

### 3. *Internet Communications Services (ICS)*

#### 3.1 **General**

3.1.1 This chapter provides guidance material on the ATN Internet Communications Service, the construction of an ATN internet and the protocols used. With respect to the addressing aspects within the ATN internetwork, reference is made to part II of the document.

#### 3.2 **ATN Internet Concept**

##### 3.2.1 **Purpose of the ATN Internetwork**

3.2.1.1 The ATN is a data communications internetwork that:

- a) provides a common communications service for all air traffic services communications (ATSC) and aeronautical industry service communications (AINSC) applications that require either ground/ground or air-ground data communications services;
- b) integrates and uses existing communications networks and infrastructure wherever possible;
- c) provides a communications service which meets the security and safety requirements of ATSC and AINSC applications, including the reliable and timely delivery of user data to its intended destination; and
- d) accommodates the different grades of service required by each ATSC and AINSC application, and the organisational policies for interconnection and routing specified by each participating organisation.

3.2.1.2 While these capabilities might, at first sight, appear ambitious, the reality is that for the ATN's users, the internetwork will be straightforward and simple to use. This is because the ATN's architecture deliberately places the responsibility for routing and maintaining an internetwork's operational status on the "routers" and therefore enables the End Systems (cf. Host Computers) to have only a minimal networking capability.

##### 3.2.2 **Technical Benefits**

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

- a) **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 local area networks (LANs), leased lines, common ICAO data interchange network (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);


- b) **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 permitting 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. aircraft communications addressing and reporting system (ACARS));
- c) **Mobile Communications.** The ATN fully supports mobile communications over a wide variety of mobile communications networks including aeronautical mobile-satellite service (AMSS), VHF digital link (VDL) and SSR Mode S. With the ATN, it is possible for a ground system to communicate with airborne avionics in any part of the world;
- d) **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. In the ATN, traffic load is balanced to the availability of communications resources;
- e) **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;
- f) **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;
- g) **Future Proofing.** The ATN’s 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;

### 3.2.3

#### Construction of the ATN Internet

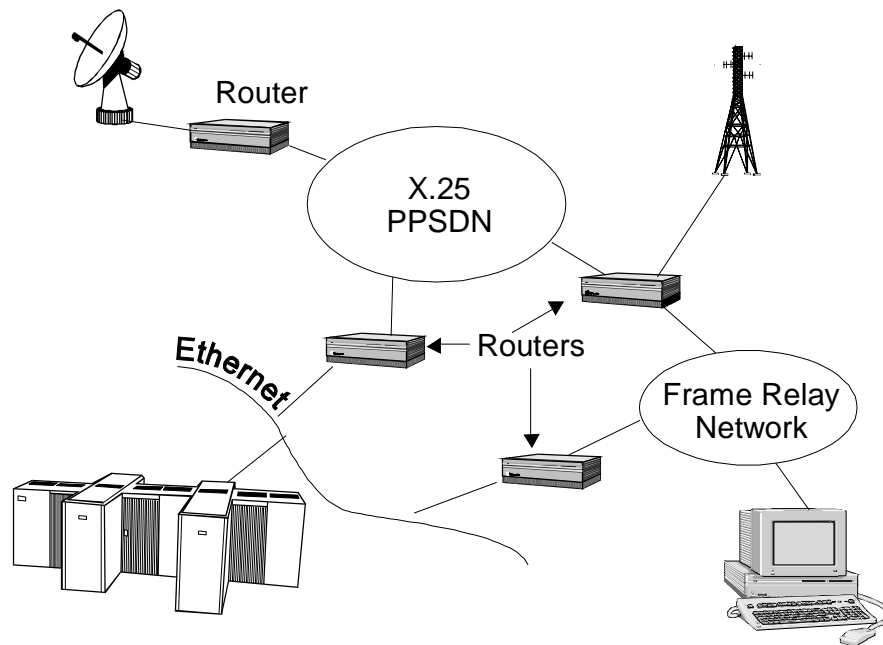
##### 3.2.3.1

A general model of the ATN is shown in Figure 3.2-1. In this model, the ATN consists of a fixed ground network which links satellite, VHF and Mode S ground stations together with ground based Host computers, including both large scale data processing engines and workstations. ATN avionics on board aircraft are then linked to the rest of the network through satellite, VDL and SSR Mode S data links, as appropriate, and may have more than one air/ground data link in use simultaneously.



**Figure 3.2-1. General Model of the ATN**

- 3.2.3.2 The fixed network is not a single entity but itself consists of many different networks all linked together, as illustrated in Figure 3.2-2. The ATN ground environment will consist of multiple networks, owned by different administrations and organisations, and implemented using many different technologies. In some cases, these will be existing networks with spare capacity made available to the ATN. Others will be new networks implemented specifically to support ATN use. There will be X.25 Private Packet Switched Data Networks (PPSDNs), Frame Relay Data Networks, Integrated Services Digital Networks (ISDNs), Local Area Networks (LANs) e.g. Ethernet, and others. These networks are then linked together through routers which provide the connectivity between the different types of data network, and to the air/ground networks. Host computers are directly connected to a nearby data network, typically a LAN.
- 3.2.3.3 User data is switched by the routers as discrete packets formatted according to the ISO/IEC 8473 Connectionless Network Protocol (CLNP). Each packet is viewed as a separate event and routed according to a “route map” of the ATN. In the ATN, each router has a portion of the full ATN route map and builds and maintains this route map dynamically using routing information passed to it by its neighbouring (adjacent) routers.



**Figure 3.2-2. The ATN Ground Environment**

- 3.2.3.4 Host computers communicate with each other either directly over a common data network, or use the services of a router to provide a communications path to a Host on another data network. It is the responsibility of the routers working together to find a suitable path through the networks which they interconnect, and data may travel through many different routers and via many different networks on its journey between two Hosts. In order to build an ATN route map for this purpose, the routers exchange, amongst themselves, information on which hosts are local to them (i.e. reachable via a single data network and with no intermediate router), and on how they relate to other routers. From such information, the routers can plot the course of data through the ATN.
- 3.2.3.5 The ATN Ground Environment permits each participating organisation to organise its networks and systems as it wishes, and form a separate “Administrative Domain”, However, it is also anticipated that Administrations and ATN regions will, wherever possible, co-ordinate their addressing plans and network topologies, such that the amount of routing information passed between organisations and regions can be kept to an absolute minimum.
- 3.2.3.6 The ATN Addressing Plan apportions a separate part of the address space to each ICAO Region, national ATS Administration, and to each IATA airline and other organisations. This allows for great flexibility in use, however, participating organisations are, as indicated above, strongly recommended to co-ordinate the allocation of addresses. For example, in Europe, Administrations should implement a co-ordinated addressing plan with a unique address prefix for Europe and address assignment that reflects the actual topology of the European ATN Internet. For similar reasons, airlines, and especially small regional airlines should consider service provider relative addresses.

### 3.2.4 Users' View of the ATN

First, it is worth considering who or what is an ATN user. In principle, the ATN User is an ATM application, or some other application supporting an aeronautical application. But, this is very much an end system view. From the networking point of view, there are many interfaces over which “a user” accesses a service. At each such interface, each user can be considered to be an ATN User. In the remainder of this section, the term “ATN User” is used in this sense, i.e. as a user of the network service or transport service, depending on context.

#### 3.2.4.1 ATN User Communications Capabilities

##### 3.2.4.1.1 General

3.2.4.1.1.1 The ATN provides its users with a robust and reliable communications service, together with the option of a datagram service. Formally, all communications aspects of a user's system are part of the ATN, but from a “user's point of view”, the ATN is out there, separate from its their own system. It is this “user's view” of the ATN that is illustrated in Figure 3.2-3. This figure shows the ATN as an abstract “cloud” which indeed is all the user needs to be aware of, with its complexity hidden from view. At this level, the ATN is a simple network that provides a datagram service to its users.

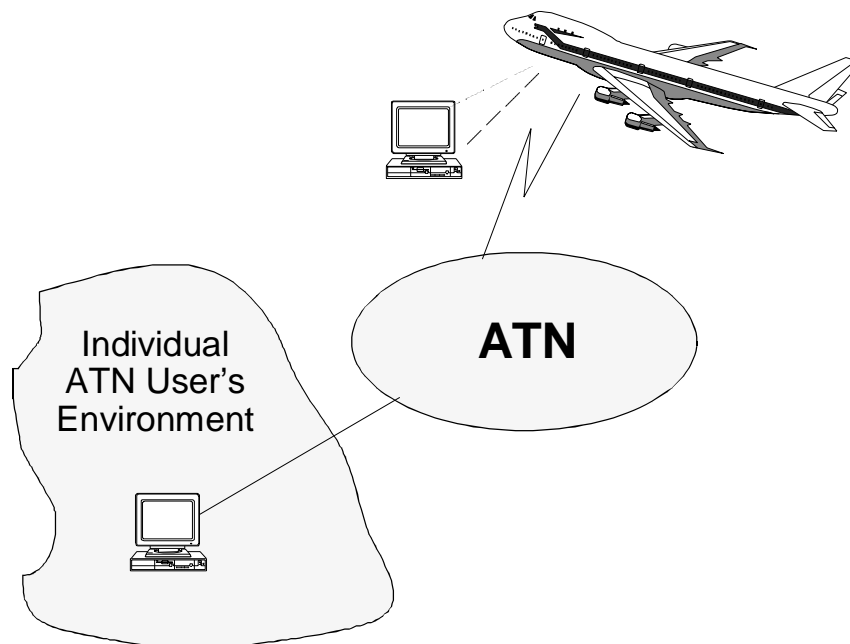


Figure 3.2-3. Individual End User's View of the ATN

- 3.2.4.1.1.2 A ground based ATN user's system, which might be anything from a complete ATC system to an entry level PC, accesses ATN services via some ATN access point. This access point is a notional socket into which the user "plugs" ~~its~~<sup>their</sup> system and thereby gains access to the ATN. However, this socket is not as tangible as an electrical power socket. A user's access to the ATN is first via an "access subnetwork", such as an Ethernet or an X.25 PSDN, and then an ATN Router. The user's system is directly connected to the access subnetwork, and this may very well involve a physical connection provided by a wall-socket, and, using the access subnetwork, the user's system communicates with the ATN Router. It is then through the ATN Router that access is gained to ATN services.
- 3.2.4.1.1.3 The communications capabilities of the user's system must obviously include the hardware and software necessary to use the "access subnetwork". Furthermore, the user's system must also support the ISO/IEC 8473 Connectionless Network Protocol in order to be part of the ATN Internet, and is recommended to support the ISO/IEC 9542 End System to Intermediate System Routing Protocol.
- 3.2.4.1.2 **ISO/IEC 8473 CLNP Protocol**
- 3.2.4.1.2.1 CLNP is a simple protocol supporting the transfer of "datagrams", i.e. packets of data transferred from sender to receiver without the need for a connection to be established in advance. Data transferred using CLNP is formatted as a block of data preceded by a protocol header containing the addresses of the sender and destination, the priority of the data, any security label associated with it, and quality of service requirements. Header and data must together not exceed 64 kilobytes.
- 3.2.4.1.2.2 An ATN user may, at any time send a CLNP formatted datagram to any valid destination address. The user does this by passing the datagram over the access subnetwork to the ATN Router. The ATN Router will inspect the protocol header, and it is then the ATN Router's responsibility to forward the datagram through the ATN to the ATN Router which provides ATN access to the addressed destination. How it does this is internal to the ATN and hence hidden from the user, although the forwarding process must respect the data priority and the Quality of Service and Security requirements identified in the protocol header. Once the datagram has arrived at the ATN Router which provides ATN access to the addressed destination, it is then transferred over the destination's access subnetwork to the destination user. If the destination user is offline (e.g., switched off), the datagram is discarded and an error report is optionally returned to the sender.
- 3.2.4.1.2.3 The operation of these processes is essentially how the user perceives the ATN. The simple CLNP is the protocol ATN users use to communicate, and permits those users to exchange information as discrete blocks of data.

### 3.2.4.1.3 **ISO/IEC 9542 ES-IS Protocol**

3.2.4.1.3.1 The other protocol that users are recommended to support - the ES-IS protocol - is really just for local administration. The user's system uses the ES-IS protocol to report its own address to the ATN Router, and this information is regularly repeated so that the ATN Router can monitor a user's online status. It is also used to report the existence and operational status of an ATN Router to its users, and enables an ATN user to have access to multiple ATN Routers, possibly over different access subnetworks, so as to provide a high availability service.

### 3.2.4.1.4 **CLNP Delivery Probability**

3.2.4.1.4.1 The ATN itself does not make any demands on the syntax or semantics of the data carried in a CLNP packet. However, the simplicity of the service does carry a penalty and this is that delivery of datagrams is not guaranteed. When a user transfers an ISO/IEC 8473 formatted packet to an ATN Router, that user is only guaranteed a probability of delivery dependent on the data priority. The probability of delivery is high, and while no targets have yet been set for delivery probability, 97% - 98% is certainly realistic. Considerations that affect this figure include:

- a) the error rates on subnetworks such as Ethernets which may lose data in transit due to line errors (although this consideration does not apply to X.25 subnetworks and similar examples, which provide a reliable transfer service);
- b) network overload which results in low priority data being discarded in order to free up congested resources; and
- c) component failures.

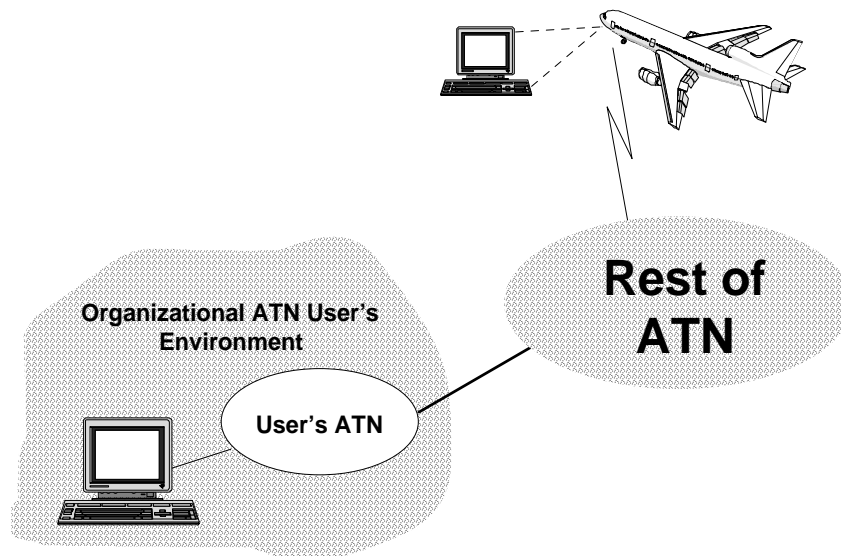
3.2.4.1.4.2 The actual delivery probability that is provided is a design issue. Once actual targets have been provided then it is possible to design a network to meet the requirement. This is achieved by minimising the use of lower reliability subnetworks, increasing overall network capacity, and through component redundancy.

3.2.4.1.4.3 However, the network can never provide a 100% delivery probability. When an ATN user does require reliable data transfer, then the end-to-end ISO/IEC 8073 Cclass 4 transport protocol is required, in addition to the CLNP. This protocol itself uses CLNP packets to convey information between ATN users. The protocol can detect data loss and recovers from it by retransmission. It can also provide end-to-end flow control and multiplexing of different data streams between the same pair of users. When this protocol is used, the impact of a comparatively low delivery probability is on mean transit delay (the average time it takes to transfer data from source to destination). This is because recovery from data loss is by retransmission, and hence the lower the delivery probability, the longer the mean transit delay. There is hence a need to offset the impact of an increased mean transit delay against the cost and design implications of higher delivery probability.

3.2.4.1.4.4 ATN users that do not require a high delivery probability (this class includes time critical applications such as radar related data transfer;) could in principle directly use the transfer service provided by CLNP, but this is not permitted in the ATN. ATM applications for which the ISO/IEC 8073 COTP Class 4 is not appropriate are instead required to use the ISO/IEC 8602 connectionless transport protocol, which specifies a format for data transferred by the CLNP. The advantage of this protocol is that it decouples the internal structure of a user's system and the applications it hosts, from network routing. This is because ISO/IEC 8602 enables multiple users to be reached through one network address rather than one per user, which would be less efficient from the network point of view, and the number of such addresses is limited.

3.2.4.1.4.5 In order to minimise the risk of mis-delivery of CLNP packets to the wrong destination, the ATN provides an end-to-end check on the integrity of the CLNP Destination NSAP Address which is assumed to be sufficient even for ATM applications with demanding requirements on the mis-delivery rate. As the CLNP header checksum is not an end-to-end checksum, because it is manipulated by Routers on its way through the ATN Internet, an extended transport layer checksum is used which includes the Source and Destination NSAP Addresses within its scope by constructing a pseudo header including these addresses. This pseudo header is never transmitted but is assumed to be part of the TPDU for checksum computation purposes. This provides the receiving transport entity with a means of detecting and rejecting TPDU's mis-delivered due to an undetected error in the Destination NSAP Address.

### 3.2.4.2 *The User as an Organisation*



**Figure 3.2-4. Organizational End User's View of the ATN**


When the ATN user is an organisation it has multiple systems and subnetworks to consider. It is no longer attaching to the ATN as an individual End System, but as an organisation,



which may access the ATN and provide transit services to other ATN users, i.e. be an ATN Service Provider. Figure 3.2-4 illustrates the changed perception. The “organisational” ATN user has both Host Computers which are individual ATN users as before, but also operates a portion of the ATN cloud, with at least one link to the rest of the ATN.

### 3.2.4.2.1 Organisations and Routing Domains

3.2.4.2.1.1 The portion of the ATN internetwork operated by the organisational user may be no more than a single ATN Router. The portion of the ATN internetwork operated by such a user is termed an administrative domain. Alternatively, a larger organisational user may operate a large number of ATN Routers, interconnected by various subnetworks also operated by the organisational user, and providing ATN access to all the systems owned by the organisational user which require ATN access. Such Routers may also be general purpose and be part of that user's own internal network. However, regardless of how many ATN Routers there are within an organisation, the ATN Routers and the Host Computers to which they provide ATN access, typically form what is known as a Routing Domain. The Routing Domain is a structure specified in the ISO standards and imposed on the ATN Internet in order to enable a structured approach to be taken to solving the routing problem. This is illustrated in [Figure 3.2-5](#).



**Figure 3.2-5. The ISO Routing Domain**

3.2.4.2.1.2 Within a Routing Domain, ATN users are recommended to use the ISO/IEC 10589 intra-domain routing protocol. This is a simple and robust routing information exchange protocol that is specified for use between systems that mutually trust each other (i.e. belong to the same user). This protocol exchanges connectivity information throughout the Routing Domain and enables each ATN Router to build up a complete topology map of the Routing Domain, so that every ATN Router knows which routers within the Routing Domain provide access to which Host Computers, and how the routers themselves are interconnected. Routes can then be plotted through the Routing Domain, and CLNP packets forwarded to their addressed destinations within the Routing Domain.

- 3.2.4.2.1.3 Within an ATN Routing Domain, there will be one or more ATN Routers that are permitted to route CLNP packets to external destinations, i.e., addressed destinations that are located in the “rest of the ATN”. These routers are known as Boundary Intermediate Systems (BISs), because they exist at the boundaries of Routing Domains.
- 3.2.4.2.1.4 The ATN simply consists of multiple Routing Domains. Each such Routing Domain is self-consistent and capable of internal routing. However, key to the ATN being a single internetwork as opposed to a collection of separate Routing Domains, is the capability of inter-domain routing between BISs.
- 3.2.4.2.1.5 In the ATN, it is mandatory for a BIS to support the ISO/IEC 10747 Inter-Domain Routing Protocol (IDRP). For inter-domain communications, this protocol requires that BISs communicate directly over a common subnetwork, which may be owned by the owner of either BIS, or by a third party. Rather than exchanging connectivity information, as is done between routers within a Routing Domain, BISs advertise routes to each other, where a route consists of the set of addresses which identifies the destinations reachable over the BIS router, and information about the route's path including the Quality of Service and Security available over the route.
- 3.2.4.2.2 **Use of Policy Based Routing by Organisations**
- 3.2.4.2.2.1 It is the BIS's responsibility to determine which routes, if any, it will advertise to an adjacent~~another~~ BIS, and the use it will make of routes which it receives. When the BISs within a Routing Domain receive alternative routes to the same destination, then they must collectively determine which is the best route and hence which of the alternatives will be used. The set of rules which determines the advertisement and use of routes is known as a Routing Policy, and each organisational user of the ATN must determine and apply its~~their~~ own Routing Policy.
- 3.2.4.2.2.2 It is the need for policy based routing between different organisations that underlies the need for the existence of Routing Domains. Policy based routing enables users to control external access to their communications resources, and to protect themselves from problems elsewhere in the internetwork. BISs may also, depending on Routing Policy, advertise to BISs in other Routing Domains routes that have been received from another Routing Domain, and thereby offer transit facilities. However, Routing Policy may also prevent such routes from being re-advertised and hence deny transit facilities.
- 3.2.4.2.2.3 Organisational ATN users must therefore ensure that they either have direct connections with the ATN Routing Domains with which communications is necessary, or that those Routing Domains with which direct connections exist also offer suitable transit facilities to the remainder. In principle, this could be done on a bilateral basis between ATN organisational users on an “as needs” basis, and this is generally what is expected for the support of ground-ground communications. However, for air-ground applications support, this is unlikely to be an efficient strategy and may actually prevent useful communications by putting too high a cost on establishing a usable path even when physical connectivity already exists.

3.2.4.2.2.4 Instead, it is intended that ATN interconnections for support of air-ground communications are coordinated on both a regional and worldwide basis, so that an ATN backbone (of Routing Domains offering general transit facilities) is created, with either a clear apportionment of costs, or a known tariff, for use of transit facilities. This way users can gain access to the full capabilities of the ATN quickly and cheaply.

3.2.4.2.2.5 Policy based routing plays a significant role in the ATN, where it is used to support user requirements for control over the users of data links, for control over the type of traffic exchanged across the data links, and for optimising the distribution of routing information concerningfor routes to mobile systems.

### 3.2.4.3 *Mobile Users*

#### 3.2.4.3.1 **General**

3.2.4.3.1.1 The ATN will incorporate many “mobile” subnetworks. Examples of such subnetworks include SSR Mode S, AMSS and VDL. If an aircraft were to attach to one mobile subnetwork only and never to any other, then even though sometimes it may be attached and at other times not attached; this has no consequence for the ATN. This is because from the point of view of the rest of the ATN, it would be no different from a fixed system that was occasionally off-line. However, that is not how mobile subnetworks are used. An aircraft will attach to many different mobile subnetworks during the course of its flight. A long haul aircraft may move between the coverage areas of different satellites; an aircraft flying over a land mass will fly between different Mode S subnetworks as it passes over different countries. And, at the same time, the applications on board the aircraft will need to maintain contact with applications on the ground. Mobile platforms thus require special routing considerations.

3.2.4.3.1.2 In the ATN, mobile “platforms” are treated in a similar manner as organisational users. That is, the systems on board an aircraft are required to form a Routing Domain and hence must include an ATN Router that is also a BIS. This is partly because the ISO/IEC 10747 inter-domain routing protocol provides a relatively efficient mechanism for the transfer of routing information over low bandwidth links, but also because aircraft are almost always organisationally separate to the ground systems with which they are in contact and the same requirements for policy based routing apply.

3.2.4.3.1.3 The existence of mobile users has a significant impact on the organisation of the ground based ATN. While the ground topology will change only slowly, each aircraft's point of contact with the ground ATN will change rapidly with a consequent impact on the volume of routing information exchanged, and the routing tables in each router. A strategy is necessary for containing this high rate of information flow, and also to avoid the problems of routing instability caused by a rapid turnover of routing information.

3.2.4.3.1.4 This is the ATN Mobile Routing Strategy and it is based on a two level concept of default route providers. The first level is provided by a default route provider to all aircraft in a given region (known as an ATN Island). This default route provider is kept informed about routes to all aircraft currently in that region and hence can always provide a path to such aircraft. Several such default route providers may exist in the same region and collectively they are said to form the ATN Island's Backbone.

3.2.4.3.1.5 The second level is provided by an aircraft's home. The "home" of an aircraft does not necessarily relate to an airline's headquarters, its maintenance facilities, or indeed any geographical concept of "home". It is simply a particular ATN Routing Domain, and, in principle, any ATN RD will do. It may be an RD belonging to an aircraft's airline, but equally it may belong to a Communications Service Provider or an Administration. Typically, all aircraft belonging to the same airline, or the General Aviation (GA) aircraft of a single country are expected to share the same home.

3.2.4.3.1.6 The ATN's default route providers in each ATN Island keep a "home" informed about the location of all of its aircraft that are current in ATN communications. Thus, if a particular default route provider needs to route a packet to an aircraft for which it does not have an explicit route (i.e. it is not in the same region), all it has to do is to route the packet to the aircraft's known home and from there it can be forwarded to the ATN Island with which it is currently in contact and thence to that aircraft.

### 3.2.4.3.2 **Route Initiation**

3.2.4.3.2.1 The establishment of a communications path between BISs in any two Routing Domains is known as "Route Initiation". These procedures apply to the establishment of both ground/ground and air/ground communications. However, as opposed to the ground/ground case, Route Initiation for mobile users is dynamic and has to follow ICAO specified procedures for which guidance is given in section ~~3.5.10~~ 3.4.10 of Part IV of this document.

### 3.2.4.4 **Routing Control**

3.2.4.4.1 An important user requirement is that ~~its~~ users can specify, on a per application basis, routing control requirements. For AOC Applications, the requirement is for control over the air/ground data link used for air/ground communications applications. For ATSC Applications, the requirement is to follow only ATSC approved routes, and further, to be able to classify routes in the range of class 'A' to class 'H', and for the user data to follow a route of the most appropriate class. Some Aadministrations may also restrict the type of traffic carried over certain air/ground data links. Such restrictions must also be taken into account.

3.2.4.4.2 The ATN internetwork meets these requirements by:

- a) permitting ~~user's~~ users to identify, in the CLNP Header, the traffic type of the data being conveyed (e.g. ATSC, AOC, General Communications, etc.); and the routing control requirements;

- b) carrying information about the air/ground data links a route traverses, and any restrictions placed upon those data links, in each IDRPsRoute; and
- c) carrying information about whether a route is approved for ATSC purposes, and the assigned ATSC Class of the route, in each IDRPsRoute.

3.2.4.4.3 When a CLNP PDU is forwarded by an ATN Router, the user requirements are matched against the available routes and the appropriate route is followed.

3.2.4.4.4 In the case of AOC traffic, the user's requirements are enforced in a "strong" manner. That is, if a route meeting the user's requirements is not available, then the data is discarded.

3.2.4.4.5 In the case of ATSC traffic, a similar "strong" interpretation is made of the requirement to follow an ATSC approved route. However, the router will then simply choose the route with the best matching ATSC Class. This is a route with the requested class, or a higher class, or if no higher class route is available, then the ATSC approved route with the highest specified class out of those available is selected.

### 3.2.5 Technical Overview of the ATN Internetwork

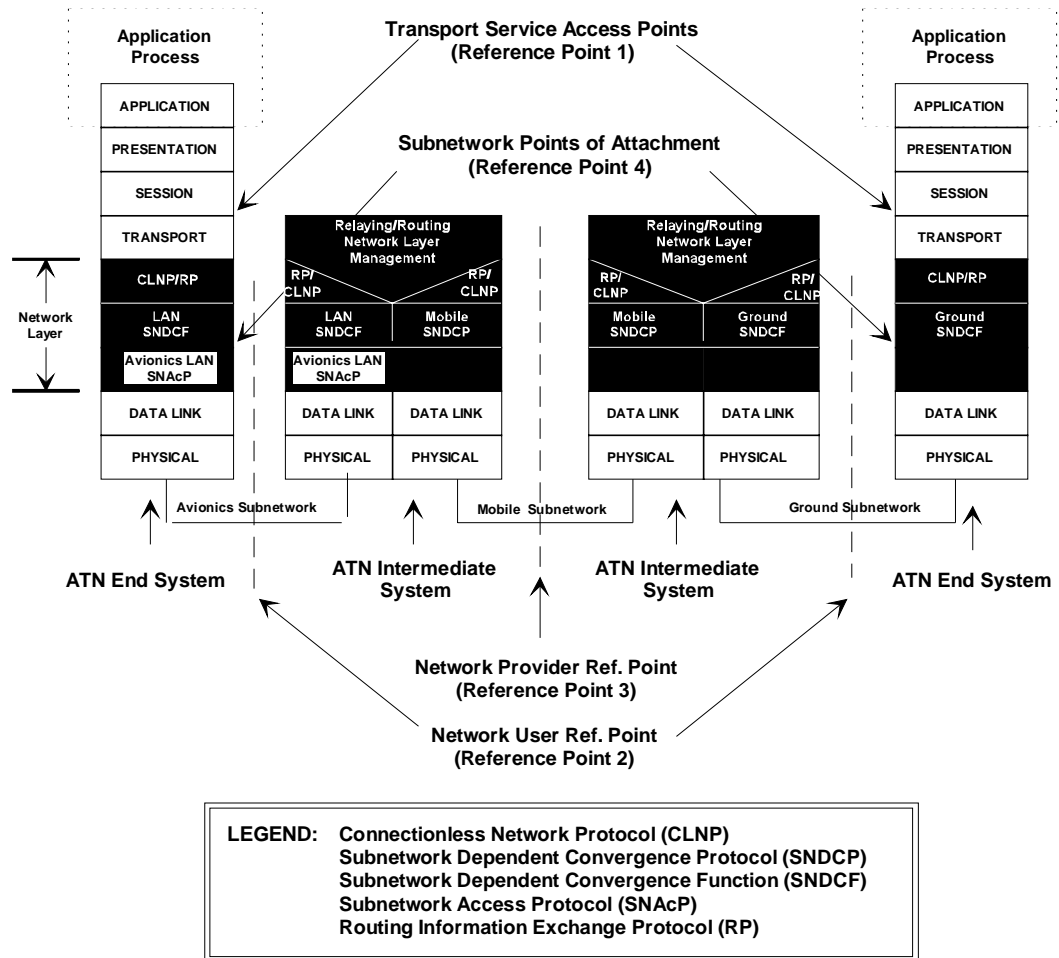
Different to the more general explanations contained in Part I of this document, the following sections go into the technical details of the ATN internetwork.

#### 3.2.5.1 Protocol Architecture

##### 3.2.5.1.1 General

3.2.5.1.1.1 The ATN Protocol Architecture is illustrated in Figure 3.2-6. This figure shows the protocols specified for two types of ATN End System and two types of ATN Router, and how those protocols relate to each other and in which OSI layer they are located, according to the OSI reference model specified by ISO in ISO/IEC 7498-1. The figure also shows a number of important interfaces.

3.2.5.1.1.2 The description of the ATN in Chapter 3.2.4 above has illustrated the fact that there is not one ATN interface, but instead there are many ATN interfaces, each of which serves a different user in a different role. In order to avoid confusion, a taxonomy of interfaces has been developed. This taxonomy identifies each significant interface as a *reference point*, and at each such reference point, there is an interface between two ATN entities, one taking on the role of a user, and the other the role of a as the service provider. Figure 3.2-76 illustrates the location of each ATN Reference Point, and the meaning of each reference point is given below.



**Figure 3.2-6. ATN Protocol Architecture**

- a) **Reference Point One**, the Transport Reference Point, is the OSI Transport Service interface and follows ISO/IEC 8072. This reference point is wholly contained within an ATN Host Computer, and represents access to the ATN Internet by an ATN Host Computer at the Transport Layer level. The user of the service provided in this example, is the OSI Session Entity; the service provider is the transport layer entity.
- b) **Reference Point Two**, the Network User Reference Point is the interface between the services of the network layers located in an ATN Host Computer and an ATN Router providing access to the ATN Internet. It comprises the OSI protocols used to access ATN Services.

*Note.— An ATN Router relays data between ATN Host Computers either directly to another ATN Host Computer or via one or more further ATN routers, which may or may not belong to other ATN Service Providers. When both Host Computer and Router are*

*owned by the same organisation, then the protocols that provide this interface are not mandated; the specification only recommends an appropriate stack.*

- c) **Reference Point Three**, the Network Provider Reference Point, is at the interface between two Routers belonging to different organisations. It comprises the OSI protocols used to support end-to-end communications via multiple ATN Routers, possibly via multiple subnetworks.

*Note.— When both Routers belong to the same ATN Service Provider then the protocols that provide this interface are not mandated; the specification only recommends an appropriate stack.*

- d) **Reference Point Four**, the Subnetwork Provider Reference Point, is at the interface between an ATN Host Computer or an ATN Router and a subnetwork. It comprises the OSI or subnetwork specific protocols used to access the service provided by that subnetwork and identifies the services provided by a real subnetwork used to connect two ATN components. This reference point identifies the lower boundary of the scope of the ATN ICS SARPs.

3.2.5.1.1.3 The provider of the service at reference point 4 is out of the scope of the ATN ICS SARPs. As indicated above, the ATN places only very limited constraints on the service provided at reference point 4, and this enables almost any subnetwork to be used as an ATN subnetwork.

### 3.2.5.1.2 **The ATN Transport Layer**

3.2.5.1.2.1 The ATN Transport Layer service provides transparent transfer of data between Transport Service users. All protocols defined in the Transport Layer have an 'End-to-End' significance, where the 'Ends' are defined as co-operating transport entities on two ATN host computers. The Transport protocols operates only between End Systems. Within the ATN, Transport Layer entities communicate over the ATN using the Network Service provided by the ATN Network Layer Entities.

3.2.5.1.2.2 There are two modes of the transport service, the Connectionless-mode Transport Service (CLTS) and the Connection-mode Transport Service (COTS). The connectionless-mode service allows two transport users to exchange individual datagrams, without flow control or the need to have previously established a connection, but with no guarantee of delivery. The connection-mode service allows two transport service users to negotiate a communications channel with a set of common characteristics, including reliable delivery of data units, and guaranteed (very high probability) order of delivery.

3.2.5.1.2.3 The two OSI protocols that provide the two modes of the transport service have separate specifications, and operate independently. Based on the higher level protocols operating within a given ATN host computer, one or both of the transport protocols may be implemented. Neither transport protocol is concerned with routing and relaying of data between End Systems, which is the responsibility of the Network Layer. The protocol in support of the CLTS is specified in ISO/IEC 8602, and the protocol in support of the COTS is specified to be ISO/IEC 8073 Class 4. The implementation of these protocols within the ATN is further described in Chapter 3.4.3.3 of Part IV of this document.

3.2.5.1.2.4 The Transport Service boundary corresponds with ATN reference point 1.

### 3.2.5.1.3 ~~The~~ ATN Network Layer

#### 3.2.5.1.3.1 **General**

3.2.5.1.3.1.1 The OSI Network ~~Layer~~ Service, like the OSI transport service is specified to provide both a connection-mode and a connectionless-mode service. However, in the ATN, the Network Layer Service is restricted to the connectionless-mode only. This is because, unlike the transport layer, the same network protocols must be implemented in every system in the ATN internetwork, if interoperability is to be guaranteed. In the case of the transport layer, the mode of the service required depends on the requirements of the users, and those End Systems that implement the same applications must also implement the same transport layer protocols. However, the internetwork itself must relay the data of all users, regardless of the mode of the transport service used. In order to provide universal connectivity, a consistent set of protocols must be implemented across the internetwork. Even if universal connectivity was ruled out, in practice, most ISs would still have to support all modes implemented by ESs, because of the tendency for data pathways to cross each other, regardless of the network service mode supported by each such data pathway.

3.2.5.1.3.1.2 It is thus cost effective to support only one mode of the network service. Implementation costs are reduced, and the complexity of validation is also reduced. Furthermore, mobile routing is not yet believed to be practicable when using the connection-mode network service.

3.2.5.1.3.1.3 The Network ~~Layer~~ Service is independent of the Transport ~~Layer~~ Service and may be used by ISO/IEC 8602 to provide the CLTS, and by ISO/IEC 8073 (class 4 procedures only) to provide the COTS.

3.2.5.1.3.1.4 The OSI Network Layer comprises three sub-layers or roles:

- a) Subnetwork Independent Convergence Role, which is responsible for providing a consistent Network Layer Service regardless of the underlying subnetwork;
- b) Subnetwork Dependent Convergence Role, which decouples the functions of the Subnetwork Independent Convergence Role from the characteristics of different subnetworks; and
- c) Subnetwork Access Role, which contains those aspects of the network layer specific to each subnetwork.



### 3.2.5.1.3.2 **The Subnetwork Independent Role**

3.2.5.1.3.2.1 In an ES, the Subnetwork Independent Role is responsible for providing the OSI Network Service independent of the real subnetwork(s) to which the ES is attached. In an IS, the Subnetwork Independent Role is responsible for the routing and relaying of user data along its route between the two communicating users. The protocols that support the exchange of routing information are also contained within this functional area.

3.2.5.1.3.2.2 In support of the connectionless-mode Network Service, it is a mandatory requirement that all ATN ESs and ISs implement the ISO/IEC 8473 ~~internetworking connectionless network~~ protocol (CLNP). This is a subnetwork independent protocol and supports the relaying of connectionless data PDUs over multiple subnetworks. By choosing such a protocol as its unifying characteristic, the ATN is cast as a subnetwork independent internetwork. CLNP supports the ISO global network addressing plan, quality of service specification, congestion control, and segmentation and reassembly of data packets. Additionally, provisions exist within CLNP for diagnostic actions such as end-to-end route recording and error reporting.

3.2.5.1.3.2.3 Three Routing Information Exchange Protocols are also specified in support of ISO/IEC 8473 within the ATN. These are:

- a) ISO/IEC 9542 — the End-System to Intermediate-System (ES-IS) protocol;
- b) ISO/IEC 10589 — the Intermediate-System to Intermediate-System (IS-IS) intra-domain routing information exchange protocol; and
- c) ISO/IEC 10747 — the Inter-Domain Routing Protocol (IDRP).

3.2.5.1.3.2.4 The use of these protocols is outlined below and described in more detail in Chapter ~~3.4.3.3~~ of Part IV of this document.

### 3.2.5.1.3.3 **End-System to Intermediate-System Routing Protocol**

3.2.5.1.3.3.1 The ISO/IEC 9542 ES-IS protocol provides a mechanism for ESs and ISs to exchange connectivity information within a local subnetwork environment. It is recommended for implementation in all ATN ESs and all ATN ISs that support ES attachment. In this role, its use applies to reference point 2.

3.2.5.1.3.3.2 The protocol enables ESs and ISs to dynamically discover each other when attached to the same subnetwork (only on broadcast subnetworks), and for ISs to inform ESs of optimal routes. In the absence of ISs (on broadcast subnetworks), ESs may also locate each other on an as needs basis.

3.2.5.1.3.3.3 The ES-IS protocol also complements the IS-IS routing protocols to support dynamic discovery of other ISs and/or their NETs, and is also used in a similar manner to support the Inter-Domain Routing Protocol over mobile subnetworks.

#### 3.2.5.1.3.4 **Intra-Domain Routing Information Exchange Protocol**

3.2.5.1.3.4.1 The ISO/IEC 10589 IS-IS intra-domain routing information exchange protocol is used by ISs within the same Routing Domain to exchange connectivity and ~~QoS~~QoS information. As the ISs within a single Routing Domain are always operated by the same organisation, this protocol is not used at any of the ATN interfaces identified by reference points.

3.2.5.1.3.4.2 The protocol works at two levels. Level 1 operates within the same Routing Area, while level 2 operates between Routing Areas. From the information exchanged by this protocol, ISs build up a topography map of the local Routing Area at level 1, or Routing Area connectivity, at level 2. From this map, optimal routes can be plotted, and the relevant information provided to each IS's Forwarding Information Base.

#### 3.2.5.1.3.5 **The Inter-Domain Routing Protocol (IDRP)**

3.2.5.1.3.5.1 The ATN has adopted the ISO/IEC 10747 Inter-domain Routing Protocol, for the exchange of dynamic routing information at the inter-domain level. IDRP is a "vector distant" routing protocol and is concerned with the distribution of routes where a route comprises a set of address prefixes for all destinations along the route and the route's path, i.e. the list of Routing Domains through which the route passes in order to reach those destinations. In addition, a route may be further characterised by various service quality metrics (e.g. transit delay).

3.2.5.1.3.5.2 Under IDRP, specialised Boundary Routers, i.e. Boundary Intermediate Systems (BISs), in each Routing Domain advertise to Boundary Routers in adjacent Routing Domains, routes to the systems contained in that Routing Domain. Typically, there is a route for each performance metric and security category supported, and the destination of these routes is the Address Prefix(es) that characterises the Routing Domain. The receiving Routing Domains then store this information and use it when they need to route packets to destinations within the other Routing Domain. A route so received may also be re-advertised to other Routing Domains adjacent to the Routing Domain that first received it, and onwards throughout the ATN Internet. Ultimately, every Routing Domain in the ATN Internet can receive a route to every other Routing Domain.

3.2.5.1.3.5.3 However, without any other functionality, IDRP would not provide a scaleable approach to routing. In order to provide such a scaleable architecture, IDRP enables the aggregation of routes to Routing Domains with common address prefixes, into a single route. It is thereby possible for the number of routes known to any one router to be kept within realistic limits without reducing connectivity within the Internetwork.

### 3.2.5.1.3.6 **Subnetwork Dependent Role**

3.2.5.1.3.6.1 The OSI Subnetwork Dependent Role is responsible for decoupling the functions of the subnetwork independent role from the characteristics of the different subnetworks and provides a consistent service to any protocols implemented by the subnetwork independent role. In doing so, it may implement a convergence protocol, implemented on a hop-by-hop basis, independently over each subnetwork. This is a Subnetwork Dependent Convergence Protocol (SNDCP).

3.2.5.1.3.6.2 ISO/IEC 8473 may be adapted to all known subnetwork types and hence a SNDCP is not specifically required. However, each subnetwork class does require a different adaptation, and each such adaptation is known as a Subnetwork Dependent Convergence Function (SNDCF). Chapter 43.6 of Part IV of this document discusses the SNDCFs that may be used to interface ISO/IEC 8473 to ATN subnetworks.

3.2.5.1.3.6.3 However, while ISO/IEC 8473 does not require an SNDCP, there is justifiable concern over the ISO/IEC 8473 protocol overhead in respect of the low bandwidth communications provided by the mobile subnetworks. For this reason, an SNDCP has been specified to provide compression of the ISO/IEC 8473 protocol header over mobile subnetworks. This is known as the LREF compression algorithm and is further described in Chapter 43.3.4.6.3 as part of the description of the Mobile SNDCF which supports LREF and other compression algorithms.

### 3.2.5.1.3.7 **Subnetwork Access Role**

The Subnetwork Access Role comprises the functions necessary to support access to a specific subnetwork. This is dependent on the specification of each subnetwork and is hence outside of the scope of this document. The service provided by the Subnetwork Access Role to the Subnetwork Dependent Role is at ATN reference point 4, which identifies the lower boundary of ~~this manual~~ the scope of the ATNI ICS SARPs.

## 3.2.5.2 **Congestion Management**

3.2.5.2.1 Congestion is a phenomenon experienced by a rRouter in an internetwork when the queuing delays through that rRouter 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.

3.2.5.2.2 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 unexpected but 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.

3.2.5.2.3 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.

3.2.5.2.4 The Congestion Avoidance pProcedures specified for the ATN represent a significant improvement on those which many will be familiar with from experience with the TCP/IP Internet. TCP Congestion Management only recognises that a network is congested when there is a need for a sender to retransmit a packet, when it is assumed that packet loss has occurred due to an overloaded router discarding packets. The sending system then first reduces the rate at which it transmits packets and then gradually speeds up again until it needs to retransmit, and so on.

3.2.5.2.5 The problem with this approach is that it constantly pushes the network into overload and, like a congested motorway, results in poor and turbulent traffic flow and a much greater loss in throughput than should have occurred.

3.2.5.2.6 The ATN takes a quite different approach. In the ATN a router that is approaching congestion indicates this by setting a “flag bit” in each packet header (this situation occurs when the outgoing packet queue is longer than one packet), and this flag is recognised by the receiving End Ssystem. If enough flags are set in a given sampling period, the receiving End Ssystem reduces the credit it offers to the sender, thus reducing the load on the network.

3.2.5.2.7 This approach has the significant advantage that it “kicks in” before network performance starts degrading thus permitting a much more stable traffic flow and enables near optimal throughput to be maintained. The ATN Congestion Management algorithm is discussed in more detail in Chapter 3.65 of Part IV of this document.

### 3.2.5.3 *Addressing*

3.2.5.3.1 Every system within a network such as the ATN, must have a unique address. This address may then be used to identify the source and destination of a packet sent through the network. ATN routers use a packet’s destination address to determine how the packet is routed to its destination.

3.2.5.3.2 An ATN address is therefore more than a unique identifier for each system, and to be truly useful, it must be possible to use an address to find out how to reach the addressed system i.e. to select the most appropriate route. That is, an address must somehow relate to a network’s topology.

- 3.2.5.3.3 Routing Domains can be viewed as being like telephone areas, and like all subscriber numbers in a telephone area, the addresses of systems within the same Routing Domain should all have a common prefix. Then a packet sent to any system in the Routing Domain, can be sent to the Routing Domain without the routers along the way having to have any knowledge of the topology of the networks and routers within that Routing Domain.
- 3.2.5.3.4 Routing Domains are, however, a more flexible concept than telephone areas. The requirement for a single common address prefix is not absolute, and it is possible to have more than one address prefix that characterises a single Routing Domain. The geographical country is also not present in either the ISO Routing Framework, or as a fixed quantity in the Address Plan. Instead, there is the very general concept of the Routing Domain.
- 3.2.5.3.5 More detailed information on allocating ATN addresses, including those for the ATN internetwork, can be found in Part II of this document.

### 3.3 **ATN Protocols and Functions**

#### 3.3.1 **Introduction**

- 3.3.1.1 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/or Airline applications. An End System is essentially any computer system that is connected to one or more ATN routers and implements the ATN communications protocols.
- 3.3.1.2 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 2 of the ATN Internet Communications Service (ICS) SARPs.
- 3.3.1.3 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/IEC 9542 ES-IS protocol. Support of ISO/IEC 9542 will be necessary if this is required by the ATN Router(s) through which the End System accesses the ATN; it is a local matter as far as the ATN ICS SARPs are concerned.
- 3.3.1.4 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 IDRPs, although the support requirements for IDRPs do differ depending on the role of the Router. Local considerations may also require support of the ISO/IEC 9542 ES-IS protocol and/or the ISO/IEC 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 5.3 of the ATN ICS SARPs.

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

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

### 3.3.2 **Transport Layer Considerations**

#### 3.3.2.1 *The ATN Transport Layer*

##### 3.3.2.1.1 **Transport Layer Model**

3.3.2.1.1.1 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 3.3-1). 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).

3.3.2.1.1.2 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:

- a) the Connection Mode Transport Service (COTS); or
- b) the Connectionless Mode Transport Service (CLTS).

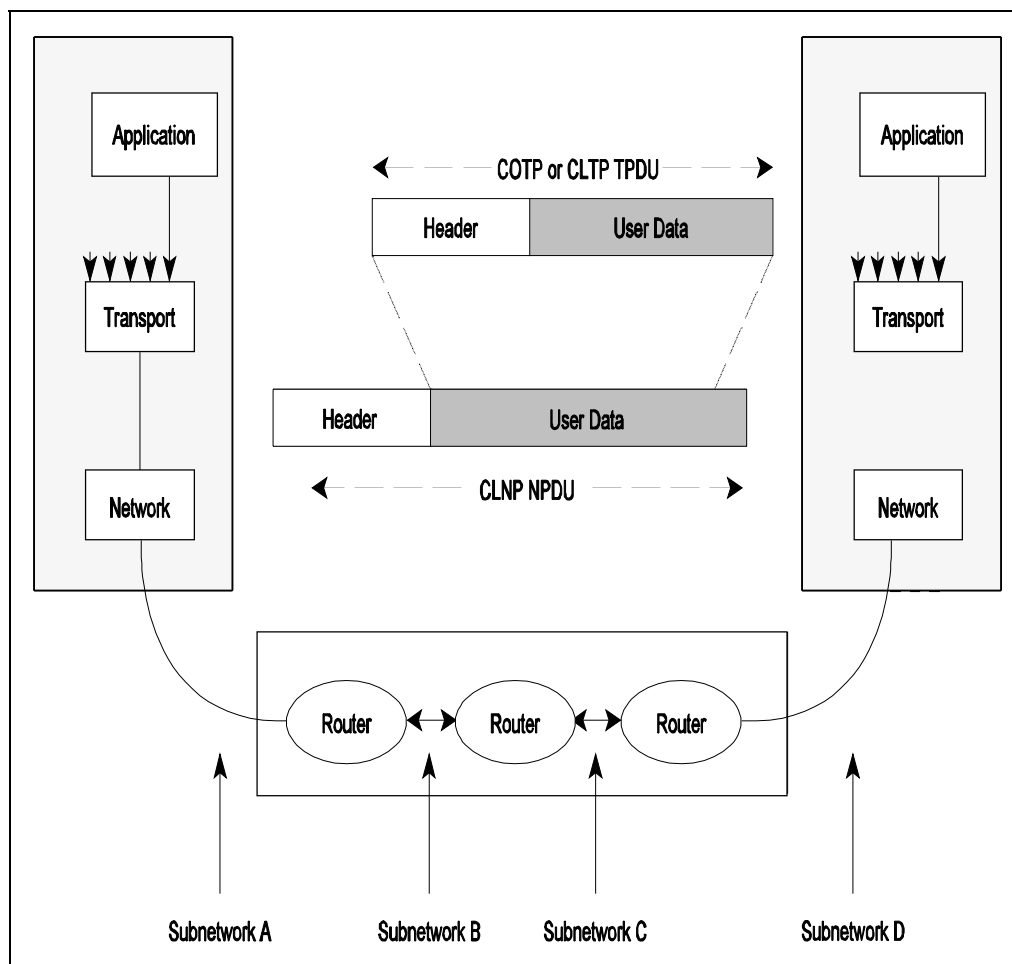


Figure 3.3-1. Scope of Transport Layer Interactions

### 3.3.2.1.2 Transport Layer Protocols

3.3.2.1.2.1 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 supported by the COTP and the CLTS by the CLTP.

3.3.2.1.2.2 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.

3.3.2.1.2.3 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 provided by the ATN Network Layer.

3.3.2.1.2.4 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/IEC 8073. This enables the reliable sequenced data transfer, where the user is guaranteed that 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.

3.3.2.1.2.5 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/IEC 8602.

### 3.3.2.1.3 **Service Provided by the ATN Transport Layer**

3.3.2.1.3.1 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.

### 3.3.2.1.3.2 **Service Provided by the COTS**

3.3.2.1.3.2.1 As far as application ~~designer's~~designers 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:

- a) 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;
- b) 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;



- c) the order of submission of TSDUs is preserved on delivery;
- d) the underlying transport protocol provides facilities to detect and recover from end-to-end transmission errors within a TSDU;
- e) 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;
- f) 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; and
- g) operation of the COTP requires system resources to maintain shared state and to monitor connection status.

### 3.3.2.1.3.3 **Service Provided by the CLTS**

3.3.2.1.3.3.1 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/IEC 8602 protocol. The characteristics of the service provided by the CLTP include the following:

- a) 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;
- b) there are no TPDUs transmitted other than those carrying user data;
- c) 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;
- d) the transport protocol can employ facilities to detect end-to-end transmission errors within a TSDU. TSDUs containing detected errors are discarded;
- e) TSDU sizes are limited to the maximum NSDU size on each end system; no segmentation is performed by the connectionless\_mode transport protocol;
- f) because there is no negotiated relationship between TS-users, the protocol does not have the capability to control the flow of TSDUs; and

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

#### 3.3.2.1.4 **Transport Addresses**

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

3.3.2.1.4.2 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, the Transmission Control Protocol), then these must use different NSAP Addresses.

#### 3.3.2.1.5 **Network Service Assumptions**

3.3.2.1.5.1 The ATN Transport Layer operates using the connectionless network service provided by the ATN network layer. All TPDU's are transmitted and received as NSDU's 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 it was sent, in duplicate, or not at all. Although it is possible for NSDU's 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 NPDU's are discarded by congested routers.

#### 3.3.2.1.6 **ATN Security and Priority**

3.3.2.1.6.1 The ATN ICS 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's 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.

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

3.3.2.1.6.3 The ATN Security Label and priority applicable to each TPDU are parameters of the ~~N-UNITDATA~~ N-UNITDATA service and are therefore encoded in the ~~NDPUNPDU~~ NDPUNPDU header, rather than in each TPDU, and are referenced by the network layer forwarding function. For such reasons, TPDU's from transport connections with different ATN Security Labels, and/or priorities, cannot be concatenated.

### 3.3.2.2 **Provision of the Connection Mode Transport Service**

#### 3.3.2.2.1 **Overview**

3.3.2.2.1.1 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.

3.3.2.2.1.2 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.

3.3.2.2.1.3 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.

3.3.2.2.1.4 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, once 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.

3.3.2.2.1.5 The COTS is realised through the implementation of the COTP.

3.3.2.2.1.6 The ATN COTP uses the ISO/IEC 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 network problems from the TS-user.

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

- a) **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.
- b) **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.

- c) **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 TS-user.

### 3.3.2.2.2 Connection\_Mode Transport Service Primitives

3.3.2.2.2.1 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.

3.3.2.2.2.2 A TS primitive issued by one TS-user will, in general, result in receipt of an indication by the other TS-user. Figure 3.3-2 gives a summary of TS-primitive time-sequence diagrams for some typical scenarios.

3.3.2.2.2.3 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 3.3-2, the flow of time is represented by the downward direction in the 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 (~).*

3.3.2.2.2.4 Figure 3.3-2 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.

3.3.2.2.3

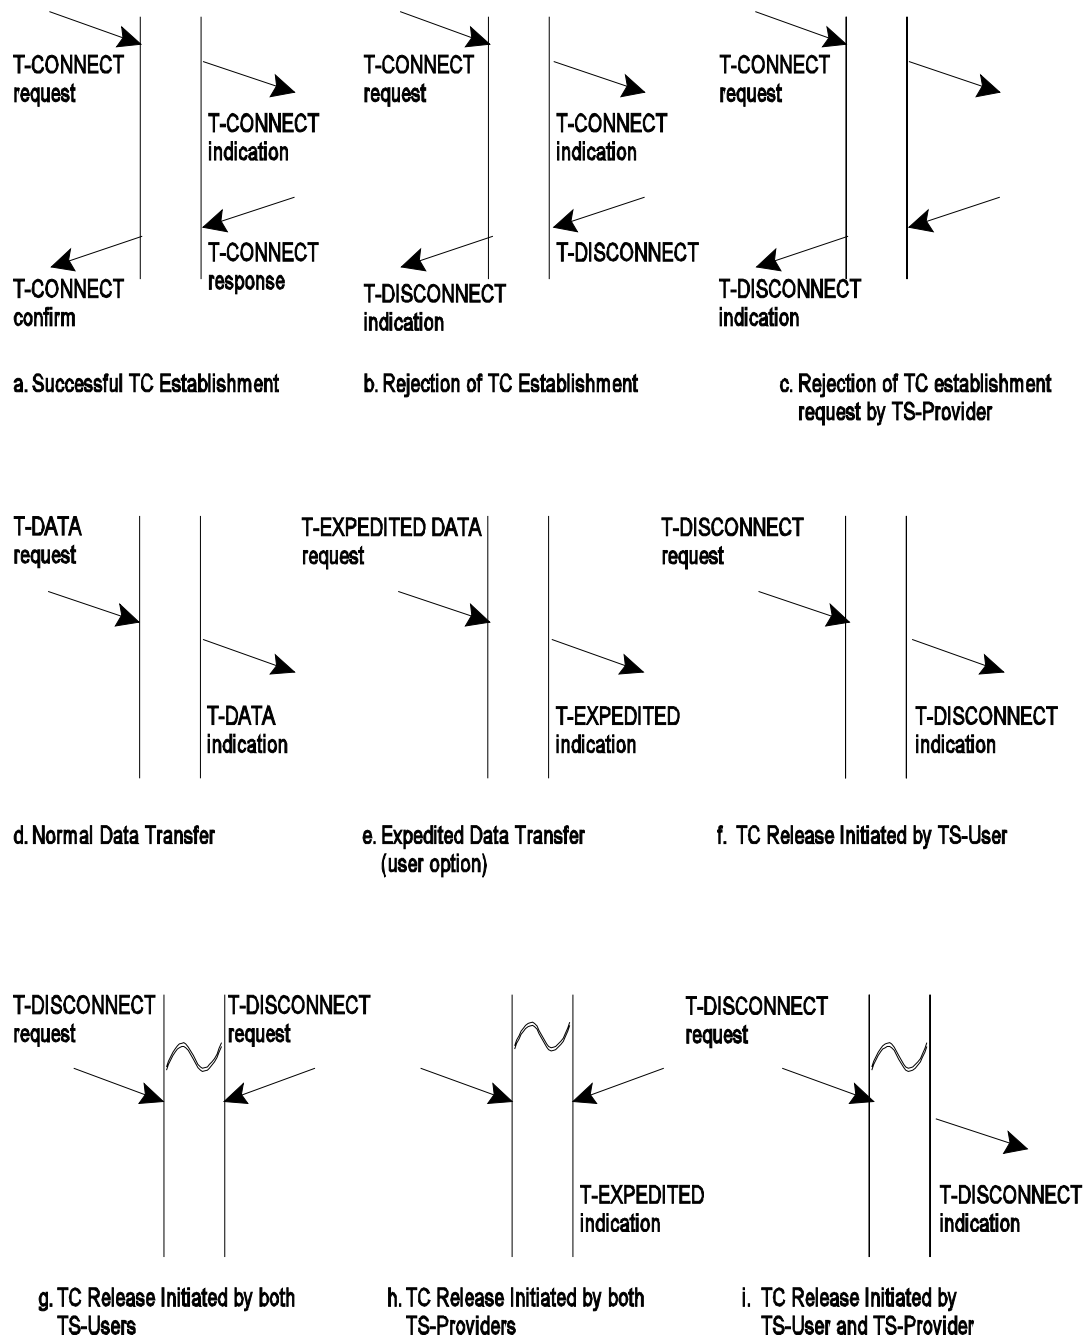


Figure 3.3-2. Transport Service Time Sequence Diagrams

## The Connection-Mode Transport Protocol (COTP)

### 3.3.2.2.3.1 Overview

3.3.2.2.3.1.1 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.

3.3.2.2.3.1.2 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.

3.3.2.2.3.1.3 The handshake and timer-based mechanisms are combined to ensure that connection identifiers are unique during the maximum time packets may remain in the system.




Figure 3.3-3. TPDUs Exchanges for Connection Establishment

### 3.3.2.2.3.2 Connection Establishment

- 3.3.2.2.3.2.1 The COTP uses a three-way handshake mechanism in combination with a timer-based mechanism to ensure connection establishment in class 4. Figure 3.3-3 illustrates a typical transport connection establishment procedure. The service TS-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 (CR) 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 TS-user.
- 3.3.2.2.3.2.2 If the TS-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 TS-user that its connection establishment request has been accepted by invoking a T-CONNECT confirm primitive.
- 3.3.2.2.3.2.3 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.
- 3.3.2.2.3.2.4 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 3.3-4.

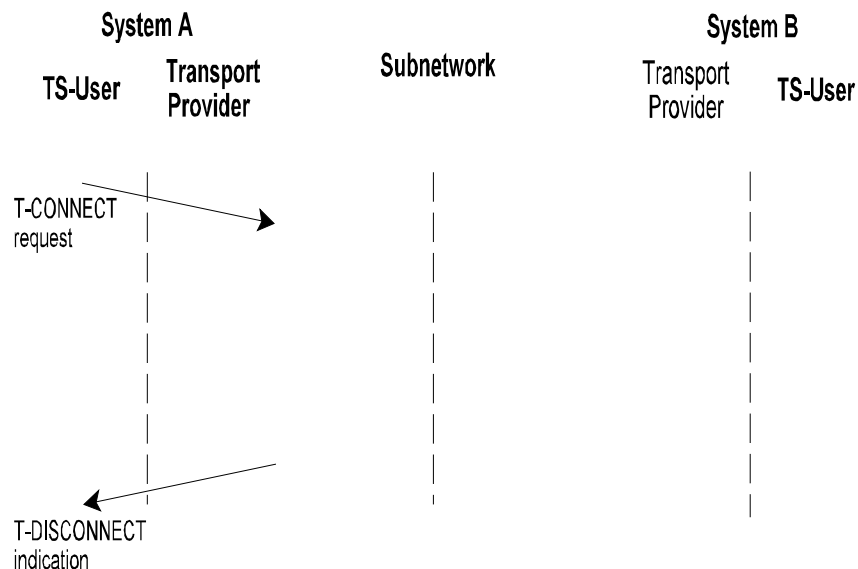


Figure 3.3-4. Connection Refusal by the TS-Provider

- 3.3.2.2.3.2.5 To initiate communications with a peer, a TS-user invokes the T-CONNECT request primitive (see Figure 3.3-2). 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 3.3-1.
- 3.3.2.2.3.2.6 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:
- a) time-out procedures, such as when a TS-user does not respond to a connection request;
  - b) rejection by the TS-provider of an attempt to establish a TC (part c of Figure 3.3-2), for reasons such as invalid or unknown called TSAP address, lack of local or remote resources of the TS-provider etc., or,
  - c) unwillingness of the called TS-user to accept the TC establishment request (part b) of Figure 3.3-2).
- 3.3.2.2.3.2.7 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.



**Table 3.3-1. TC Establishment Primitives and Parameters**

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	$\theta Q$	$\theta(=)Q(=)$		

*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

$\theta Q$  Use of this parameter is a TS-user option

### 3.3.2.2.3.2.8 Connection Request

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

- a) 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/IEC 8073);
- b) 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;
- c) 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;

- d) 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;
- e) 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 are not recommended for direct use by applications; and
- f) 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 ICS 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.*

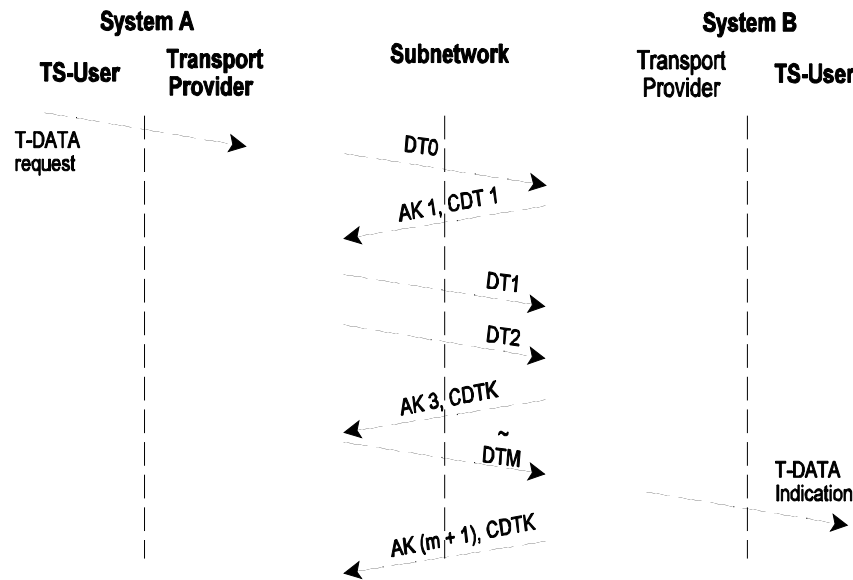
*Note 2.— In practice, not all of the parameters in a connection request must be explicitly specified, even though they exist in the service interface. For example, the invoking TS-user may only be required to specify the called transport address if the transport entity knows the calling address a priori. 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.*

#### 3.3.2.2.3.2.9 **Connection Indication**

3.3.2.2.3.2.9.1 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:

- a) calling and called address;
- b) expedited data option;
- c) TS-user-data;
- d) indicated QoS; and
- e) Security.

- 3.3.2.2.3.2.9.2 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.
- 3.3.2.2.3.2.10 **Connection Response**
- 3.3.2.2.3.2.10.1 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 3.3-4). The associated parameters and their corresponding values are the same as in the T-CONNECT request as defined in Table 3.3-1.
- 3.3.2.2.3.2.11 **Connection Confirm**
- 3.3.2.2.3.2.11.1 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.
- 3.3.2.2.3.3 **Data Transfer**
- 3.3.2.2.3.3.1 Once a connection has been successfully opened, data transfer 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.
- 3.3.2.2.3.3.2 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 3.3-5, 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.



**Figure 3.3-5. Normal Data Transfer**

3.3.2.2.3.3.3 The transport service provides for bi-directional 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 3.3-2(d) describes the primitive sequences in a successful transfer of normal data.

#### 3.3.2.2.3.3.4 **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.

#### 3.3.2.2.3.3.5 **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.

#### 3.3.2.2.3.3.6 **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 3.3-2 (e) shows the sequence of primitives in a successful transfer of expedited data.

#### 3.3.2.2.3.3.6.1 **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.

#### 3.3.2.2.3.3.6.2 **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.

#### 3.3.2.2.3.3.7 **Connection Termination**

3.3.2.2.3.3.7.1 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 3.3-2 (f) through (i).

3.3.2.2.3.3.7.2 The first scenario (f), is shown in more detail in Figure 3.3-6. 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.

3.3.2.2.3.3.7.3 In Figure 3.3-2 (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.

3.3.2.2.3.3.7.4 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 3.3-2.

Table 3.3-2. TC Release Primitives and Parameters

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.

#### 3.3.2.2.3.3.7.5 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 3.3-6. Transport Connection Termination

### 3.3.2.2.3.3.7.6 **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:

- a) **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 3.3-2), or if TS-users initiate a T-DISCONNECT request simultaneously (see part (g) of Figure 3.3-2). In these cases these data may be lost; and
- b) **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).

### 3.3.2.2.4 **The ATN Security Label**

3.3.2.2.4.1 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; ~~and~~
- d) protecting the routing information base of ATN BISs from external threats.

3.3.2.2.4.2 The ATN Internet provides mechanisms to support items (b) ~~and (c)~~ through (d) above only. ~~These mechanisms (b) and (c) 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.~~

3.3.2.2.4.3 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.

3.3.2.2.4.4 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.

3.3.2.2.4.5 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.*

### 3.3.2.2.5 **ATN Transport Layer Quality of Service**

3.3.2.2.5.1 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.

3.3.2.2.5.2 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; and



- c) the strong QoS Requirements of certain AINSC 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.

3.3.2.2.5.3 The only exception to this is the Residual Error Rate and the probability of packet mis-delivery. 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.

3.3.2.2.5.4 The risk of packet mis-delivery is mitigated by an extended transport checksum which includes the Source and Destination NSAP Adresses within its scope by constructing a pseudo header including these addresses. This pseudo header is never transmitted but is assumed to be part of the TPDU for checksum computation purposes. Through this mechanism the ATN Internet ensures that the probability of not detecting mis-delivered CLNP packets is better than 1 in  $10^8$ .

3.3.2.2.5.5 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.

### 3.3.2.2.6 **Priority**

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

3.3.2.2.6.2 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.

3.3.2.2.6.3 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.

3.3.2.2.6.4 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 connection request. TPDUs belonging to transport connections with different priorities cannot be concatenated.

3.3.2.2.6.5 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.

3.3.2.2.6.6 Transport 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.

3.3.2.2.6.7 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 3-2 of Section 1-2~~ of Subvolume 1 of the ATN 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.

3.3.2.2.6.8 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.

3.3.2.2.6.9 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 ~~Section~~Subvolume 1 of the ATN SARPs for further details on the mapping of Transport priority values to CLNP priority values).

### 3.3.2.2.7 **Negotiation of Connection Parameters**

3.3.2.2.7.1 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; and
- 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.

3.3.2.2.7.2 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 3.3-3 indicates the items that can be negotiated and their default, or mandatory, values in Class 4 operation.

**Table 3.3-3. Negotiable and Default Values for 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

3.3.2.2.7.3 **Class Negotiation — Initiator**

3.3.2.2.7.3.1 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.

3.3.2.2.7.3.2 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.

3.3.2.2.7.4 **Class Negotiation — Responder**

3.3.2.2.7.4.1 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.

3.3.2.2.7.5 **TPDU Size Negotiation**

3.3.2.2.7.5.1 All transport entities must be able to support a TPDU size of 128 octets, the default required by ISO/IEC 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.

3.3.2.2.7.5.2 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 TPDU's transmitted as NSDU's.~~

#### 3.3.2.2.7.6 **Use of Extended Format**

3.3.2.2.7.6.1 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.

3.3.2.2.7.6.2 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.

3.3.2.2.7.6.3 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/IEC 8073, the ability to support normal formats is mandatory.

#### 3.3.2.2.7.7 **Expedited Data Transport Service**

3.3.2.2.7.7.1 Support of the expedited data transport service is required by ISO/IEC 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).

#### 3.3.2.2.7.8 **Non-use of Checksum**

3.3.2.2.7.8.1 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 TPDU's, and thus all TSDUs.

3.3.2.2.7.8.2 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.

3.3.2.2.7.8.3 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.*

3.3.2.2.7.8.4 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.

### 3.3.2.2.7.9 Use of **S**elective **A**cknowledgement

3.3.2.2.7.9.1 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.

3.3.2.2.7.9.2 Because the selective acknowledgement feature reduces the need for retransmitting TPDU's, 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 [3.3.4.4.33.3.2.2.10.5.3](#) for a description of the selective acknowledgement feature.*

### 3.3.2.2.7.10 Use of Request of Acknowledgement.

3.3.2.2.7.10.1 The default for Request of Acknowledgement (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 acknowledges all TPDU's 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.

### 3.3.2.2.8 Error Handling

#### 3.3.2.2.8.1 Action on Receipt of a Protocol Error

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

- a) the transport layer can issue an ER TPDU;

- b) the transport layer can terminate the transport connection (that is, issue a DR TPDU);  
or
  - c) the transport layer can discard the TPDU (that is, ignore the error).
- 3.3.2.2.8.2 Events which qualify as a protocol error are defined in ISO/IEC 8073. It is recommended that ~~in event of a protocol error, that~~ the transport layer issues an ER TPDU in the event of a protocol error, and either discards the TPDU, or responds with a DR TPDU. This action ensures that the cause of a protocol error can be more readily identified.
- 3.3.2.2.8.3 Actions on Receipt of an Invalid or Undefined Parameter in a CR TPDU:
- 3.3.2.2.8.3.1 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.
- 3.3.2.2.8.3.2 ISO/IEC 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/IEC 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.
- 3.3.2.2.8.3.3 Another example of a parameter which may be included in a CR TPDU by the connection initiator and which will be ignored on receipt by implementations compliant with previous editions of the ATN ICS SARPs is the Extended Transport Checksum. Use of this parameter is ATN-specific and not in compliance with ISO/IEC 8073. However, it is not a protocol error to use such a parameter in a CR TPDU, and a receiving transport entity not supporting this parameter would simply pass over the option and continue to process the rest of the TPDU. A transport connection could then be established which operates without using the Extended Transport Checksum. However, if the parameter is supported by both the connection initiator and connection responder, then it will be present in the CC TPDU and thence in all TPDU's exchanged over the established transport connection.
- 3.3.2.2.8.3.4 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.
- 3.3.2.2.8.4 Actions on Receipt of an Invalid or Undefined Parameter in a TPDU other than a CR TPDU
- 3.3.2.2.8.4.1 For all other TPDU's, the decision as 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.

### 3.3.2.2.9 Timers and Protocol Parameters

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

3.3.2.2.9.2 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 are required to support configurable values for all timers and protocol parameters, rather than having fixed values. This allows modification as operational experience is gained.

*Note 1.— Refer to Table 3.3-4/5.5-1 of the ATN ICS SARPs for the complete listing of timers and parameters.*

*Note 2.— The values in this table have been derived from simulation focusing on transport entities operated over an AMSS subnetwork.*

*Note 3.— Refer also to 12.2.1.1 of ISO/IEC 8073 for more details on the timers.*

*Note 4.— In Table 3.3-4, the subscripts “R” and “L” refer to “remote” and “local”, respectively. The variable ERL, for example, refers to the maximum transit delay from the remote entity to the local entity. The variable ELR is the maximum transit delay from the local entity to the remote entity. It is assumed that these values may be different.*

*Note 5.— The values in Table 3.3-4 have been derived from simulation focusing on transport entities operated over an AMSS subnetwork.*

### **Table 3.3-4 Recommended Timer and Parameter Values and Ranges**

<i>Name</i>	<i>Description</i>	<i>Minimum value</i>	<i>Nominal value</i>	<i>Maximum value</i>
<i>MRL, MLR</i>	<i>NSDU lifetime, seconds</i>	26	400	600
<i>ERL, ELR</i>	<i>Maximum transit delay, seconds</i>	†	100	150
<i>AL, AR</i>	<i>Acknowledgement time, seconds</i>	†	20	400
<i>T1</i>	<i>Local retransmission time, seconds</i>	†2	22†	300
<i>R</i>	<i>Persistence time, seconds</i>	†	443	2710
<i>N</i>	<i>Maximum number of transmissions</i>	†	3	10
<i>L</i>	<i>Time bound on reference and/or sequence numbers, seconds</i>	160	1263	3000
<i>I</i>	<i>Inactivity time, seconds</i>	600	4500	6000
<i>W</i>	<i>Window time, seconds</i>	160	4000	6000

## 3.3.2.2.9.3

Several of the transport timers and variables listed in Table 3-3-45.5-1 of the ATN ICS SARPs are not directly configurable, but may be determined based on the values of other timers and variables. That is:

- a) the NSDU lifetime variables,  $M_{RL}$  and  $M_{LR}$ , may have a general estimate, based on the lifetime values used for NPDU. 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 transit time. The expected air-to-ground end-to-end transit time can be up to 30-40 seconds;
- b) 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;
- c) 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 TPDU so that the uncertainty about delivery is minimized. 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;
- d) 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;

This relationship is used to determine the initial, minimal and maximum value of  $T1$ . However, it does not force  $T1$  to remain fixed over time; in the ATN, all COTP



entities are required to dynamically adapt the local retransmission time individually on each transport connection in function of the round trip time measured over the connection. This issue is addressed in section 3.3.2.2.10.

- e) 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;
- f) 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 communications is no longer possible;
- g) the maximum time to receive an acknowledgement of a given TPDU, L, is bounded by ISO as:

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

- h) 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;
- i) 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; and
- j) 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.

### 3.3.2.2.10      **Dynamic Local Retransmission Time Adaptation**

#### 3.3.2.2.10.1      **Introduction**

3.3.2.2.10.1.1      The Connection-Oriented Transport Protocol (COTP) provides a reliable transport layer by retransmitting data, when the data has not been acknowledged after some period of time. A critical element of any COTP implementation is the timeout and retransmission strategy: the main issue is the determination of an appropriate timeout interval.

3.3.2.2.10.1.2      The timeout interval has important and conflicting effects on individual user throughput and overall network efficiency. To achieve optimal throughput, short timeout intervals should be used. Unfortunately, what is good for throughput is disastrous for efficient network utilization. If the timeout interval is too short, then a large number of packets may be retransmitted unnecessarily, causing a waste of the bandwidth.

3.3.2.2.10.1.3      In the ATN, the round-trip time (RTT), i.e. the time interval between sending a packet and receiving an acknowledgement for it, is expected to change over time, as routes and network

traffic load might change. For this reason an ATN COTP implementation must be designed to dynamically measure the RTT and use this measure to dynamically adapt the retransmission time out interval.

3.3.2.2.10.1.4 The following sections describe three mechanisms that have been proved efficient and suitable to the ATN environment.

### 3.3.2.2.10.2 **Dynamic Adaptation of the Retransmission Timer Value**

#### 3.3.2.2.10.2.1 **General**

3.3.2.2.10.2.1.1 The dynamic adaptation of the retransmission timer value is a key function of any COTP implementation. Ideally, the value of the retransmission timer must be continuously kept

- a) as close as possible to the experienced RTT; however,
- b) be sufficiently large so as to avoid (or limit the number of) unnecessary retransmissions, when variation of the RTT is experienced.

3.3.2.2.10.2.1.2 Clearly, the 'optimum' value of the retransmission timer is therefore related to the round-trip time experienced over the connection. Shorter round-trip times should lead to smaller values of the retransmission timer, and larger RTTs to larger values.

3.3.2.2.10.2.1.3 Hence, making a good estimation of the round-trip time is the core function for the adaptation of the retransmission timer.

3.3.2.2.10.2.1.4 An ATN COTP implementation is expected to predict future round-trip times by sampling the round-trip time of TPDU's sent over a connection and averaging those samples into a "smoothed" round-trip time estimate, SRRT. The SRRT can be thought as a prediction on the next round-trip time measurement. This RTT estimator is then used to derive a suitable value for the retransmission timer.

3.3.2.2.10.2.1.5 Different algorithms have been developed for the computation of the SRRT and of the retransmission timer values. These algorithms are discussed in the following sections.

#### 3.3.2.2.10.2.2 **Basic Algorithm (Original TCP Algorithm)**

3.3.2.2.10.2.2.1 A basic procedure for the dynamic adaptation of the retransmission timer of a transport entity was defined in the original specification of the TCP protocol. The principles of this algorithm were the following:

- a. When a TPDU is sent over a transport connection, the sender times how long it takes for it to be acknowledged, producing a sequence of round-trip time samples:  $S_1, S_2, S_3, \dots$ .

Each sample can easily be determined by the sending transport entity, if it remembers the time when each packet is transmitted. When the associated acknowledgement is received, the difference between the current time and the time when the packet was transmitted gives a new measure of the experienced RTT;

- b. With each new sample,  $S_i$ , a new average value (i.e. the SRTT) can be computed using an algorithm known as Exponential Aging. This algorithm is based on the following formula:

$$SRTT_{new} = \alpha * SRTT_{prev} + (1-\alpha) * S_i$$

where  $\alpha$  is a constant smoothing factor between 0 and 1 that controls how rapidly the SRTT adapts to change. The recommended value of  $\alpha$  is 0.9. With this value, every time a new measurement is made, ninety percent of each new SRTT estimate is from the previous SRTT estimate and ten percent is from the new measurement;

- c. Given this smoothed estimator, a new retransmission timer value has to be computed. Of course, the retransmission timer itself must be set to a value higher than the RTT estimate. The original proposal suggested using the following formula:

$$T1_{new} = \beta * SRTT_{new}$$

where  $\beta$  is a delay variance factor with a recommended value of 2, and  $T1_{new}$  is the new value of the local retransmission timer.

#### 3.3.2.2.10.2.2

Experience proved this algorithm to be inadequate, because it incorrectly assumed that the variance in RTT values would be small and constant. Van Jacobson has developed a refinement of the above algorithm that dynamically computes the variance instead of using a fixed  $\beta$ .

#### 3.3.2.2.10.2.3     **Jacobson's Algorithm**

##### 3.3.2.2.10.2.3.1

Jacobson's proposal has been to calculate the retransmission timer value based on both the mean and variance of the RTT rather than just calculating T1 as a constant multiple of the mean RTT. This led to propose the following new equations that are applied to each RTT measurement  $S_i$ :

$$\begin{aligned} Err &= S_i - SRTT_{prev} \\ SRTT_{new} &= SRTT_{prev} + g * Err \\ D_{new} &= D_{prev} + h * (|Err| - D_{prev}) \\ T1_{new} &= SRTT_{new} + 4 * D_{new} \end{aligned}$$

where  $SRTT$  is the smoothed RTT (the estimator of the average RTT) and  $D$  is the smoothed mean deviation.  $Err$  is the difference between the measured value just obtained and the current RTT estimator. Both  $SRTT_{new}$  and  $D_{new}$  are used to calculate the next retransmission timeout ( $T1_{new}$ ). The gain  $g$  is for the average and is set to 1/8

(0.125). The gain for the deviation is  $h$  and is set to 0.25. The larger gain for the deviation makes the retransmission timer value go up faster when the RTT changes.

3.3.2.2.10.2.3.2 Jacobson specifies a way to do all these calculations using integer arithmetic, and this is the implementation typically used. That is one reason  $g$ ,  $h$  and the multiplier 4 are all powers of 2, so the operations can be done using shifts instead of multiplies and divides.

3.3.2.2.10.2.3.3 Comparing the original method with Jacobson's, we see that the calculations of the smoothed average are similar ( $\alpha$  is one minus the gain  $g$ ) but a different gain is used. Also, Jacobson's calculation of the retransmission timer value depends on both the smoothed RTT and the smoothed deviation, whereas the original method used a multiple of the smoothed RTT.

3.3.2.2.10.2.3.4 This important improvement has been shown to produce dramatically improved roundtrip time estimates and is now in widespread use. Jacobson's algorithm is especially important on a low speed link, where the natural variation of packet sizes causes a large variation in RTT.

#### 3.3.2.2.10.2.4 **Recommended Algorithm for ATN COTP Implementations**

3.3.2.2.10.2.4.1 The recommended algorithm for ATN COTP implementations is derived from Jacobson's algorithm and differs only by the addition of the remote Acknowledgement timer ( $A_r$ ) in the formula used for the computation of the retransmission timer value.

3.3.2.2.10.2.4.2 With this difference, the recommended formula becomes:

$$\begin{aligned} Err &= S_i - SRTT_{prev} \\ SRTT_{new} &= SRTT_{prev} + g * Err \\ D_{new} &= D_{prev} + h * (|Err| - D_{prev}) \\ TL_{new} &= SRTT_{new} + 4 * D_{new} + A_r \end{aligned}$$

3.3.2.2.10.2.4.3 This difference is in response to the unique requirements of the aeronautical environment which may require longer acknowledgement times.

3.3.2.2.10.2.4.4 The above formula is suitable for the computation of the local retransmission timer during the steady state phase of the transport connections. However, this formula fails to compute a correct retransmission timer value when the very first RTT sample is obtained for a new transport connection. Indeed, when this formula is used on the first RTT sample, the resulting retransmission timer value is initially too short. This introduces a high risk of unnecessary retransmission on the first DT TPDUs exchanged over the transport connection. This is the reason why a different formula is recommended for the initial dynamic computation of the local retransmission time, when the first RTT sample is obtained. The recommended formula in that particular case is as follows:

$$SRTT_{init} = S_0$$

$$\begin{aligned} D_{init} &= SRTT_{init} / 4 \\ TI_{init} &= SRTT_{init} + 4 * D_{init} + A_r \end{aligned}$$

where,  $S_0$  is the first valid RTT sample.

### 3.3.2.2.10.3      **The Back-off Mechanism**

3.3.2.2.10.3.1      The retransmission timeout is used to determine when a packet has been dropped in the network. When this timeout has expired without the arrival of an ACK, the segment is retransmitted.

3.3.2.2.10.3.2      The retransmission timeout may however also expire when there is a sudden increase of transit delay between two COTP peers. This may happen when the network becomes congested or when the PDUs are suddenly routed along a slower path. In the case where PDUs are not lost, but just delayed, retransmissions are unnecessary, and must be avoided since they may contribute to further congestion. Hence, in the face of unacknowledged PDUs, it becomes necessary to readapt quickly the value of the retransmission timer to the potential increase in transit delay. Unfortunately, this cannot be done on the basis of the round-trip time sampling and estimation, since without an acknowledgement there is no new timing information to be fed into the retransmission timer value calculation.

3.3.2.2.10.3.3      A special procedure, called back-off, must be implemented to solve this problem. The back-off mechanism works as follows: whenever a timeout occurs, the COTP increases the retransmission timer value by some factor before retransmitting the unacknowledged data. Should the new, larger retransmission time expire yet again before retransmission is acknowledged, the retransmission timer is increased still further.

3.3.2.2.10.3.4      The ATN ICS SARPs prescribe the use of a binary exponential back-off. The back-off procedure is achieved by simply doubling the retransmission timer value for each consecutive attempt.

3.3.2.2.10.3.5      The back-off mechanism is essential in keeping the network stable when sudden overloads cause datagrams to be dropped. When the overload condition disappears, datagram loss stops and the COTPs reduce their retransmission timer value to their normal “mean RTT”-based values.

### 3.3.2.2.10.4      **Karn’s Algorithm**

3.3.2.2.10.4.1      When finally, after one or several retransmissions, an acknowledgement of the packet is received, a transport entity wishes to update its RTT estimator. COTP is however faced to the following dilemma: the measurement of the last RTT is normally done by computing the interval of time between the packet transmission and the acknowledgement reception; however, when retransmissions occurred, COTP cannot know whether the received acknowledgement was issued on receipt of the first packet or on receipt of one of the subsequent retransmitted packets. This is called the *retransmission ambiguity problem*.

- 3.3.2.2.10.4.2 If COTP chooses to measure the new RTT, as the interval of time from the first transmission to the receipt of the acknowledgement, it runs the risk to inflate unnecessarily the value of the RTT estimator. Inflated RTT estimates may not be a problem if the original reason for the retransmission was network congestion, because congestion tends to increase RTT anyway. However, if the path is lossy, the RTT estimator grows and throughput unnecessarily decreases to low level.
- 3.3.2.2.10.4.3 If COTP chooses to measure the new RTT, as the interval of time from the most recent retransmission to the receipt of the acknowledgement, it runs the risk to decrease inappropriately the value of the RTT estimator. The result may be dramatic: unnecessary retransmissions occur constantly, the RTT estimator stabilises at an unreasonable low estimate, useful throughput drops sharply and network bandwidth is wasted.
- 3.3.2.2.10.4.4 Another strategy is to simply ignore round trip times for packets that have been retransmitted. In that case, the RTT estimator and the derived RTO are not updated when retransmission occurs. However this method does not work if the causes of the retransmissions is a sudden increase in network round trip time (e.g. failure of a primary path that causes datagrams to be sent over a slower secondary path). This strategy may also lead to have numerous unnecessary retransmissions, and to waste the network capacity.
- 3.3.2.2.10.4.5 Karn solved the retransmission ambiguity problem by proposing the following simple rules:
- a. When an acknowledgement arrives for a packet that has been sent more than once, ignore any round-trip measurements based on this packet (this is the third strategy described above);
  - b. In addition, do not calculate a new RTO, but re-use instead the last backed-off RTO for the next transmission. Only when a packet is acknowledged without an intervening retransmission, the RTO will be recalculated from the RTT estimator.
- 3.3.2.2.10.4.6 This algorithm is expected to be implemented by any ATN COTP implementation.

### 3.3.2.2.11 **Transport Layer Protocol Conformance**

3.3.2.2.11.1 This section provides background information and notes on the ATN Protocol Requirements List (APRL)s for the connection-mode transport protocol and the encoding of TPDU's. The requirements for the connection-mode transport protocol are defined using the APRL for the ISO/IEC 8073 protocol specified in the ATN ICS SARPs, which is derived from the PICS Proforma provided with ISO/IEC 8073. ATN specific extensions are also included in the APRL.

#### 3.3.2.2.11.2 **Base Standard**

3.3.2.2.11.2.1 The base standard which applies to the ATN Transport Layer protocol is the 1992 version of ISO/IEC 8073. During the development of the APRL, an important objective was to

ensure backwards compatibility with ISO/IEC 8073: 1988, whilst permitting the use of the following features of the 1992 version which do not exist in the 1988 version:

- a) 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”;
- b) the Selective Acknowledgement option, which was added to allow a transport entity to acknowledge a non-contiguous set of TPDU’s;
- c) the Request Acknowledgement option, which was added to allow a transport entity to request that the remote entity acknowledges received TPDU’s;
- d) the inactivity time is now specified as two values, a “local” inactivity time and a “remote” inactivity time; and
- e) the values of the inactivity times can now be passed as parameters in the connection establishment phase.

### 3.3.2.2.11.3 **Caveat to Conformance with Base Standard**

3.3.2.2.11.3.1 The ISO/IEC 8073 PICS (D.6.2) identifies C4L as ISO:C2:θM reflecting that Class 4 over connectionless networks requires the implementation of class 2 for conformance purposes. However, ISO/IEC 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/IEC 8073 are therefore not necessarily in conformance with ISO/IEC 8073.

### 3.3.2.2.11.4 **Initiator/Responder Capability for Protocol Classes 0-4**

3.3.2.2.11.4.1 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.

### 3.3.2.2.11.5 **Notes on Required and Recommended Optional Functions**

#### 3.3.2.2.11.5.1 **Extended TPDU Numbering**

3.3.2.2.11.5.1.1 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 TPDU’s 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 why 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 4.2.2.93.3.2.2.9.3). If a large number of TPDU's (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.

#### 3.3.2.2.11.5.2 **Non-use of Checksum**

3.3.2.2.11.5.2.1 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 TPDU's; 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 matter. The use or non-use of checksums is dependent on the characteristics of the TS-user-data flow.*

#### 3.3.2.2.11.5.3 **Selective Acknowledgement**

3.3.2.2.11.5.3.1 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 TPDU's, 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 TPDU's, 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.

#### 3.3.2.2.11.5.4 **Request of Acknowledgement**

3.3.2.2.11.5.4.1 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 acknowledges all currently received TPDU's. 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.



3.3.2.2.11.5.5     **Reduction of Credit Window**

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

3.3.2.2.11.5.6     **Concatenation**

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

3.3.2.2.11.6     **Notes on TPDU Support**

3.3.2.2.11.6.1     **Mandatory TPDU's**

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

3.3.2.2.11.6.2     **Error TPDU Support**

3.3.2.2.11.6.2.1     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.

3.3.2.2.11.7     **Notes on TPDU Parameter Support**

3.3.2.2.11.7.1     **Optional Parameters for the CR TPDU**

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

- a) **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;

- b) **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/IEC 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 TPDU 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 ≤ n ≤ 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;~~
- c) **Preferred Maximum TPDU Size:** Support is recommended. The ~~maximum~~ preferred ~~maximum~~ 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;~~
- d) **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;
- e) **Protection Parameter:** Support is not currently recommended as no security mechanisms have been defined for the ATN Internet besides the use of the ATN Security Label and the use of ISO/IEC 10747 type 2 authentication, which ~~is both~~ are outside of the scope of this parameter. Use of this feature may be specified in later ~~versions~~ editions of the ATN ICS SARPs, if a need for lower layer protection mechanisms had been identified;
- f) **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;

- g) **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;
- h) **Priority Parameter:** Support is mandated. 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 ICS 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; and
- i) **Acknowledgement Timer and Inactivity Time Parameters:** Support is mandated 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~~loss.

#### 3.3.2.2.11.7.2 **Optional Parameters for the CC TPDU**

3.3.2.2.11.7.2.1 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.

#### 3.3.2.2.11.7.3 **Optional Parameters for a Disconnect Request TPDU**

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

#### 3.3.2.2.11.7.4 **Mandatory Parameter for a Data TPDU**

3.3.2.2.11.7.4.1 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.

### 3.3.2.2.11.7.5 **Optional Parameters for an Acknowledgement TPDU**

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

### 3.3.2.2.11.7.6 **Use of the Subsequence Number Parameter in the Acknowledgement TPDU**

3.3.2.2.11.7.6.1 Since the reduction of credit window capability is required, support of this parameter is mandatory. Support of the flow control confirmation parameter is mandated for use in congestion avoidance mechanisms.

### 3.3.2.2.11.7.7 **Use of the Selective Acknowledgement Parameter in the AK TPDU**

3.3.2.2.11.7.7.1 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.

### 3.3.2.2.11.7.8 **Optional Parameters for an Error TPDU**

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

### 3.3.2.2.11.7.9 **User Data in Class 4 TPDUs**

3.3.2.2.11.7.9.1 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.

### 3.3.2.2.12 **Use of the Network Service**

3.3.2.2.12.1 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.

### 3.3.2.2.12.2 **Use of the N-UNITDATA Request**

3.3.2.2.12.2.1 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.

3.3.2.2.12.2.2 The N-UNITDATA parameters are used as follows:

3.3.2.2.12.2.3 **NS-user-data**

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

3.3.2.2.12.2.4 **Network Service Access Point Addresses**

3.3.2.2.12.2.4.1 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.

3.3.2.2.12.2.5 **Network Quality of Service**

3.3.2.2.12.2.5.1 **Network Layer Protection**

~~3.3.2.2.12.2.5.2~~ 3.3.2.2.11.2.5.1 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.

~~3.3.2.2.12.2.5.3~~ 3.3.2.2.11.2.5.1.2 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 connection 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 TPDU) sent by both the connection initiator and the connection responder on that transport connection.

**3.3.2.2.12.2.5.4 Network Layer Transit Delay, Cost, and Residual Error Probability**

~~3.3.2.2.12.2.5.5~~ 3.3.2.2.11.2.5.2.1 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.

#### 3.3.2.2.12.2.6 **Network Layer Priority**

3.3.2.2.12.2.6.1 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 Section Subvolume 1 of the ATN SARPs.

3.3.2.2.12.2.6.2 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.

#### 3.3.2.2.12.3 **Use of the N-UNITDATA Indication**

3.3.2.2.12.3.1 The transport layer receives all TPDU's via the N-UNITDATA indication. TPDU's 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 TPDU's within a single NSDU. The parameters of an incoming N-UNITDATA indication are interpreted as follows.

#### 3.3.2.2.12.3.2 **NS-user-data**

3.3.2.2.12.3.2.1 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 TPDU's following the first one.

#### 3.3.2.2.12.3.3 **Network Service Access Point Addresses**

3.3.2.2.12.3.3.1 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.

#### 3.3.2.2.12.3.4 **Network Quality of Service**

3.3.2.2.12.3.4.1 The connection-mode transport layer does not need to interpret most of the indicated network layer QoS parameters associated with an N-UNITDATA indication, except 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.

3.3.2.2.12.3.4.2 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 NPDU 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 6 Section 3.5 of Part IV of this document](#).

3.3.2.2.12.3.4.3 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.

### 3.3.2.3 **The Connectionless-Mode Transport Layer**

3.3.2.3.1 The ATN CLTS is based on the ISO/IEC 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.

#### 3.3.2.3.2 **Overview of the Connectionless-Mode Transport Layer**

3.3.2.3.2.1 The defining characteristic of CLTS transmission is the independent nature of each invocation of the `sService`. 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 network service.

3.3.2.3.2.2 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.

3.3.2.3.2.3 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:

- a) network service selection;
- b) mapping of transport address onto Network address;
- c) TSDU delimiting (determine the beginning and end of a TSDU); and;
- d) end-to-end error detection (implying the use of a specific mechanism) and the necessary monitoring of the QoS.

3.3.2.3.2.4 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.

#### 3.3.2.3.2.5 **Service Characteristics**

3.3.2.3.2.5.1 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.

#### 3.3.2.3.2.6 **Data Transfer**

3.3.2.3.2.6.1 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.

3.3.2.3.2.6.2 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. TPDU's failing verification are discarded.

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

3.3.2.3.2.6.4 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).



### 3.3.2.3.2.7 **ATN Connectionless-Mode Transport Service Model**

3.3.2.3.2.7.1 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:

- a) discard objects;
- b) duplicate objects; and
- c) change any order of independent service requests into a different order of service indications.

3.3.2.3.2.7.2 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~~<sup>TS-user</sup>'s 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~~<sup>TS-user</sup>'s a priori knowledge of the ATN environment.

### 3.3.2.3.3 **ATN Connectionless-Mode Transport Layer Quality of Service**

#### 3.3.2.3.3.1 **Use of Transport Layer QoS**

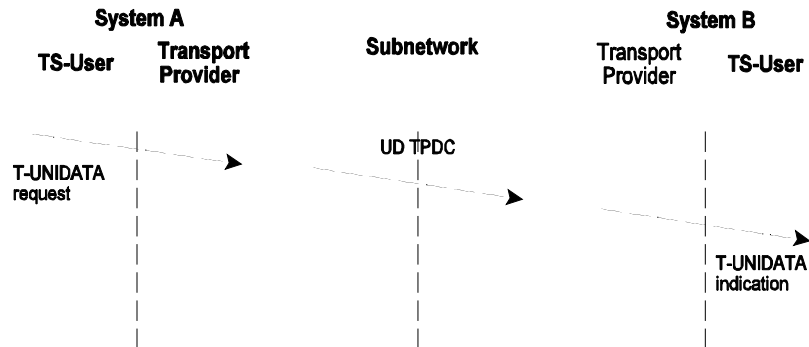
3.3.2.3.3.1.1 The use of transport layer QoS parameters for the CLTS is similar to that of the connection-mode service. However, unlike the COTS, there is no concept of negotiation of requested transport layer QoS parameters. Each invocation of the T-UNITDATA service involves a set of requested transport 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.

3.3.2.3.3.1.2 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.

#### 3.3.2.3.3.2 **Connectionless-Mode Transport Layer QoS Parameters**

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

3.3.2.3.3.2.2 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



**Figure 3.3-7. Sequence of Primitives and TPDU Exchange for Connectionless Data Transfer**

transport checksum. Both a priority and an ATN Security Label may be specified on a per TSDU basis.

### 3.3.2.3.3.3 Priority

3.3.2.3.3.3.1 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.

3.3.2.3.3.3.2 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. Subvolume Section 1 of the ATN SARPs specifies the mapping of transport layer priority values to network layer priority values.

### 3.3.2.3.3.4 ATN Security Label

3.3.2.3.3.4.1 The security parameter is used by the service TS-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 ICS SARPs.

### 3.3.2.3.3.5 Connectionless Mode Transport Layer Service Primitives

3.3.2.3.3.5.1 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 3.3-7.

### 3.3.2.3.3.6 T-UNITDATA Request

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

- a) **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;
- b) **Quality of service:** The value of the QoS is a list of subparameters. The subparameters composing the CLTS QoS are presented in 4.2.3.2.13.3.2.3.2. The TS-provider does not guarantee that it can offer the requested QoS;
- c) **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; and
- d) **Security:** The security parameter is used by the service TS-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 ICS SARPs.

3.3.2.3.3.6.2 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.

### 3.3.2.3.3.7 T-UNITDATA Indication

3.3.2.3.3.7.1 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.

3.3.2.3.3.7.2 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.

### 3.3.2.3.4      **Use of the Network Service**

*Note.— Refer to 3.3.2.2.11.2.5 for more background on selection of requested network layer QoS parameters.*

#### 3.3.2.3.4.1      **Use of the N-UNITDATA Request**

3.3.2.3.4.1.1      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:

#### 3.3.2.3.4.1.2      **NS-user-data**

3.3.2.3.4.1.2.1      The transport layer sends the UD TPDU as an NSDU.

#### 3.3.2.3.4.1.3      **Network Service Access Point Addresses**

3.3.2.3.4.1.3.1      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~~selectors~~ 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.

#### 3.3.2.3.4.1.4      **Network Quality of Service**

3.3.2.3.4.1.4.1      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.

3.3.2.3.4.1.4.2      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 ~~4.2.2.11.1.3.3.3.2.2.11.2.5.2.~~

3.3.2.3.4.1.4.3      The value of the ATN Security Label specified by the ~~service~~ TS-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.

#### 3.3.2.3.4.1.5      **Network Layer Priority**

3.3.2.3.4.1.5.1      There is no explicit priority parameter in a UD TPDU. To meet the ISO/IEC 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.

3.3.2.3.4.1.5.2 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.

#### 3.3.2.3.4.1.6 Use of the N-UNITDATA Indication

3.3.2.3.4.1.6.1 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:

##### 3.3.2.3.4.1.6.2 NS-user-data

3.3.2.3.4.1.6.3 — 3.3.2.4.1.6.2.1 The transport layer assumes that the UD TPDU begins at the first octet of the NS-user-data.

##### 3.3.2.3.4.1.6.4 Network Service Access Point Addresses

3.3.2.3.4.1.6.5 — 3.3.2.4.1.6.3.1 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.

##### 3.3.2.3.4.1.6.6 Network Quality of Service

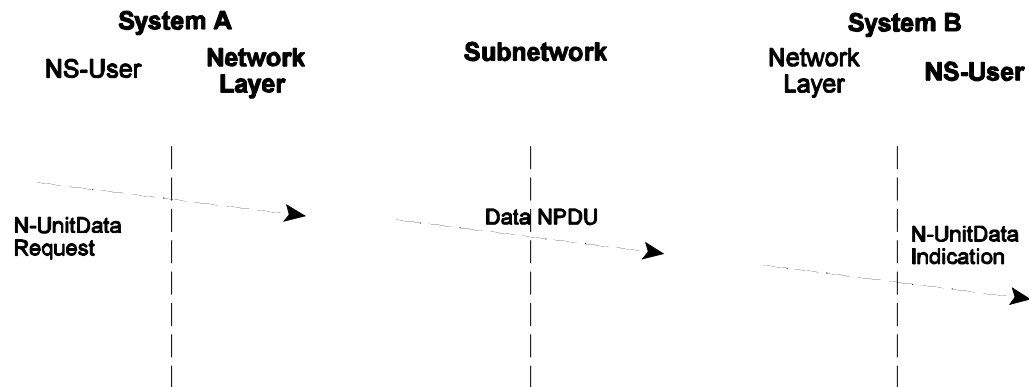
3.3.2.3.4.1.6.7 — 3.3.2.4.1.6.4.1 The connectionless mode transport layer does not need to interpret most of the indicated network QoS parameters associated with an N-UNITDATA indication, except for the network protection parameter. The relative rankings of transit delay, residual error probability, and cost are after the fact, and do not have any effect on the transport protocol.

3.3.2.3.4.1.6.8 — 3.3.2.4.1.6.4.2 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 NPDU 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.

3.3.2.3.4.1.6.9 — 3.3.2.4.1.6.4.3 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.

3.3.2.3.4.1.6.10 **Priority**

3.3.2.3.4.1.6.11 — 3.3.2.4.1.6.5.1 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 Subvolume Section 1 of the ATN SARPs for the mapping between NL priority and TL priority.

3.3.3 **CLNP Implementation Considerations**

**Figure 3.3-8. Connectionless-Mode Network Service**

3.3.3.1 ISO/IEC 8473 describes a protocol for providing the connectionless mode network service. The ISO/IEC 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.

3.3.3.2 **The Connectionless-Mode Network Service**

3.3.3.2.1 The OSI connectionless-mode network service is the service provided to a network service user when the ISO/IEC 8473 connectionless network protocol is used as a SNICP. The operation of the connectionless network service is illustrated in Figure 3.3-8. It consists of a single end-to-end primitive, i.e. the N-UNITDATA service.

3.3.3.2.2 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, i In the ATN, the service parameters also includes the Network Priority of the data and the associated ATN Security Label.

3.3.3.2.3 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.

### 3.3.3.3 **The Connectionless Network Protocol**

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

#### 3.3.3.3.2 **Satisfying the N-UNITDATA Request**

3.3.3.3.2.1 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 the ATN Security Label.

3.3.3.3.2.2 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.

3.3.3.3.2.3 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/IEC 9542 ES-IS Routing Information Exchange Protocol (see 3.3.4.24.4.1).

3.3.3.3.2.4 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.

3.3.3.3.2.5 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/IEC 8473. Two special SNDCFs have, however, been specified for ATN Mobile Subnetworks (see 3.3.4.6.34.4.4.1).

### 3.3.3.3.3 **NPDU Forwarding**

3.3.3.3.3.1 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~~these were separate Data PDUs and may even be further fragmented.

3.3.3.3.2 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.

3.3.3.3.3 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/IEC 9542 ES-IS (see [3.3.4.24.4.1](#)), the ISO/IEC 10589 IS-IS Intra-Domain Routing Information Exchange Protocol (see [3.3.4.4.54.4.3.2](#)), and the ISO/IEC 10747 Inter-Domain Routing Protocol (see [3.4.44.4.3.2.3](#)).

#### 3.3.3.3.4 **At the Destination End System**

3.3.3.3.4.1 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.

3.3.3.3.4.2 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.

#### 3.3.3.4 **Addressing Consideration**

3.3.3.4.1 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/IEC 8348. The ATN Addressing Plan specified in the ATN ICS 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/IEC 8473 protocol to convey NSAP Addresses is the preferred binary encoding specified in ISO/IEC 8348.

#### 3.3.3.4.2 **Network Entity Titles**

3.3.3.4.2.1 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).



3.3.3.4.2.2 A NET identifies a Network Entity in an End-s System or Intermediate-s 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.

3.3.3.4.2.3 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.

### 3.3.3.5 ***Other NS-User Services***

3.3.3.5.1 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 may be given.

3.3.3.5.2 The service characteristics information that may be made available includes:

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

3.3.3.5.3 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 a standard mechanism available to support the dynamic distribution of such Quality of Service Information.

### 3.3.3.6 ***Error Reports***

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

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

3.3.3.6.3 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 done as an extension to the service interface or through a local management mechanism.

### 3.3.3.7 ***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 ICS 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.

### 3.3.3.8 *Priority*

3.3.3.8.1 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.

3.3.3.8.2 In the ATN ~~Network~~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.

3.3.3.8.3 ATN ~~Network~~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.

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

### 3.3.3.9 *ATN Security*

3.3.3.9.1 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 ICS SARPs.

3.3.3.9.2 The ATN Security Label conveys the following information:

- a) **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. The following traffic types are defined in the ATN-:

- 1) ATN Operational Communications:

- i) ATSC;
  - ii) AOC;
  - 2) ATN Administrative Communications;
  - 3) General Communications; and
  - 4) ATN Systems Management Communications.
- b) An **ATSC Class**: this is valid for ATSC Traffic Types only 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 t~~The ATN Internet aims to send data with a declared ATSC Class over a subnetwork supporting that ATSC Class or higher. If no such subnetwork is available, then the next lowest available class subnetwork is chosen.
- c) An **Air/Ground Subnetwork Preference**: this is valid for ~~AOC~~<sup>AINSC</sup> Traffic Types only and identifies the Air/Ground subnetworks over which the data may be forwarded and the relative preference of such subnetworks.

3.3.3.9.3 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.

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

3.3.3.9.5 In order to conserve limited bandwidth on air/ground subnetworks, ATN Air/Ground Routers may omit the Security Information field from routes advertised to ATN Airborne Routers which have received relevant information in the uplinked ISH PDU during route initiation.

3.3.3.9.6 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.

### 3.3.3.10 **ISO/IEC 8473 Mandatory Internetwork Protocol Functions**

3.3.3.10.1 This section describes the functions which are performed as part of the ATN Internetwork Protocol within all Network entities conforming to ISO/IEC 8473. These are listed in Table 3.3-54, which also classifies these functions according to their conformance requirement and by protocol subset. ATN Intermediate Systems have to be able to support both the full and the non-segmenting subset. ATN End Systems are required to support the segmenting subset only. 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.

#### 3.3.3.10.2 **PDU Composition Function**

3.3.3.10.2.1 The PDU COMPOSITION function is responsible for the construction of a Network Protocol Data Unit (NPDU) according to the rules governing the encoding of NPDUs. Protocol Control Information (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.

3.3.3.10.2.2 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.

3.3.3.10.2.3 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 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.

**Table 3.3-4. ISO/IEC 8473 Protocol Functions**

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

3.3.3.10.2.4 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:

- a) the total length of the NPDU in octets is determined by the originator and placed in the Segment Length field of the NPDU header;
- b) the segmentation field is not changed for the lifetime of the NPDU; and
- c) no Data Unit Identification is provided.

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

### 3.3.3.10.3 **PDU Decomposition Function**

3.3.3.10.3.1 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:

- a) 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;
- b) 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; and
- c) information relating to the ~~QoS~~**QoS** provided during the transmission of the NPDU is determined from the ~~QoS~~**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. |

### 3.3.3.10.4 **Header Format Analysis Function**

3.3.3.10.4.1 The HEADER FORMAT ANALYSIS function determines whether the full protocol described in this section is employed, or one of the defined subsets thereof. If the ~~NPDU~~**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. |

3.3.3.10.4.2 If the ~~NPDU~~ 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.

### 3.3.3.10.5 **PDU Lifetime Control Function**

3.3.3.10.5.1 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.

3.3.3.10.5.2 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 NPDU<sub>s</sub>.

3.3.3.10.5.3 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:

- a) the transit delay in the underlying service from which the NPDU was received; and
- b) 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.

3.3.3.10.5.4 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 NPDU<sub>(s)</sub>.

### 3.3.3.10.6 **Route PDU Function**

3.3.3.10.6.1 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, IS-IS and IDRP routing information exchange protocols.

3.3.3.10.6.2 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 (including the Security Label) of the **N-UNITDATA** Request, which specifies the **QoSQoS** requested by the sending ATN NS user. The ROUTE PDU function determines whether this **QoSQoS** 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 **QoSQoS** 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.

### 3.3.3.10.7 **Forward PDU Function**

3.3.3.10.7.1 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 **QoSQoS** 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.

### 3.3.3.10.8 **Segmentation Function**

3.3.3.10.8.1 For an ATN Network Entity implementing the full protocol, segmentation is performed when the size of the NPDU 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.

3.3.3.10.8.2 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.

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

- a) the source address,
- b) the destination address, and
- c) the data unit identifier.

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

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

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

3.3.3.10.8.6 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.

### 3.3.3.10.9 **Reassembly Function**

3.3.3.10.9.1 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).

3.3.3.10.9.2 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 PDU 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.

3.3.3.10.9.3 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:

- a) interact with the routing algorithm to favour paths on which fewer segments are generated;
- b) generate more segments than absolutely required in order to avoid additional segmentation at some subsequent point; or
- c) 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.

3.3.3.10.9.4 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.

#### 3.3.3.10.10 **Discard PDU Function**

3.3.3.10.10.1 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:

- a) a violation of protocol procedure has occurred;
- b) an NPDU is received whose checksum is inconsistent with its contents;
- c) an NPDU is received, but due to local congestion, it cannot be processed;
- d) an NPDU is received with a correct header checksum, but whose header contents are invalid;

- e) 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;
- f) an NPDU is received whose destination address is unreachable or unknown;
- g) 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;
- h) an NPDU is received whose NPDU lifetime has expired or a segmented NPDU is received whose lifetime expires during reassembly; and
- i) an NPDU is received which contains an unsupported Type 2 option.

### 3.3.3.10.11 **Error Reporting Function**

3.3.3.10.11.1 The ERRORREPORTING 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.

3.3.3.10.11.2 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<sub>user</sub>.

3.3.3.10.11.3 The ERROR REPORTING function performs as follows:

- a) an ER NPDU is not generated to report the discard of an ER NPDU;
- b) 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;
- c) 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; and
- d) if a DT NPDU is discarded for one of the reasons in listed in section 3.3.3.10.10.1 ~~Paragraph 4.3.9.9~~, and the ER flag has been set to allow Error Reports, an ER NPDU is generated.

3.3.3.10.11.4 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).

#### 3.3.3.10.11.5 **Initiation of Error Reports**

3.3.3.10.11.5.1 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.

3.3.3.10.11.5.2 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.

3.3.3.10.11.5.3 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.

#### 3.3.3.10.11.6 **Processing of Received Error Reports**

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

- a) 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;
- b) the reason for generating the Error Report is extracted from the Options Part of the NPDU header; and
- c) 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.

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

#### 3.3.3.10.11.7 **Relationship of Data NPDU Options to Error Report NPDUs**

3.3.3.10.11.7.1 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.

3.3.3.10.11.7.2 The processing of an Error Report is controlled by options which are present in the corresponding DT NPDU. In particular, op

tions selected for the original DT NPDU affect which options are included in the corresponding ER NPDU.

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

- a) if the Priority Option or the ~~QoS~~ 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;
- b) 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; and
- c) the Record Route Option, if selected in the DT NPDU, is specified in the ER NPDU.

3.3.3.10.11.7.4 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.

### 3.3.3.10.12 PDU Header Error Detection

3.3.3.10.12.1 The PDU HEADER ERROR DETECTION function protects against failure of ATN IS or ES network entities due to the processing of erroneous information in the NPDU header. The PDUHEADERERRORDETECTION 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.

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

$$\sum a_i (1 \leq i \leq L) \pmod{255} = 0$$

$$\sum (L - i + 1)a_i (1 \leq i \leq L) \pmod{255} = 0$$

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.

3.3.3.10.12.3 An efficient algorithm for calculating and checking the checksum octets is provided in Annex D of ISO/IEC 8073 and ISO/IEC 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.

### 3.3.3.11 ISO/IEC 8473 Optional Internetwork Protocol Functions

3.3.3.11.1 ISO/IEC 8473 internetwork protocol options are selected by the ATN ES Network entity which originates ISO/IEC 8473 NPDUs. As a part of the ISO/IEC 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/IEC 8473 CLNP header is delivered unchanged to each successive ATN network entity along the end-to-end path between source and destination ES.

#### 3.3.3.11.2 Padding Function

3.3.3.11.2.1 The PADDING Function allows extending the length of ISO/IEC 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.

#### 3.3.3.11.3 Security Function

3.3.3.11.3.1 The SECURITY function supports imposition of Network Layer security provisions by way of an options field conveyed within the ISO/IEC 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.

3.3.3.11.3.2 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 ATN ICS SARPs does not mandate support of the source or destination specific formats for the ISO/IEC 8473 security parameter, and hence to avoid service irregularities, neither format should be used.

3.3.3.11.3.3 The Security options field is included in the ISO/IEC 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, which is also the format used for ~~and this is~~ the encoding of the ATN Security Label provided on the N-UNITDATA-Request. As discussed in ~~4.3.83.3.3.9.2~~, 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.

- 3.3.3.11.3.4 When the NPDU reaches its destination, the value of the Security options field is provided to the destination NS\_user as the Security Label associated with the NSDU.
- 3.3.3.11.4 **Source Routing Function**
- 3.3.3.11.4.1 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.
- 3.3.3.11.5 **Record Route Function**
- 3.3.3.11.5.1 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.
- 3.3.3.11.5.2 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:
- a) 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;
  - b) 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; and
  - c) 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.
- 3.3.3.11.5.3 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 2b) 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.

3.3.3.11.5.4 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.

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

### 3.3.3.11.6 **Quality of Service Maintenance Function**

3.3.3.11.6.1 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.

3.3.3.11.6.2 The ISO/IEC 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.

3.3.3.11.6.3 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.



### 3.3.3.11.7 **Priority Function**

3.3.3.11.7.1 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.

3.3.3.11.7.2 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.

3.3.3.11.7.3 ATN use of the Priority Function is discussed in [3.3.3.84.3.7](#).

### 3.3.3.11.8 **Congestion Notification Function**

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

3.3.3.11.8.2 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).

3.3.3.11.8.3 ATN Intermediate-~~s~~ 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-~~s~~ System will start to discard NPDUs; at this time, the ATN Intermediate-~~s~~ 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.

3.3.3.11.8.4 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-~~s~~ System which processes the NPDU to indicate that that ATN Intermediate-~~s~~ System is experiencing congestion. The method of initiating Congestion Notification is discussed in ~~Chapter 6~~ [Chapter 3.5 of Part IV of this document](#).

### 3.3.3.11.9 **Echo Request and Response**

3.3.3.11.9.1 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.

3.3.3.11.9.2 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.

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

### 3.3.3.11.10 **Notes on the CLNP APRLs**

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

#### 3.3.3.11.10.2 **Security**

3.3.3.11.10.2.1 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 if connectivity is to be maintained. However, within a Routing Domain, it is acceptable for the actual value of this parameter to be ignored.

#### 3.3.3.11.10.3 **Complete Route Recording**

3.3.3.11.10.3.1 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.

#### 3.3.3.11.10.4 **Source Routing**

3.3.3.11.10.4.1 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.

#### 3.3.3.11.10.5 **Priority**

3.3.3.11.10.5.1 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.

#### 3.3.3.11.10.6 **Padding**

3.3.3.11.10.6.1 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.

### 3.3.4 **The Implementation of the Routing Information Exchange Protocols**

#### 3.3.4.1 **General**

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

3.3.4.1.2 ISO/IEC 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/IEC 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.

3.3.4.1.3 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 communications between RDs. This is specified in ISO/IEC 10747, and is known as the Inter-Domain Routing Protocol (IDRP).

3.3.4.1.4 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.

### 3.3.4.2 ***ES-IS Implementation Considerations***

#### 3.3.4.2.1 **Overview**

3.3.4.2.1.1 ISO/IEC 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.

3.3.4.2.1.2 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.

3.3.4.2.1.3 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.

3.3.4.2.1.4 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.

3.3.4.2.1.5 The protocol can also support routing in the absence of an ~~End~~ Intermediate 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 communications can proceed.

3.3.4.2.1.6 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.

3.3.4.2.1.7 In the ATN, the Configuration Information function is also used with mobile networks. The majority of 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

communications take place. An exchange of ISO/IEC 9542 Configuration Information is required as part of this procedure.

#### 3.3.4.2.2 **ATN Use of ISO/IEC 9542**

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

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

3.3.4.2.3.1 ISO/IEC 9542 is also required when ISO/IEC 10589 is implemented.

#### 3.3.4.3 **The ES-IS Protocol**

##### 3.3.4.3.1 **General**

##### 3.3.4.3.1.1 **PDU Formats and Use**

3.3.4.3.1.1.1 ISO/IEC 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).

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

3.3.4.3.1.1.3 For each type of ISO/IEC 9542 PDU, Table 3.3-56, Table 3.3-67, Table 3.3-78 and Table 3.3-89 respectively indicate:

- a) the main contents of the PDU;
- b) the type of systems which generates this PDU;
- c) the event which triggers its generation, and the destination systems of this PDU; and
- d) its functional role.

**Table 3.3-5. ISO/IEC 9542 PDU Types**

ISO/IEC 9542 PDUs	Main Contents
ESH	<p><b>Source address parameter:</b></p> <p>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 unit size, but in the end, the ES must have reported information about all its NSAPs, via one or several ESHs)</p>
ISH	<p><b>Source address parameter:</b></p> <p>NET of the IS sending the ISH PDU (the protocol allows only one NET in each ISH)</p>
RD	<p><b>Source address parameter:</b></p> <p>NET of the IS sending the RD PDU (only one NET);</p> <p><b>Destination address parameter:</b></p> <p>Destination NSAP address of the PDUs affected by the redirection (and possibly a mask selecting a “class” of NSAPs);</p> <p>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 destinations).</p>

**Table 3.3-6. Generation of ISO/IEC 9542 PDUs**

ISO/IEC 9542 PDUs	Generation of PDUs
ESH	By each ES: on timer expiry or 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 3.3-7. Propagation of ISO/IEC 9542 PDUs**

ISO/IEC 9542 PDUs	Propagation of PDUs
ESH	<ul style="list-style-type: none"> <li>• Transmitted on each SNPA the ES is attached to (the transmitted PDUs may be different but they must provide the same information)</li> <li>• Transmitted from an ES in response to a query configuration</li> <li>• Transmitted to all the ISs on the subnetwork</li> </ul>
ISH	<ul style="list-style-type: none"> <li>• Transmitted on each SNPA the IS is attached to</li> <li>• to all the ESs on each subnetwork the IS is attached to</li> </ul>
RD	<ul style="list-style-type: none"> <li>• Transmitted by any IS</li> <li>• Transmitted to the ES originating the PDU when the IS knows a better path</li> </ul>

**Table 3.3-8. Role of ISO/IEC 9542 PDUs**

ISO/IEC 9542 PDUs	Functional Role
ESH	<p><b>CONFIGURATION</b></p> <ul style="list-style-type: none"> <li>• 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</li> <li>• Allows the ESs to discover the existence and reachability of another ES on the same subnetwork, along with the NSAPs this ES supports</li> </ul>
ISH	<p><b>CONFIGURATION</b></p> <ul style="list-style-type: none"> <li>• Allows all the ISs to discover the existence and reachability (SNPA) of an IS on the same subnetwork, along with the NET of that ES (when ISO/IEC 9542 is used between ISs)</li> <li>• Allows the ESs to discover the existence and reachability (SNPA) of an IS on the same subnetwork, along with the NET of this IS</li> </ul>

ISO/IEC 9542 PDUs	Functional Role
RD	<p data-bbox="553 302 873 331"><b>ROUTE REDIRECTION</b></p> <ul data-bbox="553 373 1338 579" style="list-style-type: none"> <li data-bbox="553 373 1338 506">• 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)</li> <li data-bbox="553 548 1338 579">• It may also relate to a “class” of NSAPs (using Address Masks)</li> </ul>

3.3.4.3.1.1.4 The basic transmission mechanism for ISO/IEC 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.

3.3.4.3.1.1.5 Two broadcast subnetwork destination addresses are possible:

a) “All ESs network entities”; or

b) “All ISs network entities”.

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

### 3.3.4.3.1.2 Main protocol functions

3.3.4.3.1.2.1 ISO/IEC 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.

#### 3.3.4.3.1.2.2 Report Configuration Function

3.3.4.3.1.2.2.1 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.

#### 3.3.4.3.1.2.3 Record Configuration Function

3.3.4.3.1.2.3.1 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.



3.3.4.3.1.2.4     **Flush Old Configuration Function**

3.3.4.3.1.2.4.1    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 (e.g. SNPA re-initialisation).

3.3.4.3.1.2.5     **Query Configuration Function**

3.3.4.3.1.2.5.1    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.

3.3.4.3.1.2.5.2    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”.

3.3.4.3.1.2.5.3    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.

3.3.4.3.1.2.6     **Configuration Response Function**

3.3.4.3.1.2.6.1    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.

3.3.4.3.1.2.7     **Configuration Notification Function**

3.3.4.3.1.2.7.1    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.

3.3.4.3.1.2.8     **Request Redirect Function**

3.3.4.3.1.2.8.1    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.

3.3.4.3.1.2.8.2    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.

**3.3.4.3.1.2.9 Record Redirect Function**

3.3.4.3.1.2.9.1 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.

**3.3.4.3.1.2.10 Refresh Redirect Function**

3.3.4.3.1.2.10.1 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.

**3.3.4.3.1.2.11 Flush Old Redirect Function**

3.3.4.3.1.2.11.1 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 (e.g. SNPA re-initialisation).

**3.3.4.3.1.2.12 PDU Header Error Detection**

3.3.4.3.1.2.12.1 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.

**3.3.4.3.1.2.13 Protocol Error Processing Function**

3.3.4.3.1.2.13.1 An ISO/IEC 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/IEC 9542 protocol.

**3.3.4.3.1.3 ISO/IEC 9542 Operation Among ESs**

3.3.4.3.1.3.1 When ISO/IEC 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.

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

3.3.4.3.1.3.3 The operation of ISO/IEC 9542 among the ESs generates no additional information transmission (compared with the “standard operation”).

#### 3.3.4.3.1.4 **ISO/IEC 9542 Operation as an Initiation Phase for the Routing Protocols**

3.3.4.3.1.4.1 In the same way, when ISO/IEC 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”).

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

#### 3.3.4.3.2 **ISO/IEC 9542 Operation over Fixed Ground Subnetworks**

##### 3.3.4.3.2.1 **Overview**

3.3.4.3.2.1.1 The use of ISO/IEC 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/IEC 9542.

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

- a) each system must be able to discover the existence of neighbour systems attached to the same subnetwork;
- b) 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; and
- c) each IS must be able to dynamically monitor connectivity changes over the local subnetwork.

##### 3.3.4.3.2.2 **General Topology Subnetworks**

3.3.4.3.2.2.1 In the case ISO/IEC 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, in terms of bandwidth, it can be very costly to simulate broadcast over non-broadcast subnetworks. However, in some cases (high-bandwidth subnetworks), this solution can be chosen.

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

##### 3.3.4.3.2.3 **Broadcast Subnetworks**

3.3.4.3.2.3.1 As far as broadcast subnetworks are concerned, the full use of ISO/IEC 9542 is recommended, since this protocol was designed for operation over this kind of subnetwork.

The use of ISO/IEC 9542 over broadcast subnetworks is not too costly and allows to dynamically ascertain local configuration changes.

#### 3.3.4.3.2.4 **Point-to-Point Subnetworks**

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

#### 3.3.4.3.3 **ISO/IEC 9542 Operation over Air-ground Mobile Subnetworks**

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

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

3.3.4.3.3.3 This action is performed via the exchange of an ISO/IEC 9542 ISH PDU, and is discussed in more detail in section [3.4.105-10](#), which deals with the Route Initiation procedure.

#### 3.3.4.3.4 **Notes on the ISO/IEC 9542 APRLs**

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

#### 3.3.4.3.4.2 **Route Redirection Information**

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

#### 3.3.4.3.4.3 **Configuration Notification**

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

#### 3.3.4.4 ***Intra-Domain Routing Implementation Considerations***

3.3.4.4.1 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.

3.3.4.4.2 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.

3.3.4.4.3 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.

#### 3.3.4.4.4 **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; and
- 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; and
  - 4) ATN Intra-Domain Routing must support error control and diagnosis.

#### 3.3.4.4.4.1 **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; and
  - 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 Exchange 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/IEC 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; and
- l) Intra-Domain Routing must be integrated within the ATN Routing Framework described in Chapter 3.4 of Part IV of this document ~~Chapter 2~~. 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/IEC 9542 ES\_to-IS Protocol), in order to meet the ATN Intra-Domain Routing Goals defined in 3.3.4.4.4.1.

#### 3.3.4.4.4.2 **Intra-Domain Requirements relevant to Inter-Domain Routing**

- a) 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; and
- b) 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.

#### 3.3.4.4.5 **Overview of ISO/IEC 10589**

3.3.4.4.5.1 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.

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

3.3.4.4.5.3 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.

3.3.4.4.5.4 Under ISO/IEC 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.

- 3.3.4.4.5.5 In general terms, an LSP identifies the generating IS's neighbour ISs (i.e. those with which it has active communications links), the End Systems (ESs) to which the IS also has links; (discovered by ISO/IEC 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".
- 3.3.4.4.5.6 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.
- 3.3.4.4.5.6.1 **Routing Areas**
- 3.3.4.4.5.6.1.1 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 addresses. When two ISs discover each other, they will determine whether or not they are in the same Routing Area.
- 3.3.4.4.5.6.1.2 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:
- a) the address prefixes of their local End Systems (and of the IS itself);
  - b) the identity of adjacent ISs (i.e. those ISs in the local routing area with which the IS is in communications and can exchange ISO/IEC 8473 PDUs) and the associated quality of service parameters; and
  - c) the identity of adjacent End Systems (i.e. those ESs in the local routing area with which the IS is in communications and can exchange NPDU) and the associated quality of service parameters.
- 3.3.4.4.5.6.1.3 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.
- 3.3.4.4.5.6.1.4 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.
- 3.3.4.4.5.6.1.5 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:
- a) the set of area addresses that characterise the local routing area;



- b) the identity of adjacent level 2 ISs (i.e. the level 2 ISs in the routing domain with which the IS is in communications and can exchange NPDUs) and the associated quality of service parameters; and
- c) 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.

3.3.4.4.5.6.1.6 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.

3.3.4.4.5.6.1.7 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.

#### 3.3.4.4.5.6.2 **Partition Repair**

3.3.4.4.5.6.2.1 Level 1 routing areas may become disjoint, either due to failures or mis-configuration, and ISO/IEC 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 communications when a path exists and the only barrier to communications is a purely artificial constraint.

3.3.4.4.5.6.2.2 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.

3.3.4.4.5.6.2.3 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.

#### 3.3.4.4.5.6.3 **Support for Inter-Domain Routing**

3.3.4.4.5.6.3.1 ISO/IEC 10589 also recognises that some ISs may also be Boundary ISs, that is they are at the periphery of a 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.

### 3.3.4.5 ***IDRP Implementation Considerations***

#### 3.3.4.5.1 **Overview**

3.3.4.5.1.1 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 communications between different Administrative Domains is necessary, then a different approach is required.

3.3.4.5.1.1.1 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.

3.3.4.5.1.1.2 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 it supports.

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

3.3.4.5.1.1.4 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 [section 3.4](#) ~~Chapter 5~~.

#### 3.3.4.5.1.2 **Routing Policy**

3.3.4.5.1.2.1 Within an OSI RD, in general routing decisions are made on the basis of performance, taking into account the ~~QoS~~ QoS available over a given subnetwork connection and the ~~QoS~~ 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:

- a) 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; and
- b) the internal NSAP Addresses for which routes are advertised to adjacent RDs, and the scope of any further distribution.

3.3.4.5.1.2.2 Routing Policy is necessary because even when connectivity exists, 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~~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.

3.3.4.5.1.2.3 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.

### 3.3.4.5.1.3 **BIS-BIS Communications**

3.3.4.5.1.3.1 BISs exchange routing information in a pair-wise fashion. They use the services of ISO/IEC 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/IEC 8473 Data PDU. The BIS-BIS protocol is thus connection mode in operation and has similar features to the ISO/IEC 8073 Class 4 transport protocol.

3.3.4.5.1.3.2 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.

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

- 3.3.4.5.1.3.4 The BIS-BIS protocol uses the UPDATE PDU to copy routes from the Aadj-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 Aadj-RIB-out and Aadj-RIB-in.
- 3.3.4.5.1.3.5 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.
- 3.3.4.5.1.3.6 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.
- 3.3.4.5.1.3.7 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.
- 3.3.4.6 ***SNDCF Implementation Considerations***
- 3.3.4.6.1 The ATN specification is predicated on the use of the Connectionless Network Protocol (CLNP) specified in ISO/IEC 8473. CLNP provides the unifying end to end internetwork protocol. However, it is necessary to provide an adaptation 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 concern 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.
- 3.3.4.6.2 **SNDCFs for Fixed Data Networks**
- 3.3.4.6.2.1 ISO/IEC 8473 provides several standard SNDCFs for use with common subnetwork types including IEEE 802 compatible LANs and ISO/IEC 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.
- 3.3.4.6.3 **~~The~~ Mobile SNDCFs**
- 3.3.4.6.3.1 The mobile networks are a key component of the ATN. Some 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. 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.

3.3.4.6.3.2 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; and
- 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 data stream mode compression using the Deflate algorithm. This is a two-stage compression mechanism which first removes redundancy in the data stream by replacing re-occurring strings by backward references to previous occurrences of such strings, and in the second stage replaces these backward references as well as strings for which no backward reference exists with shorter symbols which are elements of a Huffman Code set.~~

3.3.4.6.3.3 ~~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 occurring strings with shorter symbols.~~

### 3.3.4.6.3.3 **Overview of LREF Compression Algorithm**

3.3.4.6.3.3.1 LREF compression is specified for use over a reliable virtual circuit. It is a directory 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 circuits. The compression algorithm is described as follows.

3.3.4.6.3.3.2 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;
- b) the inward NSAP address is identical to the packet's source address;
- c) the protocol version number is the same as that contained 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.

- 3.3.4.6.3.3.3 If the above conditions ~~are~~<sup>is</sup> 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 header may be replaced by a compressed header.
- 3.3.4.6.3.3.4 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.
- 3.3.4.6.3.3.5 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.
- 3.3.4.6.3.3.6 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.
- 3.3.4.6.3.3.7 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 SNDCEF by the local management protocol. An SNDCEF Error PDU is specified for this purpose.
- 3.3.4.6.3.3.8 **Creating Local Directory Entries**
- 3.3.4.6.3.3.8.1 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.
- 3.3.4.6.3.3.8.2 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.
- 3.3.4.6.3.3.8.3 Whenever an uncompressed CLNP packet is received over a virtual circuit supporting the Mobile SNDCEF, 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.

3.3.4.6.3.3.8.4 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. Communications 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.

3.3.4.6.3.3.8.5 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.

#### 3.3.4.6.3.3.9 **Re-use of Directory Entries**

3.3.4.6.3.3.9.1 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.

3.3.4.6.3.3.9.2 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.

#### 3.3.4.6.3.4 **Overview of the Delate Compression Algorithm**

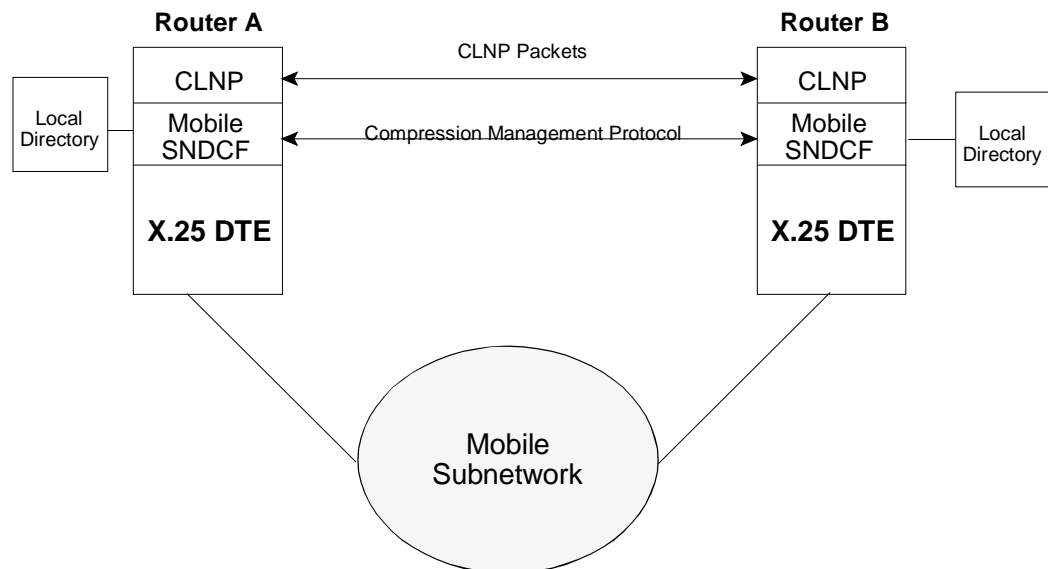
3.3.4.6.3.4.1 >To be developped<

#### 3.3.4.6.4 **Implementation Model**

3.3.4.6.4.1 The current generation of ICAO Mobile Networks all provide a network access service compliant with ISO/IEC 8208 (ITU-T recommendation X.25), although the VDL Mode 3 Subnetwork offers a frame mode service interface in addition. The CLNP specification already provides a set of procedures for passing CLNP packets over X.25 virtual circuits; ISO/IEC 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.

3.3.4.6.4.2 The implementation model for the Mobile SNDCF is illustrated in Figure 3.3-9 Implementation Model of the Mobile SNDCF. Note that the specification is not necessarily restricted to X.25. In principle, this specification may be readily adapted to any connection mode data link.

- 3.3.4.6.4.3 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.
- 3.3.4.6.4.4 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.
- 3.3.4.6.4.5 Note that the format of the compressed headers is such that they can be distinguished from normal CLNP packets, and as well as IS-IS, ES-IS and NLSP packets, and the local management protocol.
- 3.3.4.6.4.6 In each router, the Mobile SNDCF maintains a local directory for use by the LREF compression algorithm. A separate local directory is maintained 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 Rrouter or End System 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.



**Figure 3.3-9. Implementation Model of the Mobile SNDCF**



- 3.3.4.6.4.7 The local directory consists of entries numbered from zero to a maximum of 32767, each entry consisting of:
- a) a pair of NSAP Addresses, known as the inward and outward NSAP Addresses respectively;
  - b) the ISO/IEC 8473 protocol version number; and
  - c) the value of the security options parameter (see ISO/IEC 8473 Clause 7.5.3), which may be empty; ~~and~~
- 3.3.4.6.4.8 The directory is initially empty. The minimum directory size that may be supported is 128 entries.
- 3.3.4.6.4.9 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.

## 3.4 **ATN Routing**

### 3.4.1 **Introduction**

- 3.4.1.1 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 implementors 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 Domain will typically be configured to ignore the CLNP Security Parameter and hence the ATN Security Label.
- 3.4.1.2 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 IDRPs 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.
- 3.4.1.3 The following sections are ~~This Chapter is~~ concerned with describing how IDRPs work, how it is used in the ATN to support routing to both fixed and mobile systems, and the routing policies that have been adopted.

### 3.4.2 **Background to IDRP**

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

- a) there is a need for the communication of routing information between End Systems and Intermediate Systems. This requirement is satisfied by ISO/IEC 9542;
- b) 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; and
- c) there is a need for the communication of routing information between Intermediate Systems in different Routing Domains. ~~It is to satisfy~~ This requirement is satisfied by ISO/IEC 10747 that (IDRP) was developed.

3.4.2.2 In fulfilling the role of an inter-domain routing protocol, IDRP has to exchange routing information in what is described as a domain of limited 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.

3.4.2.3 Scalability 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.

3.4.2.4 In addition to meeting the requirements of ISO/IEC TR 9575, the ISO/IEC 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.

### 3.4.3 **Choice of IDRP for the ATN**

3.4.3.1 IDRP was chosen for ATN use early on in the development of the ATN ICS 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.

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

3.4.3.3 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 needs to be exchanged.

3.4.3.4 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.

#### 3.4.4 **IDRP Overview**

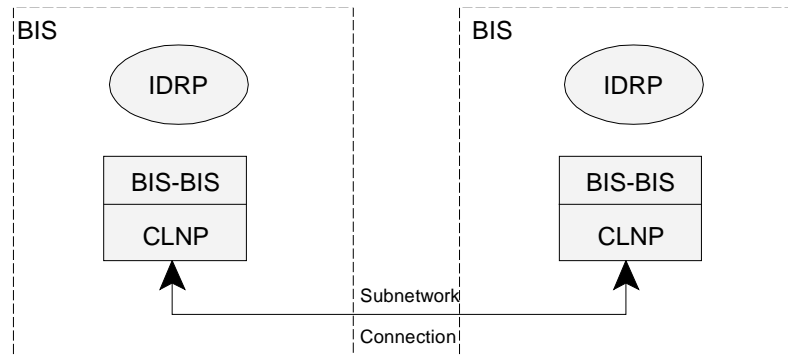
##### 3.4.4.1 *General*

3.4.4.1.1 IDRP is a routing information exchange protocol that supports:

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

3.4.4.1.2 IDRP is architecturally described by the protocol and process models shown respectively in Figure 3.4-1 and Figure 3.4-2, respectively.

- 3.4.4.1.3 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 communications follows the connection\_mode, i.e. the reliable exchange or routing information is supported 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/IEC 8073 Class 4 Transport Protocol, and uses the services of the Connectionless Network Protocol (CLNP) for data transfer between two adjacent BISs.
- 3.4.4.1.4 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.
- 3.4.4.1.5 The CLNP forwarding information necessary for BIS-BIS communications is typically



**Figure 3.4-1. IDRP Protocol Model**

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.

- 3.4.4.1.6 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; and
- b) information describing the path to such destinations.

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

3.4.4.1.8 IDRPs process the routes it receives from adjacent BISs (and locally provided routes) according to the process model shown in Figure 3.4-2.

#### 3.4.4.2 ***The Adj-RIB-in***

3.4.4.2.1 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 communications. 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 3.4.5). In such cases, there is a separate set of routes for each set of distinguishing path attributes.

3.4.4.2.2 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.

3.4.4.2.3 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 IDRPs. 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.

#### 3.4.4.3 ***The Loc-RIB***

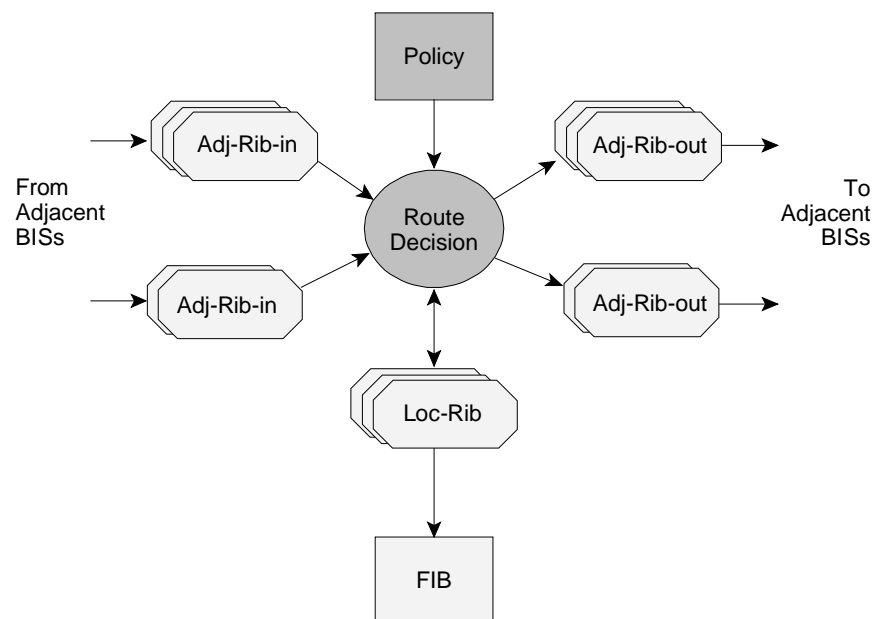
3.4.4.3.1 The Routing Decision process then selects routes for local use. This includes routes received from BISs in the local Routing Domain, external RDs and local routes provided either by a System Administrator or by intra-domain routing. This decision process is much like that described above where 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.

3.4.4.3.2 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.

#### 3.4.4.4 *The Adj-RIB-out*

3.4.4.4.1 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(s). 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.

3.4.4.4.2 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



**Figure 3.4-2. IDRP Process Model**

Domain is known as a Transit Routing Domain (TRD).

- 3.4.4.4.3 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.
- 3.4.4.5 ***Route Aggregation***
- 3.4.4.5.1 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 ISO/IEC 10747 sStandard.
- 3.4.4.5.2 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 scalability. 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.
- 3.4.4.5.3 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 ICS 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.
- 3.4.4.6 ***Route Information Reduction***
- 3.4.4.6.1 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(s) 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 scalability.

### 3.4.4.7 ***Routing Domain Confederations***

3.4.4.7.1 The last important feature of IDRPs 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.

3.4.4.7.2 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 IDRPs append 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.

### 3.4.5 **The ATN Security Path Attribute**

#### 3.4.5.1 ***General***

3.4.5.1.1 In IDRPs, the information that describes a route, in addition to a route's destination(s), 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(s), are still regarded as different routes.

3.4.5.1.2 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 NPDU's 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 IDRPs 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.

3.4.5.1.3 The ATN does not make use of the Quality of Service path attributes. However, it does use the IDRPs 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-Att's: 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.



3.4.5.1.4 The ATN ICS 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:

- a) the air/ground subnetwork type security tag; and
- b) the ATSC class security tag.

3.4.5.2 ***The Air/Ground Subnetwork Type Security Tag***

3.4.5.2.1 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.

3.4.5.2.2 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; and
- 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.

3.4.5.2.3 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.

3.4.5.3 ***The ATSC Class Security Tag***

3.4.5.3.1 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.

3.4.5.3.2 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; ~~and~~
- c) supports a lower ATSC Class.

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

- a) the route is available to both ATSC and non-ATSC data; and
- b) the route is available to ATSC data only.

3.4.5.3.4 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 are 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.

## 3.4.6 ~~The BIS-BIS Protocol~~

### 3.4.6.1 *General*

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

3.4.6.1.2 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/IEC 8473 DT PDU. Routes once advertised may also later be withdrawn by another UPDATE BISPDU.

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

### 3.4.6.2 *BIS-BIS Connections*

3.4.6.2.1 UPDATE BISPDU 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 BISPDU (again as the contents of a single ISO/IEC 8473 DT PDU) initialises the connection.

- 3.4.6.2.2 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 simultaneously issue an OPEN BISPDU.
- 3.4.6.2.3 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.
- 3.4.6.2.4 As long as routes are being exchanged in both directions then all the protocol information necessary to maintain reliable communications 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).
- 3.4.6.2.5 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.
- 3.4.6.3 ***RIB Refresh***
- 3.4.6.3.1 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.
- 3.4.6.3.2 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 local RIB, and rectify any discrepancies.
- 3.4.6.3.3 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.
- 3.4.6.4 ***Route Combination***
- 3.4.6.4.1 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.

3.4.6.4.2 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.

### 3.4.6.5 *Authentication and Security*

3.4.6.5.1 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.

3.4.6.5.2 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.

3.4.6.5.3 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.

3.4.6.5.4 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.

3.4.6.5.5 The ATN ICS SARPs currently require type 1 authentication. However, it should be noted that this may not 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 as 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.

3.4.6.5.6 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 be 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.

### 3.4.7 **The Route Decision Process**

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

- a) the selection of routes for Internal Distribution (Phase 1 decision);
- b) the selection of routes for Local Use (Phase 2 decision); and
- c) the selection and update of routes for External Distribution (Phase 3 decision).

Each of these three phases is described below.

#### 3.4.7.2 **The Phase One Decision Process**

3.4.7.2.1 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 from BISs in the local Routing Domain, because all such BISs are assumed to be in direct communications and will receive such a route direct from the local BIS from which it came.

3.4.7.2.2 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.

3.4.7.2.3 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.

3.4.7.2.4 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 off 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 Atttribute for internal distribution.

### 3.4.7.3 *The Phase Two Decision Process*

3.4.7.3.1 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, this will be the preferred route for a given value of the security path attribute.

3.4.7.3.2 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.

### 3.4.7.4 *The Phase Three Decision Process*

3.4.7.4.1 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 3.4-3. This illustrates the data structures and processes needed to implement the Route Decision process.

3.4.7.4.2 Two policy information base (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.

3.4.7.4.3 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) 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;

- b) 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
- c) 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).

3.4.7.4.4 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.

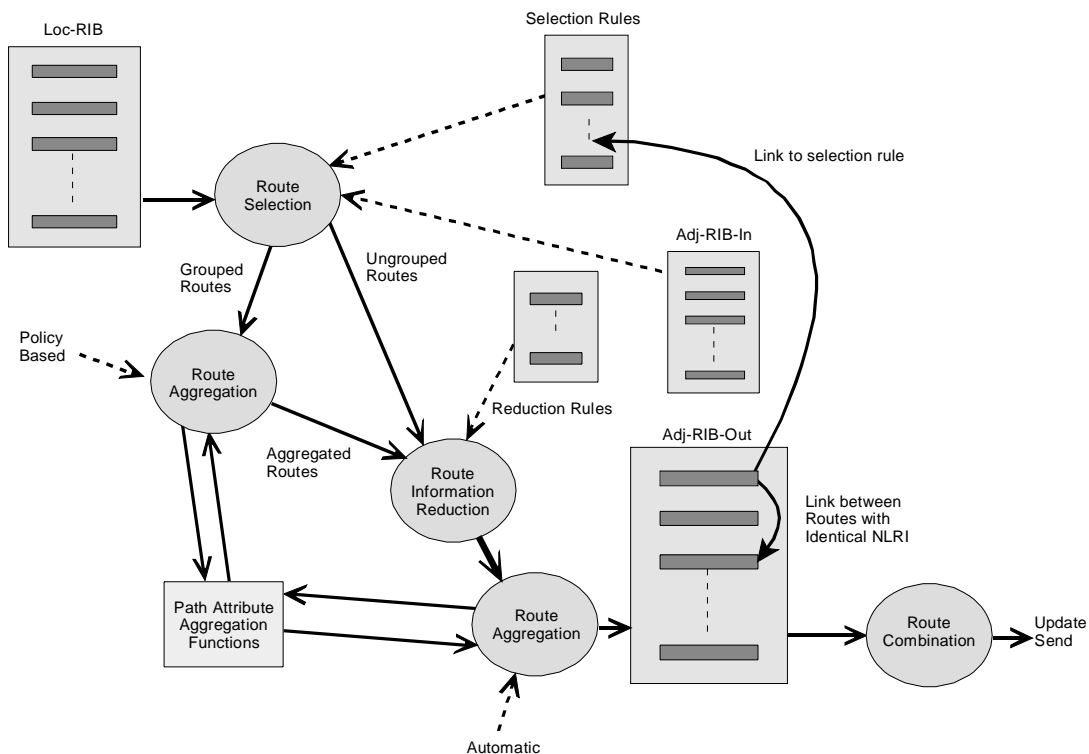
3.4.7.4.5 Type 2b rules are defined in response to specific ATN requirements for supporting routes to mobile systems. In order to optimise route information distribution, it is necessary to formulate rules that advertise a route to a given BIS, only if that BIS is advertising the selected route to a particular destination. The Type 2b rule is a class of rule that meets this requirement.

3.4.7.4.6 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:

- a) DIST\_LIST\_INCL/EXCL path attributes;
- b) different values of NEXT\_HOP; and
- c) different values of MULTI\_EXIT\_DISC.

3.4.7.4.7 The outcome, in such cases, is a local matter. However, it is recommended that a deterministic outcome is always ensured.

- 3.4.7.4.8 The remaining routes selected by Ttype 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.
- 3.4.7.4.9 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 of 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 ICS SARPs as a simplified Route Aggregation procedure.
- 3.4.7.4.10 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



**Figure 3.4-3. Generic Approach to Route Selection, Aggregation and Information Reduction**



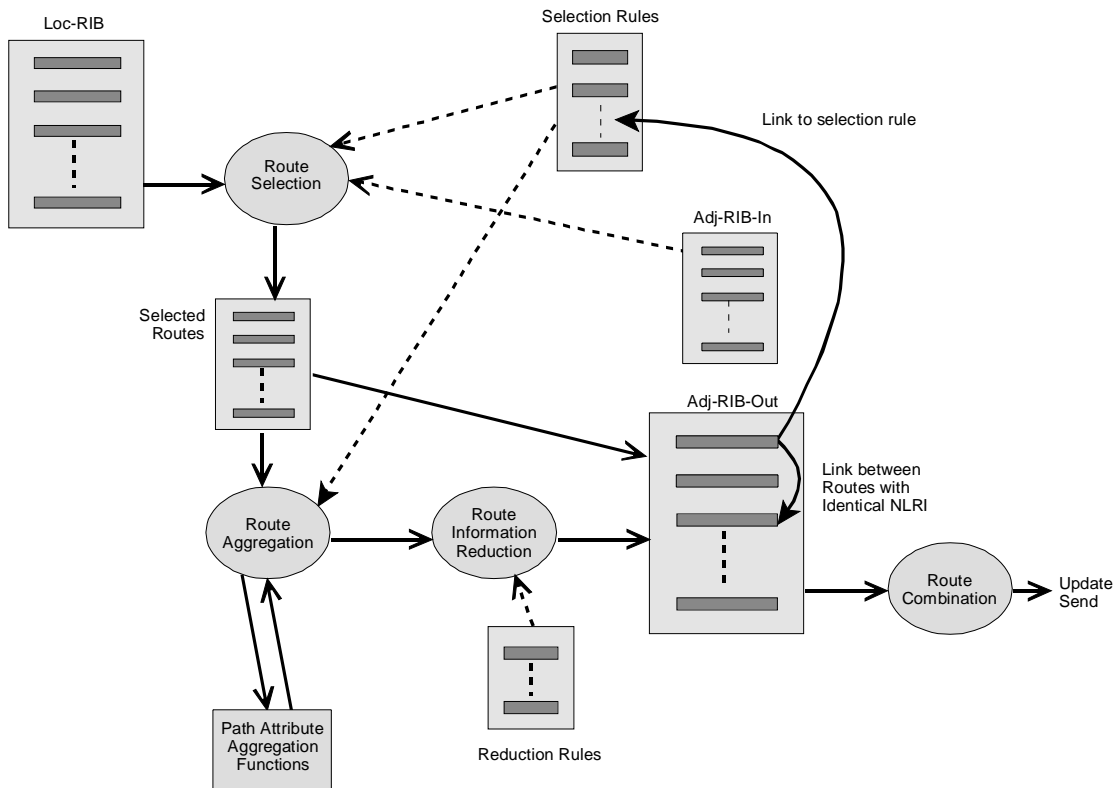
route.

- 3.4.7.4.11 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 Ttype 1 rule, then the new route is marked as replacing the route linked to that Selection Rule. If it is a Ttype 2a or Ttype 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.
- 3.4.7.4.12 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.
- 3.4.7.4.13 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.
- 3.4.7.4.14 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 3.4-4.
- 3.4.7.4.15 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 more 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.
- 3.4.7.4.16 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:
- a) route Information Reduction is only applied to aggregated routes (i.e. routes selected by Ttype 1 rules);

- b) 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; and
- c) routes selected by different Ttype 1 rules cannot, as a result of Route Information Reduction, have identical NLRI.

3.4.7.4.17 With these assumptions in place, the process model illustrated in Figure 3.4-4 can be considered.

3.4.7.4.18 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 (Ttype 1, Ttype 2a and Ttype 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 3.4-4. Simplified Model for Route Selection, Aggregation and Information Reduction**

3.4.7.4.19 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.

3.4.7.4.20 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.

3.4.7.4.21 The second assumption ensures that a route copied directly to an Aadj-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.

3.4.7.4.22 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.

### 3.4.8 Relationship to Intra-Domain Routing

3.4.8.1 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.

3.4.8.2 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 BISs 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 scalability 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~~RDs.

3.4.8.3 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.

3.4.8.4 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:

- a) 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;
- b) whenever the intra-domain routing function passes a PDU to the inter-domain routing function which is either;
  - 1) decapsulated and then routed to another Routing Domain; or
  - 2) 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.

- c) 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;
- d) 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; and
- e) 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.

#### 3.4.8.5

The consequence of the above approach is that BISs learn about the external destinations that systems inside the Routing Domain want to reach and, provided 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 is thereby avoided. This is advantageous because encapsulation always carries the risk of unnecessary segmentation, with the overheads that that implies.

### 3.4.9 Route Selection, Aggregation and Information Reduction

3.4.9.1 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 scalability of any internet network, this role has not yet been fully explained. The purpose of this section is to illustrate how these mechanisms are used to implement a scalable internet network. The approach taken is deliberately informal, in order to present a complex subject in an accessible manner.

#### 3.4.9.2 What is Route Aggregation?

3.4.9.2.1 Firstly, look at the signpost alongside in Figure 3.4-5, 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?

3.4.9.2.2 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.




Figure 3.4-5. Signposting the Way

3.4.9.2.3 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 Petersburg, and so on, until there is no further room on the signpost to hang another sign. What does our poor Signpost Manager do now?

3.4.9.2.4 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.


3.4.9.2.5 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 3.4-6.

3.4.9.2.6 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.

3.4.9.2.7 OK, so this is how road signs work, but is it really relevant to network routing?

3.4.9.2.8 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.



**Figure 3.4-6. The rationalised Signpost**

3.4.9.2.9 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 replaced the list of individual destinations, by a single common destination name. This procedure is properly known as *Route Information Reduction*.

### 3.4.9.3 ***Structured Addresses and Routing***

3.4.9.3.1 From this you may conclude that routers adopt a principle similar to that illustrated in Figure 3.4-6, and minimise the amount of routing information by collecting routes together and signposting routes to appropriate group addresses. Unfortunately, you would not always be right in making such a conclusion.

3.4.9.3.2 For example, in the TCP/IP Internet, the routers implemented by the Internet Service Providers are much more like the signpost in Figure 3.4-5. 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.


3.4.9.3.3 The reason why this is so is twofold. Firstly, the network addresses used in the TCP/IP Version 4 Internet Protocol 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.

- 3.4.9.3.4 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.
- 3.4.9.3.5 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:
- a) if the signpost is in London, then we start looking for a sign first to Baker Street;
  - b) otherwise, if the signpost is in England, we look for London;
  - c) otherwise, if the signpost is in Europe, we look for England; and
  - d) finally, if the signpost is not even in Europe, we look for a sign for Europe.
- 3.4.9.3.6 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.
- 3.4.9.3.7 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.
- 3.4.9.3.8 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 3.4-6. 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.


- 3.4.9.3.9 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 ~~of~~ 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.
- 3.4.9.4 **The Allocation of Structured Addresses**
- 3.4.9.4.1 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.
- 3.4.9.4.2 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.
- 3.4.9.4.3 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 ~~s~~Systems within a given organisation would then 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.
- 3.4.9.4.4 With this approach, similar network addresses, as illustrated in Figure 3.4-7, 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.





**Figure 3.4-7. Similar Network Addresses**

- 3.4.9.4.5 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.
- 3.4.9.4.6 Furthermore, as routing is done by such a simple prefix matching rule, the Routers do not themselves 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.
- 3.4.9.5 *Towards a Scaleable Routing Concept*
- 3.4.9.5.1 Our signpost analogy is really only one part of the routing concept. As illustrated in Figure 3.4-8, 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.
- 3.4.9.5.2 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 a 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.



**Figure 3.4-8. The Route - A Path Between A and B**

- 3.4.9.5.3 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.
- 3.4.9.5.4 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 IDRPs, 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.
- 3.4.9.5.5 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 scalable internetwork. This is because for an internetwork to be scalable, 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.
- 3.4.9.5.6 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.
- 3.4.9.5.7 For example, consider the example in Figure 3.4-9. 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.

3.4.9.5.8 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.

3.4.9.5.9 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.

3.4.9.5.10 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 scalability is all about.

3.4.9.5.11 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.

3.4.9.5.12 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




Figure 3.4-9. Aggregating Routes Together

Administration maintains. Thus only a single route needs to 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.

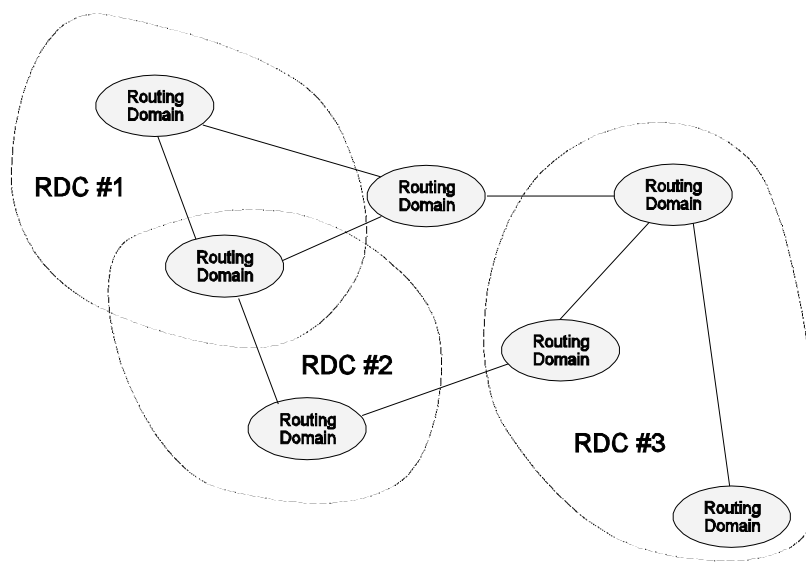
3.4.9.5.13 Looking ahead to Chapter 3.4.11.2, this principle is further exploited by the ATN Island concept. An ATN Island is essentially a regional grouping of Administrations with co-ordinated addressing plans. In such a situation, it is possible to form a single route to “the ATN Island”, and, indeed, 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.

### 3.4.9.6 *Containment Boundaries and Routing Domain Confederations*

3.4.9.6.1 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 3.4.4.7). 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.

3.4.9.6.2 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.

3.4.9.6.3 In the example in 3.4.9.5 above, there is clearly some sort of Containment Boundary enclosing the Service Provider and its users. This can simply be a conventional boundary.



**Figure 3.4-10. Routing Domain Confederations**

However, IDRP does provide a means to make this more concrete in the shape of a Routing Domain Confederation (RDC).

3.4.9.6.4 An RDC is no more than a group of Routing Domains, as illustrated in Figure 3.4-10, 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.

3.4.9.6.5 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.

3.4.9.6.6 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.

3.4.9.6.7 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 3.5.1+3.4.11, RDCs, Address Allocation and Route Aggregation are used together to create a scaleable ATN supporting mobile routing.

### 3.4.10 **Route Initiation**

#### 3.4.10.1 *The Purpose of Route Initiation*

3.4.10.1.1 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/IEC 10747 and used and disseminated according to local routing policies specified in accordance with the ATN ICS 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”.

3.4.10.1.2 Route Initiation procedures are required whenever two ATN RDs need to be interconnected. Since the ATN ICS 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.

3.4.10.1.3 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/IEC 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.*

### 3.4.10.2 **Ground-Ground Route Initiation**

#### 3.4.10.2.1 **The Communications Environment**

3.4.10.2.1.1 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 ICS SARPs, but can be assumed to include:

- a) X.25 Public and Private Data Networks;
- b) leased lines;
- c) integrated services digital networks (ISDNs);
- d) frame relay services; and
- e) the Public Switched Telephone Network (PSTN).

3.4.10.2.1.2 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.~~

3.4.10.2.1.3 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/IEC 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 communications 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.*

### 3.4.10.2.2 **Summary of Procedures**

3.4.10.2.2.1 The sequence of procedures for a typical ground-ground Routing Initiation is illustrated in Figure 3.4-11, and summarised below. They are described in greater depth in the following sections. This figure illustrates the co-ordination of two sSystems (“A” and “B”) interconnecting over a common network. The procedures are:

- a) adjacent BIS MOs are established in both sSystems. 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;
- b) a logical communications path is established over the subnetwork; typically one system is initiator and the other responder;
- c) establishment of the logical communications path is notified to the Systems Manager; and
- d) in response, the Systems Manager for each system:
  - 1) adds a route to the local FIB and to the remote sSystem; and
  - 2) invokes the IDRP “Start Event” action, or re-runs the decision process if a BIS-BIS connection already exists with the remote system.

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.*

### 3.4.10.2.3 **Initial Route Initiation**

3.4.10.2.3.1 Route Initiation begins with the decision to establish a communications path between a pair of BISs, including the decision on which communications network(s) to use. The first procedure is to establish the underlying communications circuit between the BISs and hence to establish the logical communications path.

3.4.10.2.3.2 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).

- 3.4.10.2.3.3 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.
- 3.4.10.2.3.4 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. |






Figure 3.4-11. Ground-Ground Route Initiation Sequence

### 3.4.10.2.4 **Route Initiation in CLNP**

3.4.10.2.4.1 The ATN ICS SARPs specify the use of the Connectionless Network Protocol (CLNP) specified in ISO/IEC 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:

- a) assign an appropriate Subnetwork Dependent Convergence Function (SNDCF) to interface the data link to the Network Entity; and
- b) update the Forwarding Information Base (FIB) to record statically known routes available over the data link and via the remote BIS.

3.4.10.2.4.2 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.

3.4.10.2.4.3 The SNDCF is typically specified for a network type and associated at system configuration time with a physical communications 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 communications path~~data link~~. However, for some network types there may be alternatives chosen at connection establishment time.

3.4.10.2.4.4 The FIB may be updated with any statically known routes that are known a priori to exist via the newly established communications path~~data link~~, where a route consists of an NSAP Address prefix paired with an identifier for a communications path~~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 communications path~~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 communications path~~data link~~ has been established.

3.4.10.2.4.5 Therefore, once a communications path~~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 communications path~~data link~~. As a minimum, this prefix may be the complete NET; and

- b) the communications path~~data link~~ to that remote BIS.

*Note 1.— ~~T~~he reverse must also take place when the data link is terminated, i.e. the above route must be removed from the FIB.*

*Note 2.— ~~A~~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.*

### 3.4.10.2.5 **Route Initiation in IDRP**

- 3.4.10.2.5.1 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.
- 3.4.10.2.5.2 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.
- 3.4.10.2.5.3 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 in the MO 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.
- 3.4.10.2.5.4 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.
- 3.4.10.2.5.5 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 ATN ICS 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 theto risk of both being configured in the passive role by mistake.

- 3.4.10.2.5.6 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.
- 3.4.10.2.5.7 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.
- 3.4.10.2.5.8 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 ATN ICS SARPs require that the IDRP Decision Process be aware of the aggregated QoS and Security Restrictions over the communications paths to a given remote BIS (Adjacent BIS). The ATN ICS 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(s) and to use this updated QoS when determining the degree of preference of the route and when re-advertising it.
- 3.4.10.2.5.9 The ATN ICS 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(s). 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~~ Adj-RIB-out for that Router (and hence advertised to it).
- 3.4.10.2.5.10 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.

### 3.4.10.3 **Air-Ground Route Initiation**

- 3.4.10.3.1 Air-Ground Route Initiation is similar to ground-ground Route Initiation, but differs for the following reasons:
- a) ICAO specified subnetworks are used for air-ground communications with their procedures for use mandated by SARPs rather than subject to bilateral negotiation;
  - b) route initiation typically starts as soon as communications 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);

- c) 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;
- d) 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; and
- e) 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.

### 3.4.10.3.2 **Communications Environment**

3.4.10.3.2.1 The following ICAO Air/=Ground data networks are expected to be used to support the ATN:

- a) the aeronautical mobile satellite service (AMSS);
- b) the VHF data link (VDL); ~~and~~
- c) the Mode S data link; and
- d) the HF data link.

3.4.10.3.2.2 In each case, ITU recommendation X.25 provides the data network access procedures (with VDL Mode 3 offering an additional frame-mode network access), and the responsible ICAO Panels have required that:

- a) AMSS communications are “air initiated”, that is the aircraft is responsible for initiating communications with the ground;
- b) VDL communications are similarly air initiated; and
- c) Mode S communications are “ground initiated”, that is a ground ATN Router attached to a Mode S data link is responsible for initiating communications with an aircraft.

### 3.4.10.3.3 **Summary of Procedures**

3.4.10.3.3.1 The Air-Ground Route Initiation procedures are illustrated in Figure 3.4-12, and summarised below. They are described in greater depth in the following sections. This figure illustrates the case where a Joint Event is generated by the air-ground subnetwork. If the subnetwork cannot generate a Joint 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.



3.4.10.3.3.2 The Route Initiation Procedures are:

- a) 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;
- b) system “A” either:
  - 1) 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
  - 2) if polling, System “A” issues a Call Request to the next address on its poll list.
- c) 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 and information on the capabilities of the air-ground subnetwork, in the air-initiated case, i.e. if System “B” is an ATN Ground Router;
- d) 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; and
- e) on receipt of the ISH;
  - 1) 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~~ BISPDU, while System “B” generates one and sends it to System “A”. System “A” responds to the OPEN BISPDU, with its own OPEN BISPDU; or
  - 2) alternatively, if a BIS-BIS connection already exists with the remote system, then the IDRPs Decision Process is re-run.

Once the ~~BIS-Open~~ OPEN BISPDUs have been exchanged, the Route Initiation procedures have been completed.

### 3.4.10.3.4 Initial Route Initiation

#### 3.4.10.3.4.1 General

3.4.10.3.4.1.1 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 ICS SARPs to be implemented in each airborne Router and each ground Router with air-ground connectivity. This is 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.*

3.4.10.3.4.1.2 The IS-SME may have to handle two different classes of air-ground subnetwork:

- a) 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; and
- b) air-ground subnetworks which have no mechanism for recognising the above event and/or reporting it.

3.4.10.3.4.1.3 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.*

#### 3.4.10.3.4.2 The Join Event

3.4.10.3.4.2.1 Ideally, the Join Event should be an OSI Systems Management Notification sent to the IS-SME from a Management Entity in the subnetwork itself. This notification should provide the following information:

- a) a subnetwork identifier allowing the BIS to associate the event with an air-ground subnetwork to which the Router is connected;
- b) the address on that subnetwork of the remote airborne or ground Router; and
- c) the expected lifetime of the adjacency, i.e. how long a communications path is expected to be available.



3.4.10.3.4.2.2 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 subnetwork at the time it comes into contact with that air-ground subnetwork.

3.4.10.3.4.2.3 On receipt of a Join Event, an ATN Ground Router will, if communications 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.

3.4.10.3.4.2.4 Likewise, on receipt of a Join Event, an ATN Airborne Router will, if communications 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.

3.4.10.3.4.2.5 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.

#### 3.4.10.3.4.3 **The Join Event for Subnetworks that do not support ATN Systems Management**

3.4.10.3.4.3.1 It is anticipated that not all ICAO air-ground subnetworks will support the OSI 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:

- a) 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; and
- b) 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.

*Note.*— An example of Join Event packet is provided in Table 3.4-1.

#### 3.4.10.3.4.4 **Procedures for Air-Ground Subnetworks that do not Provide a Join Event**

3.4.10.3.4.4.1 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.

- 3.4.10.3.4.4.2 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.

**Table 3.4-1. Joint Event Format**

Field	Size, in 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

**Note.**— **Notes:** *The following conventions apply:*

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 (SP). 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 - 1-200 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~~ Join 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 ISO/IEC 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”; ~~and~~*
  - b) *a one-octet length is set to the length of the remote Router's SNPA address; and*
  - ~~c) The variable-length value contains the actual DTE address of the remote Router.~~*

- 5.4. Multiple SNPA fields may be included within a single routing initiation event to report the reachability of several Routers simultaneously.
- 6.5. The Version~~VER~~ field should be set to '1'.
- 7.6. The value of the Message ID~~type~~ field identifying the following data to be of type 'SNPA' should be set to '1'.

3.4.10.3.4.4.3 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.*

### 3.4.10.3.5 **Route Initiation in CLNP**

3.4.10.3.5.1 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 ICS 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.

3.4.10.3.5.2 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 ICS SARPs.

3.4.10.3.5.3 An ISO/IEC 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 and to uplink capability information on the air/ground subnetwork for use by the airborne NPDU forwarding process.*

### 3.4.10.3.6 **Route Initiation in IDRP**

3.4.10.3.6.1 Route Initiation in IDRP in the air-ground case is then almost identical to the ground-ground case, except that the ATN ICS 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~~ subnetworks are used.

3.4.10.3.6.2 The ATN ICS 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 ATN ICS 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.

3.4.10.3.6.3 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.

3.4.10.3.6.4 An IDRP “Start Event” is then invoked by the IS-SME, provided that a BIS-BIS connection does not already exist 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.

### 3.4.10.4 ***Air-Ground Route Initiation without IDRP***

3.4.10.4.1 Due to certain avionics limitations, the ATN ICS 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 ICS 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 non-support of IDRP by an Airborne Router ~~use 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.*

#### 3.4.10.4.2 **Summary of Procedures**

3.4.10.4.2.1 The procedures for Air-Ground Route Initiation without IDRP are illustrated in Figure 3.4-13, 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:

- a) 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;
- b) 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
- c) 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;
- d) 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;
- e) on receipt of the ISH PDU, both systems update their local FIB to include the routing information received on the PDU;
- f) system “A” generates the derived routes using the NET of System “B”, inserts them into the IDRP RIB, and invokes the IDRP Decision Process; and
- g) 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.

#### 3.4.10.4.3 **Initial Route Initiation**

3.4.10.4.3.1 There is no difference in the initial Route Initiation procedures when IDRP is not used over the air-ground data link.




Figure 3.4-13. Air-Ground Route Initiation without IDRPs

#### 3.4.10.4.4 **Route Initiation in CLNP**

- 3.4.10.4.4.1 The ATN ICS 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/IEC 10589. The ATN ICS 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.*— ~~Support~~*that support of a NET with an NSEL of zero is necessary in such Airborne Routers when, for example, they also support ISO/IEC 10589 within the aircraft.*

- 3.4.10.4.4.2 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.

#### 3.4.10.4.5 **IS-SME Procedures without the use of IDRP**

##### 3.4.10.4.5.1 **In the Ground Router**

- 3.4.10.4.5.1.1 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.

- 3.4.10.4.5.1.2 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 received 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.

- 3.4.10.4.5.1.3 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.

- 3.4.10.4.5.1.4 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 routes to the Airborne Router.

3.4.10.4.5.1.5 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-Attr) 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. 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.

3.4.10.4.5.1.6 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.— In that in contrast with the use of IDRP over an air-ground data link, when the ATNICS 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.*

#### 3.4.10.4.5.2 In the Airborne Router

3.4.10.4.5.2.1 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; and
- 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/IEC 10589 is implemented, the generated routes are used to generate Reachable Address MOs and the ISO/IEC 10589 entity is used to update the FIB.

3.4.10.4.5.2.2 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:



- a) restrictions on Security Types for such destinations that are additional to ITU restrictions imposed by the Air-Ground Subnetwork; and
- b) 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.

3.4.10.4.5.2.3 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.

3.4.10.4.5.2.4 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 from capability information for~~ on each active Air-Ground Subnetwork supported. ~~This latter information on the characteristics of each individual subnetwork is uplinked to the Airborne Router during the route initiation process over each such subnetwork and is used in selecting the appropriate subnetwork for downlinking NPDU's determining processed to determine .~~ 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.

3.4.10.4.5.2.5 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 and the subnetwork characteristics.

3.4.10.4.5.2.6 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/IEC 10589 is implemented, then the IS-SME may make such routes available to ISO/IEC 10589 by creating a Reachable Address MO for each such route, and removing the MO when the route ceases to be available. The ISO/IEC 10589 implementation may be relied upon to maintain the FIB with this routing information.

#### 3.4.10.4.6 **Management of the ISH PDU Holding Time**

3.4.10.4.6.1 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:

- a) the ISH PDU exchange is first used to negotiate the use or non-use of IDRP;
- b) 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; ~~and~~

- c) the initial ISH PDU exchange is also used to inform the peer ATN Router about those extended capabilities which are supported by the sending ATN Router over the air/ground adjacency (such as generation or respectively reception of routes without security information, or ATN security services) and about the characteristics of the air/ground subnetwork used to exchange the ISH PDU. This information is included in the options part of the ISH PDU. To ensure backwards compatibility, it will be ignored on receipt by those ATN Routers which are not compliant with the most current edition of the ATN ICS SARPs;
- d) the ISH PDU may also be used by an Air/Ground Router to request the public-key certificate from the peer Airborne Router in support of the ATN security services; and
- e) 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.

- 3.4.10.4.6.2 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.
- 3.4.10.4.6.3 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.
- 3.4.10.4.6.4 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 communications 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

PDU's 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.

3.4.10.4.6.5 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.

3.4.10.4.6.6 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.

3.4.10.4.6.7 Implementors should also note that existing implementations of ISO/IEC 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.

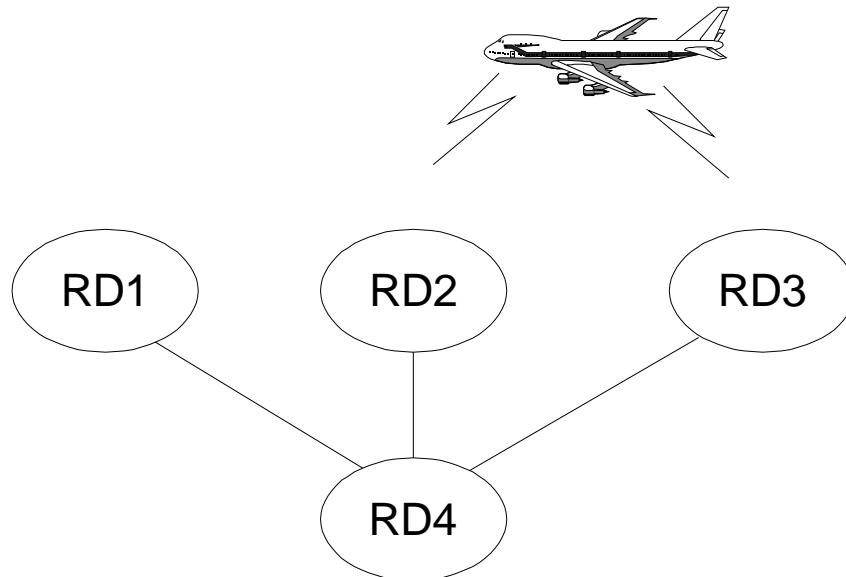
### 3.4.11 **Support for Mobile Systems**

#### 3.4.11.1 *Mobility and Routing Domains*

3.4.11.1.1 The scalability 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 IDRPs, 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 scalability.

3.4.11.1.2 This property of the routing architecture and of IDRPs; is exploited by the ATN to support Mobile Routing.

- 3.4.11.1.3 In the ATN, the systems onboard an aircraft form a Routing Domain unique to that aircraft which is characterised by one address prefix for ATSC systems, and another for AINSC 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



**Figure 3.4-14. Mobile Routing Example**

the airline's routing policy. Routing Information is then exchanged between ground Routing Domains, using IDRPs, so that all ground Routing Domains are aware of the current route to that aircraft. This is illustrated in Figure 3.4-14.

- 3.4.11.1.4 In this example, there are four ground based Routing Domains RD1 through to RD4. RD1, RD2 and RD3 all support air/ground data links, while RD4 depends on the other three for air/ground communications. The aircraft currently has communications over air/ground data links with both RD2 and RD3.
- 3.4.11.1.5 Using IDRPs, 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.
- 3.4.11.1.6 As the aircraft continues on its journey, it may lose communications with RD3. For example, it goes out of range of the VHF data link it was using to communicate with RD3. RD3 informs RD4 of this situation by issuing the appropriate IDRPs 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 communications with RD3 and must now route all traffic via RD2.

- 3.4.11.1.7 Further on the journey, the aircraft comes into contact with an air/ground data\_link offering communications with RD1. A data\_link 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, 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.
- 3.4.11.1.8 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.
- 3.4.11.2 ***Containing the Impact of Mobility***
- 3.4.11.2.1 While the principles of mobile routing outlined in the previous section are straightforward they are not scaleable using the existing IDRPs associated with Route Aggregation and RDCs. The problem is that even if an aircraft is given an NSAP 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.
- 3.4.11.2.2 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 3.4.5.11.4 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.
- 3.4.11.3 ***Routing to Mobiles within an ATN Island***
- 3.4.11.3.1 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 data\_links. These Routing Domains form an RDC, as illustrated in Figure 3.4-15, and an ATN Island is essentially an RDC in which certain Routing Policy rules are followed. All ATN Routing Domains that have air/ground data\_link are members of an ATN Island and, although most ATN Routing Domains which do not have air/ground data\_link 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.
- 3.4.11.3.2 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.

3.4.11.3.3 Within the ATN Island, the Backbone RDC provides a default route to *all aircraft*. As illustrated in Figure 3.4-145, this is advertised to all other Routing Domains within the Island as a route to the common address prefix for all aircraft.

3.4.11.3.4 Routing Domains with explicit routes to aircraft then have a simple routing policy rule to determine to which adjacent Routing Domain they must advertise such a route. 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 `data_links` available to the Island's Routing Domains, and can thus act as default route providers for packets addressed to airborne systems.

*Note.*— 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.

3.4.11.3.5 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 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.

- 3.4.11.3.6 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 communications 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 3.4.5.11.6 below).
- 3.4.11.3.7 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 providers to any aircraft currently in communications 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.
- 3.4.11.3.8 The ATN Island is only the first part of achieving a scalable 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 data link needs only the capacity to handle the

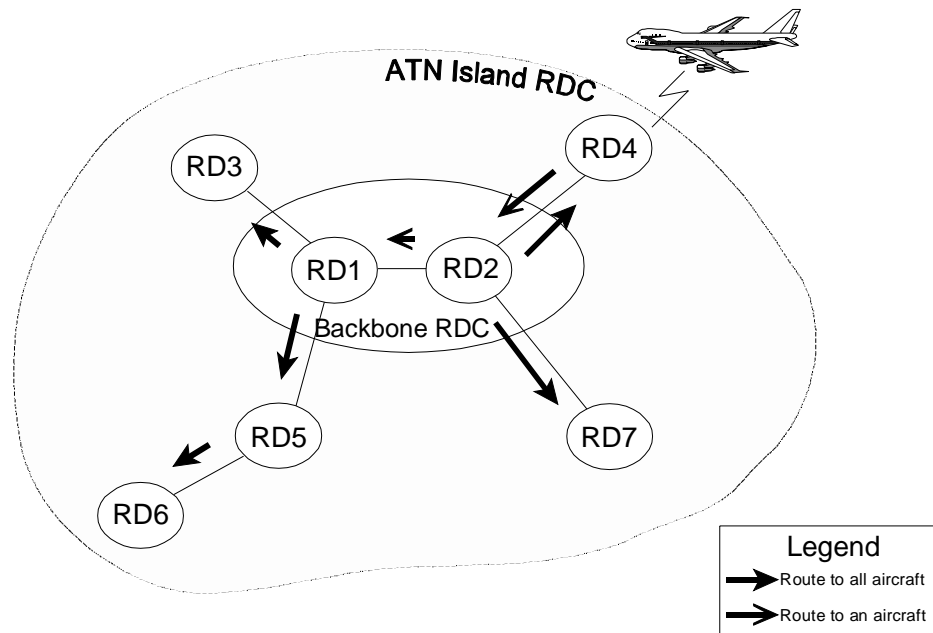


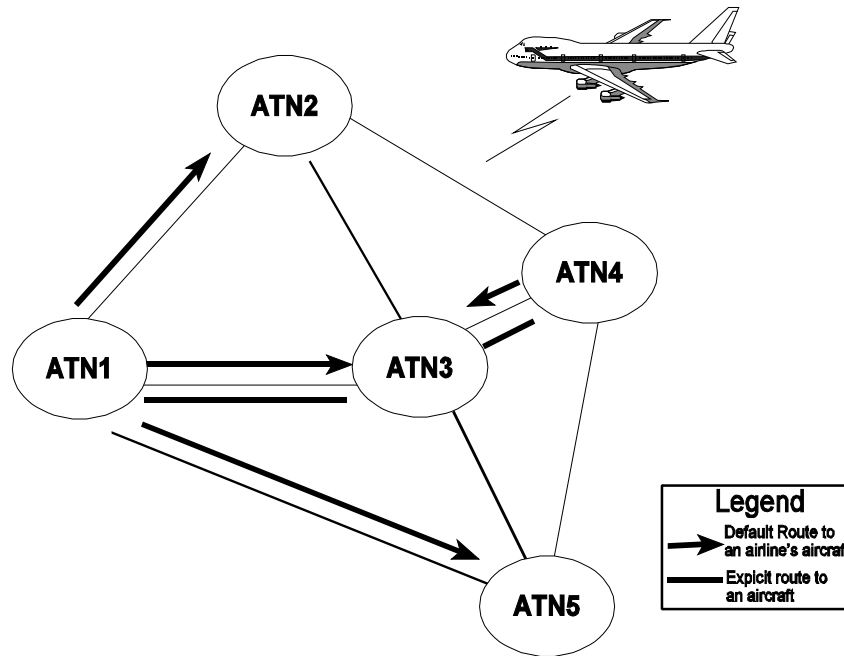
Figure 3.4-15. Mobile Routing Within an ATN Island

aircraft supported by its data link(s), and there is a similar impact on Routing Domains that are Transit Routing Domains providing a route between the backbone and an air/ground data link equipped Routing Domain. For all other Routing Domains on the Island, there is no impact on routing overhead due to aircraft.

- 3.4.11.3.9 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.
- 3.4.11.4 ***Routing to Mobiles between ATN Islands***
- 3.4.11.4.1 ATN Islands can be set up such that their geographical spread matches Air Traffic Control communications 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 scalability.
- 3.4.11.4.2 The mechanism used to achieve this derives from the concept of the “Home” domain.
- 3.4.11.4.3 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 communications 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.
- 3.4.11.4.4 The operation of the “Home” domain is illustrated in Figure 3.4-16. 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.



3.4.11.4.5 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 airline or country of registration exists then the actual route to the aircraft is advertised to the Routing Domain advertising that default route. Otherwise, the explicit



**Figure 3.4-16. Inter-Island Routing**

route is not advertised outside of the Island. In Figure 3.4-16, 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.

*Note.— 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.*

3.4.11.4.6 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).

3.4.11.4.7 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.

- 3.4.11.4.8 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”.
- 3.4.11.4.9 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.
- 3.4.11.5 ***Impact on Air/Ground Data Links***
- 3.4.11.5.1 A final limiting factor on the ATN is the capacity of the air/ground data links. At present, these are low bandwidth communications channels and only the minimum routing information can be transferred over them.
- 3.4.11.5.2 IDRPs are 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.
- 3.4.11.5.3 In addition, IDRPs are 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.
- 3.4.11.5.4 The ATN uses these properties of IDRPs to keep the transfer of routing information over an air/ground data link to a minimum. When the data link 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.
- 3.4.11.5.5 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.
- 3.4.11.5.6 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:
- a) the local Routing Domain (typically that providing Air Traffic Services); and
  - b) the ATN as a whole:

in addition to other routes as determined by local policy.

3.4.11.5.7 The airborne router will then be able to choose between the alternative routes (via different ground routers to these destinations.

### 3.4.11.6 *The Impact of Routing Updates*

#### 3.4.11.6.1 **General**

3.4.11.6.1.1 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.

3.4.11.6.1.2 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.

3.4.11.6.1.3 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.

3.4.11.6.1.4 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.


3.4.11.6.1.5 The KeepAlive timer is used within IDRP to determine the health of a link. This directly controls the frequency which IDRP KeepAlive PDUs are sent on BIS-BIS connections. There is a trade-off concerning the setting of this timer. A small value of this timer will more accurately determine a change in link status, however this will increase the protocol overhead of an already bandwidth limited air/ground resource. The setting of this timer to a small value will also increase the financial cost of the resource. A large value of the Keepalive timer will be less responsive to determine a change in link status, however this will decrease the protocol overhead across the air/ground resource. The setting of this timer to a large value will also decrease the financial cost of the resource. It is recommended that this value be based on operational experience between the various States and Organizations.

#### 3.4.11.6.2 **“Hold Down” Timer Use**

3.4.11.6.2.1 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.

3.4.11.6.2.2 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 3.4-17.

3.4.11.6.2.3 In this example, RD4 provides a route to the aircraft, to RD5. When the aircraft loses



**Figure 3.4-17. Impact of a Hold Down Timer**

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.

3.4.11.6.2.4 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.

- 3.4.11.6.2.5 This has important implications for the design of an ATN Island. If an Island's air/ground data\_links 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 data links, then the update frequency can be significantly reduced, without affecting the responsiveness to real connectivity changes.
- 3.4.11.6.2.6 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 data link 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.
- 3.4.11.6.2.7 Routes advertised to an aircraft's "Home" are also affected by the hold down timer and, in this case, RDCs and the Hold Down Timer work together to keep the routing overhead to an absolute minimum.
- 3.4.11.6.2.8 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.
- 3.4.11.6.2.9 Results from simulation work have shown that the "ideal" setting for the **minRouteAdvertisementInterval** is under one minute (typically 30 seconds). Furthermore, simulation has shown that complex topologies for the ATN Island Backbone should be avoided as they significantly increase the convergence time. Two independent studies have shown that an hierarchical arrangement of ATN Island, each with a small number of Backbone BISs, both reduces the volume of IDRP Update traffic and promotes a scaleable architecture. Figure 3.4-18 illustrates such an architecture.
- 3.4.11.6.2.10 Simulations have also shown that the optimal interconnection of ATN Islands is a single direct adjacency between each pair of ATN Islands.

- 3.4.11.6.2.11 Having a small number of BISs on a backbone has been demonstrated to be an optimal arrangement as increasing the number of BISs increases the number of IDRP updates and peer relationships each BIS must handle. However, it may not be possible to produce a topology that satisfies this. Under the circumstances that an RDC is formed from a number of fully meshed BISs on a common subnetwork, the use of a Route Server may improve the route convergence.


**Figure 3.4-18. A Hierarchical Structure of ATN Islands**

- 3.4.11.6.2.12 A Route Server is a system that participates in IDRP, but does not participate in the actual CLNP packet forwarding. A Route Server is a BIS dedicated to the processing of routes: it acquires routing information from all the BISs connected to a common WAN, performs the IDRP decision process over this information, and then redistributes the results to the routers. Figure 3.4-19 illustrates the use of a Route Server in a fully meshed RDC.
- 3.4.11.6.2.13 The Route Server approach relies on the use of an optional feature of the ISO/IEC 10747 (IDRP) standard:
- a) the BISs which are clients to the route server would need to support “the generation of the NEXT\_HOP attribute in support of route servers” options;
  - b) the Router Server is a BIS which does support the “propagation of the NEXT\_HOP attribute in support of route servers”; and

- c) the support on receipt of the NEXT\_HOP attribute is a mandatory IDRPs function and is therefore assumed to be supported by every standard ATN IDRPs implementation.

3.4.11.6.2.14 The Route Server approach relies therefore on standard mechanisms and can be used in the ATN provided that the options mentioned above are implemented.

3.4.11.6.2.15 Note that this approach delegates the Routing Policy decisions to the Route Server. However, this can be appropriate as long as the Routing Policies among all the BISs



**Figure 3.4-19. The Use of a Route Server in a Fully Meshed RDC**

connected to the common subnetwork are consistent among themselves.

### 3.4.11.7 Failure Modes

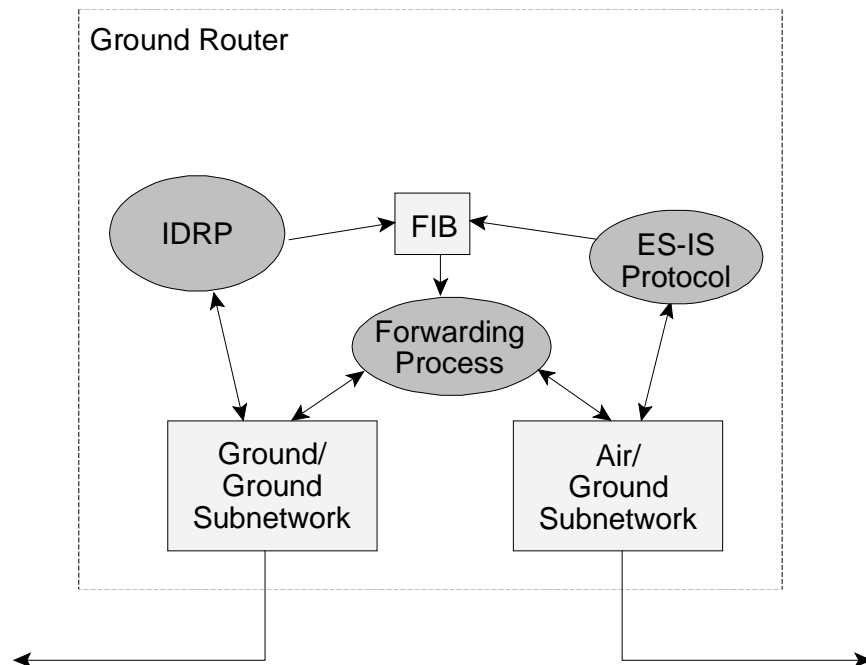
3.4.11.7.1 In the pure ground-ground environment, loss of a router or a communications path can be readily recovered from provided an alternative route exists and routing policy permits its use. However, the situation is not so straightforward with the policy rules that support mobile routing. The ATN Mobile Routing Concept depends upon two default route providers, the ATN Island Backbone and the “Home”. Failure of either of these or loss of access to them will impact mobile routing.

#### 3.4.11.7.2 Loss of the “Home”

3.4.11.7.2.1 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.

- 3.4.11.7.2.2 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.
- 3.4.11.7.2.3 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.
- 3.4.11.7.3 **Failure of an ATN Island Backbone**
- 3.4.11.7.3.1 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 data links 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.
- 3.4.11.7.3.2 For similar reasons to those already detailed in ~~5.11.7.13.4.11.7.2~~, there is no need for loss of an ATN 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.
- 3.4.11.8 **Optional non-Use of IDRP**
- 3.4.11.8.1 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 in all cases.
- 3.4.11.8.2 An alternative approach is possible using provisions in the ISO/IEC 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





**Figure 3.4-20. Architecture of an Initial Ground Network Router**

operation of ISO/IEC 9542) can be used to keep the link between ground and airborne routes “live” in the absence of IDRP.

- 3.4.11.8.3 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.
- 3.4.11.8.4 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.
- 3.4.11.8.5 The architecture of a ground router implementing such functionality is illustrated in Figure 3.4-20. 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.



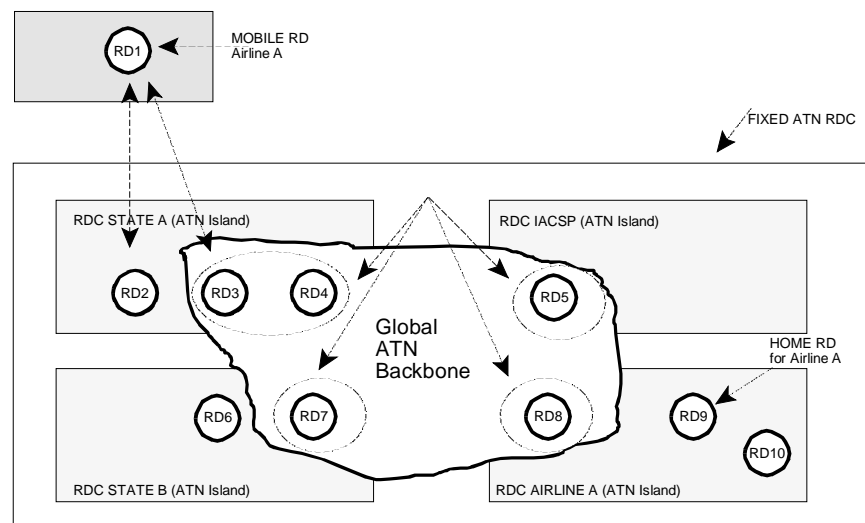
3.4.11.8.6 As the ISH PDU mechanism is also used for route initiation in the full ATN, some convention for distinguishing between its use in this scenario and in the full ATN is necessary. This can be readily achieved by addressing conventions. A non-zero value in the NET's "SEL" field (254 decimal) is used to signal use of the above procedures.

3.4.11.8.7 Routing information learnt in this way by the Air/Ground Router may then be disseminated throughout the ATN Ground Environment using normal IDRP procedures.

### 3.4.11.9 **Routing Policies in Support of Mobile Routing**

3.4.11.9.1 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 5.8.35.3.7 of the ATN ICS 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.

3.4.11.9.2 Figure 3.4-21 defines the example Routing Architecture scenario that has been used as the basis for the guidance provided in this section.



**Figure 3.4-21. Example Routing Policy Scenario**

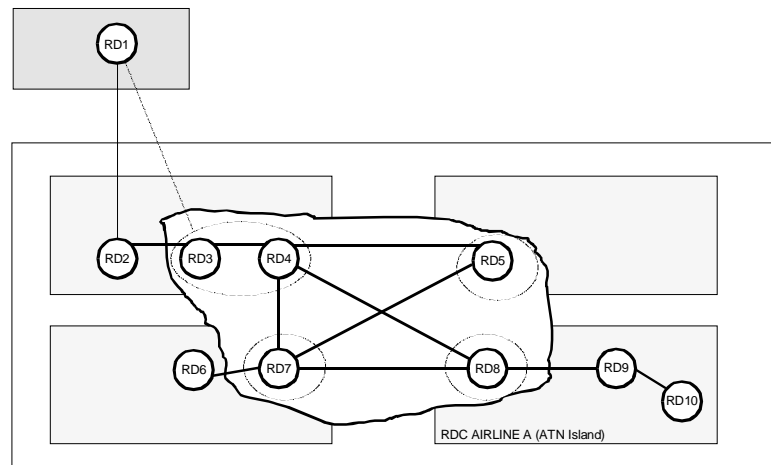
3.4.11.9.3 The following are the key components of the example network:

- a) the scenario defines at the highest level the "Fixed ATN RDC" which the ATN ICS SARPs define to comprise of all fixed ATN RDs;
- b) within the Fixed ATN RDC are defined four organisational RDCs:

- 1) an RDC for State “A”;
- 2) an RDC for State “B”;
- 3) an RDC for an International Aeronautical Communications Service Provider (IACSP); and
- 4) an RDC for an airline (airline “A”).

*Note.— The term “RDC” in this context is synonymous with the term “ATN Island”.*

- c) 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; and



**Figure 3.4-22. Routing Policy Example Connectivity**

- d) the connectivity between the RDs is illustrated in Figure 3.4-22.

#### 3.4.11.9.4

The following RDCs are defined:

- a) State “A” RDC:

- 1) The State A RDC comprises three RDs (RD2, RD3 and RD4); and
- 2) The State A RDC includes a Backbone RDC that comprises RDs 3 & 4 and a TRD (RD2) off the Backbone.

- b) State “B” RDC:

- 1) The State B RDC comprises two RDs (RD6 and RD7); and

- 2) One RD (RD7) is the only member of the State B Backbone RDC.
- c) IACSP RDC:
  - 1) The IACSP RDC comprises one RD (RD5) which is the only member of the IACSP RDCs-Backbone RDC.
- d) Airline A RDC:
  - 1) The Airline A RDC comprises three RDs (RD8, RD9 and RD10);
  - 2) The Backbone RDC comprises only RD (RD8);
  - 3) RD9 is a TRD and is designated as the “Home RD” for the airline; and
  - 4) RD10 is defined to be an ERD.

3.4.11.9.5 RD1 is a Mobile RD belonging to Airline A.

3.4.11.9.6 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.

3.4.11.9.7 It should be noted that an overriding requirement in the ATN ICS 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 ATN ICS 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 ATN ICS SARPs.

Table 3.4-2. Routing Policy Requirements for Members of an ATN Island Backbone RDC (ATN ICS SARPs Ref. 5.3.7.2)

ATN ICS SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
5.3.7.1.2	<b>Adjacent ATN RD's within the ATN Island Backbone RDC</b>	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	<p>Each Router in RD3 is required to advertise the following routes to each adjacent Router in RD4:</p> <ul style="list-style-type: none"> <li>● a route to NSAPs &amp; NETs contained within RD3;</li> <li>● 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;</li> <li>● the selected route to every Fixed ATN RD in the same Island i.e. a Route to RD2.</li> </ul>
5.3.7.1.3	<b>All other ATN RDs within the ATN Island</b>	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	<p>In this case RD8 will advertise <u>the following routes</u> to RD9:</p> <ul style="list-style-type: none"> <li>● a route to NSAPs &amp; NETs contained within RD8;</li> <li>● the selected Route to every Fixed ATN RD in the same ATN Island for which a Route is available (not applicable in this example);</li> <li>● a Route to all Mobile RDs thereby providing a default Route to all Mobiles;</li> <li>● 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.</li> </ul>

<b>ATN ICS SARPs Reference</b>	<b>Category</b>	<b>Description</b>	<b>Applicable RDs from Scenario</b>	<b>Example</b>
5.3.7.1.4	<b>Mobile RDs</b>	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 Mobile RD.	RD3–RD1	<p>In this case RD3 will advertise to the Mobile RD1:</p> <ul style="list-style-type: none"> <li>● a Route to NSAPs &amp; NETs contained within RD3.</li> </ul> <p>The ATN ICS SARPs additionally recommend that RD3 should advertise to the Mobile RD1:</p> <ul style="list-style-type: none"> <li>● an aggregated Route to NSAPs &amp; NETs contained within the State A RDC (i.e. the local RDC) and;</li> <li>● an aggregated Route to NSAPs &amp; 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).</li> </ul>

ATN ICS SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
5.3.7.1.5	<b>ATN RDs in other ATN Islands</b>	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 to 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 <u>ATN</u> Island Backbone RDCs: <ul style="list-style-type: none"> <li>• an aggregated Route to NSAPs and NETs contained within the Airline <u>A</u> RDC;</li> <li>• a Route to all Mobile RDs assigned to Airline A since the Home RD (RD9) belongs to the same Island as RD8;</li> <li>• 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.</li> <li>• 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 &amp; 7.</li> </ul>



Table 3.4-3. Routing Policy Requirements for a Mobile RD (ATN ICS SARPs Ref. 5.3.7.2)

ATN ICS SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
5.3.7.2.1	<b>Mobile RD</b>	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 3.4-4. Routing Policy Requirements for an ATN TRD that is not a member of the ATN Island Backbone RDC (ATN ICS SARPs Ref. 5.3.7.3)

ATN ICS SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
5.3.7.3.2	<b>Adjacent ATN RDs that are members of the ATN Island's Backbone RDC</b>	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: <ul style="list-style-type: none"> <li>● a Route to NSAPs &amp; NETs contained within RD9;</li> <li>● 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.</li> <li>● the selected Route to every Fixed ATN RD in the Airline Island i.e. a Route to RD10;</li> <li>● a Route to each Home that the TRD itself (i.e. RD9) provides for Mobile RDs (e.g. for Mobile RD1).</li> </ul>

ATN ICS SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
5.3.7.3.3	<b>Adjacent ATN RDs within the same ATN Island and which are not members of the ATN Island's Backbone RDC</b>	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: <ul style="list-style-type: none"> <li>● a Route to NSAPs and NETs contained within RD9;</li> <li>● the selected Route to every Fixed RD in the Airline Island for which a Route is available i.e. a Route to RD8;</li> <li>● 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;</li> <li>● the preferred Route to all Mobiles i.e. via RD8;</li> <li>● a Route to each Mobile RD for which RD10 is advertising the preferred Route to the Mobile RDs Home (not applicable in this example);</li> <li>● a Route to the Home of all Mobile RDs assigned to Airline A since RD9 is the Home RD for Airline A.</li> </ul>

ATN ICS SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
5.3.7.3.4	<b>Mobile RDs</b>	The policy requirements are applicable to the routes advertised by a Fixed TRD which is not a member of its <u>ATN Islands Backbone</u> to an adjacent Mobile RD.	RD2→RD1	<p>In this case RD2 will advertise to the Mobile RD1:</p> <ul style="list-style-type: none"> <li>● a Route to NSAPs &amp; NETs contained within RD2.</li> </ul> <p>The ATN ICS SARPs additionally recommend that RD2 should advertise to the Mobile RD1:</p> <ul style="list-style-type: none"> <li>● an aggregated Route to NSAPs &amp; NETs contained within the State A RDC (i.e. the local RDC) and;</li> <li>● an aggregated Route to NSAPs &amp; 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).</li> </ul>

**Table 3.4-5. The Routing Policy for a Fixed ATN ERD (ATN ICS SARPs Ref. 5.3.7.4)**

ATN ICS SARPs Reference	Category	Description	Applicable RDs from Scenario	Example
5.3.7.4.1	<b>Fixed ATN ERD</b>	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 3.4-6. 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
<b>RD1 Mobile RD</b>			5.3.7.3.4	5.3.7.1.4							
<b>RD2 Fixed TRD off Backbone</b>	5.3.7.2.1			5.3.7.1.3							
<b>RD3 Fixed RD on Backbone</b>	5.3.7.2.1	5.3.7.3.2			5.3.7.1.2						
<b>RD4 Fixed RD on Backbone</b>				5.3.7.1.2		5.3.7.1.5		5.3.7.1.5	5.3.7.1.5		
<b>RD5 Fixed RD on Backbone</b>					5.3.7.1.5		5.3.7.1.5	5.3.7.1.5	5.3.7.1.5		
<b>RD6 Fixed TRD off Backbone</b>								5.3.7.1.3			
<b>RD7 Fixed RD on Backbone</b>					5.3.7.1.5	5.3.7.1.5	5.3.7.3.2		5.3.7.1.5		
<b>RD8 Fixed RD on Backbone</b>					5.3.7.1.5	5.3.7.1.5		5.3.7.1.5		5.3.7.3.2	
<b>RD9 Fixed TRD off Backbone</b>									5.3.7.1.3		5.3.7.4.1

RD10 Fixed ERD off Backbone									5.
-----------------------------------	--	--	--	--	--	--	--	--	----

### 3.5 Congestion Avoidance in the ATN Internetwork

#### 3.5.1 Network Congestion

3.5.1.1 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.

3.5.1.2 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.

3.5.1.3 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.

#### 3.5.2 Possible Techniques

3.5.2.1 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.

3.5.2.2 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:

- a) high throughput (in bit/s), together with a small end-to-end transit delay, should be experienced by network users;
- b) a small buffer load within the traversed routers should be achieved; and

c) the probability of packet loss should be minimal;

3.5.2.3 In pursuit of these goals, two candidate algorithms were initially investigated during the development of the ATN ICS 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.

3.5.2.4 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.

3.5.2.5 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.

### 3.5.3 Receiving Transport Layer Congestion Avoidance

#### 3.5.3.1 Overview

3.5.3.1.1 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.

3.5.3.1.2 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.

3.5.3.1.3 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!

3.5.3.1.4 When the receiving transport entity gets an NPDUNSDU 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 NPDUNSDUs 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 ~~NPDU~~NSDUs 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.

3.5.3.1.5 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  $n$  TPDU. Normally, the objective is to acknowledge received DT TPDU 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 TPDU (i.e. that it runs out of credit). The sender is then able to transmit data as fast as it can.

3.5.3.1.6 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. ~~T~~there will be times when it cannot send any DT TPDU, 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  $n$ .

3.5.3.1.7 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. The mentioned ~~Our~~ goals for a Congestion Avoidance Algorithm are thereby met.

### 3.5.3.2 **Determining the Onset of Congestion**

3.5.3.2.1 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<sub>2</sub> then the CE-bit should be set to indicate that the queuing delay is excessive<sub>2</sub> 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?


3.5.3.2.2 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.

- 3.5.3.2.3 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

$$\text{Power} := \frac{\text{Throughput}}{\text{Delay}}$$

- 3.5.3.2.4 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.

- 3.5.3.2.5 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).



**Figure 3.5-1. A Queuing System**

- 3.5.3.2.6 Packets arrive at a server with arrival rate  $\lambda$ , 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  $\mu$ . 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  $\lambda < \mu$ .

- 3.5.3.2.7 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

$$E(T) = \frac{1}{\mu - \lambda} \quad (1)$$

- 3.5.3.2.8 The throughput of an M/M/1 system is equal to  $\lambda$ , if the system is operating in a stable state (i.e. one can never receive a higher throughput than the server forwarding rate  $\mu$ , but the server is also not able to forward packets faster than they arrive).

- 3.5.3.2.9 From the above, the Power of a M/M/1 system thus can be derived to be:

$$\text{Power} := \frac{\text{Throughput}}{\text{Delay}} = \frac{\lambda}{1 / (\mu - \lambda)} = \lambda \cdot (\mu - \lambda) \quad (2)$$

3.5.3.2.10 This measure is maximised if the following condition holds:

$$\lambda = \frac{\mu}{2} \quad (3)$$

3.5.3.2.11 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  $\lambda$  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.

3.5.3.2.12 Although it has been suggested that a number greater than one is appropriate for 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.).

3.5.3.2.13 However, this is not to say that special considerations do not apply to air/ground data links. The queuing model 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 communications 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.

### 3.5.3.3 **Reporting Congestion Experienced to the NS-User**

3.5.3.3.1 Congestion is experienced by an NPDUs, 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.

3.5.3.3.2 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:


- a) to indicate to the NS-User both the total number of NPDU's received for a single NSDU, and the number of NPDU's received having the CE flag set to the transport layer;
- b) 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;
- c) to merge the CE flags received by bitwise ANDing all values. Thus, only if all NPDU's had the CE flag set, congestion will be indicated to the NS-User; or
- d) to only forward the CE flag setting of the last NPDU received during reassembly of an NSDU to the local transport layer.

Strategy (a) 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.

### 3.5.3.4 Credit Window Management by the Receiving Transport Entity

- 3.5.3.4.1 The receiving transport entity monitors incoming TPDU's 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 TPDU's, then the effective credit window is reduced by multiplying it by a reduction factor  $\beta$ . Otherwise, if the credit window is currently less than the  $\alpha$  value which will permit maximum throughput, then it may be increased by adding an integral value  $\delta$ . 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 TPDU's for each DT TPDU transferred.
- 3.5.3.4.2 Only DT TPDU's are monitored during a sampling period. This is because only DT TPDU's are subject to credit management. Other TPDU's 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.
- 3.5.3.4.3 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 TPDU's 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 TPDU's are subject to the new credit regime and hence its effect on the network state can be reasonably determined.
- 3.5.3.4.4 The reason why a "freeze period" is necessary can be readily seen from the following example.

- 3.5.3.4.5 Figure 3.5-23.4-2 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).
- 3.5.3.4.6 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.
- 3.5.3.4.7 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 Figure 3.5-2).
- 3.5.3.4.8 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  $\beta = \frac{1}{2}$ , to better visualise the operation. Thus, the window  $W$  is reduced to  $W_1 = 8 \times \frac{1}{2} = 4$ . Afterwards, the sender is transmitting with a smaller load, indicated by a greater spacing of the packets.
- 3.5.3.4.9 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.
- 3.5.3.4.10 During the next round trip time (starting at time  $t_2$ ), DT TPDU's 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 3.5-2. Window Adaptation over Time**

### 3.5.3.5 *The Congestion Avoidance Algorithm*

3.5.3.5.1 In the ATN ICS 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.

3.5.3.5.2 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.*— A new connection starts advertising an initial window size  $W_0$  (as defined in section 5.5.2.5.2 of the ATN ICS SARPs text 5.2.6.3) to its peer. This is reflected in the initialization of variable 'w\_new'.

- 3.5.3.5.3 Whenever a ~~NSDU~~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 NPDU's forming that TPDU
// nCE         -      number of NPDU's forming that TPDU that had their CE
                    flag set on
//
                    reception
if (dt_TPDU) n_DT++; //      count total # DT-TPDU's 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-TPDU's, 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 ‘W<sub>0</sub>’ are parameters which are defined in section 5.5.2.5.4 of the ATN ICS SARP's text; see section 5.5.2.5.4: “Recommended algorithm values”.

### 3.5.3.6 ***Sending Transport Entity Procedures***

- 3.5.3.6.1 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 TPDU's and the utilisation of the received credit window.

3.5.3.6.2 ~~However~~<sup>4 December 1998</sup>However, implementors should note that commercial implementations of the transport protocol may often include “transport layer backoff” procedures similar to the ~~Van Jacobson~~<sup>van Jacobson</sup> 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.

### 3.5.3.7 Known Limitations

#### 3.5.3.7.1 Fairness

3.5.3.7.1.1 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 Figure 3.5-3.

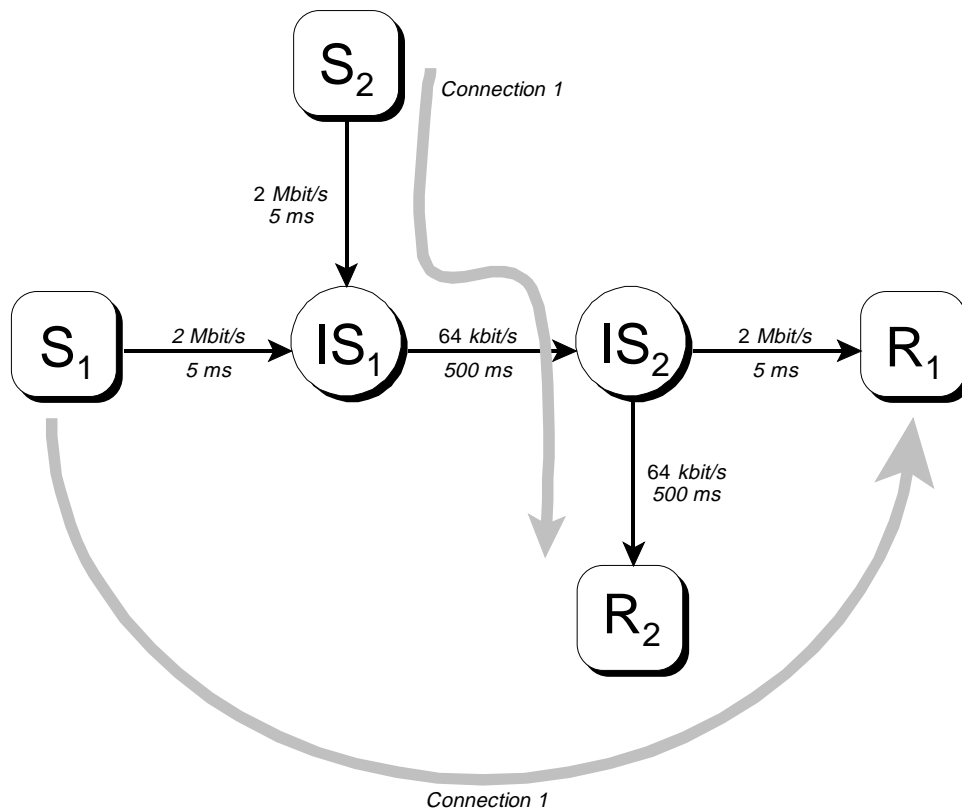


Figure 3.5-3. Fairness Among Competing Users



3.5.3.7.1.2 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 experience a lower throughput than another with the same credit window and a shorter Round Trip Time.

3.5.3.7.1.3 It may be possible to balance throughput by varying the value of  $\beta$  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.

### 3.5.3.7.2 Two-Way Traffic

3.5.3.7.2.1 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).

3.5.3.7.2.2 Figure 3.5-4 depicts this scenario.

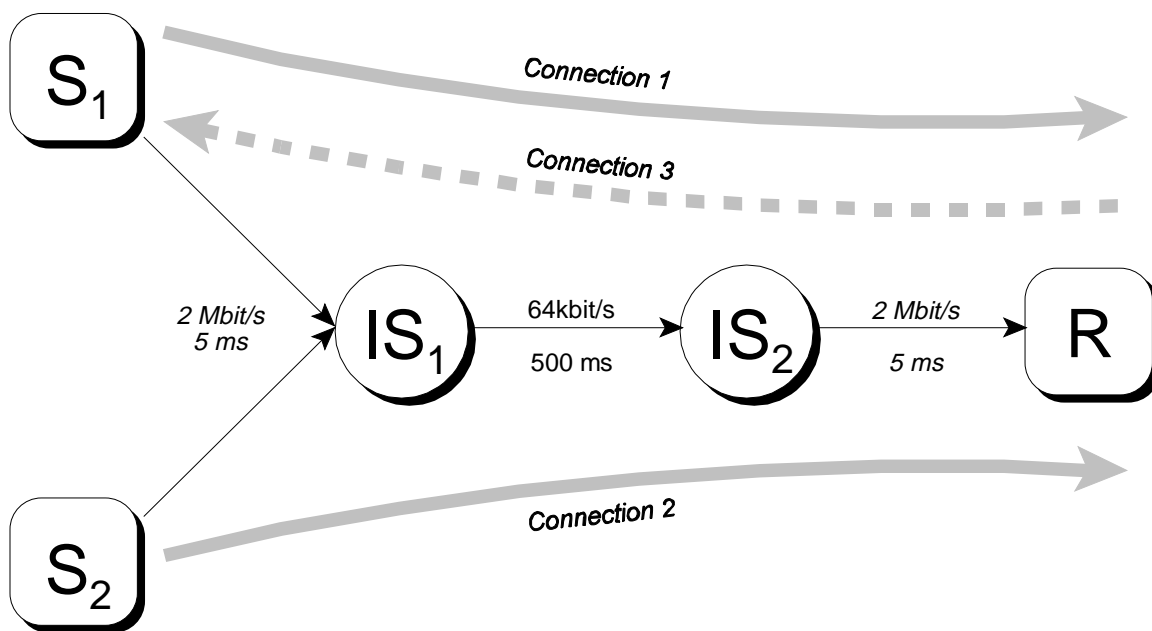


Figure 3.5-4. Two-Way Traffic

### 3.5.3.7.3 **A-Credit Window of One is the Minimum**

3.5.3.7.3.1 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.

### 3.5.3.8 **Conclusion**

3.5.3.8.1 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 importance 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.

## 3.6 **ATN Subnetworks**

### 3.6.1 **Introduction**

3.6.1.1 The ATN ICS SARPs specify requirements for the Subnetwork Dependent Convergence Function (SND CF) and require that Subnetwork (SN)-Service (SNS) primitives or equivalent mechanisms be provided.

3.6.1.2 This chapter provides guidance on the necessary features of the SND CF to support the ISO/IEC 8473+ Connectionless Network Layer Protocol (CLNP) over these various subnetworks. First it describes the ATN requirements which are common to all subnetworks and thereafter, it is further broken down into mobile (air-ground) and ground subnetworks. The list of subnetworks is not exhaustive, and future subnetworks may well be capable of serving as ATN subnetworks.

### 3.6.2 **General Characteristics of ATN-suitable Subnetworks**

3.6.2.1 It is true to say that almost any data communications network can be an ATN Subnetwork. The ATN is based on internetworking techniques — it is an internetwork — and has the facilities to integrate both existing and new networking technologies.

3.6.2.2 The minimum requirements for an ATN Subnetwork are that:

- a) it supports packet mode communications;
- b) except for point-to-point data links, each system attached to the network must be individually addressable; and

- c) it must provide subnetwork connectivity information to attached ATN Routers in the form of “Join Events” and “Leave Events” respectively; and
- d) it must support the transparent communications of octet aligned data, i.e. there must be no “special” characters that are interpreted by the underlying network.

In practice only the first ~~three~~<sup>two</sup> requirements are absolute, the ~~fourth~~<sup>third</sup> can always be handled by a special adaptation layer.

3.6.2.3 Other features may also be regarded as desirable for ATN Subnetworks. For example, subnetwork prioritisation of data may be important when the network will share both safety related and routine data transfer between different pairs of communicating systems. Some networks can prioritise virtual circuits on an a priori basis prioritising the data of different users, while others can support this on a per packet or a per virtual circuit basis.

### 3.6.3 Subnetwork Adaptation for the ATN

3.6.3.1 The SNDCF has already been introduced in section ~~3.3.4.6.34.5~~, and an SNDCF is essentially an adaptation layer that interfaces the service provided by an individual subnetwork to the (SN)-Service expected by CLNP. Provided a suitable SNDCF exists for a given subnetwork and is implemented in both systems (End Systems or Routers) that will use the subnetwork to communicate, then that subnetwork can be used as an ATN Subnetwork.

3.6.3.2 For Ground Subnetworks, ISO/IEC 8473 already defines SNDCFs for most common subnetwork types. Industry standards exist for other common requirements (e.g. encapsulation of CLNP over IP). For CIDIN, the ATN ICS SARPs define a special purpose SNDCF.

3.6.3.3 For Air/Ground Subnetworks, the ATN ICS SARPs have specified ~~two~~<sup>a</sup> special purpose SNDCFs, called the Mobile SNDCF and the Frame-Mode SNDCF respectively. ~~This~~<sup>The</sup> Mobile SNDCF is based upon the SNDCF for ISO/IEC 8208 subnetworks specified in ISO/IEC 8473-3 and additionally supports:

- a) identification of the NETs of the communicating systems by ISH PDUs included in call setup user data; and
- b) the negotiation and use of data compression procedures.

3.6.3.4 The data compression procedures specified include:

- a) mandatory Local Reference (LREF) compression for the compression of CLNP Header information (see ~~3.3.4.6.344.5.2.1~~); and
- b) optional data stream mode compression using the Deflate algorithm.

~~c) optional Address Compression - a compression algorithm that searches for bytes sequences that appear to be ICAO NSAP Addresses and which then removes redundant information in those addresses; and~~

~~d) optional ITU-T V.42bis style Data Compression.~~

3.6.3.5 ~~Air/Ground~~ Mobile Subnetworks also require specific local management procedures for reacting to the presence of a new mobile system on the subnetwork (the Join Event), and reacting to a mobile system leaving the subnetwork (the Lease Event). Formally, these procedures are not part of the Mobile SNDCF, although they are supported by the service provided by the Mobile SNDCF (see 3.3.4.6.35-10.3).

### 3.6.4 **Air/Ground Subnetworks**

The following sections briefly summarise the features of individual mobile subnetworks. Where necessary reference is made to the appropriate ICAO Annex 10 material.

#### 3.6.4.1 ***VDL Mode 1 and Mode 2 Subnetworks***

##### 3.6.4.1.1 **Introduction**

3.6.4.1.1.1 The VHF digital link (VDL) is a mobile subnetwork of the ATN, operating in the VHF aeronautical mobile frequency band. This subnetwork can be run in either one of the following four~~two~~ modes:

- a) **Mode 1.** A minimum shift keying modulation scheme;~~and~~
- b) **Mode 2.** A differentially encoded phase shift keying modulation scheme;
- c) **Mode 3.**
- d) **Mode 4.**

~~Note. — As of 1997, further modes of operation may be defined.~~

3.6.4.1.1.2 The VDL airborne unit uses a VHF Data Radio (VDR) which communicates with a Remote Ground Station (RGS) interconnected via Ground WANs with Air/Ground ATN routers. An RGS is a ground station equipment with radio antenna and WAN access capabilities.

##### 3.6.4.1.2 **General VDL Mode 1 and Mode 2 Characteristics**

###### 3.6.4.1.2.1 **Route Initiation**

3.6.4.1.2.1.1 In the Route Initiation phase, the Air/Ground Router learns of the existence of an Airborne Router which can be accessed through the VDL Subnetwork, and the Airborne Router learns of Air/Ground ATN Routers which are able to forward air-initiated traffic.

3.6.4.1.2.1.2 When entering the coverage of an RGS, the mobile ATN router learns of Air/Ground ATN Routers interconnected with the RGS when the RGS transmits the DTE Address of each Air/Ground router via the Ground Station Information (GSIF) or Link Establishment XIDs (Exchange ID) frames. The receipt of such information is formally described as a Jjoin Eevent.

3.6.4.1.2.1.3 Once the Jjoin Eevent is received, the Airborne Router proceeds with the Route Initiation procedures described in 3.4.10.3, following the “air-initiated” model, and using the DTE Address(es) received in the Jjoin Eevent to determine with which Air/Ground Routers, calls are initiated. Note that if more than one DTE Address is received, then it is a local policy decision as to which one to use, or whether to attempt multiple simultaneous connections.

### 3.6.4.1.2.2 **VDL Handovers**

3.6.4.1.2.2.1 With VDL, the subnetwork connections between an Airborne Router and an Air/Ground Router are typically short lived. This is because they last only as long as an aircraft is within range of the RGS that supports the connection. It would prove to be very costly if the Route Initiation procedures were invoked every time the RGS changed, and instead a VDL Handover procedure is used to avoid this overhead.

3.6.4.1.2.2.2 Firstly, it is considered good practice that an Air/Ground Router is attached to many RGSs within the same region. Thus even when the aircraft moves out of range of one RGS and into the area of coverage of another, it can still remain in contact with the same Air/Ground Router. Although a new virtual circuit has to be established, this is transparent to IDRP and there is no need to exchange additional routing information.

3.6.4.1.2.2.3 Secondly, when a new virtual circuit is established between an Airborne Router and an Air/Ground Router a procedure known as the M/I bit procedure is used to signal that the LREF compression directory is to be shared with the existing virtual circuit. This further ensures that there is no need to re-create the directory used for CLNP header compression.

3.6.4.1.2.2.4 Thus as an aircraft moves between the RGSs of a given Air/Ground Router, the only impact is the establishment of a new virtual circuit and the addition of this new route to the local Forwarding Table. Otherwise, the data compression context is preserved and no new route initiation is required.

3.6.4.1.2.2.5 Whenever multiple virtual circuits exist between the Air/Ground Router and a given aircraft, through the VDL subnetwork, then the last established call is the *active* virtual circuit. Through the active virtual circuit, ISO/IEC 8473 NPDU's are sent and received. Through the remaining connections, with the same aircraft, ISO/IEC 8473 NPDU's are never sent, once a new active connection is established but can still be received.

### 3.6.4.1.2.3 **Route Termination**

3.6.4.1.2.3.1 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 virtual circuits, active or not, between the Air/Ground Router and a given aircraft via that RGS are cleared.

- 3.6.4.1.2.3.2 When the loss of all VDL virtual circuits with a given aircraft is detected by the Air/Ground Router, it then activates the Route Termination phase.
- 3.6.4.1.2.3.3 The Route Termination phase consists of the issue of a Leave Event to local management functions, which results in appropriate updates to the routing tables in the Air/Ground Router reflecting the loss of the route via VDL.
- 3.6.4.1.3 **Use of X.25 Facilities**
- 3.6.4.1.3.1 **Fast Select with no Restriction on Response**
- 3.6.4.1.3.1.1 The X.25 facility “Fast Select with no restriction on response” is used, in Route initiation, to allow a responder to pass user data information in the Call Confirm packet.
- 3.6.4.1.3.2 **Priority**
- 3.6.4.1.3.2.1 The VDL Mode 1 and 2 have no X.25 priority capability.
- 3.6.4.1.4 **Use of Compression Algorithm**
- 3.6.4.1.4.1 **LREF Compression**
- 3.6.4.1.4.1.1 **Use of the M/I Bit Procedure**
- 3.6.4.1.4.1.1.1 Use of the LREF compression algorithm is mandated by the ATN ICS SARPs for all air/ground subnetworks. In the case of VDL, the M/I bit management procedure described in the ATN ICS SARPs must be used during VDL Handover.
- 3.6.4.1.4.1.2 **LREF and Handovers Procedure**
- 3.6.4.1.4.1.2.1 With regard to the specification provided by the ATN ICS SARPs, it may happen that, during a VDL handover, 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 are 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 is generated to the sending SNDCF.
- 3.6.4.1.4.1.2.2 In this case the transport protocol will eventually re-transmit the packet.
- 3.6.4.1.4.1.3 **LREF and Call Clearing**
- 3.6.4.1.4.1.3.1 When the virtual circuit has been cleared for other reasons than handover and the LREF compression was used for the cleared virtual circuit, the internal resources used to handle the Local Reference Directory are released.

3.6.4.1.4.1.3.2 When the virtual circuit has been cleared due to handover and the M/I bit was not set or refused for the newly established associated virtual circuit, the Local Reference Directory is released.

3.6.4.1.4.1.3.3 In all other circumstances the Local Reference Directory is maintained.

## 3.6.4.2 **AMSS Subnetwork**

### 3.6.4.2.1 **Introduction**

3.6.4.2.1.1 The AMSS satellite subnetwork components are briefly summarised below:

**AES:** Aircraft Earth Station, encompasses airborne equipment from the Satellite Data Unit (SDU), High Power Amplifier (HPA), Radio Frequency Unit (RFU), Antenna.

**GES:** Ground Earth Station, encompasses ground equipment for the satellite subnetwork interface with a ground WAN.

3.6.4.2.1.2 Two types of architecture currently exist to access a mobile system from the DTE using satellite links; the DATA-2 mode, and the DATA-3 mode.

3.6.4.2.1.3 DATA-2 is an implementation of the link layer; (OSI layer 2) with non-standard access and relay layers. DATA-3 communications is an evolution of DATA-2, integrating the OSI standards for layer 3; (ISO/IEC 8208), with routing and relay functions. As the ATN ICS SARPs recommend require the use of an ISO/IEC 8208 mobile subnetwork, the AMSS mode to be used in conjunction with ATN is the DATA-3 mode.

3.6.4.2.1.4 The GES DATA-3 Interworking Function (IWF) is defined in the AMSS SARPs. This provides a mapping between the ISO/IEC 8208 compatible protocol elements used on the Air/Ground Data Link with the ground X.25 network. The IWF implements a set of minimal relay functions to support the operation of the ATN mobile SNDCEF, including fast select facility and the management of priorities.

3.6.4.2.1.5 Not all ISO/IEC 8208 facilities can be used for the ATN access to the satellite subnetwork. Restrictions may arise depending on the implementation of the IWF function, and the ground X.25 network. As of ~~2000~~1996, for example, some Public X.25 Service Providers are unable to support Fast Select or priority in their X.25 data networks and such facilities may therefore not be available between an Air/Ground Router and an Airborne Router using AMSS.

### 3.6.4.2.2 **General Characteristics of the AMSS Subnetwork**

#### 3.6.4.2.2.1 **Route Initiation**

3.6.4.2.2.1.1 In the Route Initiation phase, the Air/Ground Router learns of the existence of an Airborne Router which can be accessed through the Satellite Subnetwork, and the Airborne Router similarly learns about Air/Ground ATN Routers. This information is reported as a Join Event.

3.6.4.2.2.1.2 Once the Join Event is received, the Airborne Router proceeds with the Route Initiation procedures described in 3.4.10.3, following the “air-initiated” model, and using the DTE Address(es) received in the Join Event to determine with which Air/Ground Routers calls are initiated. Note that if more than one DTE Address is received, then it is a local policy decision as to which one to use, or whether to attempt multiple simultaneous connections.

*Note.— As of ~~2000~~1996, GESs are unable to provide aircraft with the list of ATN Air/Ground routers to which they are interconnected~~to~~. The airborne router has to maintain a table indicating, for each GES, the list of Air/Ground routers with which a virtual circuit can be established.*

### 3.6.4.2.2.1.3 Procedure for the Establishment of Connections

3.6.4.2.2.1.3.1 Following successful logon to a GES, the SDU provides a Join Event with the GES identification to the airborne router, and this router will then try to establish one virtual circuit with one or several ATN Air/Ground routers associated with the GES.

3.6.4.2.2.1.3.2 AMSS is air-initiated, and it is the responsibility of the Airborne Router to establish the first connection with an Air/Ground Router. The Airborne Router periodically tries to initiate a call with each known Air/Ground Router’s DTE Address, and, if Fast Select is available, includes an ISH PDU in the call user data in order to identify its own NET to the Air/Ground Router. Currently, the DTE addresses of reachable Air/Ground Routers via a given GES are fixed addresses, pre-defined in a table.

3.6.4.2.2.1.3.3 Once a connection has been established with one Air/Ground Router, it is a local policy matter as to whether the Airborne Router attempts to establish further connections with other Air/Ground Routers.

### 3.6.4.2.2.2 Route Termination

3.6.4.2.2.2.1 Connection loss can be due to:

- a) the aircraft leaving the satellite coverage area;
- b) the handover of the AES connection to another GES;
- c) a ground communications subnetwork system management procedure;
- d) air/ground Router system management procedure; and
- e) the expiration of the ISO/IEC 8208 mobile SNDCF idle timer.

3.6.4.2.2.2.2 When a disconnection takes place, the GES will clear all terrestrial connections supporting the exchange of traffic with this aircraft.

*Note.— As of ~~2000~~1996, no specific cause and diagnostic have been defined for the clearing of a virtual circuit by the GES due to the loss of the Satellite media. Therefore,*



*upon detection of the loss of all connections with an airborne SND CF, the Air/Ground router mobile SND CF enters the Route Termination phase.*

3.6.4.2.2.2.3 This Route Termination Phase causes a Leave event to be passed to local management and the airborne router's routing tables to be updated to reflect the loss of the AMSS route to the aircraft.

3.6.4.2.2.3 **Leaving / Entering GES Coverage**

3.6.4.2.2.3.1 The AMSS SARPs constrain the SDU to be connected to only one GES at any time. Therefore when an aircraft leaves the coverage of one GES, the connectivity is interrupted before the aircraft establishes a new virtual circuit to the same or another Air/Ground router via another GES.

3.6.4.2.2.4 **ISO/IEC 8208 Services Supported by the IWF**

3.6.4.2.2.4.1 For the definition of the ATN convergence function over ISO/IEC 8208 it is important to note that the IWF function implements the following services:

- a) connection establishment/release;
- b) extended address management;
- c) data transfer and expedited data transfer;
- d) receipt confirmation;
- e) interrupt;
- f) reset;
- g) transparent mapping of the QoS;
  - 1) throughput class negotiation;
  - 2) minimum throughput class negotiation;
  - 3) subnetwork transit delay negotiation;
  - 4) end to End transit delay negotiation;
- h) transparent mapping of cause/diagnostic codes;
- i) fast select facility; and
- j) priority.

### 3.6.4.2.2.5 **Use of X.25 Facilities**

#### 3.6.4.2.2.5.1 **Fast Select with no Restriction on Response**

3.6.4.2.2.5.1.1 The use of Fast Select is not mandatory for the AMSS subnetwork but it should be used (with no restriction on response) if the GES IWF and ground X.25 network support Fast Select. As Route initiation is performed by the airborne systems, the airborne router has to maintain, in the Air/Ground Router address table, information indicating whether Fast Select calls can be used for the corresponding Air/Ground Router and GES.

#### 3.6.4.2.2.5.2 **AMSS Priority**

3.6.4.2.2.5.2.1 As satellite communications services are used by applications other than ATN applications, the use of AMSS subnetwork priorities is critical to ensure that the traffic relating to ATM safety applications is not delayed by non-ATN applications traffic. Ten AMSS priority levels are available; passenger communications use priorities 0 to 3.

3.6.4.2.2.5.2.2 It should be noted that each virtual circuit has a different priority so there are as many open virtual circuits as X.25 priorities in use between a mobile SNDCF and a remote peer. This means that the mobile SNDCF must be able to handle ten virtual circuits for each remote peer when used over the AMSS subnetwork.

### 3.6.4.2.3 **Compression Procedures**

3.6.4.2.3.1 Use of the LREF compression procedures is mandated by the ATN ICS SARPs. Use of the M/I bit management procedure is not required as the transition from one GES to another will always cause the loss of all virtual circuits before new ones can be established.

## 3.6.4.3 **Mode S**

### 3.6.4.3.1 **Introduction**

3.6.4.3.1.1 The Mode S air/ground subnetwork is a mobile subnetwork of the ATN. It is an extension of the SSR Mode S surveillance system providing air/ground data communications through a connection-oriented communications 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/IEC 8208, and is entirely conformant with the ATN architecture. In addition to its function as an ATN subnetwork, Mode S offers specific services.

3.6.4.3.1.2 It is likely that a number of Mode S interrogators are connected to a single Ground Data Link Processor (GDLP), thus extending the coverage of the subnetwork. Unlike VDL or AMSS, transfer from one Mode S interrogator to another is handled by the subnetwork and thus completely invisible to the Internetwork.

### 3.6.4.3.2 **General Mode S Characteristics**

#### 3.6.4.3.2.1 **Route Initiation**

3.6.4.3.2.1.1 When an aircraft enters the coverage of a Mode S subnetwork, a “Join Event” is generated. The “Join Event” is always generated by the ground part of the Mode S subnetwork, i.e. the GDLP.

3.6.4.3.2.1.2 Similarly, “Leave Events” are generated by both the aircraft side (ADLP) and the ground side as soon as the aircraft leaves the coverage of a particular cluster of Mode S interrogators. In addition, refresh cycles may be performed.

3.6.4.3.2.1.3 “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.

#### 3.6.4.3.2.2 **Route Termination**

3.6.4.3.2.2.1 When none of the interrogators connected to a GDLP has a particular aircraft in its coverage, the GDLP activates the Route Termination phase by sending a Leave Event to local management functions, which result in appropriate updates to the routing tables of the air/ground router reflecting the loss of that particular route via Mode S.

### 3.6.4.3.3 **Use of X.25 Facilities**

3.6.4.3.3.1 In accordance with ISO/IEC 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, 1024, 2048 and 4096 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.

#### 3.6.4.3.4 **Fast Select with no Restriction on Response**

3.6.4.3.4.1 The Mode S Subnetwork provides the X.25 “Fast Select” capability. This facility is to be used typically during route initiation.

#### 3.6.4.3.5 **Priority**

3.6.4.3.5.1 The Mode S subnetwork provides means for distinguishing two priorities ('high' and 'low').

### 3.6.4.3.6 Use of Compression Algorithms

3.6.4.3.6.1 Use of the LREF compression procedures is mandated by the ATN ICS SARPs. Use of the M/I bit management procedure is not required for the Mode S subnetwork.

## 3.6.5 Ground/Ground Subnetworks

3.6.5.1 This section presents guidance to States and Organisations wanting to implement new or already existing networks as ATN subnetworks inside their respective boundaries. Figure 3.6-1 shows various ground-ground networking technologies that may be used as subnetworks to support the ATN Internet Communications Service. Some of these technologies may also be applicable within an aircraft.

### 3.6.5.2 Subnetwork Addressing

3.6.5.2.1 A subnetwork point-of-attachment (SNPA) address is needed for each point of attachment between an End System or an Intermediate System and a subnetwork.

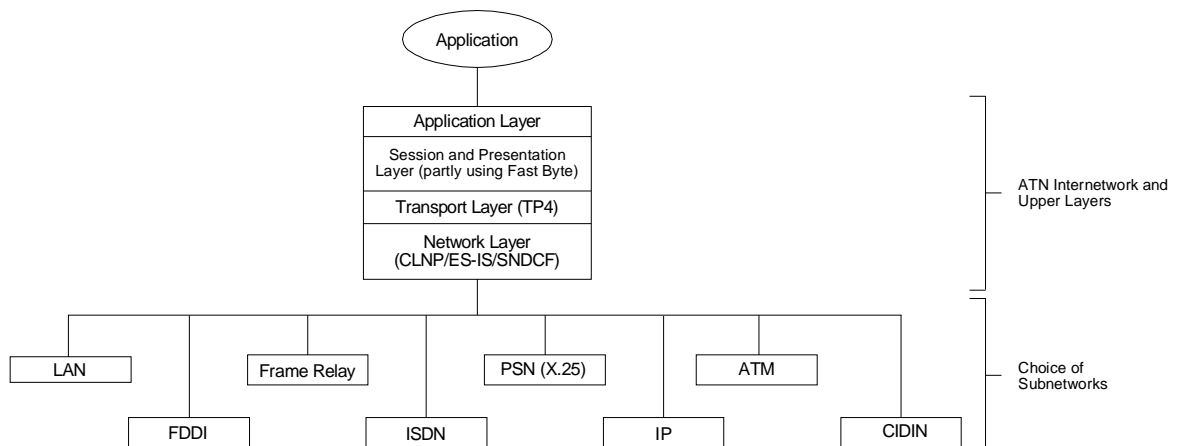


Figure 3.6-1. ATN Use of Various Ground-g/Ground Ssubnetworks

- 3.6.5.2.2 The routing function of the Network layer manages the correspondence between NSAP addresses and SNPA addresses, which may be complex. There is no need for an NSAP address to incorporate a corresponding SNPA address, although this may facilitate routing. The use of the SYS field in the ATN NSAP address structure for this purpose is specified in section 5.4.3.8.6 of the ATN ICS SARPs.
- 3.6.5.2.3 Guidance on the use of the ISO/IEC 9542 ES-IS routing protocol over ground-ground subnetworks is given in section 3.3.4.24.4.2.2.
- 3.6.5.3 ***Mapping CLNP over an ISO/IEC 8802 Subnetwork***
- 3.6.5.3.1 The ATN ICS SARPs (section 5.7.3) specify that the subnetwork service over ISO/IEC LANs is provided as specified in 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 SNDCEF. Each such request is mapped in turn to an ISO/IEC 8802-2 Logical Link Control Type 1 Unnumbered Information frame, with the DSAP and the SSAP fields set to the assigned Service Access Point hexadecimal value for CLNP [0FE].
- 3.6.5.3.2 The SNDCEF provisions for ISO/IEC 8802 LANs apply equally to the Fibre Distributed Data Interface (ISO/IEC 9314) and other LANs that support the ISO/IEC 8802-2 Logical Link Control service.
- 3.6.5.3.3 **ISO/IEC 8802 LAN Addressing**
- 3.6.5.3.3.1 The structure of Local Area Network (LAN) subnetwork addresses is defined in the ISO/IEC 8802 series of standards (including the associated Technical Report series ISO/IEC 11802), and applies to FDDI (ISO/IEC 9314) in addition to the ISO/IEC 8802 LAN types. There are address parameters in both the Logical Link Control (LLC) and the Medium Access Control (MAC) service.
- 3.6.5.3.3.2 LLC addresses have a small number of fixed values.
- 3.6.5.3.3.3 MAC addresses have to be unique within each extended LAN (~~ie~~, i.e. a group of LANs connected by MAC bridges), and one is required for each LAN SNPA. System configuration becomes easier if MAC addresses are in fact globally unique; in practice this is not a major issue because LAN interfaces are supplied with globally unique addresses, allocated originally by an agreement between the manufacturers and now administered by the IEEE as the International Registration Authority for ISO/IEC 8802.
- 3.6.5.3.3.4 The ISO/IEC 9542 ES-IS protocol supports the selection of the appropriate MAC address for each SN-UNITDATA transmission.

### 3.6.5.4 **Mapping CLNP over a Frame Relay Network**

3.6.5.4.1 The Frame Relay access protocol is based on High-level Data Link Control (HDLC/Q.921), and the link access protocol was developed for signalling over the D<sub>2</sub> channel of narrow-band Integrated Services Digital Network (ISDN) (ITU-T Recommendation Q.922). 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.

3.6.5.4.2 The ATN ICS SARPs do not specify the SNDCF for use over a Frame Relay network. ISO/IEC 8473-27 (*Information Technology — Protocol for Providing the Connectionless-Mode Network Service — Part 7: Provision of the Underlying Service by Frame Relay Subnetworks*) provides an appropriate specification.

*Note.— As of 1996, this specification was at Committee Draft status. States and Organizations who wish to make use of draft versions of ISO documents (e.g., for trial implementation) are advised to contact the relevant national ISO member body.*

### 3.6.5.4.3 **Frame Relay Subnetwork Addressing**

3.6.5.4.3.1 Frame relay uses the same address formats as ISDN. See section 7.4.4.13.6.5.5.3.

### 3.6.5.5 **Mapping CLNP over ISDN**

3.6.5.5.1 Where an ISDN service is available, a dynamically established ISDN connection may provide a useful backup to a permanent X.25 wide area link.

3.6.5.5.2 The ATN ICS SARPs do not specify the SNDCF for use over an ISDN network. ISO/IEC 8473-5 (*Information Technology — Protocol for Providing the Connectionless-Mode Network Service — Part 5: Provision of the Underlying Service for Operation over ISDN Circuit-switched B-channels*) provides an appropriate specification.

*Note.— As of 1996, this specification was at Committee Draft status. States and Organizations who wish to make use of draft versions of ISO documents (e.g., for trial implementation) are advised to contact the relevant national ISO member body.*

### 3.6.5.5.3 **ISDN Subnetwork Addressing**

3.6.5.5.3.1 The structure of addresses for use with public ISDN subnetworks is defined in ITU-T Recommendation E.164. There is little practical experience with OSI networking over ISDN, and further specification may be needed.

### 3.6.5.6 **Mapping CLNP over an ISO/IEC 8208 Network**

3.6.5.6.1 The ATN ICS SARPs (section 5.7.5) specify that the subnetwork service over ground-ground subnetworks using ISO/IEC 8208 is provided as specified in ISO/IEC 8473-3 (*Information Technology — Protocol for Providing the Connectionless-Mode Network Service — Part 3: Provision of the Underlying Service by ISO 8208 Subnetworks*).

3.6.5.6.2 Management of virtual circuits established to support the SNDCF is discussed in detail in ISO/IEC 8473-3.

### 3.6.5.6.3 **ISO/IEC 8208 Subnetwork Addressing**

3.6.5.6.3.1 The structure of SNPA addresses for use in access via ISO/IEC 8208 to public packet-switched data networks is defined in ITU-T Recommendation X.121. Address formats for private packet-switched data networks are a matter for the network operator but are generally based on the specification of X.121. One SNPA address is needed for each End System or Intermediate System connected to a subnetwork via ISO/IEC 8208.

3.6.5.6.3.2 There is a need for Link layer addresses in the ISO/IEC 7776 protocol, but these have fixed values depending on the DTE/DCE roles of the systems.

### 3.6.5.7 **Mapping CLNP over IP**

#### 3.6.5.7.1 **General**

3.6.5.7.1.1 There are two approaches that will allow the ATN Internet Communications Service to be tunneled across an internetwork using the Internet Protocol (STD0005).

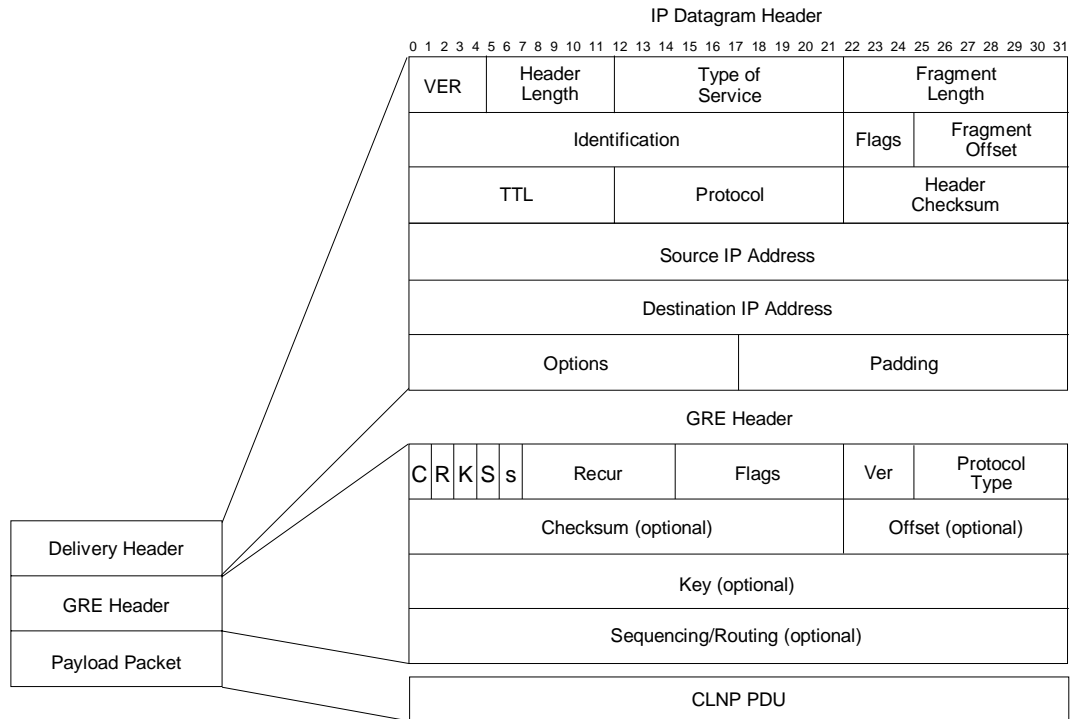
3.6.5.7.1.2 Recent IETF RFCs have been developed to allow the encapsulation of CLNP PDUs over IP. Alternatively, current commercial off-the-shelf (COTS) routers will encapsulate the subnetwork PDUs (for example, X.25 packets) into IP datagrams, and decapsulate the datagrams and forward them to the peer OSI application.

3.6.5.7.1.3 If there is a need for direct encapsulation of CLNP PDUs over IP, then IETF RFCs 1701, 1702 and 1070 define a 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.

*Note.— These RFCs are categorized as ‘Informational’ by the IESG. According to RFC 1602, Informational RFCs are specifications “published for the general information of the Internet community, and [do] not represent an Internet community consensus or recommendation. The Informational designation is intended to provide for the timely publication of a very broad range of responsible informational documents from many*

sources, subject only to editorial considerations and to verification that there has been adequate coordination with the standards process”.

3.6.5.7.1.4 The Delivery Header for IP will consist of the fields shown in Figure 3.6-2.



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 3.6-2. Delivery Header for IP**

3.6.5.7.1.5 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 3.6-1.

3.6.5.7.1.6 In this case, the mapping to the SNS parameters is as follows:

- a) SN-Source-Address: this field should contain a source IP address;
- b) SN-Destination Address: this field should contain a destination IP address;
- c) 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;



- d) SN-Quality-of-Service: if supported, this field should contain the Type of Service value; and
- e) SNS-User Data: this field should contain the CLNP NPDU.

**Table 3.6-1. Example Protocol Type Values**

Protocol Family	Protocol Type Value (Hex)
Reserved	0000
OSI network layer	00FE
IP	800
Frame Relay	0808
Raw Frame Relay	6559
IP Autonomous Systems	876C
Secure Data	876D
Reserved	FFFF

### 3.6.5.7.2 **IP Addressing**

3.6.5.7.2.1 Addressing for networks using the Internet Protocol is specified in STD0005 and various supporting RFCs. ATN NSAP addressing is specified in the ATN ICS SARPs, section 5.4. Although IP addresses may be mapped into NSAP addresses, the reverse is only possible for addresses used with the new IP version 6. For the predominant IP version 4, the incompatible addressing structures must be accommodated using an encapsulation or conversion technique.

### 3.6.5.8 **Mapping CLNP over Asynchronous Transfer Mode (ATM)**

#### 3.6.5.8.1 **General**

3.6.5.8.1.1 The Multiprotocol over ATM (MPOA) specification ~~of currently being developed by the~~ ATM Forum should provide the basis for this specification. However, early implementations could take a similar approach to that developed for the encapsulation of IP over ATM. This is described below.

3.6.5.8.1.2 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 traffic, routed and bridged PDUs, over an ATM network. The first method (called “LLC encapsulation”) allows multiplexing of multiple protocols over a single ATM virtual circuit.

The protocol of a carried PDU is identified by prefixing the PDU by an ISO/IEC 8802-2 LLC header. The second method (called “VC-based Multiplexing”) performs higher-layer protocol multiplexing implicitly using ATM Virtual Circuits (VCs).

*Note.— This RFC is categorized as ‘Informational’ by the IESG. According to RFC 1602, Informational RFCs are specifications “published for the general information of the Internet community, and [do] not represent an Internet community consensus or recommendation. The Informational designation is intended to provide for the timely publication of a very broad range of responsible informational documents from many sources, subject only to editorial considerations and to verification that there has been adequate coordination with the standards process”.*

- 3.6.5.8.1.3 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 3.6-3.
- 3.6.5.8.1.4 The Payload field contains user information up to  $(2^{16}-1)$  octets.
- 3.6.5.8.1.5 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.

CPCS-PDU Payload
PAD
CPCS-UU
CPI
CPCS-PDU Trailer
Length
CRC

**Figure 3.6-3. AAL5 CPCS-PDU Format**


- 3.6.5.8.1.6 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 RFC 1483 and can be set to any value.
- 3.6.5.8.1.7 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 coded as the hexadecimal value [0x00]
- 3.6.5.8.1.8 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 the hexadecimal value [0x00] is used for the abort function.

3.6.5.8.1.9 The Cyclical Redundancy Check (CRC) field protects the entire CPCS-PDU except the CRC field itself.

3.6.5.8.1.10 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.


### 3.6.5.8.2 **LLC Encapsulation**

3.6.5.8.2.1 In LLC Encapsulation the protocol of the routed PDU is identified by prefixing the PDU by an ISO/IEC 8802-2 LLC header, which may be followed by an Subnetwork Attachment Point (SNAP) header. In LLC Type 1 operation, the LLC header consists of three 1-octet fields as shown in Figure 3.6-4.



**Figure 3.6-4. LLC Header Format**

3.6.5.8.2.2 The LLC header value 0xFE-FE-03 identifies that a CLNP PDU follows. The Control field value 0x03 specifies Unnumbered Information Command PDU. For CLNP PDUs, the format of the AAL5 CPCS-PDU Payload field shall thus be as follows as shown in Figure 3.6-5.



**Figure 3.6-5. AAL5 CPCS-PDU Payload Field Format for Routed ISO PDUsx**

3.6.5.8.3 The CLNP 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:

- a) SN-Source-Address: This field should contain the hexadecimal value [0xFE];
- b) SN-Destination Address: This field should contain the hexadecimal value [0xFE];
- c) SN-Priority: This field does not map to AAL-5 fields;
- d) SN-Quality-of-Service: This field does not map to AAL-5 fields; and

- e) SNS-User Data: This field should contain the ISO network layer PDU.

#### 3.6.5.8.4 **ATM Addressing**

- 3.6.5.8.4.1 Addressing support for ATM is defined in the ATM Forum specification User Network Interfaces 3.0/3.1. The address format is based on the OSI syntax for NSAP addresses, but despite the similar structure, these 20-byte ATM addresses are better described as private ATM SNPA addresses. There are three different formats: NSAP Encoded E.164, Data Country Code (DCC) Format, and International Code Designator (ICD) Format. Implementation of ATM subnetworks will require an address conversion process in order to map from the ATN NSAP address to the ATM address.

#### 3.6.5.9 **Mapping CLNP over CIDIN**

*Note.— The Common ICAO Data Interchange Network (CIDIN) is specified in Annex 10, Volume III. In addition, ongoing working group activities in the European region are concerned with protocol refinements, profile specifications, network management and provision of guidance material for the CIDIN. The outcome of these activities is published in the EUR CIDIN Manual (ICAO EUR DOC 005).*

#### 3.6.5.9.1 **General Characteristics of CIDIN**

##### 3.6.5.9.1.1 **Provided Communications Service**

- 3.6.5.9.1.1.1 CIDIN, at the time of definition conceived as a general purpose data network providing a code and byte independent connectionless transport service for AFS applications, makes use of packet switching techniques according to the CCITT Recommendation X.25. CIDIN protocols are defined at four levels: data link protocol (level 2), X.25 packet protocol (3a), CIDIN packet protocol (3b) and transport protocol (4). The level 1 is related to the physical interface to the transmission media. Routing and multiple dissemination is performed at the level of the CIDIN packet protocol. The user interface is provided at the level 4. The X.25 packet protocol may be performed in the DTE-DTE mode on leased lines (using permanent virtual circuits) or at the DTE-DTE interface to packet switched data networks (using switched or permanent virtual circuits).

*Note.— For CIDIN use of packet switched data networks see EUR CIDIN Manual.*

- 3.6.5.9.1.1.2 In the CIDIN concept a user of the CIDIN service is represented by an abstract functional unit called application entity. An application entity invokes the CIDIN transport service for user data and provides the parameters needed to specify the requested service (request to send a CIDIN message). In the opposite direction, control information and transported user data are accepted (reception of a CIDIN message). Individual types of application entities are distinguished by the assigned Message Code and Format (MCF) value. Only application entities of the same type (MCF value) are allowed to communicate across the CIDIN.

*Note.— Presently, application entities and the corresponding MCF values are specified for the transport of AFTN formatted messages, OPMET data, and CIDIN management information (EUR CIDIN Manual).*

3.6.5.9.1.1.3 The access point of an application entity to the CIDIN transport service is identified by the CIDIN entry address point (point of CIDIN message submission) and exit address (point of CIDIN message delivery).

*Note.— Special structures for the 8-letter CIDIN entry/exit addresses may be established on a regional basis.*

3.6.5.9.1.1.4 When sending a CIDIN message, the application entry can indicate by a service parameter whether the message transport should be acknowledged end-to-end within the CIDIN. Using this acknowledgement option, the CIDIN provides information on successful or non-successful message delivery per exit address delivery confirmation.

3.6.5.9.1.2 **The CIDIN Transport Interface**

3.6.5.9.1.2.1 The interactions with the users of the CIDIN transport service (level 4) are a local matter, i.e. not specified in Annex 10, Volume III. In accordance with the EUR CIDIN manual, Table 3.6-2 provides some guidance to the use of service parameters at this interface when sending or receiving a CIDIN message respectively.

**Table 3.6-2. Service Parameters Used for Sending and Receiving CIDIN Messages**

Service Parameter	Sending a CIDIN Message	Receiving a CIDIN Message
Exit Address (Ax)	Mandatory	Optional <sup>1</sup>
Entry Address (Ae)	Mandatory	Mandatory
Message Code and Format (MCF) indicator	Mandatory	Mandatory
Message priority (MP) indicator	Mandatory	Mandatory
Network Acknowledgement (NA) Indicator	Optional	Optional
User Data (CIDIN message)	Mandatory	Mandatory

<sup>1</sup> Is known to the addressed application entity

3.6.5.9.1.2.2 In the following some explanations are given to the service parameters listed in Table 3.6-2.

- a) Exit Address(es) (Ax): Identification of the receiving application entity (entities);

*Note.— In the European Region, a maximum number of 16 exit addresses may be associated with a CIDIN message (EUR CIDIN Manual).*

- b) Entry Address (Ae): Identification of the sending application entity;
- c) Message Code and Format (MCF) Indicator: Identifies the type of the communicating application entities; and
- d) Message Priority (MP) Indicator: Eight levels of priorities are defined. The highest priority (level 1) is reserved for CIDIN network management messages. The remaining priorities are available for user messages.

*Note.— For the transport of AFTN-formatted message the following correspondences between AFTN priority indicators and CIDIN priorities have been agreed: SS = 2, DD = 4, FF = 5, GG = 6, and KK = 7: (see EUR CIDIN Manual).*

- e) Network Acknowledgement (NA) Indicator: NA = 0 (no acknowledgement required) or NA = 1 (acknowledgement required).
- f) User Data (CIDIN message): The coding of the user data is code and byte independent. According to the CIDIN SARPs user data may have unlimited length.

*Note.— There is an agreement between States in the European Region to restrict the maximum length of user data to 64 kilobytes (see EUR CIDIN Manual).*

### 3.6.5.9.2 **Integration of CIDIN as ATN Subnetwork**

3.6.5.9.2.1 As illustrated in the section above, CIDIN has been specified as a general purpose transport system between peer CIDIN entry/exit centers. Thus the concept of “underlying subnetworks” as applied by the ATN architecture is not obvious in the CIDIN context.

3.6.5.9.2.2 However, CIDIN can be integrated in the ATN as an ATN subnetwork in which the subnetwork service is provided by the CIDIN transport service. In this configuration, the CIDIN transport protocol operates as subnetworks access protocol (SNAcP) according to the structure of the OSI network layer. The service provided by the CIDIN transport protocol is raised to the level required by the ATN internetwork protocol (CLNP) by means of a suitable SNDCF. This CIDIN SNDCF is specified in section 5.7.4 of the ATN ICS SARPs and described in more detail in section 3.6.5.9.37-5.8.3 below. The following Ffigure 3.6-6 illustrates how the CIDIN transport service is accessed by the ATN internetwork layer.

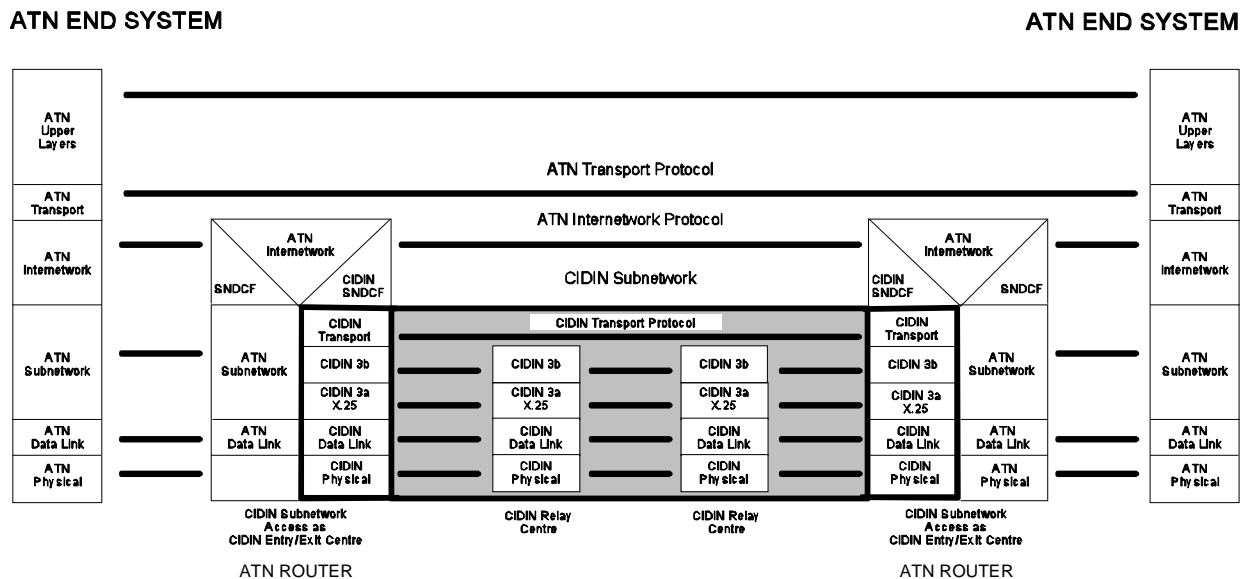


Figure 3.6-6. CIDIN as ATN subnetwork

- 3.6.5.9.2.3 In this configuration where the ATN internetwork protocol operates over the CIDIN transport protocol, there is a considerable degree of functional overlap between the SNIDP (i.e. the CIDIN transport protocol and packet protocol). For example, the CIDIN transport protocol also provides segmenting and reassembling functions and the CIDIN packet protocol performs routing on permanent or switched virtual circuits. The CIDIN transport and packet headers have to be carried in addition to the CLNP headers within the CIDIN subnetwork.
- 3.6.5.9.2.4 Furthermore, some functions provided by the CIDIN transport service together with the CIDIN packet protocol responsible for the handling of the 256-octet CIDIN packets are not used by the CLNP in this configuration. This includes the multiple dissemination of messages and the acknowledgement of messages between CIDIN entry and CIDIN exit centres.
- 3.6.5.9.2.5 However it is important to note that when CLNP operates over the CIDIN transport protocol, CIDIN maintains its integrity, i.e. it could simultaneously serve as an ATN subnetwork and as an end-to-end data network providing service to other CIDIN applications, such as the transport of AFTN-formatted messages.
- 3.6.5.9.2.6 Because of the almost unlimited length of CIDIN messages, the non-segmenting subset of the CLNP is sufficient when operating over CIDIN.

### 3.6.5.9.3 **CIDIN SNDCF**

3.6.5.9.3.1 The CIDIN SNCDF performs a mapping between the SN-Service required by the ATN internetwork protocol (CLNP) and the CIDIN (transport) service. The ATN ICS SARPs specify the relationship between the SN-UNITDATA service primitives and the actions at the CIDIN (transport) interface:

- a) a SN-UNITDATA Request corresponds to a request to send a CIDIN message; and
- b) a CIDIN message received at a CIDIN exit center translates into a SN-UNITDATA indication.

3.6.5.9.3.2 The acknowledgement option of the CIDIN is not invoked by the CIDIN SNDCF. This means that CIDIN will not provide a delivery confirmation, when used as ATN subnetwork. No segmentation of the SNS-Userdata is needed.

3.6.5.9.3.3 The parameters of the SN-UNITDATA service primitive, i.e. SN-Source-Address, SN-Destination-Address, SN-Priority and SN-Userdata have equivalents handled by the CIDIN transport service. Table 3.6-3 indicates the correspondence between these SN-UNITDATA service parameters and the CIDIN service parameters.

3.6.5.9.3.4 The SN-Quality-of-Service parameter of the SN-UNITDATA service primitive can be assumed to have a constant (a-priori) value for a CIDIN subnetwork and is entered e.g. as management data in the ATN router. It is ignored by the CIDIN SNDCF when receiving a SN-UNITDATA request and pre-set by the CIDIN SNDCF when generating a SN-UNITDATA indication.

**Table 3.6-3. Correspondence between SN-Service and CIDIN Service Parameters**

SN Service Parameter	CIDIN Transport Parameters
SN-Source-Address	Entry Address (Ae)
SN-Destination-Address	Exit Address (Ax)
SN-Priority	Message Priority (MP) Indicator
SNS-Userdata	CIDIN Message

3.6.5.9.3.5 Except for the above mapping to and from CIDIN transport parameters, the CIDIN SNDCF has to assign and MCF value identifying the User Data as ATN traffic.

*Note.*— The currently (January 2001+997) unassigned MCF value of 4 may be used for ATN communications traffic over CIDIN. Corresponding allocation will be initiated on a regional basis.



3.6.5.9.3.6 The correspondence between ATN CLNP priority and CIDIN priority as recommended by ASPP/3 is shown in Table 3.6-4:

**Table 3.6-4. Mapping of Priorities**

CLNP priority	CIDIN priority
35778	2
35739	5
35549	7

3.6.5.9.3.7 ~~A priori~~ ~~A-priority~~ values for transit delay, protection against unauthorized access, cost determinants and residual error probability have to be entered as management data into the ATN router.

#### 3.6.5.9.4 **Synopsis**

3.6.5.9.4.1 CIDIN is in wide use in certain regions as communications service for AFS applications, and may provide the only means of data communications to remote facilities. However, the use of CIDIN as an ATN subnetwork can often not be regarded as a technically straightforward solution.

3.6.5.9.4.2 If ATN-compliant X.25 services are available for the whole communications path (i.e. Svs) which are accessible on X.25 level, it is more advisable to use the underlying X.25 service directly as an ATN subnetwork in order to reduce overhead from the encapsulation of CLNP. In this case, a “standard” X.25 SNDCF can be used in the ATN router.

3.6.5.9.4.3 However, if the X.25 protocol cannot be accessed directly, then the use of a CIDIN SNDCF provides the possibility to make use of the already existing infrastructure for ATN purposes.

### 3.7 **Interoperability with Previous Editions of the ATN ICS SARPs**

3.7.1 The ATN ICS SARPs have been amended and upgraded according to the principle of preserving interoperability with previous editions to the utmost extent. Consequently, the ATN ICS SARPs contain a number of provisions which are complementary to the specification of new technical capabilities and have been included to ensure backward compatibility with implementations compliant with older editions of the ATN ICS SARPs.

3.7.2 The key mechanisms used to ensure backward compatibility when introducing additional or enhanced technical features in the ATN ICS SARPs are:

- a) use of protocol options in ISO/IEC protocols which, according to the existing protocol specifications, have to be ignored by the receiver in the case it does not support or understand the option(s). This mechanism is used to convey additional

configuration data between peer ATN Routers and an example of this mechanism is the Mobile Subnetwork Capabilities Parameter (see 3.7.3);

- b) negotiation of new technical features during connection initiation by signalling enhanced capabilities in the initial PDU exchange and invoking and using only those capabilities for which the peer system also indicates support in the responding PDU. Examples of this mechanism are the ATN Data Link Capabilities Parameter (see 3.7.4) and the Extended Transport Checksum (see 3.7.5); and
- c) upgrade of the version number of an existing ATN protocol to indicate an enhanced version of this protocol and fall back to an older version of this protocol, if the peer system indicates that it does not support this enhanced version of the protocol. An example of this mechanism is the Mobile SNDCF Protocol (see 3.7.6).

### 3.7.3            **Mobile Subnetwork Capabilities Parameter**

3.7.3.1            The Mobile Subnetwork Capabilities Parameter is included by an ATN Air/Ground Router in the options part of an ISO/IEC 9542 ISH PDU uplinked during the air/ground route initiation process. This parameter informs the receiving ATN Airborne Router about the traffic type restriction(s) and the supported ATSC Class(es) of the air/ground data link over which the ISH PDU is received. This allows the ATN Airborne Router to make appropriate routing decisions when downlinking CLNP packets.

3.7.3.2            As indicated above, this parameter will be ignored by ATN Airborne Routers which are implementing an older edition of the ATN ICS SARPs; these Airborne Routers are required to use a priori configuration information to respect ITU restrictions on traffic types and ATSC Class qualifications of data links when forwarding CLNP packets.

3.7.3.3            ATN Airborne Routers compliant with the most recent edition of the ATN ICS SARPs will use the received subnetwork capability information to update its local configuration data and will use this updated information in re-building its Loc RIB and FIB.

### 3.7.4            **ATN Data Link Capabilities Parameter**

3.7.4.1            The ATN Data Link Capabilities Parameter is included by ATN Air/Ground Routers and Airborne Routers respectively in the options part of an ISO/IEC 9542 ISH PDU during the air/ground route initiation process. In the uplink, this parameter signals whether the Air/Ground Router

- a) supports the generation of UPDATA BISPDU's without a security path attribute (in order to efficiently use scarce air/ground resources);
- b) intends to perform IDRPs Type 2 authentication in subsequent BISPDU exchanges with the receiving Airborne Router; and
- c) requires the receiving Airborne Router to downlink its Public-Key Certificate for subsequent validation.

- 3.7.4.2            In the downlink, this parameter allows an Airborne Router to indicate its capability of receiving and processing UPDATE BISPDUs which do not contain a security path attribute.
- 3.7.4.3            From the value of the ATN Data Link Capability Parameter the receiving ATN Router determines the extended capabilities, if any, of the sending ATN Router and invokes and uses in further communications with this ATN Router only those extended capabilities which are supported by both ATN Routers.
- 3.7.4.4            According to ISO/IEC 9542, unknown parameters contained in the options part of received ISH PDUs are ignored by the receiving system but the PDUs are not discarded. Consequently, an ATN Router compliant with an older edition of the ATN ICS SARPs will ignore the ATN Data Link Capability Parameter, but will be in a position to communicate with an ATN Router implementing the current specification of the ATN ICS SARPs. However, the latter router has to "downgrade" its feature set and is not allowed to use extended capabilities. The extended capabilities may only be used, if both ATN Routers implement the most current specification of the ATN ICS SARPs and indicate support for the extended capabilities in the ATN Data Link Capabilities Parameter.
- 3.7.5                **Extended Transport Checksum**
- 3.7.5.1            The Extended Transport Checksum has been specified as an ATN-specific parameter in the variable part of COTP and CLTP PDUs to ensure a reasonably low probability of packet misdelivery (i.e. less than 10<sup>-8</sup>). Use of this parameter is optional if the COTP is operated and mandatory if the CLTP is operated. In both cases, this parameter is not in compliance with the relevant ISO/IEC standards.
- 3.7.5.2            However, it is not a protocol error to use it in a Connect Request (CR) PDU. If present in the CR PDU, the connection initiator signals that it wants a high level of protection against packet mis-delivery. If the receiving transport entity does not support this parameter (because it implements an earlier version of the ATN ICS SARPs), it will not include it in the responding Connect Confirm (CC) TPDU. To ensure backward compatibility, the procedure requires that the Extended Transport Checksum will not be present in any further TPDU exchanged over the established transport connection.
- 3.7.5.3            Otherwise, if the responding transport entity also supports the Extended Transport Checksum, this parameter will be included in the CC TPDU and thereafter in all other TPDUs exchanged over the established transport connection.
- 3.7.5.4            To ensure interoperability in the case that the transport connection (TC) responder implements a later edition of the ATN ICS SARPs than the TC initiator, the TC responder is not allowed to include the Extended Transport Checksum parameter in any TPDU, if it was not present in the received CR TPDU.
- 3.7.5.5            Use of the Extended Transport Checksum parameter has also been specified for the CLTP, however on a mandatory basis as there is no negotiation phase for the connectionless transport protocol. Interoperability between implementations compliant with different editions of the ATN ICS SARPs is provided as long as the sending transport entity knows

a priori or through some directory that the receiving entity also supports this parameter and that there is a need to use the Extended Transport Checksum.

### 3.7.6 Enhanced Mobile SNDCF Protocol

3.7.6.1 The Mobile SNDCF specification has been enhanced to include the transfer of a Deflate compression dictionary from an existing virtual circuit to a newly established virtual circuit, and to negotiate the use of pre-stored Deflate compression dictionaries (see xxx) . To support these enhancements the Mobile SNDCF Header contained in the User Data of a Call Request or Call Accept packet respectively has been amended to include a variable length SNDCF Parameter Extension Block.

3.7.6.2 The presence of this SNDCF Parameter Extension Block in the Call Request User Data is indicated by setting the Version Number in the Mobile SNDCF Header to binary [0000 0010]. A receiving system which also supports this enhanced version (i.e. Version 2) of the Mobile SNDCF Protocol will understand and process the information received in the SNDCF Parameter Extension Block and will respond with a Call Accept packet, if the call is acceptable in general. The Call Accept User Data may or may not include a SNDCF Parameter Extension Block depending on the proposed options in the received SNDCF Parameter Extension Block. In the first case, the most significant bit (i.e. the PEXT bit) of the first octet of the User Data is used to signal the presence of this Extension Block. The initiating Mobile SNDCF will be able to understand and process this SNDCF Parameter Extension Block, as it has indicated in the previous Call Request that it supports Version 2 of the Mobile SNDCF Protocol.

3.7.6.3 In the case that the Mobile SNDCF Header in the User Data of the incoming Call Request packet contains a Version Number of binary [0000 0001], indicating the first version of the Mobile SNDCF Protocol, the responding SNDCF will neither set the PEXT bit nor include a SNDCF Parameter Extension Block in the Call Accept, even if it supports a later version (i.e. Version 2). This maintains backward compatibility if the called Mobile SNDCF is compliant with a more recent edition of the ATN ICS SARPs than the calling Mobile SNDCF.

3.7.6.4 Interoperability is also ensured in the case that the calling Mobile SNDCF issues a Call Request with User Data formatted according to Version 2 of the Mobile SNDCF Protocol and the called Mobile SNDCF implements an earlier edition of the ATN ICS SARPs. In this case, the receiving Mobile SNDCF will reject the call with a diagnostic code indicating "Version number not supported". The specified provisions require the call initiator to re-attempt the call establishment with a Call Request User Data formatted according to Version 1 of the Mobile SNDCF Protocol (i.e. without an SNDCF Parameter Extension Block). The consequence of this procedure is that two attempts may be required to set up a virtual circuit between the Air/Ground and the Airborne Router, but communications will be established in any case, if the call is acceptable in general.

-----