold down timer for withdrawn routes. However, while RD4 recognises this event and switches to the route provided by RD2, it does not necessarily inform RD5 of this now minor change to the route immediately (the route still exists, only the detail of the path is different), and anyway, the update must be sent not less than the period **minRouteAdvertisementInterval** since any previous update. In this example, it should be noted that the minor change will not affect RD5's routing decision, as it has no alternatives available.

Sometime later, the aircraft comes into contact with RD1. RD4 is immediately informed as this is a new route. However, even if RD4 switches to this new route, it does not inform RD5 of the change until the **minRouteAdvertisementInterval** has again expired.

This has important implications for the design of an ATN Island. If an Island's air/ground datalinks are all connected to Routing Domains which are themselves adjacent to the Backbone RDC, all connectivity changes will be immediately reported to the Backbone giving a high route update rate. On the other hand, if there are intermediate Routing Domains between the backbone and the Routing Domains connected to air/ground datalinks, then the update frequency can be significantly reduced, without affecting the responsiveness to real connectivity changes.

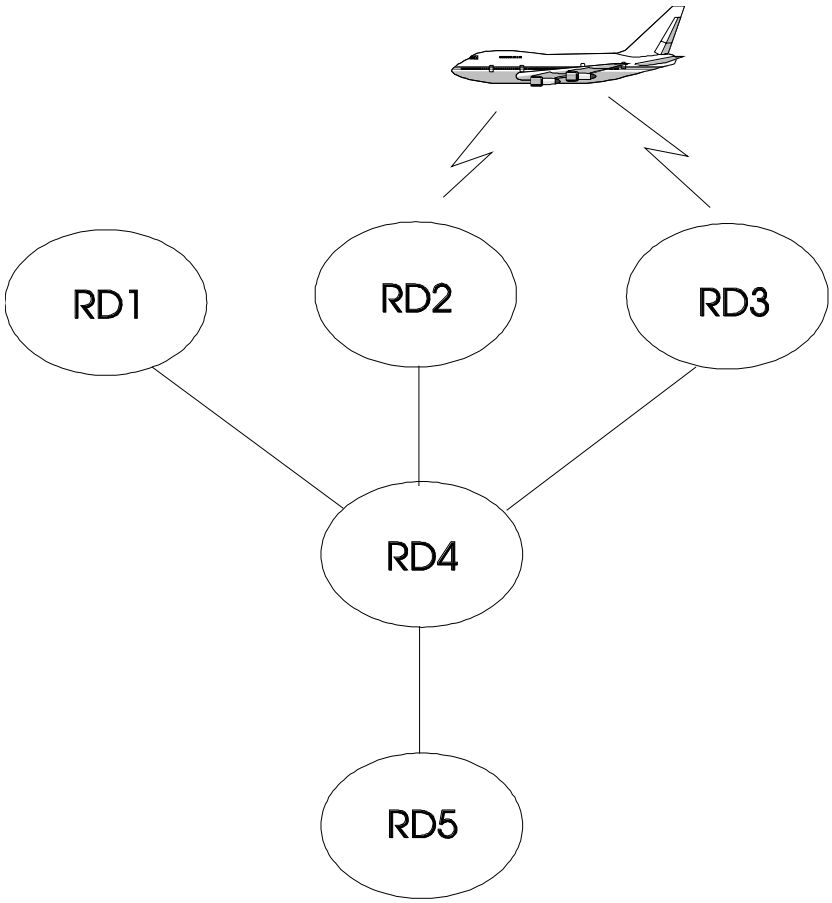

Figure 5-17 Impact of a Hold Down Timer

This is an important benefit derived from using IDRP to support mobile routing compared with, for example, a directory based approach to mobile routing. Under a directory based approach, there would be a central directory server on each ATN Island (c.f. the Backbone), updates on the position of aircraft would be sent direct to the directory, and other routers would consult the directory in order to determine the current location of a specific aircraft. In terms of overhead, this situation is analogous to an ATN Backbone Routing Domain directly connected to each Island Routing Domain with air/ground datalink capability, and the directory has to be able to take the full update rate. IDRP can, however, distribute the update load throughout the ATN Island.

Routes advertised to an aircraft's "Home" are also affected by the hold down timer and, in this case, RDCs and the Hold Timer work together to keep the routing overhead to an absolute minimum.

As an ATN Island is an RDC, routes advertised to other Islands have their path information for the transit through the RDC replaced by a single RDC identifier, and therefore, in many cases, changes in the route will not even be visible to another ATN Island. When changes are visible (e.g. a change in hop count or QoS metric), and such changes can be kept to a minimum by careful network design, then the Hold Timer limits the rate at which such changes can be advertised and prevents minor changes which are also short lived, being exported outside of the Island.

Results from simulation work have shown that the "ideal" setting for the **minRouteAdvertisementInternval** is approximately one minute. Furthermore, complex topologies for the ATN Island Backbone should be avoided as they significantly increase the convergence time. Typically, an ATN Island Backbone should consist of a small number of routers linked as a chain with ring shaped topologies avoided.

5.11.7 Failure Modes

In the pure ground-ground environment, loss of a router or a communications path can be readily recovered from provided an alternative route exists and routing policy permits its use. However, the situation is not so straightforward with the policy rules that support mobile routing. The ATN Mobile Routing Concept depends upon two default route providers, the Island Backbone and the "Home". Failure of either of these or loss of access to them will impact mobile routing.

5.11.7.1 Loss of the "Home"

Loss of the "Home" may come about from either the loss of the Routing Domain advertising a route to the "Home" for a given set of aircraft, or the loss of the communications path to it. The consequence of either failure is clear: the affected aircraft are now only reachable from systems on the ATN Island to which they are currently adjacent.

In practice, there should not be a single point of failure related to the "Home" Routing Domain. A Routing Domain may comprise many BISs, each of which may advertise the route to the "Home". Only loss of all of these BISs will result in the complete loss of the route to the "Home". Furthermore, there may be many communications paths, using different network technologies, linking two adjacent Routing Domains. Such concurrent links may be between the same pair of BISs, or between different pairs. Only if all such links are lost, will total loss of communications occur.

Therefore, it will always be possible to design a network topology that will avoid the loss of the "Home" being due to any single failure, and which can ensure that the probability of loss of the "Home" is kept within acceptable limits. Where inter-Island communications are required in support of air safety, then the design of the Inter-Island ATN topology must be supported by an appropriate failure mode analysis to ensure that safety limits are maintained.

5.11.7.2 Failure of an ATN Island Backbone

Failure of an ATN Island may also result from the failure of the Routing Domain(s) that comprise an Island's Backbone, or of communications paths with an Island's backbone. The consequence of such a failure is that the aircraft currently adjacent to the Island are only reachable from the Routing Domains supporting air/ground datalinks with those aircraft, and any other Routing Domains on the Island to which routing information to those aircraft is advertised according to explicit policy rules.

For similar reasons to those already detailed in 5.11.7.1, there is no need for loss of an Island Backbone to be due to a single point of failure, and an appropriate network design should be developed for each ATN Island to ensure that the probability of the loss of the backbone is within acceptable limits.

5.11.8 Optional non-Use of IDRP

Simple networks can often avoid dynamic routing mechanisms in favour of statically defined routing tables, initialised by a System Manager. However, even in the early ATN, the existence of Mobile Systems does not permit the general use of static routing techniques. Aircraft may join and leave the air/ground subnetwork(s) at any time and this dynamic behaviour must be recognised by the routers and reflected in the routing tables. Some dynamic adaptive routing protocol is needed to support this requirement. IDRP is specified for this purpose. However, implementing IDRP functionality on an airborne router may not be practicable in the early stages of ATN implementation.

An alternative approach is possible using provisions in the ISO 9542 ES-IS protocol. An exchange of Intermediate System Hello (ISH) PDUs is already required as part of the route initiation process, and, in a limited topology, an exchange of ISH PDUs can be sufficient to provide the exchange of dynamic routing information necessary to support mobile routing. Furthermore, a regular exchange of ISH PDUs (part of the normal operation of ISO 9542) can be used to keep the link between ground and airborne routes "live" in the absence of IDRP.

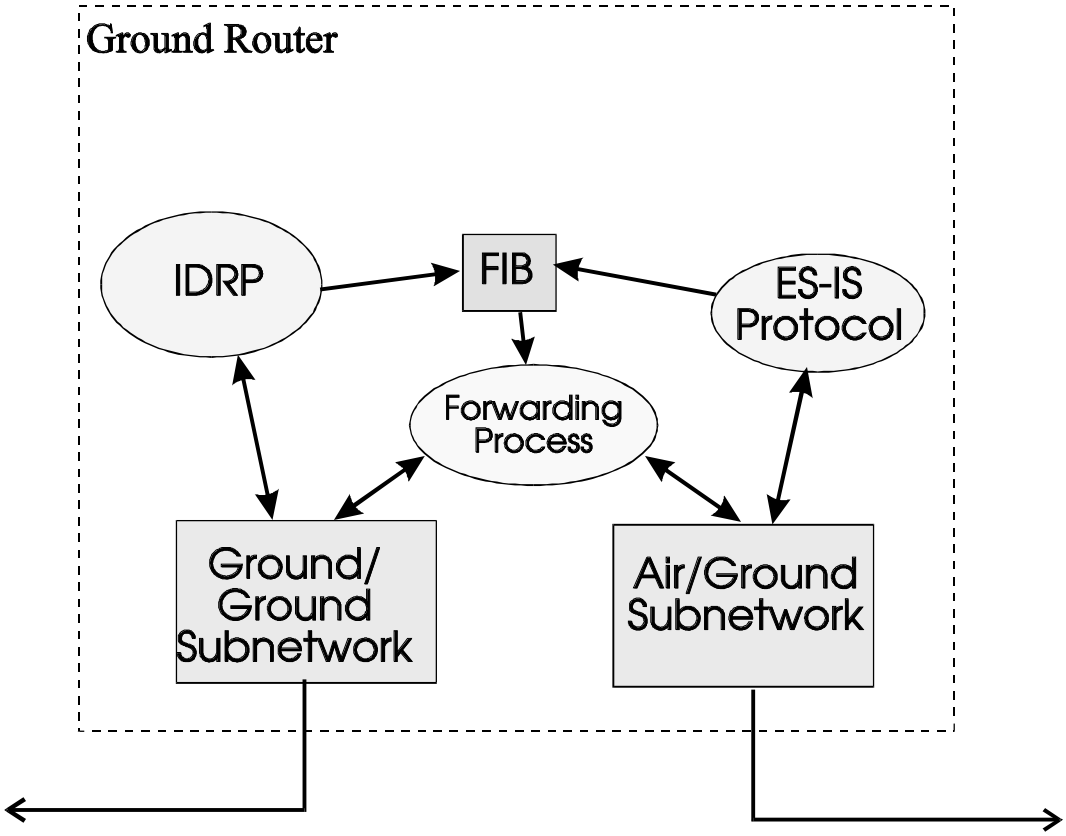

Figure 5-18 Architecture of an Initial Ground Network Router

Such a use of the ISH PDUs depends upon an assumed relationship between the Network Entity Title (NET) of each router - which is essentially the router's address - and the NSAP Addresses in the ground and airborne End Systems. The NET is exchanged as part of the ISH PDU. When the Air/Ground router receives an ISH PDU from an airborne router, it may infer from the ATN Addressing Plan the common NSAP address prefix of all NSAPs onboard that aircraft. This being the first eleven octets of the NET. This NSAP Address Prefix may then be used as the destination of a route to the NSAPs onboard that aircraft and the route entered into the ground router's Forwarding Information Base. It is then possible for the ground End Systems to send data to airborne End Systems on that aircraft.

The same process may also take place on the Airborne Router, on the receipt of an ISH PDU from the Air/Ground router, enabling airborne End Systems to send data to ground End Systems. The routing information remains current until either a regular exchange of ISH PDUs ceases, or the subnetwork connection is cleared, when the ground and airborne routers remove the associated routes from their forwarding information bases.

The architecture of a ground router implementing such functionality is illustrated in Figure 5-18. The architecture is straightforward enough with the ES-IS protocol active on both subnetworks. Both protocol entities update the Forwarding Information Base (FIB) which is, in turn, used by the Forwarding process to route packets.

As the ISH PDU mechanism is also used for route initiation in the full ATN, some convention for distinguishing between its use in this scenario and in the full ATN is necessary. This can be readily achieved by addressing conventions. A non-zero value in the NET's "SEL" field (254 decimal) is used to signal use of the above procedures.

Routing information learnt in this way by the Air/Ground Router may then be disseminated throughout the ATN Ground Environment using normal IDRP procedures.

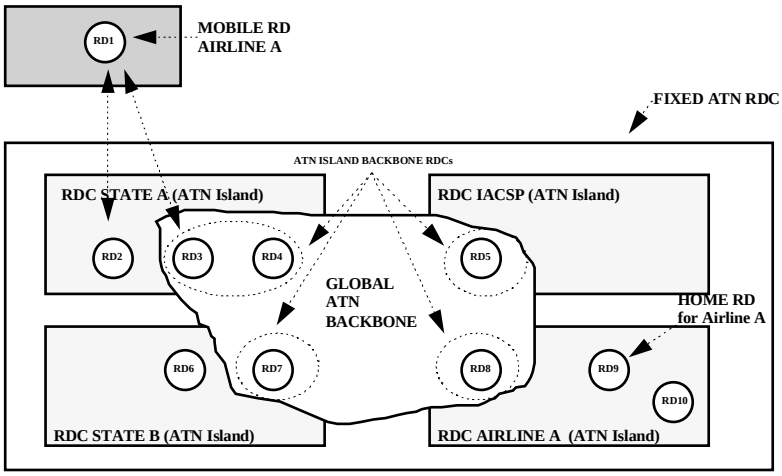

Figure 5-19 Example Routing Policy Scenario

5.11.9 Routing Policies in Support of Mobile Routing

No special features of IDRP are required to implement the mobile routing strategy described above, other than the ATN specific use of the Security Path Attribute. Instead a prescribed set of Routing Policies are used to provide this functionality. These rules are fully specified in section 5.7.3 of the ATN Internet SARPs, and it should be noted that different sets of rules apply to ATN Routers in different roles. This section attempts to illustrate the application of those rules by describing an example network of routers and discussing the application of the rules to this example network.

Figure 5-19 below defines the example Routing Architecture scenario that has been used as the basis for the guidance provided in this section.

The following are the key components of the example network:

1) The scenario defines at the highest level the "Fixed ATN RDC" which the SARPs define to comprise of all fixed ATN RDs;

2) Within the Fixed ATN RDC are defined four organisational RDCs[3]:

   - an RDC for State "A";
   - an RDC for State "B",;
   - an RDC for an International Aeronautical Communications Service Provider (IACSP); and
   - an RDC for an airline (airline "A").

1) The scenario additionally defines a Mobile RD (RD1) belonging to airline "A" that is currently connected with two RDs (RDs 2 & 3) within the State A RDC.

2) The connectivity between the RDs is illustrated in Figure 5-20.

The Following RDCs are defined:

1) State A RDC

[3] The term "RDC" is synonymous with the term "ATN Island".

- The State A RDC comprises three RDs (RD2, RD3 and RD4).
- The State A RDC includes a Backbone RDC that comprises RDs 3 & 4 and a TRD (RD2) off the Backbone.

2) State B RDC

- The State B RDC comprises two RDs (RD6 and RD7).
- One RD (RD7) is the only member of the State B Backbone RDC.

3) IACSP RDC

- The IACSP RDC comprises one RD (RD5) which is the only member of the IACSP RDCs Backbone RDC.

4) Airline RDC

- The Airline A RDC comprises three RDs (RD8, RD9 and RD10).

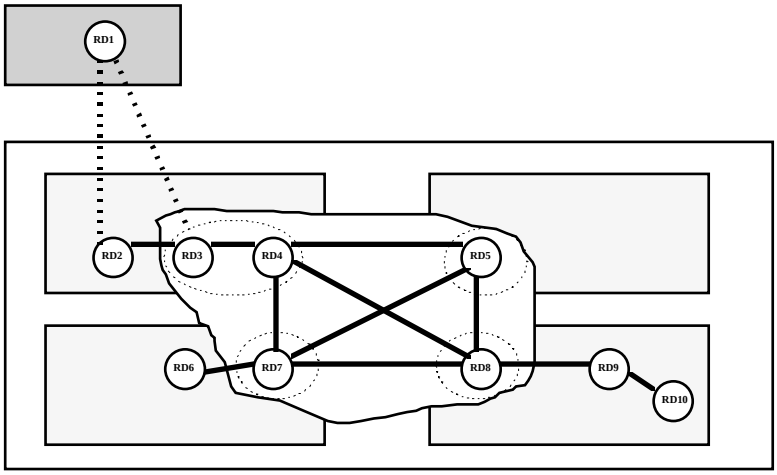

Figure 5-20 Routing Policy Example Connectivity

- The Backbone RDC comprises on RD (RD8).
- RD9 is a TRD and is designated as the "Home RD" for the airline.
- RD10 is defined to be an ERD.

Mobile RD

- RD1 is a Mobile RD belonging to Airline A.

Global ATN Backbone

- The "Global ATN Backbone" comprises all RDs that are members of the Backbone RDCs of each of the 4 organisational RDCs i.e. RDs 3, 4, 5, 7 & 8.

It should be noted that an overriding requirement in the SARPs is that all Routers within the same RD are required to implement the same Routing Policy. With respect to the Routing Policy rules defined in the SARPs, and explained in the following sections, it should be noted that rules have only been defined in support of air/ground routing. Routing Policy rules for ground/ground routing have been considered to be a local matter and are therefore outside the scope of the SARPs.

Table 5-2 Routing Policy Requirements for Members of an ATN Island Backbone RDC (SARPs Ref. 5.3.7.2)

SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
5.3.7.1.2	**Adjacent ATN RD's within the ATN Island Backbone RDC**	The policy requirements are applicable to the exchange of routing information between adjacent routing domains both of which are members of the ATN Island Backbone RDC.	RD3→RD4 RD4→RD3	Each Router in RD3 is required to advertise the following routes to each adjacent Router in RD4: • a route to NSAPs & NETs contained within RD3; • the selected Route to every Mobile for which a route is available i.e. either direct to the mobile RD1 from RD3 or a route via RD2 to the mobile RD1; • the selected route to every Fixed ATN RD in the same Island i.e. a Route to RD2.
5.3.7.1.3	**All other ATN RDs within the ATN Island**	The policy requirements are applicable to the advertisement of routing information from an RD that is a member of an ATN Island Backbone RDC and an RD that is not a member of the ATN Island RDC but belongs to the same ATN Island.	RD3 → RD2 RD7→RD6 RD8→RD9	In this case RD8 will advertise Routes to RD9: • a route to NSAPs & NETs contained within RD8; • the selected Route to every Fixed ATN RD in the same ATN Island for which a Route is available (not applicable in this example); • a Route to all Mobile RDs thereby providing a default Route to all Mobiles; • a Route to each Mobile RD (i.e. to Mobile RD1) for which the adjacent RD (RD9) is advertising a Route to the Mobile RDs Home.
5.3.7.1.4	**Mobile RDs**	The policy requirements are applicable to the advertisement of routing information between a Router in an RD that is a member of an ATN Island Backbone RDC and a Router in an adjacent	RD3→RD1	In this case RD3 will advertise to the Mobile RD1; • a Route to NSAPs & NETs contained within RD3. The SARPs additionally recommend that RD3 should advertise to the Mobile RD1: • an aggregated Route to NSAPs & NETs contained within the State A

SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
		Mobile RD.		RDC (i.e. the local RDC) and; • an aggregated Route to NSAPs & NETs contained within the State B RDC, the IACSP RDC and the Airline RDC (i.e. all other Island RDCs for which a Route is available).
5.3.7.1.5	**ATN RDs in other ATN Islands**	The policy requirements are applicable to the advertisement of routing information by a Router in an RD that is a member of an ATN Island Backbone RDC and a Router in a RD that belongs an adjacent ATN Island Backbone RDC.	RD4→RD5 RD4→RD8 RD4→RD7 RD5→RD4 RD5→RD8 RD5→RD7 RD7→RD4 RD7→RD5 RD7→RD8 RD8→RD4 RD8→RD5 RD8→RD7	For example RD8 will advertise the following Routes to all adjacent Routers (RD4, RD5, RD7) in adjacent Island Backbone RDCs: • an aggregated Route to NSAPs and NETs contained within the Airline RDC; • a Route to all Mobile RDs assigned to Airline A since the Home RD (RD9) belongs to the same Island as RD8; • a Route to each Mobile RD for which the adjacent RDs are advertising a route to the Mobile RD's Home (not applicable in this example). However, RD4 would advertise a Route to Mobile RD1 to RD8 since RD8 would be advertising a Route to the Home for the Mobile RD1. • a Route to each Mobile RD for which there is no home (not applicable in this example). However, if a Mobile RD was connected to either RD8, RD9 or RD 10 then RD8 would advertise this route to RDs 4, 5 & 7.

Table 5-3 Routing Policy Requirements for a Mobile RD (SARPs Ref. 5.3.7.2)

SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
5.3.7.2.1	**Mobile RD**	The policy requirements relate to the advertisement of routing information between a Router in a Mobile RD and all ground Router (irrespective of whether or not they belong to one or more RDs) to which it is connected.	RD1→RD2 RD1→RD3	For example the Mobile RD1 will advertise to RDs 2 & 3 a Route to NSAPs and NETs contained within mobile RD1.

Table 5-4 Routing Policy Requirements for an ATN TRD that is not a member of the ATN Island Backbone RDC (SARPs Ref. 5.3.7.3)

SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
5.3.7.3.2	**Adjacent ATN RDs that are members of the ATN Island's Backbone RDC**	The policy requirements are applicable to the advertisement of routing information from Routers in a TRD that do not belong to the ATN Islands Backbone RDC to adjacent Routers that are members of the ATN Island's Backbone RDC.	RD6→RD7 RD2→RD3 RD9→RD8	For example RD9 (TRD) will advertise to RD8: • a Route to NSAPs & NETs contained within RD9; • the selected Route to every Mobile RD for which a Route is available (not applicable in this example). However, the rule is applicable to RD2 which would advertise to RD3 a Route to Mobile RD1. • the selected Route to every Fixed ATN RD in the Airline Island i.e. a Route to RD10; • a Route to each Home that the TRD itself (i.e. RD9) provides for Mobile RDs (e.g. for Mobile RD1).
5.3.7.3.3	**Adjacent ATN RDs within the same ATN Island and which are not members of the ATN Island's Backbone RDC**	The policy requirements are applicable to the advertisement of routing information from routers in a TRD that do not belong to the ATN Islands Backbone to a router in an adjacent RD which also does not belong to the ATN Islands Backbone.	RD9→RD10	In this example RD9 would advertise to RD10: • a Route to NSAPs and NETs contained within RD9; • the selected Route to every Fixed RD in the Airline Island for which a Route is available i.e. a Route to RD8; • if RD9 is currently advertising the preferred Route to all Mobile RDs (which is must be since there is no alternative available) then every known Route to a Mobile is advertised to RD10 from RD9; • the preferred Route to all Mobiles i.e. via RD8; • a Route to each Mobile RD for which RD10 is advertising the preferred Route to the Mobile RDs Home (not applicable in this example); • a Route to the Home of all Mobile RDs assigned to Airline A since RD9 is the Home RD for Airline A.

SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
5.3.7.3.4	**Mobile RDs**	The policy requirements are applicable to the routes advertised by a Fixed TRD which is not a member of its Islands Backbone to an adjacent Mobile RD.	RD2→RD1	In this case RD2 will advertise to the Mobile RD1; • a Route to NSAPs & NETs contained within RD2. The SARPs additionally recommend that RD2 should advertise to the Mobile RD1: • an aggregated Route to NSAPs & NETs contained within the State A RDC (i.e. the local RDC) and; • an aggregated Route to NSAPs & NETs contained within the State B RDC, the IACSP RDC and the Airline RDC (i.e. all other Island RDCs for which a Route is available).

Table 5-5 The Routing Policy for a Fixed ATN ERD (SARPs Ref. 5.3.7.4)

SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
5.3.7.4.1	**Fixed ATN ERD**	The policy requirements are applicable to the routes advertised by a Fixed ERD to adjacent RDs to which it is connected.	RD10→RD9	For example RD 10 will advertise to RD9 a Route to NSAPs and NETs contained within RD10.

Table 5-6 RD Matrix

To / From	RD1 Mobile RD	RD2 Fixed TRD off Backbone	RD3 Fixed RD on Backbone	RD4 Fixed RD on Backbone	RD5 Fixed RD on Backbone	RD6 Fixed TRD off Backbone	RD7 Fixed RD on Backbone	RD8 Fixed RD on Backbone	RD9 Fixed TRD off Backbone	RD10 Fixed ERD off Backbone
RD1 Mobile RD		5.3.7.3.4	5.3.7.1.4							
RD2 Fixed TRD off Backbone	5.3.7.2.1		5.3.7.1.3							
RD3 Fixed RD on Backbone	5.3.7.2.1	5.3.7.3.2		5.3.7.1.2						
RD4 Fixed RD on Backbone			5.3.7.1.2		5.3.7.1.5		5.3.7.1.5	5.3.7.1.5		
RD5 Fixed RD on Backbone				5.3.7.1.5		5.3.7.1.5	5.3.7.1.5	5.3.7.1.5		
RD6 Fixed TRD off Backbone							5.3.7.1.3			
RD7 Fixed RD on Backbone				5.3.7.1.5	5.3.7.1.5	5.3.7.3.2		5.3.7.1.5		
RD8 Fixed RD on Backbone				5.3.7.1.5	5.3.7.1.5		5.3.7.1.5		5.3.7.3.2	
RD9 Fixed TRD off Backbone								5.3.7.1.3		5.3.7.4.1
RD10 Fixed ERD off Backbone									5.3.7.3.3	

6. Congestion Avoidance in the ATN Internetwork

6.1 Network Congestion

Congestion is a phenomenon experienced by a Router in an Internetwork when the queuing delays through that Router exceed the maximum acceptable limit. In such a situation, the end-to-end transit delay is likely to exceed the maximum acceptable for the internetwork's users. In the extreme case, a congested router, due to lack of buffer space, may not be able to accept incoming NPDUs at the rate that an adjacent router is trying to send them, and is hence forced to discard lower priority NPDUs, or those near the expiry of their lifetime, in order to make way for higher priority NPDUs.

Congestion is not a problem for an internetwork. Congested routers can simply discard NPDUs when they start running out of buffers. However, it is a serious problem for the users of the internetwork. Congestion first results in an acceptably long transit delay. However, if network users assume that the lack of arrival of an end-to-end acknowledgement is due to packet loss, rather than simply an unexpectedly long delay in the network, then they can retransmit such unacknowledged packets, thus adding to the load on the network.

In fact, a catastrophic degradation in transit delay and throughput can be observed in a congested network. First the network becomes congested, then users start retransmitting, making the network even more congested, resulting in more retransmissions, and so on, until the point is reached where only insignificant amounts of data can be transferred. It is therefore vital that Congestion Avoidance mechanisms are put in place in any internetwork, if it is not to be perceived as unstable and unreliable.

6.2 Possible Techniques

In a connectionless internetwork, Congestion Avoidance has to be a co-operative activity in which a major part is played by the users of the network. Successful operation of the network depends on its users being "good citizens" and reducing the load placed upon the network once the onset of congestion has been determined.

In general, any suitable Congestion Avoidance technique must be able to control overload situations in the underlying network in such a way that data transfer is performed as efficiently as possible. To be acceptable, the adopted Congestion Avoidance technique must satisfy the following goals:

1. High throughput (in bit/s), together with a small end-to-end transit delay, should be experienced by network users.

2. A small buffer load within the traversed routers should be achieved.

3. The probability of packet loss should be minimal.

In pursuit of these goals, two candidate algorithms were initially investigated during the development of the ATN Internet SARPs. These were a sending transport entity back-off algorithm, similar to the

Van Jacobsen Slow-Start algorithm that is widely used in the TCP/IP Internet, and a Receiving Transport Entity Congestion Avoidance algorithm.

Although widely used, the former was rejected. The Slow-Start algorithm probes the network until congestion occurs, when the transport entity backs off and then proceeds to probe again. It is effective when congestion is a rare event, and avoids catastrophic congestion occurring, but is inefficient on a heavily loaded network, as that network is regularly forced into a congested state during the regular "probes". In a mobile network, such as the ATN, there is also considerable scope for the Slow-Start algorithm to be confused by a mobile system changing its point of attachment. The resulting packet losses will be interpreted by the sending transport entity as an indication of a congested network, forcing a back-off state and hence a resulting in a lowering of throughput.

On the other hand, the chosen algorithm relies upon indications received from the network layer (i.e. the CE-bit in an NPDU Header) in order to determine when the network is approaching a congested state, and adjusts the advertised credit window in response. This has the advantages of avoiding the continued probing that is characteristic of the Slow-Start algorithm, and of remaining unaffected by a mobile system changing its point of attachment. It therefore appears to give a significantly better throughput in the aeronautical environment.

6.3 Receiving Transport Layer Congestion Avoidance

6.3.1 Overview

The Receiving Transport Layer Congestion Avoidance algorithm depends on the "Congestion Experienced" (CE) bit that may be included in an NPDU Header. This bit is set initially to zero by the End System that creates the NPDU. Should the NPDU pass through a Router, on its journey through the internetwork, that is either congested, or is nearing the point of congestion, then the CE-bit is set to one by that Router.

When an NPDU is received by the destination End System, it can therefore readily inspect the CE-bit and determine if the NPDU experienced congestion anywhere on its route.

This is a simple mechanism for determining the congested state of an internetwork, and does so without generating additional network traffic. This is important, as a network reaching the point of congestion should not suffer additional traffic, just because it is congested!

When the receiving transport entity gets an NPDU with a CE-bit set to one, it is not required to take immediate action - indeed, it should not do so, as such an isolated event may well be transitory. However, if enough NPDUs are received with a CE-bit set to one, during a suitable sampling period, then it must take action to tell the sending transport entity to slow down and reduce the load it is putting on the network. This is because when NPDUs start to be regularly received with the CE-bit set, then it is indicative of a network that if not already congested, is starting to become so.

The way the receiving transport entity tells the sending transport entity to slow down, is to reduce the advertised credit window. Whenever a received TPDU is acknowledged, an AK TPDU is sent back to the sender that includes the sequence number of the most recently received TPDU and gives permission for the sender to send another *n* TPDUs. Normally, the objective is to acknowledge received DT TPDUs in a timely manner to ensure that the sender never gets into a situation whereby it no longer has permission to send any more DT TPDUs (i.e. that it runs out of credit). The sender is then able to transmit data as fast as it can.

Normally, *n* is set large enough for this to be the case. However, when congestion is detected, if the receiving transport entity sets *n* to a smaller value, the sender will start to occasionally run out of credit, there will be times when it cannot send any DT TPDUs, and hence the load on the network is reduced. Thus the sending transport entity can be readily told to slow down simply by reducing the value of *n*.

Later on, if a smaller proportion of packets are received with the CE-bit set to one, then *n* can be safely increased again, until congestion is once again determined. On a congested network, this

algorithm results in a small oscillation around the ideal data transfer rate, while never pushing the network into a congested state. A high throughput with the minimum of transit delay is thereby achieved, without forcing the routers to discard packets as part of a “probing” process. Our goals for a Congestion Avoidance Algorithm are thereby met.

6.3.2 Determining the Onset of Congestion

In line with the definition given earlier, a router can be considered congested when the queuing delays imposed by a transit through it exceed a certain threshold. A useful metric for congestion can therefore be gained from a simple inspection of the length of the outgoing queue when a forwarded packet is queued for transfer to another system. If the queue exceeds a certain length then the CE-bit should be set to indicate that the queuing delay is excessive i.e. congestion has been experienced. However, what is an appropriate queue length (i.e. the threshold) to determine when the CE-bit is to be set?

When specifying a queue threshold, it is necessary to take into account what it is intended to do with this signal. The final goal is to achieve a data transfer service with fairly good user visible performance (i.e. low end-to-end delay, high throughput), without producing too high a buffer load (so the global network is operating stable). If the buffer load found in any output queue is very large, it runs the risk of packet loss, which will trigger packet retransmissions. These in turn will increase the end-to-end delay of the data transmission (since packet loss first has to be detected and recovered, before normal data transmission can continue) and thus will also reduce the throughput visible to the user. Finally, packet losses put an additional burden on the network, since the lost packet will already have (uselessly) traversed part of the network, before it gets lost.

Since high throughput and low end-to-end delay are competing goals, L. Kleinrock proposed in his standard work on queuing systems to optimise the "Power" of a connection, which he defined as

$$\text{Power} := \frac{\text{Throughput}}{\text{Delay}}$$

This measure has served well since its introduction, and is widely used within network optimisation. By adapting that goal to the problem considered here, we have to derive a threshold value for the output queue load such that the Power of the system is maximised.

To derive an appropriate queue threshold value, we consider an output queue together with its outgoing link as a M/M/1 queuing system (exponentially distributed inter-arrival times, exponentially distributed service times; see also **Error! Reference source not found.**).

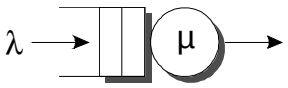

Figure 6-1 A queuing system

Packets arrive at a server with arrival rate λ, where they eventually get queued if the server is currently busy. Packets are fetched from the queue by the server, which forwards packets at a rate of μ. The system is in a stable state only if packets do not arrive faster than the server can forward them, i.e. if and only if $\lambda < \mu$.

Such a queuing system is referred to as an M/M/1 system, and for such an M/M/1 system, the average time a packet spends in the system is given by

$$E[T] = \frac{1}{\mu - \lambda} \qquad (1)$$

The throughput of an M/M/1 system is equal to λ, if the system is operating in a stable state (i.e. one can never receive a higher throughput than the server forwarding rate μ, but the server is also not able to forward packets faster than they arrive).

From the above, the Power of a M/M/1 system thus can be derived to be

$$\text{Power:} = \frac{\text{Throughput}}{\text{Delay}} = \frac{\lambda}{1/(\mu-\lambda)} = \lambda\cdot(\mu-\lambda) \quad (2)$$

This measure is maximised if the following condition holds:

$$\lambda \equiv \frac{\mu}{2} \quad (3)$$

The average number of customers found in a M/M/1 system is given by

$$E[N] = \frac{\lambda}{\mu-\lambda} \quad (4)$$

which, for λ as given in (3) to optimise the power, finally evaluates to a value of 1 packet. If the output queue threshold is thus set to 1 packet, every system will try to operate at a point of maximum power, i.e. offering a high throughput to the user, while also making sure that the end-to-end delay (e.g. for short messages exchanged between communicating entities) is kept reasonably small.

Although it has been suggested that a number greater than one is appropriate low bandwidth data links, consideration of the above shows that there is no justification for this. A value larger than one simply implies that longer queuing delays are tolerated with clear downside implications for throughput (e.g. by requiring a longer retransmission timer, reducing the rate at which AK TPDUs can be sent, etc.).

However, this is not to say that special considerations do not apply to air/ground data links. The queuing model presented in **Error! Reference source not found.**, and the associated argument assumes that all outgoing queues from a given router are independent. This is not true for a network such as AMSS, where a single transponder is shared for communication with all aircraft. Although the Air/Ground Router servicing an AMSS data link will see a separate outgoing queue for each aircraft, the reality is that they are all constrained by a common uplink queue. In such cases, the number of packets on the outgoing queues should be summed up and the CE-bit set when the total packets queued for uplink over the same transponder is greater than one.

6.3.3 Reporting Congestion Experienced to the NS User

Congestion is experienced by an NPDU, while it is an NSDU that is passed to the NS-User as part of an N-UNITDATA.indication. In many, if not most cases, there will be a one to one relationship between NPDUs and NSDUs. In such a case, there is little problem in reporting Congestion Experienced, and, as additional information to the N-UNITDATA.indication, the Network Layer can pass an indication that the NSDU reported congestion experienced on its route from the sender.

However, this leaves open what happens when an NSDU is segmented into two or more NPDUs, some of which may experience congestion, while others do not. Possible strategies for the network layer are to:

1. to indicate to the NS-User both the total number of NPDUs received for a single NSDU, and the number of NPDUs received having the CE flag set to the transport layer;

2. to merge the CE flags received by bitwise ORing all values. Thus, if a single NPDU had the CE flag set, congestion will be indicated to the NS-User.

3. to merge the CE flags received by bitwise ANDing all values. Thus, only if all NPDUs had the CE flag set, congestion will be indicated to the NS-User.

4. to only forward the CE flag setting of the last NPDU received during reassembly of an NSDU to the local transport layer.

Strategy (1) is the preferred strategy. This is because it gives the NS-User the maximum amount of information on which to base a decision. All the alternatives hide information from the NS-User, and there is little value in doing so.

6.3.4 Credit Window Management by the Receiving Transport Entity

The receiving transport entity monitors incoming TPDUs and determines whether or not congestion was experienced by the TPDU during its transit through the internetwork. If, during some sampling period, congestion was experienced by enough TPDUs, then the effective credit window is reduced by multiplying it by a reduction factor β. Otherwise, if the credit window is currently less than the a value which will permit maximum throughput, then it may be increased by adding an integral value δ. Initially, the credit window is set to a low value (e.g. two). The algorithm then ensures that it increases until either maximum throughput is achieved or, congestion starts to be experienced, when the credit window oscillates about the optimal value. Note that starting from a lower value (i.e. one) has a downside in that a credit of one demands two AK TPDUs for each DT TPDU transferred.

Only DT TPDUs are monitored during a sampling period. This is because only DT TPDUs are subject to credit management. Other TPDUs types, such as expedited data or acknowledgements are not subject to credit management, and therefore no feedback can be gained by monitoring them to see if any restrictions on the credit window are working in respect of reducing network congestion.

Furthermore, once a sampling period has been completed, and a new credit window determined, no more sampling should be undertaken for a period equal to the estimated Round Trip Time (RTT). This is because any DT TPDUs received during this period will have been subject to the previous credit window. Only once the RTT has elapsed, can it be assumed that the received DT TPDUs are subject to the new credit regime and hence its effect on the network state can be reasonably determined.

The reason why a “freeze period” is necessary can be readily seen from the following example.

4th depicts a sender transmitting data towards a receiver. Each packet is indicated by a line going from the sender to the receiver. Transmission of a packet through the network takes a certain amount of time, represented by the slope of the line (time proceeds from top to bottom).

Initially, the transmission is performed in this example with a window size of 8. It is also assumed that the network is currently overloaded, so the receiver will see CE flags being set, and reported to the Transport Entity by the Network Layer.

At time t_1, the receiver decides to ask the sender to reduce its load, in order to remove the overload found in the network. The sender is informed about this decision using an AK TPDU, transmitted from the receiver to the sender (indicated by the dashed line in 4th).

Once this indication is received by the sender, the sender reduces its window to the value advertised by the receiver. For the scenario considered here, it is assumed that $\beta = ½$, to better visualise the operation. Thus, the window W is reduced to $W_1 = 8 \times ½ = 4$. Afterwards, the sender is transmitting with a smaller load, indicated by a greater spacing of the packets.

As can be seen from the figure, immediately after the decision to reduce the advertised window, the receiver will continue to get packets still transmitted with the old window. These may also have their CE flag set, since the sender is not yet aware of the decision to reduce the window, and still is transmitting with the large window. It takes approximately one round trip time, until the first DT TPDU transmitted at the lower load (i.e. with the reduced window size) arrives at the receiver. Note that within this time interval, W_0 packets will be received that still have been transmitted using the old window size.

During the next round trip time (starting at time t_2), DT TPDUs being sent using the reduced window size, arrive at the receiver. Assuming that the load is now small enough, there will be no more congestion within the network. In consequence, these packets will not have their CE-flag set. The receiver will thus see another 4 packets without the CE flag set. After the second RTT (i.e. at time t_3), the receiver will make a new decision how to modify the advertised window size.

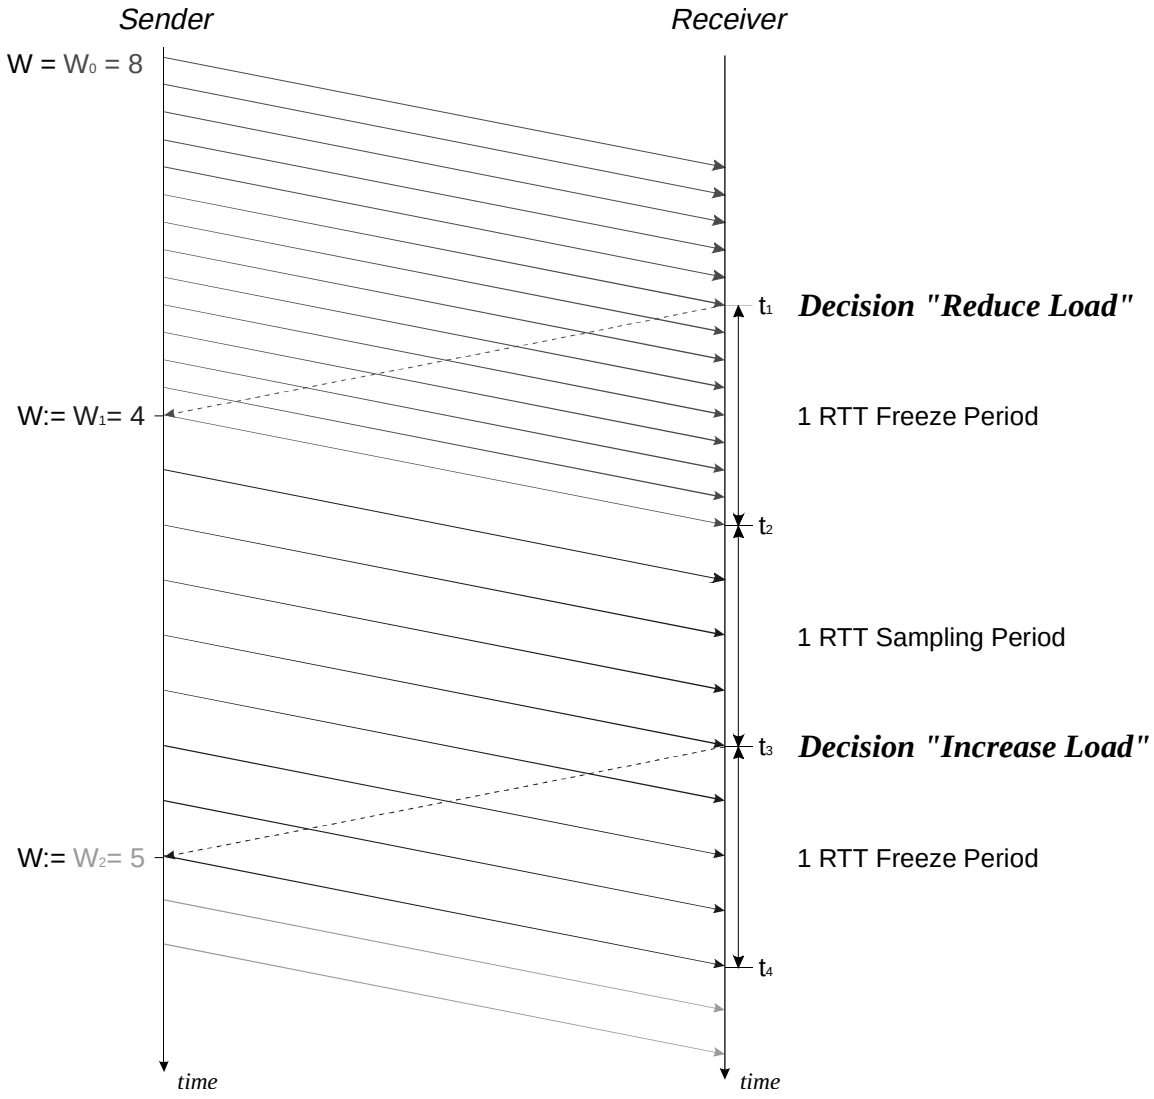

Figure 6-2 Window Adaptation over Time

6.3.5 The Congestion Avoidance Algorithm

In the ATN Internet SARPs, the Congestion Avoidance algorithm is presented as a set of requirements, following the normal style for SARPs. It is represented here in a 'C' code format, in order to make the algorithm more readily understandable to implementors.

Firstly, to support the Congestion Avoidance algorithm, each connection keeps a number of state variables, defined and initialised as follows:

```
int  n_DT    = 0;         // number of DT-TPDUs received
int  n_total = 0;         // total number of CE signals received
int  n_CE    = 0;         // number of active CE signals received
int  W_old   = 0;         // previously advertised window size
int  W_new   = W0;        // newly advertised window size
bool sampling= TRUE;      // are we currently sampling CE-flags?
```

Note that a new connection starts advertising an initial window size W_0 (as defined in SARPs text 5.2.6.3) to its peer. This is reflected in the initialization of variable 'w_new'.

Whenever a TPDU is received from the network layer, the routine *CongestionAvoidance()* is called with the congestion information received from the network layer forwarded to it. This routine performs the congestion avoidance algorithm, and updates the state variables as follows:

```
CongestionAvoidance(bool dt_TPDU, int nTotal, int nCE)
{
  // dt_TPDU  - flag indicating whether a DT-TPDU had been received
  // nTotal   - total number of NPDUs forming that TPDU
  // nCE      - number of NPDUs forming that TPDU that had their CE flag set on
  //            reception
  if (dt_TPDU) n_DT++; // count total # DT-TPDUs received
  n_total+= nTotal;    // count total # signals received so far
  n_CE   += nCE;       // count # active signals received so far
  if (n_DT > W_old) {
         // received enough DT-TPDUs, phase is completed
      if (sampling) {
             // was in sampling phase; compute new window and advertise
          if (n_CE > lambda * n_total) {
              W_new *= beta;
          } else {
              W_new++;
          }
          AdvertiseWindow(W_new);
          sampling= FALSE;
      } else {
             // was not sampling; just switch to sampling phase
          W_old   = W_new;
          sampling= TRUE;          // now entering sample phase
          n_total = n_CE= n_DT= 0; // reset counts
      }
  }
}
```

Note.: 'lambda', 'beta' and 'W0' are parameters defined in the SARPs text; see section 5.5.2.5.4: "Recommended algorithm values".

6.3.6 Sending Transport Entity Procedures

No specific features are required of the sending transport entity, in order to support this Congestion Management algorithm. It is only required to implement normal behaviour with respect to the handling of AK TPDUs and the utilisation of the received credit window.

However, implementors should note that commercial implementation of the transport protocol may often include "transport layer backoff" procedures similar to the van Jacobson Slow-Start algorithms. Implementors are strongly advised to remove such a feature from the implementation prior to it being deployed on the ATN. The backoff procedure is not required for congestion management and is likely to detect false indications of network congestion when a mobile system moves its point of attachment. This will result in reduced throughput, and implementations that include the backoff procedure will be perceived as being slower and giving poorer performance than those that do not.

6.3.7 Known Limitations

6.3.7.1 Fairness

It is known from previous research in the area of Congestion Management algorithms, that the adaptation of a window (instead of the transmission rate) is likely to cause problems if competing users have different path lengths (i.e. round trip times). Such a situation is shown in 3rd.

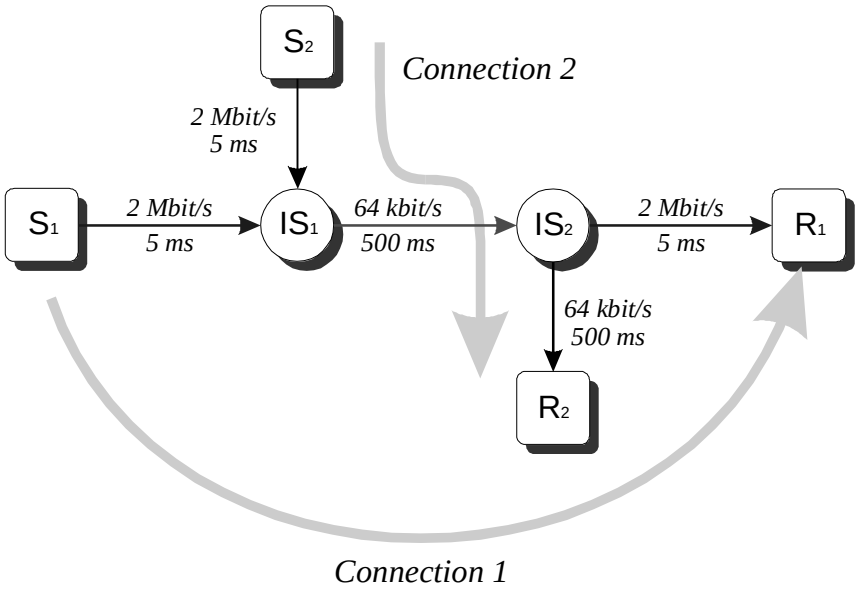

Figure 6-3 Fairness among competing Users

The problem is that the specified Congestion Avoidance will tend to result in approximately equal credit windows for all transport connections through the congested node. However, throughput depends not just on credit window, but also on the Round Trip Time. Once credit windows become restricted below the point at which greatest throughput is achieved, a transport connection will experienced a lower throughput than another with the same credit window and a shorter Round Trip Time.

It may be possible to balance throughput by varying the value of β taking into account the Round Trip Time. However, this requires network wide co-ordination to be effective and is only then useful with large window sizes. This limitation therefore appears to be a feature which has to be accepted.

6.3.7.2 Two-Way Traffic

Another well-known problem of many Congestion Control algorithms is caused by traffic along the reverse path. If data packets are transmitted along the reverse path, they will keep the intermediate system busy for some time. Acknowledgements arriving during that time will get queued, waiting for the IS to become available again. As soon as the system becomes free, these acknowledgements are transmitted back-to-back. This can have some adverse influence on the operation of the Congestion Control algorithm (e.g. leading to bursts of data packets emitted by one of the senders).

2nd depicts this scenario.

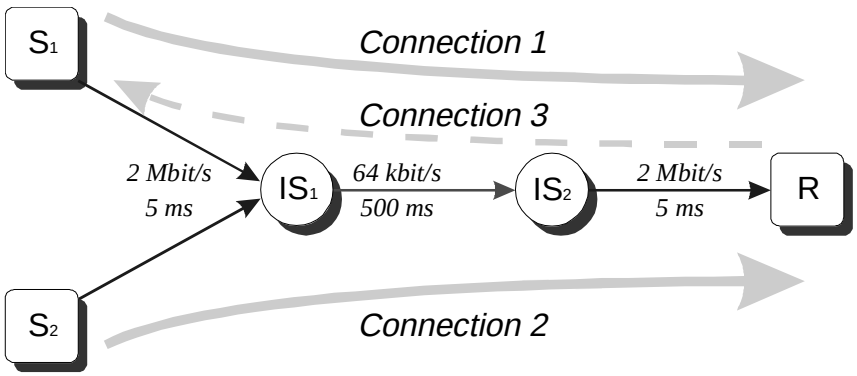

Figure 6-4 Two-way Traffic

6.3.7.3 A Credit Window of One is the Minimum

Like all Congestion Management algorithms, there is a point beyond which the algorithm cannot stop the network getting into a state of catastrophic congestion. In this case, this point is reached once all transport connections through a congested node have had their credit window reduced to one. After this point, the algorithm cannot reduce the load on the network any more and any increase in the load results in congestion, packet discards, re-transmissions and the network will become congested.

6.3.8 Conclusion

Congestion Avoidance is an essential feature for any internetwork. The specified algorithm appears to be the best for the ATN and achieves the best throughput while avoiding congesting the network as part of its own operation. There is, however, a limit to its effectiveness. This limit point is well beyond the point at which the algorithm starts to give useful benefits. However, it still underlines the important of good network design and capacity planning in respect of ensuring that network performance is maintained. A good Congestion Avoidance algorithm is an essential defence mechanism. However, it cannot give you network capacity that does not exist.

7. ATN Subnetworks

7.1 Introduction

The ATNP Internetwork SARPs specify requirements for the Subnetwork Dependent Convergence Function (SNDCF) and require that Subnetwork (SN)-Service (SNS) primitives or equivalent mechanisms be provided.

This chapter provides guidance on the necessary features of the SNDCF to support the ISO/IEC 8473-1/2/3/4 Connectionless Network Layer Protocol (CLNP) over these various subnetworks. Fur this purpose, it is firstly describing ATN requirements which are common to all subnetworks; thereafter, it is further broken down into mobile (air-ground) and ground subnetworks. The list of subnetworks listed is not exhaustive, and in particular future subnetworks may well be capabable to serve as ATN subnetworks.

7.2 General Characteristics of ATN-suitable subnetworks

ATN subnetworks are connected to the ATN internetwork via subnetwork dependent conversion functions (SNDCFs). An SNDCF, where each individual subnetwork requires one of its own, is part of the ATN router and provides a number of functions:

- Address Compression and De-compression;
- Priority Mapping, if supported by the subnetwork;
- Assembling and re-assembling of data packets;
- processing of signals generated by the subnetwork (e.g. join/leave events).

It is not the purpose of this document to provide guidance on the implementation of these subnetworks itself, since this is done in the relevant subnetwork specific material. Rather, guidance on the implementation of ATN using these subnetworks is given. In order to enable and optimize the SNDCF's operation over these subnetworks, the characteristics and features which are common to all networks which are used as ATN subnetworks are as summarized as follows:

- firstly, the subnetworks have to provide a byte and code indepentent communication service; in particula, this implies that character-oriented communication systems cannot (directly) be used as ATN subnetworks.
- mobile, i.e. air/ground, subnetworks have to provide join and leave events in order to enable routing initiation by the internetwork layer;
- although not being a strict requirement, ATN subnetworks should provide a priority feature, allowing to map the priority indicated by the network layer to a subnetwork priority and enabling to distinguish higher priority data from lower priority data;

- similarly, Fast Select capability, although not required generally for subnetworks, is a desirable feature;

- in the same way, "Fast Select" capability should be available from X.25-type subnetworks in order to optimize bandwidth;

Editors note: is this last point correctly understood, or do we have "Fast Select" also with other subnetworks??

7.3 Air/Ground Subnetworks

The following sections briefly summarize the features of individual mobile subnetworks. Reference is made to the appropriate ICAO Annex 10 material and the material of the relevant panels.

7.3.1 AMSS

In accordance with Draft ICAO AMSS SARPs (§ 4.7.1), the AMSS satellite subnetwork provides connection-oriented packet data service by establishing subnetwork connection between subnetwork service (SNS) users.

The packet data interface allows AMSS to function as the satellite subnetwork of the ATN. The satellite subnetwork transfers data packets from air to ground and from ground to air. Packet data transfers are provided in the form of connection-mode service, using ISO 8208 as a subnetwork access protocol.

7.3.1.1 Subnetwork Access

Access to the Satellite SubNetwork is normally provided by the SubNetwork Access (SNAc) function. ISO 8208, second edition, is the recommended International Standard to be used as SubNetwork Access Protocol (SNAcP) between the Aircraft Earth Station (AES) or the Ground Earth Station (GES) and an attached DTE (i.e. for example an ATN router), exchanging ISO 8208 Packets. The SNAc function performs the ISO 8208 DCE protocol function in the AES and the GES.

7.3.1.2 Expedited Data Negotiation

The Expedited Data Negotiation (EDN) Facility of the ISO 8208 INCOMING CALL/CALL CONNECTED Packet is mapped from the Expedited Data Selection parameter of the corresponding SN-CONNECT primitive.

7.3.1.3 Fast Select

Within AMSS, the fast select facility of the ISO 8208 INCOMING CALL packet is mapped from the Fast Select parameter of the SN-CONNECT indication primitive.

7.3.1.4 Priority

Within AMSS, the Priority facility of the ISO 8208 INCOMING CALL/CALL CONNECTED packet is mapped from the Priority parameter of the corresponding SN-CONNECT primitive.

If the parameter is present, then the IWF should encode a Priority Facility with a Facility Parameter Field of length 1. The Priority parameter value should be placed into the one-octet Facility Parameter field. If the Priority parameter is absent from the SN-CONNECT primitive, then the Priority facility should be omitted from the corresponding packet.

7.3.2 VDL

This chapter provides guidance on the development of the VDL SNDCF based on a development for initial VDL validation. This solution uses VDL station interconnected as DTE to a ground X.25 network. Other Solutions are not precluded.

7.3.2.1 General

The VDL Convergence functions provide a set of routing services to ISO 9542 and ISO 10747 routing protocols, which handle the Routing Initiation and Termination mechanism, which takes place in two phases:

1. The *Routing Initiation* or *Peer Discovery* phase during which the ground ATN Router learns of the existence of a mobile ATN Router which can be joined through the VDL Subnetwork and the mobile ATN Router learns of ground ATN Routers which are able to forward air initiated traffic.

   The Peer Discovery phase is always initiated by airborne mobile ATN Routers and uses the ISO 9542 protocol to exchange network layer reachability information and to update the corresponding routing tables.

2. The *Routing Termination* phase during which the VDL Convergence Functions notify network layer entities (ISO 9542 and ISO 10747), about the loss of the VDL subnetwork connectivity with a mobile ATN Router.

That information is used by the above network layer routing entities to update their corresponding routing tables.

The VDL Convergence Functions also handle VDL subnetwork hand-offs so that network layer routing tables remain unchanged. Hand-off is the term used to describe a process executed by the airborne ATN/VDL entities to identify and use a different RGS, other than the one currently being used, to relay VDL traffic.

The Header Compression (**LREF**) mechanism described in the ATN SARPs, section 5.7.3 Convergence Provisions for ISO 8208 Mobile Subnetworks, is part of the VDL convergence functions.

7.3.2.2 VDL Hand-offs

When the airborne ATN/VDL entities determine (by means which are outside the scope of this document) that a hand-off is required, the ATN/VDL entities may place new calls to all ground ATN Routers with which they are currently communicating via the VDL subnetwork.

The VDL Convergence Functions will accept the new call and clear the previous call only after a period of time determined by the value of the runtime-configurable G-TG5 timer.

If multiple virtual calls exist between the ground ATN Router and a given aircraft, through the VDL subnetwork, then the last established call is the *active VC*.

Through the active VC, ISO 8473 NPDUs can be either sent or received. Through the remaining connections, with the same aircraft, ISO 8473 NPDUs are never sent but can be received.

Cause Hex80 and diagnostic 0x84 is used by the SNDCF for the VC clearance, due to the TG5 timer expiry. A CLEAR REQUEST packet with the above cause shall not be forwarded by the ground station via the RF, towards the aircraft.

7.3.2.3 Routing Mechanisms over the VDL SN

7.3.2.3.1 Introduction

In order to establish and close ATN communications over the VDL subnetwork, both the airborne and the ground routers must be advised of modifications in the subnetwork connectivity.

The *Routing Initiation* is always initiated by the airborne ISO *9542* sub-layer.

The *Routing Termination* shall be invoked by the *SNDCF* sub-layer whenever the subnetwork connectivity is lost.

Note:- Through the VDL subnetwork only the airborne ATN Routers are allowed to initiate virtual calls.

7.3.2.3.2 The Routing Initiation mechanism

7.3.2.3.2.1 Call Request

The peer discovery mechanism is initiated by airborne ATN Routers in a way which is outside the scope of this document.

During the peer discovery phase the airborne ATN Router issues a CALL REQUEST towards the ground ATN Router. The INCOMING CALL received by the ground ATN Router has the following parameters:

1. **Calling Address:** The SNPA address of the RGS through which the call is established. Because an RGS acts as DTEs with regard to the X.25 networks to which it is connected, the RGS SNPA received by the ground ATN Router usually does not contain a DNIC, but only a DNA and optionally a sub-address.

2. **Called Address:** A local X.121 address of the ground ATN Router, managed by the VDL Convergence Functions.

3. **Facilities:** the fast select without restriction on response.

4. **Non X.25 Facilities:** a non-X.25 facility is used to convey the aircraft SNPA, which contains the BCD encoding of the octal representation of the 24-bits ICAO binary address of the aircraft.

   If the aircraft SNPA length is “len”, the encoding of non-X.25 facility field is :

0x00	0xFF	0xFE	len+2	0xFF	SNPA

5. **Call User Data:** contains the "parameter block" which can be followed by an ISO 9542 ISH PDU, which conveys the NET of the airborne mobile ATN Router.

   The layout of the Call User Data is shown in Table 7-1

Octet	Value	Meaning
1	1100.0001	Mobile SNDCF
2	0000.0101	Parameter block length indicator: 5 octets
3	0000.0001	First version
4	SNCR low	Number of VC currently established between the aircraft
5	SNCR high	and the ground router, except this one (hand-offs)
6	0AV0.M0LC	A = 1 ACA compression requested, V = 1 V42bis compression requested M=1 maintenance of SNDCF context requested L=1 LREF requested, C=1 Cancellation Option requested
7	xxxx.xxxx	Low octet of Maximum number of directory entries
8	xxxx.xxxx	High octet. Bits 7 and 8 are present only if L=1.

Table 7-1 Mobile SNDCF Call User Data Layout

Note 1.– In case SNCR is greater than 1, the connection is cleared with diagnostic INVALID SNCR.

Note 2.– In the case L=0 and the length of the call user data is greater than 6 and if the 7th octet contains the value 1000.0010 then the Satellite SNDCF assumes that an ISO 9542 ISH PDU has been sent by the peer entity together with the CALL REQUEST.

Note 3.– In the case L=1, LOCAL reference compression requested, see [1], the SNDCF looks for the ISH starting at octet number 9.

7.3.2.3.2.2 ISH in the INCOMING CALL

The VDL Convergence Functions perform the following actions upon receipt of an INCOMING CALL indication carrying an ISO 9542 ISH:

1. Does not immediately confirm the call.

2. Extracts the NET from the received ISH PDU and associates it with the new VC (whose set-up is still in progress).

3. For local management purposes, extracts the RGS DNA from the Calling Address field of the INCOMING CALL and associates it with the new VC.

4. Issues an SN-UNITDATA Indication primitive in order to relay the ISO 9542 ISH PDU up to the local 9542 entity

Upon receipt of the ISH PDU the local ISO 9542 sub-layer shall update the routing tables and shall operate an ISO 9542 *Configuration Notification Function* (respond with an ISH PDU containing the NET of the local network entity

Then the VDL Convergence Function will confirm the corresponding VC (same source and destination addresses) whose set-up is in progress. The layout of the Call User Data of the CALL ACCEPTED packet is shown in Table 7-2.

Octet	Value	Meaning
1	0AV0.M0LC	A = 1 ACA compression accepted V = 1 V42bis compression accepted M=1 maintenance of SNDCF context accepted L=1 LREF accepted, C=1 Cancellation Option accepted
2	1000.0001	9542 ISH protocol identifier. The ISH response begins here.

Table 7-2 Mobile SNDCF Call Accept User Data Layout

The establishment of this ISO 8208 connection completes the peer discovery phase. This ISO 8208 connection becomes the active VC between the ground ATN Router and the aircraft and it is the only VC used by the ground ATN Router to forward ISO 8473 NPDUs towards the aircraft, through the VDL subnetwork (see section 5.2. VDL Hand-offs).

7.3.2.3.3 The Routing Termination mechanism

When the RGS detects the loss of coverage for a given aircraft, it clears all the appropriate calls within the terrestrial network. As a consequence, all VCs, active or not, between the ground ATN Router and a given aircraft get cleared.

The loss of all VCs with a given aircraft is detected by the ground VDL SNDCF, which then activates the Routing Termination phase.

The Routing Termination phase consists in issuing a *Leave Event* towards the network layer routing entities.

If between the ground and airborne ATN Routers there is only one VC and if the set-up of this single VC is in progress, then in the case that VC gets cleared a *Leave Event* is issued in order to update the routing tables (in the case an ISH, sent by the airborne BIS, has already been relayed up to the network layer entities).

7.3.2.4 LREF Compression

7.3.2.4.1 Connection Context

The LREF mechanism is implemented in the VDL SNDCF as specified in the ATN SARPs, though the connection context is associated to a pair aircraft Id/Ground DTE rather than a pair DTE/DTE, as in this implementation case the remote DTE is the DTE of the Radio Ground station.

7.3.2.4.2 LREF and Handoffs procedure

With regard to the ATN SARPs, it may happen that during a Handoff procedure, a packet requesting the creation of a directory entry is sent on the old virtual circuit, and subsequent packets for the same source/destination NSAPs be sent on the new virtual circuits in compressed mode. As VDL does not ensure that the packet sent on the old virtual circuit will be received by the adjacent router before the compressed one, compressed packets arriving first will be discarded and an error report be generated to the sending SNDCF.

In this case the transport protocol should ensure retransmission of the packet.

7.3.2.4.3 LREF and Call clearing

In the case where the VC has been cleared for other reasons than GT5-Timer expiration and in the case the LREF compression was used for the cleared VC the internal resources used to handle the Local Reference Directory are freed.

In the case where the VC has been cleared due to GT5-Timer expiration and the maintenance bit was not set or refused for the newly established associated CV, the Local Reference directory is freed.

Otherwise the Local reference Directory is maintained.

7.3.2.4.4 LREF and Call Reset

Otherwise a DTE generated RESET also results in the total clearing of the LREF directory.

7.3.3 Mode S

The Mode S air/ground subnetwork primarily provides a connection-oriented communication service between two subnetwork points of attachment (SNPA), one in the aircraft and the other on the ground. This service may be accessed by means of the protocol defined in ISO 8208, and is entirely conformant with the aeronautical telecommunication network (ATN) architecture.

7.3.3.1 Subnetwork Access

The ADLP and GDLP have a standard ISO 8208 interface between the output of the subnetwork and the ATN router(s). ISO 8208 DATA packets are received from the ATN router via the XDLP data terminating equipment/data circuit terminating equipment (DTE/DCE) interface.

Although DTE address assignment is a local issue, it is necessary to assign DTE addresses in an unambiguous manner in order to ensure the correct operation of the Mode S subnetwork.

In accordance with ISO 8208 the default maximum user data field length is 128 bytes. In addition, other (non-standard) default maximum user data field lengths may be available from the following list: 16, 32, 64, 256, 512, 1 024, 2 048 and 4 096 bytes. The selection of a non-standard default value is a local issue at a DTE/DCE interface and has no influence on the Mode S packet layer protocol, because the exact length of the user data field can be extracted from the data link layer information field of the DTE/DCE interface.

The interface between the GDLP and a ground DTE (including an ATN router) is required to be functionally conformant with the ISO 8208 DTE-DCE interface. However, the form of the physical connection between the GDLP and a ground DTE remains an option at the discretion of implementors. Alternative approaches are discussed in the following sub-paragraphs, although other options are also possible.

The simplest form of physical connection which may be used is a simple point-to-point link directly between the DTE and the DCE. In this case, only a suitable link layer protocol (e.g. ISO 9676) must be provided below the DTE, and the connection to the physical medium established.

7.3.3.2 Fast Select

The Mode S Subnetwork provides Fast Select Capability.

7.3.3.3 Priority

The Mode S subnetwork provides means for distinguishing two priorities ('high' and 'low')

7.3.3.4 Route Initiation

In order to enable the initiation of routing table updates, mobile subnetworks have to identify to the ATN internetwork when new routes to mobile systems are available. When an aircraft is newly

entering the coverage of a Mode S subnetwork, a 'join event' is generated by that particular subnetwork. This 'join event' is always generated by the ground part of the Mode S subnetwork, i.e. the GDLP.

Similarly, 'leave events' are generated by both, the aicraft side (ADLP) and the ground side. In addition, refresh cycles may be performed.

'Join event' and 'leave event' messages contain at least the following fields:

a) message type;

b) message length;

c) aircraft address; and

d) optionally, time and position of aircraft entry or exit.

7.4 Ground/Ground Subnetworks

This section presents a guidance for States which want to implement new or already existing networks as ATN subnetworks inside their respective boundaries.

7.4.1 Mapping CLNP over An ISO/IEC 8802 Subnetwork

This mapping is provided in the ATNP Internetwork Communications Service SARPs. The subnetwork service is provided as specified in Clause ISO/IEC 8473-2 *Information Technology—Protocol for Providing the Connectionless-Mode Network Service Part 2: Provision of the Underlying Service by an ISO/IEC 8802 Subnetwork*. In this case, the generation of an SN-UNITDATA request by CLNP results in a Data Link Layer (DL)-UNITDATA request (as described in ISO/IEC 8802-2) being generated by the SNDCF.

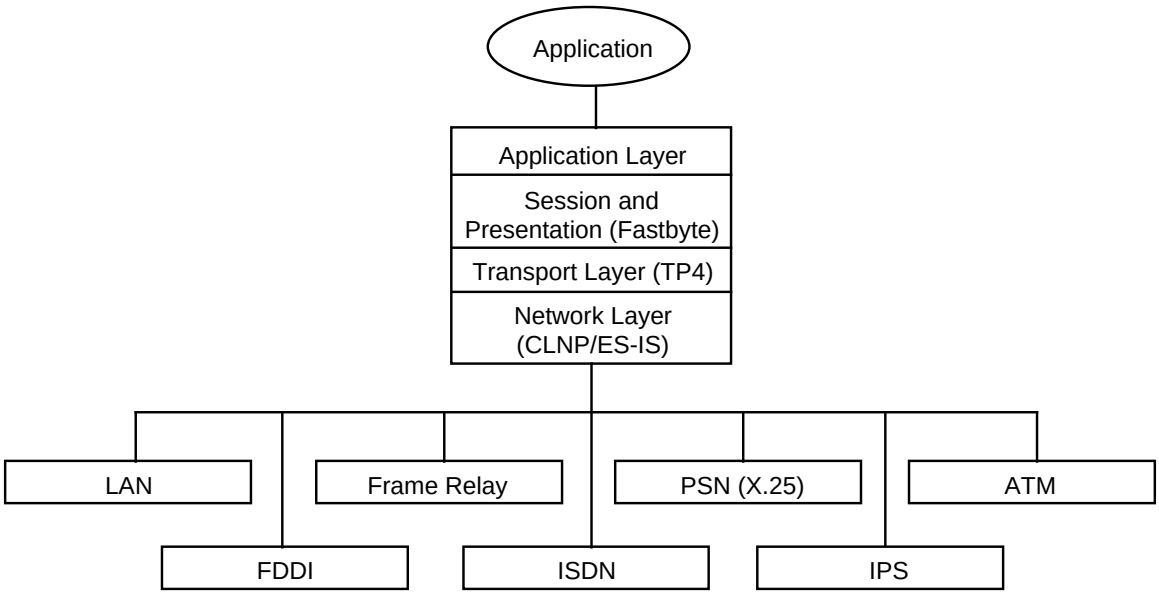

Figure 7-1 Expanded Subnetwork

7.4.2 Mapping CLNP over An FDDI Network

FDDI provides a high bandwidth LAN consisting of a physical layer and Media Access Control (MAC) sublayer as defined in ISO 9314-2. Encapsulation of CLNP in FDDI frames can be performed in a manner similar to the Internet Protocol (IP) encapsulation over FDDI (see Internet Engineering Task Force [IETF] Request for Comments [RFC] 1188. As per RFC 1188, CLNP can

be encapsulated via the IEEE 802.2 Logical Link Control (LLC) service as shown in **Error! Reference source not found.**.

The Destination Service Access Point (DSAP) and Source Service Access Point (SSAP) should be set to the hexadecimal value [0xFE] to indicate that the Information field contains an OSI network layer Protocol Data Unit (PDU).

In this case, the mapping to the SNS parameters should be the following:

- SN-Source-Address: This address should be set to the assigned value [0xFE] to indicate that the information field contains an OSI network layer PDU.
- SN-Destination Address: This address should be set to the assigned value [0xFE] to indicate that the information field contains an OSI network layer PDU.
- SN-Priority: If supported, this field should contain the mapped LLC priority based on the received SN-priority.
- SN-Quality-of-Service: If supported, this field should contain the mapped LLC quality-of-service based on the received SN-quality-of-service.
- SNS-User Data: This field should contain the ISO network layer PDU.

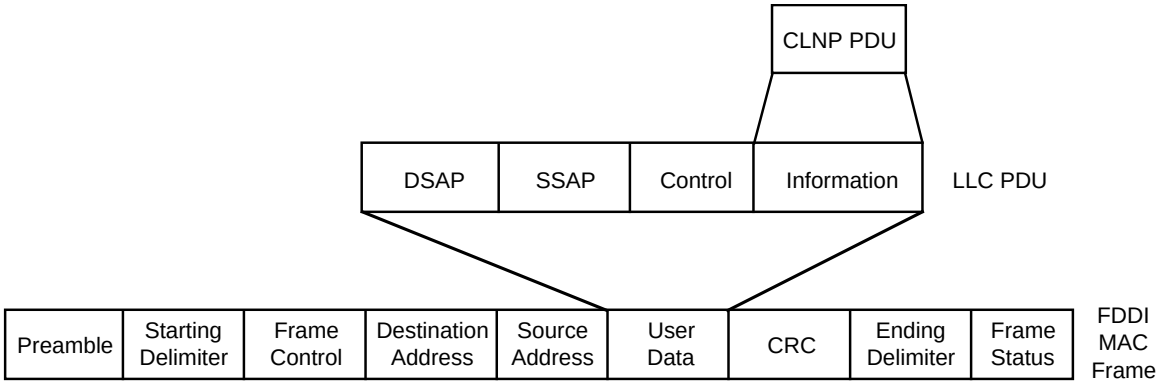

Figure 7-2 Mapping of CLNP into FDDI MAC Frames

IEEE 802.2 LLC Type 1 communication should be used exclusively. All frames should be transmitted in standard IEEE 802.2 LLC Type 1 Unnumbered Information format, with the DSAP and the SSAP fields of the 802.2 header set to the assigned Service Access Point value for ISO-IP (0xFE).

The IEEE 802.2 LLC subnetwork may then map the LLC PDU to the FDDI Frame as shown in **Error! Reference source not found.**.

7.4.3 Mapping CLNP over Frame Relay ISDN Network

While encapsulation of CLNP in Frame Relay networks can be accommodated in a manner similar to IP encapsulation over Frame Relay as outlined in IETF RFC 1490. The Frame Relay access protocol is based on High-level Data Link Control (MDLC/Q.921), and the link access protocol was developed for signaling over the D channel of narrow band Integrated Services Digital Network (ISDN) (International Telephone and Telephone Consultative Committee (CCITT) Recommendation Q.922). As discussed in IETF RFC 1490, the Frame Relay network provides a number of virtual circuits that form the basis for connections between stations attached to the same Frame Relay network. The resulting set of interconnected devices form a private Frame Relay group which may be either fully interconnected with a complete “mesh” of virtual circuits or only partially

interconnected. In either case, each virtual circuit is uniquely identified at each Frame Relay interface by a Data Link Connection Identifier (DLCI). In most circumstances, DLCIs have strictly local significance at each Frame Relay interface.

All protocols encapsulate their packets within a CCITT Q.922 Annex A frame or ATNSI T1.618. Additionally, the frames contain information necessary to identify the protocol carried within the PDU, thus allowing the receiver to properly process the incoming packet. The frame format is shown in Figure 7-1.

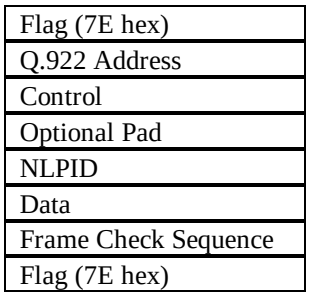

Figure 7-3 Frame Format

The Q.922 address is 2 octets and contains a 10-bit DLCI. In some networks Q.922 addresses may optionally be increased to 3 or 4 octets.

The control field is the Q.922 control field.

The Pad field is used to align the remainder of the frame to a 2-octet boundary. There may be zero or one Pad octet within the Pad field and, if present, must have a value of zero.

The Network Level Protocol ID (NLPID) field is administered by ISO and CCITT. It contains values for many different protocols including IP, CLNP, and IEEE Subnetwork Access Protocol (SNAP). This field tells the receiver what encapsulation or what protocol follows. Values for this field are defined in ISO/IEC Technical Report (TR) 9577. Some commonly used NLPIDs are defined below:

[0x00]	Null Network Layer or Inactive Set
[0x80]	SNAP
[0x81]	ISO CLNP
[0x82]	ISO ES-to-IS Protocol
[0x83]	ISO IS-to-IS
[0xCC]	Internet IP
[0x08]	ISDN

Table 7-3 Common Network Layer Protocol Ids

In the case of ISO protocols, the NLPID is considered to be the first octet of the protocol data. The single octet serves both as the demultiplexing value and as part of the protocol data. ISO/IEC 8473-1 has a NLPID value of [0x81]. For CLNP, PDUs, the frame would consist of the following fields as shown in Figure 7-4.

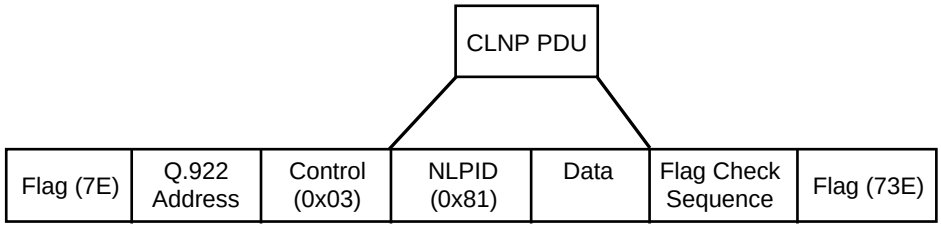

Figure 7-4 Frame Format for Transfer of CLNP PDUs

In this case, the mapping to the SNS parameters should be the following:

- SN-Source-Address: No source address field is provided in the frame relay header.
- SN-Destination Address: This address should be set to the destination Q.922 address.
- SN-Priority: No priority field is provided in the Frame Relay header.
- SN-Quality-of-Service: This field should contain the mapped Frame Relay quality-of-service. (Note that forward or backward explicit congestion notification may be indicated in the frame relay header.)
- SNS-User Data: This field should contain the ISO network layer PDU.

7.4.4 Mapping CLNP over An ISO 8208 Network

This information is provided in the ATNP Internet Communications Service SARPs. The subnetwork service is provided using ISO/IEC 8473-3.

Management of X.25 connections is discussed in detail in ISO/IEC 8473-3.

7.4.5 Mapping CLNP over IP

Recent IETF RFCs have been developed to allow the encapsulation of CLNP PDUs over IP. However, for OSI applications that may need to communicate via an IP internet, current commercial off-the-shelf (COTS) routers will encapsulate the ISO subnetwork PDUs (for example, X.25 PDUs) into IP datagrams, and decapsulate the datagrams and forward to the peer OSI application.

If there is a need for direct encapsulation of CLNP PDUs over IP, then IETF RFCs 1701, 1702 and 1070 define an Generic Routing Encapsulation (GRE) protocol to allow a number of different protocols to be encapsulated over IP. As defined in these RFCs, the packet to be encapsulated and routed is called a payload packet. The payload is first encapsulated in a GRE packet. The resulting GRE packet can then be encapsulated in some other protocol (such as IP) and then forwarded. This outer protocol is called the delivery protocol.

The Delivery Header for IP will consist of the fields shown in Figure 7-5.

Within the GRE Header, the Protocol Type field contains the protocol type of the payload packet. Example protocol types are listed below as shown in Table 7-4.

In this case, the mapping to the SNS parameters should be the following:

- SN-Source-Address: This field should contain a source IP address.
- SN-Destination Address: This field should contain a destination IP address.

- SN-Priority: If supported, the priority can be indicated in IP datagrams via the precedence bits in the Type of Service field. This field should indicate the IP priority.

- SN-Quality-of-Service: If supported, this field should contain the Type of Service value.

- SNS-User Data: This field should contain the ISO network layer PDU.

The other method using RFC 1070 permits ISO Systems using ES-IS to change their topological relationship to the IP format.

7.4.6 Mapping CLNP over Asynchronous Transfer Mode (ATM)

While encapsulation of CLNP over ATM networks has not been standardized, it could be accommodated in a similar manner that IP is encapsulated over ATM (IETF RFC 1483).

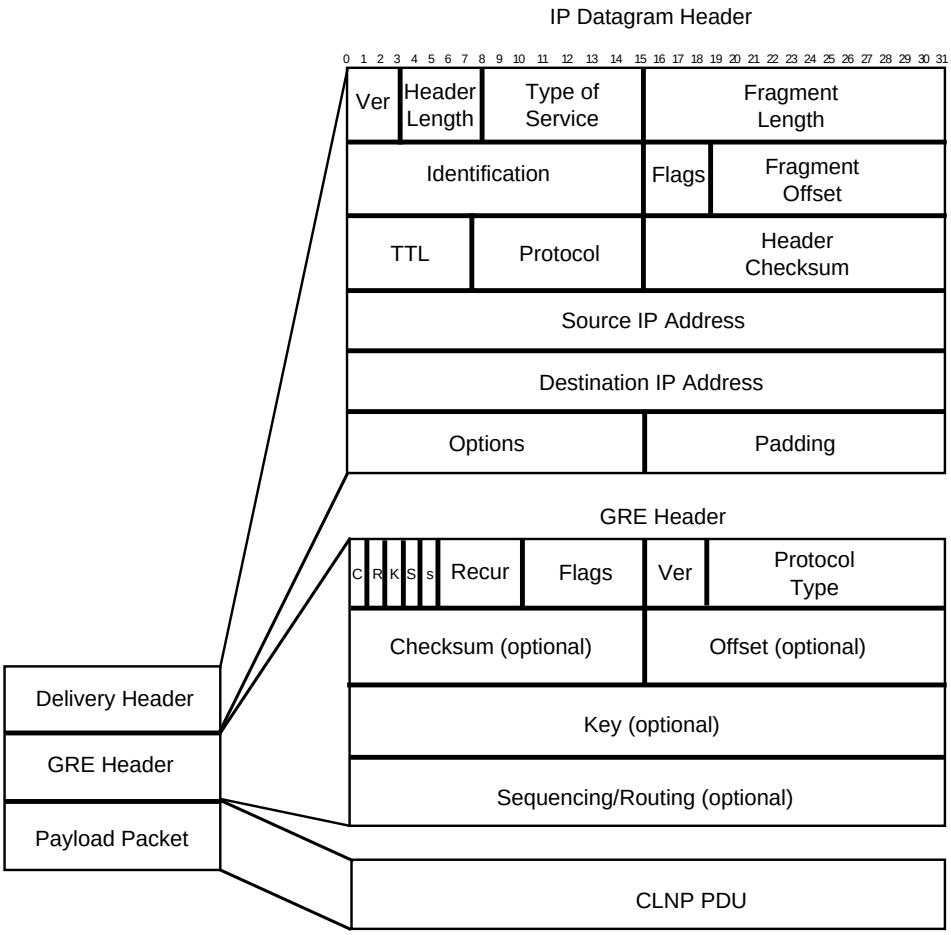

Legend:		
C	=	Checksum Present (1)
R	=	Routing Present (1)
K	=	Key Present (2)
S	=	Sequence Number Present (1)
s	=	Strict Source Route (1)
Recur	=	Recursion Control (3)
Ver	=	Version Number (3)

Figure 7-5 Delivery Header for IP

Protocol Family	Protocol Type Value (Hex)
Reserved	0000
OSI network layer	00FE
IP	0800
Frame Relay	0808
Raw Frame Relay	6559
IP Autonomous Systems	876C
Secure Data	876D
Reserved	FFFF

Table 7-4 Example Protocol Type Values

As described in RFC 1483 *Multiprotocol Encapsulation over ATM Adaptation Layer 5 (AAL5),* ATM-based networks are of increasing interest for both local and wide area applications. There are

two different methods for carrying connectionless network interconnect traffic, routed and bridged PDUs, over an ATM network. The first method allows multiplexing of multiple protocols over a single ATM virtual circuit (called "LLC encapsulation"). The protocol of a carried PDU is identified by prefixing the PDU by an IEEE 802.2 LLC header. The second method performs higher-layer protocol multiplexing implicitly by ATM Virtual Circuits (VCs) (called "VC-based Multiplexing").

No matter which multiplexing method is selected, routed and bridged PDUs are encapsulated within the Payload field of AAL5 Common Part Convergence Sublayer (CPCS)-PDU. The format of the AAL5 CPCS-PDU is shown in Figure 6.

The Payload field contains user information up to (2^16–1) octets.

The Padding (PAD) field pads the CPCS-PDU to fit exactly into the ATM cells such that the last 48-octet cell payload created by the new Segmentation and Reassembly sublayer will have the CPCS-PDU Trailer right justified in the cell.

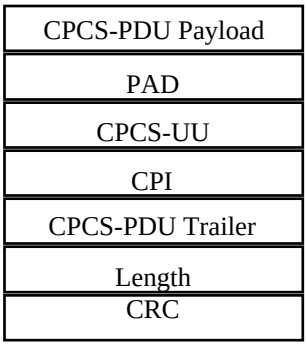

Figure 7-6 AAL5 CPCS-PDU Format

The CPCS-User-to-User (UU) field is used to transparently transfer CPCS-UU information. The field has no function under the multiprotocol ATM encapsulation described in this memo and can be set to any value.

The Common Part Indicator (CPI) field aligns the CPCS-PDU trailer to 64 bits. Possible additional functions are for further study in CCITT. When only the 64 bit alignment function is used, this field shall be codes as 0x00.

The Length field indicates the length, in octets, of the Payload field. The maximum value for the Length field is 65,535 octets. A Length field coded as 0x00 is used for the abort function.

The Cyclical Redundancy Check (CRC) field protects the entire CPCS-PDU except the CRC field itself.

RFC 1483 describes the use of LLC encapsulation for CLNP PDUs which is described below. For additional information concerning VC-based multiplexing, the reader is referred to the RFC.

7.4.6.1 LLC Encapsulation

In LLC Encapsulation the protocol of the routed PDU is identified by prefixing the PDU by an IEEE 802.2 LLC header, which may be followed by an IEEE 802.2 SubNetwork Attachment Point (SNAP) header. In LLC Type 1 operation, the LLC header consists of three 1-octet fields as shown inFigure 7-7.

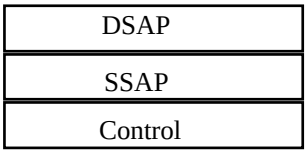

Figure 7-7 LLC Header Format

The LLC header value 0xFE-FE-03 identifies that a routed ISO PDU follows. The Control field value 0x03 specifies Unnumbered Information Command PDU. For routed ISO PDUs, the format of the AAL5 CPCS-PDU Payload field shall thus be as follows as shown in Figure 7-8.

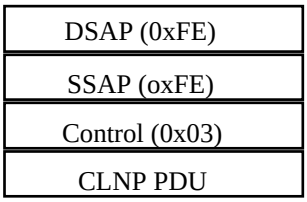

Figure 7-8 AAL5 CPCS-PDU Payload Field Format for Routed ISO PDUs

The routed ISO protocol is identified by a 1 octet NLPID field that is part of Protocol Data.

In this case, the mapping to the SNS parameters should be the following:

- SN-Source-Address: This field should contain the value [0xFE].
- SN-Destination Address: This field should contain the value [0xFE].
- SN-Priority: This field does not map to AAL-5 fields.
- SN-Quality-of-Service: This field does not map to AAL-5 fields.
- SNS-User Data: This field should contain the ISO network layer PDU.
-
-
-