AERONAUTICAL NETWORK PANEL

Working Group 2

Proposed Guidance Material

Section 2

Brisbane 5-9 February 1996

Prepared and Presented by: Ron Cossa

SUMMARY

This paper presents proposed editing of the Chapter 2 Guidance Material for the CNS/ARM-1 Package SARPs and Guidance Material.

~~Introduction~~

1. The ATN CONCEPT

1.1 The ATN Specification

The ATN is a CLNP Internetwork ~~that interconnects for the purpose of data interchange~~designed to interconnect:

- Existing, planned and future ATS and Aeronautical Industry Networks
- ICAO Mobile Data Communications Networks
- ATS and Aeronautical Industry ATM Systems, including those onward aircraft in flight.

~~T~~While the specification of these interconnected networks and systems is outside of the scope of the ATN Internet,~~ while~~ the specification of the ATN Internet SARPs is concerned with the architecture of the ATN, how data packets are routed through the ATN, and the specification of the ATN Routers, which provide the essential packet interchange and routing facilities that enable ATM Systems to communicate across a variety of networking technologies.

The ATN Internet SARPs are therefore concerned with the following functional areas:

- The selection and specification of the data communications protocols used to format and pass user data packets throughout the ATN;
- The selection and specification of the Routing Information Exchange Protocols used to support the proper operation of ATN Routers;
- Security and Quality of Service Maintenance;
- The specification of ATN Routers;
- The specification of the End System Components necessary for ATN Access;
- The specification of the procedures for supporting ATN Mobile Routing.

ATN Mobile Routing itself comprises several functional areas:

- The Management of Routes to Mobile Systems in the Ground ATN;
- The use of Mobile Networks including procedures for Route Initiation (i.e. the establishment of communications with a mobile system);
- Data Compression over low bandwidth Mobile Networks.

Where possible, the ATN uses communications protocols that have already been specified by the International Standards Organisation (ISO) or the International Telecommunications Union (ITU). The ATN Internet SARPs are therefore concerned only with the identification of these protocols and their adaptation (i.e. profiling) for aeronautical use.

The main area of original specification in the ATN Internet SARPs is in the specification of the procedures for supporting ATN Mobile Routing. In particular, the routing policies that support mobile routing in line user and application requirements, the Route Initiation procedures, and the Data Compression procedures for use over low bandwidth networks, are all unique to the ATN Internet SARPs.

1.2 An ATN Overview

This section is concerned with the ATN Architecture and the Service it provides to its users. It provides introductory material to the ATN Concept, especially targeted at the reader that requires only limited information on the internal operation of the ATN.

~~However, 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 chapter, the term "ATN User" is used in this sense, i.e. as a user of the network service or transport service, depending on context.~~

~~The ATN: a User's Perspective~~

The ATN is a data communications internetwork that:

1. provides a common communications service for all ATSC and AINSC applications that require either ground/ground or air-ground data communications services.
2. integrates and uses existing communications networks and infrastructure wherever possible.
3. provides a communications service which meets the security and safety requirements of ATSC and AINSC applications.
4. accommodates the different grades of service required by each ATSC and AINSC application.

While the~~ above~~se 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 OSI 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.

1.2.1 The ~~Individual ATN User~~ATN: a User's Perspective

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 chapter, the term "ATN User" is used in this sense, i.e. as a user of the network service or transport service, depending on context.

1.2.1.1 ATN User Communications Capabilities

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 their own system. It is this "user's view" of the ATN that is illustrated in Figure 1-1. This figure shows the

Figure 1-1 Individual End User's View of the ATN

ATN as an abstract "cloud" which indeed is all the user need 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.

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" their 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.

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 8473 Connectionless Network Protocol in order to be part of the ATN Internet, and is recommended to support the ISO 9542 End System to Intermediate System Routing Protocol.

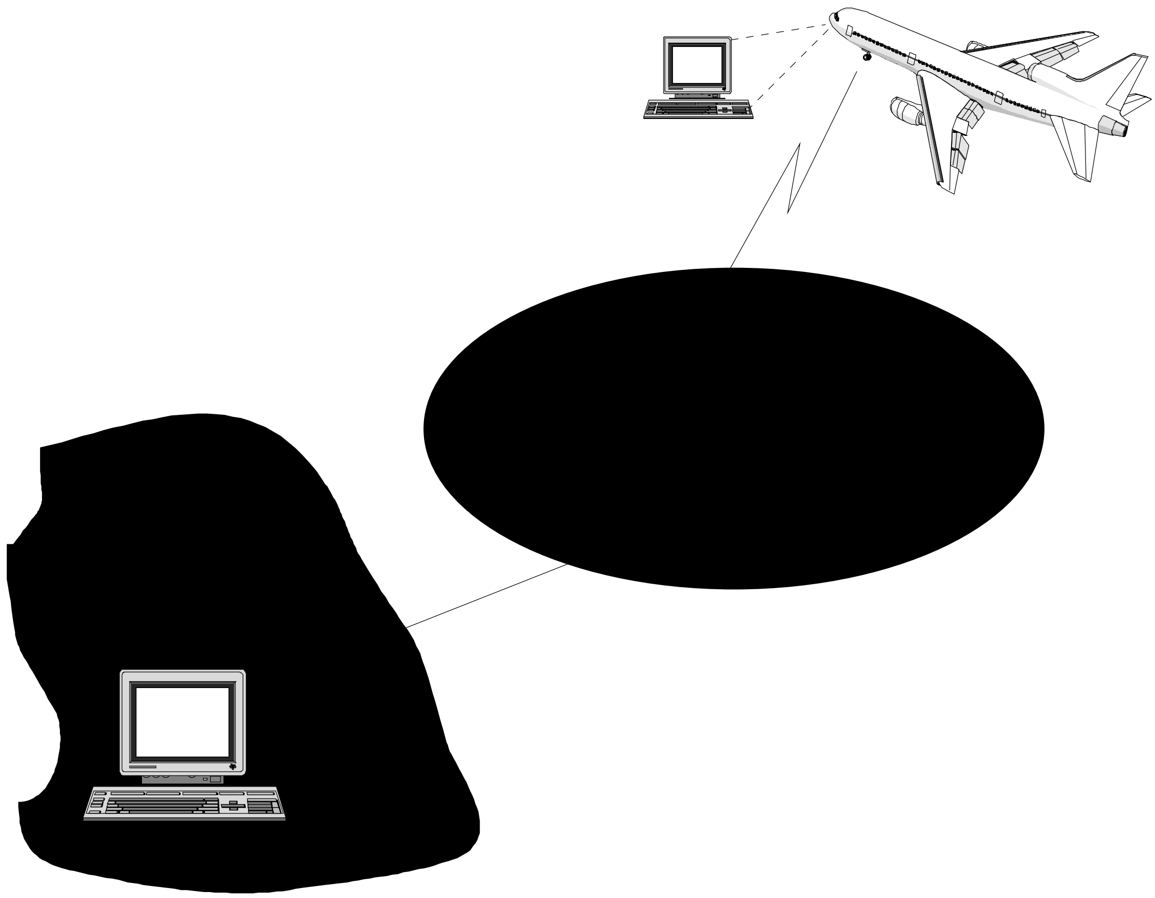

1.2.1.1.1 ISO 8473 CLNP Protocol

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.

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.

The ~~above~~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.

1.2.1.1.2 ISO 9542 ES/IS Protocol

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.

1.2.1.1.3 CLNP Delivery Probability

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 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:

1. 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)

2. network overload which results in low priority data being discarded in order to free up congested resources

3. component failures.

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 minimizing the use of lower reliability subnetworks, increasing overall network capacity, and through component redundancy.

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 8073 class 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.

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 8073 COTP Class 4 is not appropriate are instead required to use the ISO 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 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.

1.2.1.2 The User as an Organization

Wh~~ile the previous description is satisfactory from the individual user's point, it does assume that there is some external ATN Service Provider, that provides the ATN Router and hides ATN complexity from the end user. However, while this may indeed be true for many ATN users, wh~~en the ATN user is an organization with multiple systems to consider; <u>when it </u>has many subnetworks of its own; and <u>when it </u>also wishes to operate ATN Routers, then the organization will also be an ATN Service Provider and hence needs to be aware of some of the internal complexity.

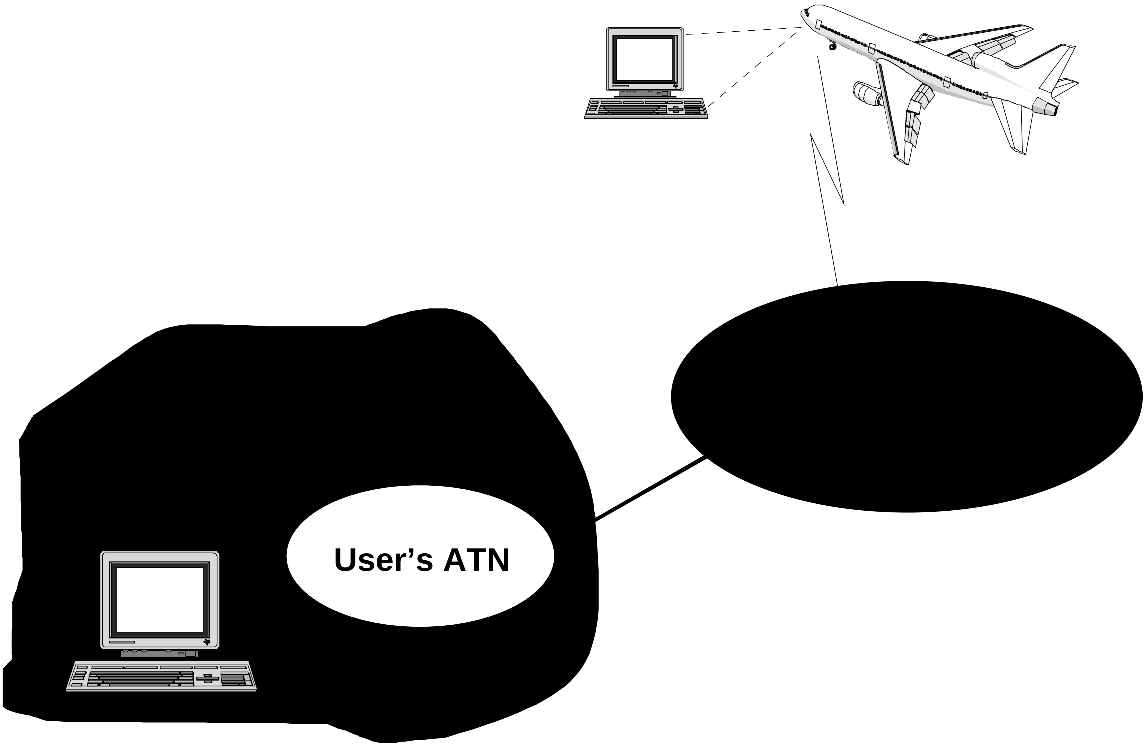

Figure 1-2 Organizational End User's View of the ATN

Figure 1-2 illustrates the changed perception. The "organizational" 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.

1.2.1.2.1 Organizations and Routing Domains

The portion of the ATN operated by the organizational user may be no more than a single ATN Router. Alternatively, a larger organizational user may operate a large number of ATN Routers, interconnected by various subnetworks also operated by the organizational user, and providing ATN access to all the systems owned by the organizational 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 organization, 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.

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 compete 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.

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.

The ATN simply consists of multiple Routing Domains, each operated by an ATN participant. 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.

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 router, and information about the route's path including the Quality of Service and Security available over the route.

1.2.1.2.2 Use of Policy Based Routing by Organizations

It is the BIS's responsibility to determine which routes, if any, it will advertise to 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 organizational user of the ATN must determine and apply their own Routing Policy.

It is the need for policy based routing between different organizations 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.

Organizational ATN users must therefore ensure that they either have direct connections with the ATN Routing Domains with which communication 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 organizational users on an "as needs" basis. However, in practice, this is unlikely to be an efficient strategy and may actually prevent useful communication by putting too high a cost on establishing a usable path even when connectivity already exists.

Instead, it is intended that ATN interconnections are coordinated on both a regional and worldwide basis, so that ATN backbones (of Routing Domains offering general transit facilities) are 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.

1.2.1.3 Mobile Users

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. [c5 t 0115] 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.

In the ATN, mobile "platforms" are treated in a similar manner as organizational 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 routing protocol provides a relatively efficient mechanism for the transfer of routing information over low bandwidth links, but also because aircraft are almost always organizationally separate to the ground systems with which they are in contact and the same requirements for policy based routing apply.

The existence of mobile users has a significant impact on the organization 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.

1.2.1.3.1 Mobile "Home" Concept

This strategy is based on the notion of 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 Service Provider or an Administration. Typically, all aircraft belonging to the same airline, or the General Aviation (GA) aircraft of a single country share the same home. Through the ISO/IEC 10747 routing protocol, a route to an aircraft's home is known throughout the fixed ATN.

As regards routes to the aircraft themselves, the routing policies used in the ATN constrain the distribution of routing information about a given aircraft. The constraints ensure that the routing information is propagated to the aircraft's home and those RDs along the path to the home only. [c5 t 0145]

With ISO/IEC 10747, a CLNP packet is routed not by destination address as such, but by address prefixes. Packets are routed along a route which provides an address prefix which is best (i.e., the longest match) with the destination address. In the ATN, both the route to the home and to the aircraft itself are characterized by address prefixes to the actual addresses contained within the aircraft. However, the route to the home is characterized by a shorter prefix than is the route to the aircraft itself. Typically, the address prefix that characterizes the route to the home is a prefix to all an airline's aircraft.

Thus, when a CLNP packet is sent to an address in an aircraft from an arbitrary point within the ATN, it will typically follow the route to the home. This is because this route is known throughout the ATN and it provides a prefix for the destination address. However, as soon as the CLNP packet reaches the home RD, or indeed, any RD on the path between the home RD and the aircraft, the route to the aircraft provides a longer match in respect of the address prefix, and hence the packet now follows the route to the aircraft and its destination.

By this technique an almost optimal route is followed while constraining the distribution of routing information to mobiles to a limited set of RDs, and hence minimizing the impact of mobiles on the ATN.

1.2.1.4 Routing Control

The ATN may be used for Operational Communications, Administrative Communications, Systems Management Communications and for General Communications. Although some parts of the ATN may allow the sharing of facilities between these different classes of data, there will usually be a strict separation of communications resources between these different classes of data. The ATN supports this separation by specifying a separate traffic type for each class of data, and enabling each route through the ATN to be labelled according to the class of data which it may convey. ATN Routers are required to relay user data along only those routes available for its class of data. The requirement for strict separation is thus implemented and enforced by the ATN routing procedures.

1.2.2 ATN Users and Service Providers

The ~~above discuss~~description in the previous section 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 as the service provider. Figure 1-3 illustrates the location of each ATN Reference Point, where the layers of the OSI reference model are represented as numbered rectangles, with the numbers 1 to 7 corresponding to the physical, data link, network, transport, session, presentation and application layers respectively. Figure 1-1(a) illustrates the reference points by abstraction from Figure 1-2, while Figure 1-1(b) provides a more formal model of the ATN.

1. Reference Point One, the Transport Reference Point, is the OSI Transport Service interface and follows ISO 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.

2. 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 futher ATN routers, which may or may not belong to other ATN Service Providers. When both Host Computer and Router are owned by the same organization, then the protocols that provide this interface are not mandated; the specification only recommends an appropriate stack.

3. Reference Point Three, the Network Provider Reference Point, is at the interface between two Routers belonging to different organizations. 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.*

4. 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 Internet SARPs.

The provider of the service at reference point 4 is out of the scope of this document. As ~~discuss~~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.

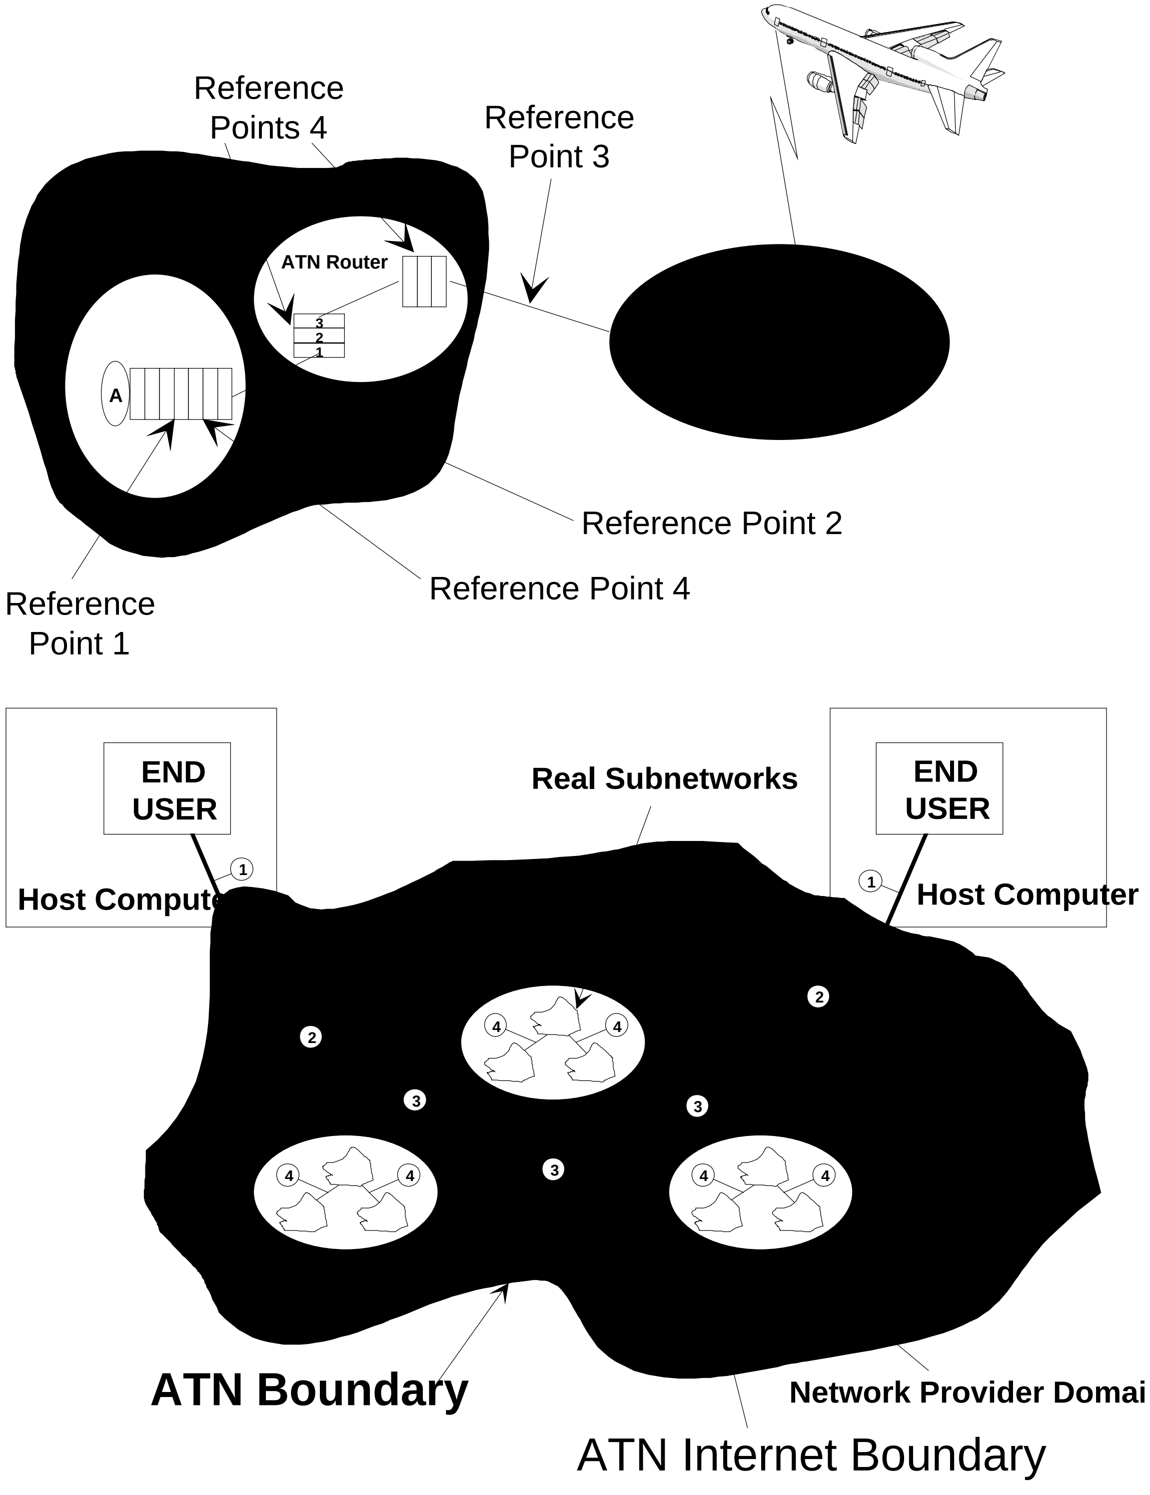

Figure 1-3 ATN User and Service Provider Reference Points

1.3 ATN Architectural Components

1.3.1 The OSI Reference Model

The ATN is an implementation of OSI and frequently refers the OSI Reference Model. It is therefore useful to consider what the OSI Reference Model means in practice and introduce some of the concepts it defines. The “seven layers” of the reference model are illustrated in the figure opposite.

The reference model is a layered model building on the layered architecture introduced with X.25, and identifies two types of Systems:

- **End Systems:** which are Host Computers, the users of network services, and which comprise seven protocol layers, providing communications services to applications.
- **Intermediate Systems:** which are either Routers or Packet switches, and comprise only the three layers appropriate to network communications.

End Systems may either communicate directly, using the services of a physical communications medium, or communicate via one or more Intermediate Systems. The definition of each protocol layer is:

- **Application Layer:** contains all the information (or semantics) that is exchanged between End Systems. In particular, it contains all user information that is exchanged. It also provides the means to allow the End Systems to agree to the semantics of the information exchanged.
- **Presentation Layer:** provides the means to represent the information exchanged (i.e. the Syntax) between the End Systems without changing the semantics of the information.
- **Session Layer:** provides the means to mark significant part of the information exchanged between systems: for example, a Unit or Word, a page, or a chapter.
- **Transport Layer:** provides end-to-end control and information interchange with the level of reliability that is needed for the using application. The services provided to the upper

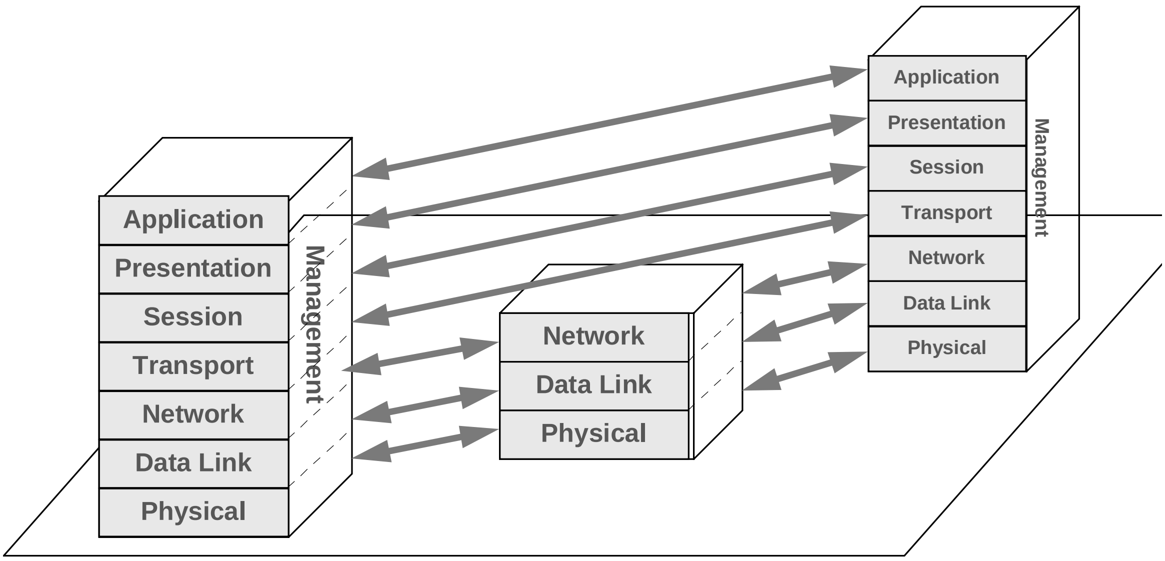

Figure 1-4 The OSI Reference Model

layers are independent of the underlying network implementation. The Transport Layer is therefore the “user’s liaison, acting as the go-between for the user and the network, enhancing the network’s service to that required by the application.

- **Network Layer:** provides the means to establish, maintain and terminate the switched connections between End Systems, or to transfer datagrams between two End Systems. Addressing and routing functions are included in the Network Layer.

- **Data Link Layer:** provides the synchronisation and error control for the information transmitted over the physical link.

- **Physical Layer:** provides the functional and procedural characteristics to activate, maintain and deactivate the physical connection. It also includes the electrical and mechanical characteristics of the physical interface to the external transmission media.

The seven layer model has proved remarkably resilient to change since it was first introduced, although it has been considerably enhanced over the years. The original reference model considered only connection mode communications. It was later amended to include connectionless communications, and now also includes a Security Model, a Naming and Addressing Model, and a Systems Management Model. A Quality of Service (QoS) Management Framework is a likely future addition, as this is now becoming a subject generating major interest.

OSI Reference Model terminology is used extensively in both the ATN SARPs and Guidance Material.

1.3.2 ATN Functional Components

The ATN comprises the following functional components:

- End Systems (ESs) (i.e. Host Computers where aeronautical applications may reside)

- Intermediate Systems (ISs) (i.e. ATN Routers)

The relationships between these components are illustrated in Figure 1-5.

The OSI layer 1-4 aspects of ESs and all ATN ISs comprise the ATN Internet. The ATN Internet provides data transfer services to ATN ESs. Through ESs, ATM applications gain access to ATN services and are hence able to exchange application specific data.

The ATN also recognizes the existence of subnetworks, and such subnetworks, when used to provide communication paths between ESs and ISs, and ISs, may also be viewed as being part of the ATN. Subnetworks are characterized by the ATN as being:

- Ground subnetworks

- Mobile subnetworks

- Aircraft subnetworks

Ground, Mobile and Aircraft subnetworks are used to interconnect the components described above.

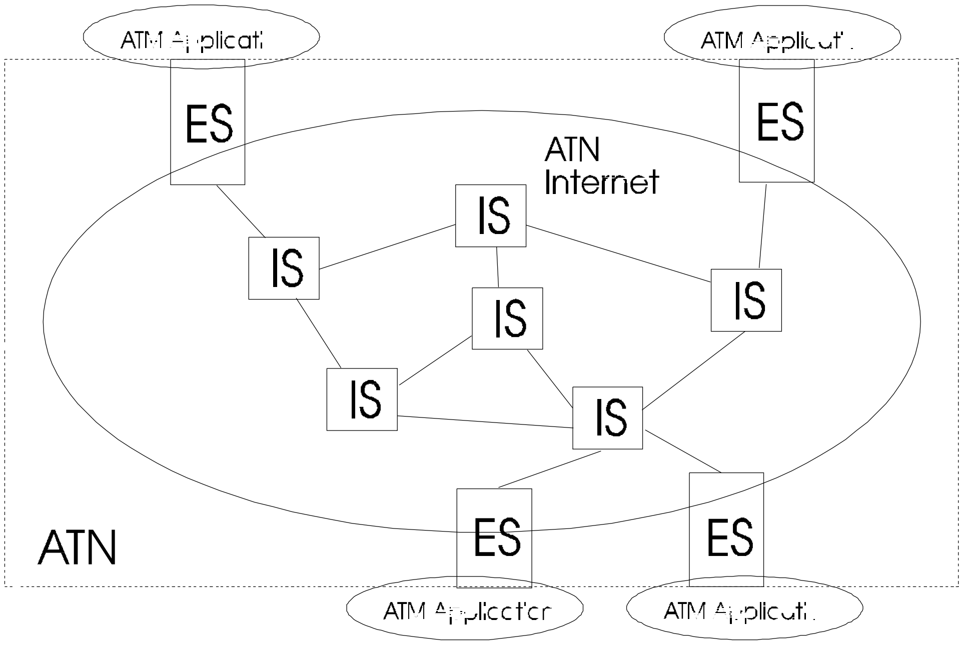

Figure 1-5 ATN Functional Components

1.3.2.1 End System Model

The function of an ATN End System is to provide the end-user applications with an OSI compliant communications interface to enable them to communicate with remote end-user applications. An ATN End System is the representation in ATN Architecture of a Host Computer supporting one or more CNS/ATM applications.

ATN End System implementation of the protocols required for Layers 1 and 2 (i.e. Physical and Data Link), and subnetwork access functions in layer 3, is purely a local issue and wholly dependent on the subnetwork to which the particular End System is attached, as such definition of these protocols is outside the scope of the SARPs.

ATN End System protocols and functions in the transport layer, and those in the Network Layer which are subnetwork independent or are concerned with the convergence of subnetwork dependencies with subnetwork independent protocols and functions, are within the scope of the SARPs.

ATN End System higher layer protocols and functions in support of ATM applications are defined in Parts 2, 3 and 4 of the SARPs.

1.3.2.2 Intermediate System Model

The functions of the Intermediate System are to perform the relaying and routing of ISO 8473 NPDUs between other Intermediate Systems and End Systems. An ATN Intermediate System is the representation in ATN Architecture of a Router interconnecting two or more ATN Networks for the purpose of data interchange between ATN End Systems connected directly or indirectly to those ATN Networks.

The protocol support required for the Physical and Data Link layers of an ATN Intermediate System are purely dependent on the subnetwork to which they are attached and are therefore a local issue and outside the scope of the SARPs.

ATN Intermediate System protocols and functions in the Network Layer which are subnetwork independent or are concerned with the convergence of subnetwork dependencies with subnetwork independent protocols and functions, are within the scope of the SARPs.

The functions and data structures in support of routing and relaying are d~~iscussed below~~etailed in the next section.

1.3.3 Administrative Domains and Routing Domains

In order to develop a scaleable routing architecture, the ATN has adopted the ISO Routing Framework, presented in ISO TR 9575 and illustrated in Figure 1-6. This provides the structure necessary to support routing in the complex ATN ground environment and to Mobile Systems.

The ISO Routing Framework first recognises that Host computers, routers and networks are owned and operated by different organisations, and therefore defines the *Administrative Domain*. An Administrative Domain comprises the Hosts computers, routers and networks operated by the same organisation. The purpose of the Administrative Domain is to clearly indicate the domain of an organisation's responsibility and to differentiate communication within an organisation from communication between organisations.

However, the most appropriate structures for routing control do not necessarily always follow organisational boundaries, and this is recognised by the definition of the *Routing Domain.* A Routing Domain simply comprises a set of Host Computers and Routers owned by the same organisation, and which implement a common routing algorithm. There may be many Routing Domains within a single Administrative Domain.

The routing framework also classifies routers according to their role. Those routers that operate solely within a Routing Domain are termed Intermediate Systems, while those that support inter-domain routing are termed Boundary Intermediate Systems (also referred to as Boundary Routers).

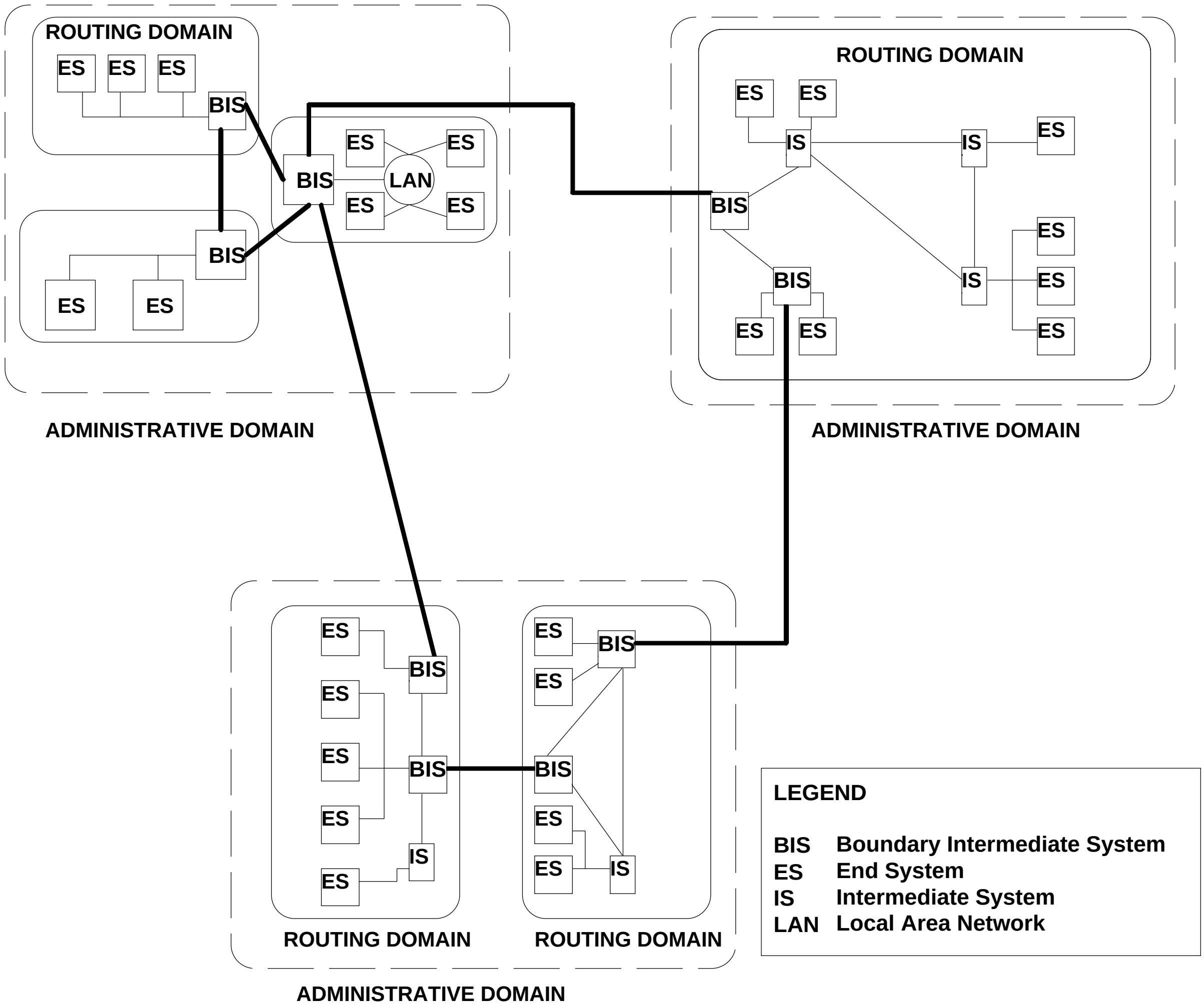

Figure 1-6 ISO Routing Framework

1.3.4 Addressing

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.

An 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.

A useful example of this concept is provided by telephone numbers. The ITU has published a global numbering plan for telephone number allocation and, in principle, each telephone number consists of a "country code", an "area code", an "exchange code", and a "subscriber number". Telephone numbers are closely related to the topology of the telephone network. Given a subscriber's telephone number it is possible to identify the country and area in which their local exchange is located, and by using this close relationship between topology and telephone numbers, the routing of telephone calls can be readily accomplished. As long as, for example, each inter-area telephone exchange has routing tables that identify the routes to each other area within the same country, telephone calls between different areas within the same country can be readily made. Only within an area does the connectivity between exchanges inside the area, need to be known.

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.

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 Confederation (see 1.6 below).

There is also no requirement in the ISO Routing Framework for the address prefixes that characterise adjacent (i.e. linked by a common network) Routing Domains, to have any similarity (i.e. for there to be another (shorter) address prefix common to each Routing Domain's address prefix).

If all inter-Domain interconnections are simply developed on an ad hoc basis with no aim to create a Global ATN Internet, then any lack of similarity between the address prefixes assigned to adjacent Routing Domains is not an issue. However, if a scaleable routing architecture (i.e. one which permits effectively unlimited growth) is to result then there does need to be some similarity between the address prefixes characterising adjacent Routing Domains. Then it will be possible to group Routing Domains together and advertise routes to a group of Routing Domains, rather than to each individually. This is similar to telephone networks grouping of all the areas in one country together and treating them as a whole from other countries. With such a strategy it is possible to develop a scaleable routing architecture such that the further away a router is from a packet's destination, the less detailed the routing information needs be to successfully route the packet.

This is a very important feature of a scaleable architecture, because if the amount of routing information required by at least one router is in proportion to the size of the ATN Internet then the maximum size that router can be, places a limit on the size of the network as a whole.

~~ATN Islands~~

~~The "Home" of an ATN Mobile~~

1.4 ATN Protocol Architecture

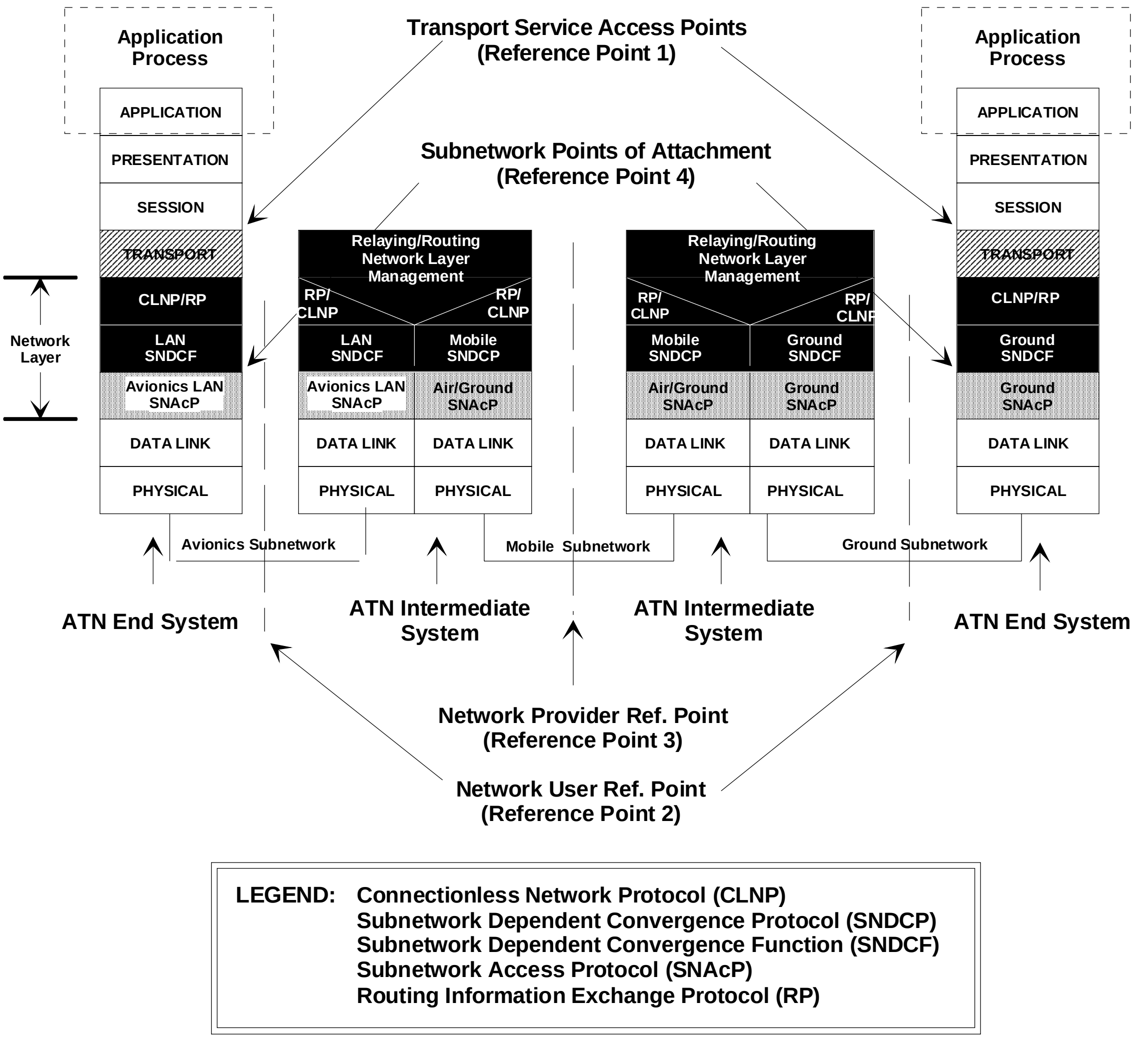

Figure 1-7 ATN Protocol Architecture

1.4.1 The ATN Transport Layer

The OSI 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 protocol 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.

The ATN protocols follow the reference model specified by ISO in ISO 7498-1.

The ATN protocol architecture is illustrated in Figure 1-7, and specifically illustrates the use of a mobile subnetwork for air-ground communications. This protocol model has been derived directly from the OSI Reference Model.

There are two modes of the transport service, the Connectionless mode Transport Service and the Connection-mode Transport Service. 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.

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 8602, and the protocol in support of the COTS is specified ~~in ISO 8073~~to be ISO 8073 Class 4. The implementation of these protocols within the ATN is further described in Chapter 8.

The Transport Service boundary corresponds with ATN reference point 1.

1.4.2 The ATN Network Layer

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 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.

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.

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

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

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

1.4.2.1 The Subnetwork Independent Role

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.

In support of the connectionless mode Network Service, it is a mandatory requirement that all ATN ESs and ISs implement the ISO 8473 internetworking protocol. 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.

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

- ISO 9542 - the End-System to Intermediate-System (ES-IS) protocol
- ISO 10589 - the Intermediate-System to Intermediate-System (IS-IS) intra-domain routing information exchange protocol
- ISO/IEC 10747 - the Inter-Domain Routing Protocol (IDRP)

The use of these protocols is d~~iscuss~~escribed in more detail in Chapter 11.

1.4.2.1.1 End-System to Intermediate System Routing Protocol

The ISO 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.

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.

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.

1.4.2.1.2 Intra-Domain Routing Information Exchange Protocol

The ISO 10589 IS-IS intra-domain routing information exchange protocol is used by ISs within the same Routing Domain to exchange connectivity and QOS information. It is recommended for implementation in all ATN ISs. As the ISs within a single Routing Domain are always operated by the same organization, this protocol is not used at any of the ATN interfaces identified by reference points.

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.

1.4.2.1.3 The Inter-Domain Routing Protocol (IDRP)

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).

Under IDRP, specialised Boundary Routers 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.

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.

1.4.2.2 Subnetwork Dependent Role

The OSI Subnetwork Dependent Role is responsible for decoupling the functions of the subnetwork independent role from the characteristics of 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 .

ISO 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. Chapter 10 discusses the SNDCFs that may be used to interface ISO 8473 to ATN subnetworks.

However, while ISO 8473 does not require an SNDCP, there is justifiable concern over the ISO 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 8473 protocol header over mobile subnetworks. This SNDCP is further d~~iscuss~~escribed in Chapter 10.

1.4.2.3 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.

1.5 Policy Based Routing

The intuitive view of routing in a packet based network is generally termed “performance based routing”. In this mode of operation, all communications paths are available, and a router’s objective is to choose the best out of those available, using metrics such as “hop count”, “capacity”, “transit delay”, “cost”, etc. in order to determine which is “best”.

However, while this may be the most appropriate strategy within an organisation, when packets are routed between organisations, or over commercial networks, the fact that a route is available (i.e. connectivity exists) may not always be the only reason to consider its use, other policy based criteria may apply. The application of such policy criteria to routing is know as “policy based routing”, and is another feature of IDRP.

Policy based routing has always been applied informally, using static configuration of routing tables. However, IDRP formalises policy based rules for route selection within the context of a dynamic routing framework.

Policy is applied at two points. Firstly, when a route is received from another Routing Domain, a policy decision is taken on whether to use it, either at all, or in preference to alternative routes to the same destination. And, secondly, a policy based decision is made when a route is considered for onward advertisement to an adjacent Routing Domain. Through routing policy, a network manager can choose both the received routes that are accepted for use, and those which it is prepared to offer for the use of other Routing Domains. For example, through the implementation of appropriate policy rules, a Routing Domain connected to many other Routing Domains, can be a Transit Routing Domain i.e. relaying between those Routing Domains, or an End Routing Domain, i.e. only accepting packets addressed to local destinations.

1.6 Routing Domain Confederations

Although the structuring of an internetwork into Administrative and Routing Domains enables a structured approach to routing to be developed, this is not in itself readily scaleable. Once there exists a large number of Routing Domains, the structuring problem re-asserts itself, and there is a need to provide another level to the Routing Framework~~, and so on. This problem is resolved in a recursive fashion by the Routing Domain Confederation~~. This problem can be resolved by the Routing Domain Confederation concept.

A Routing Domain Confederation (RDC) is a set of Routing Domains and/or RDCs which have agreed to join together and form a Routing Domain Confederation. The formation of a RDC is done by private arrangement between its members without any need for global co-ordination. From the outside, an RDC appears exactly like a single Routing Domain in the sense that the routes that CLNP PDUs can follow cannot re-enter an RDC, no more than they can re-enter a Routing Domain. All Routing Domains within an RDC must also be reachable from each other without the route passing through a Routing Domain that is outside of the RDC; this is a simple consequence of a route not being able to re-enter an RDC.

Routing Policies can refer to entire RDCs in the same way that single RDs are referred to, which enables the straightforward specification of routing policy rules that apply to whole classes of RD. There is no requirement for there to be co-ordination of routing strategies or the adoption of any common routing policy rules. However, efficiencies can result from the co-ordination of Addressing Plans and policies.

Figure 1-8 illustrates the RDC concept. RDCs are simply groupings of Routing Domains. A Routing Domain may be a member of zero, one or more RDCs, and hence RDCs may overlap, may be nested, and may be disjoint. RDCs are first a shorthand way of referring to communities of Routing Domains, but are at their most powerful when they are closely related to Address assignment and when combined with IDRP's features for *route information reduction* and *route aggregation* (see 1.6.2 below). RDCs are also essential for ensuring that the size of a route's path information does not ~~itself become a limit of the size (or more specifically the diameter) of~~become so large as to be a limiting factor in the Internet.

1.6.1 Limiting the Size of Path Information

In complex Internet topologies, it is possible that routes may loop back on themselves if some mechanism is not introduced to detect and suppress looping routes. This is simply achieved in IDRP by including the unique identifier of each Routing Domain that a route passes through, as part of a route's path information. A Routing Domain adds its identifier to every route that it advertises, and a simple loop test may then be introduced for each received route. However, a consequence of this is that the path information associated with each route grows each time that the route is re-advertised.

The IDRP protocol limits the length of the message that conveys each route and, anyway, routers will need to impose a limit on message length for practical implementation reasons. There is thus a limit on the number of Routing Domain Identifiers that the path information can contain, and, without RDCs, this limit will provide an upper bound on the number of Routing Domains through which a route may pass. If all Routing Domains in an Internet are to be able to communicate with each other, then this limit translates into a limit on the "diameter" of the Internet, and hence a point beyond which the Internet cannot grow.

However, in IDRP, path information reduction is also a feature the RDC.

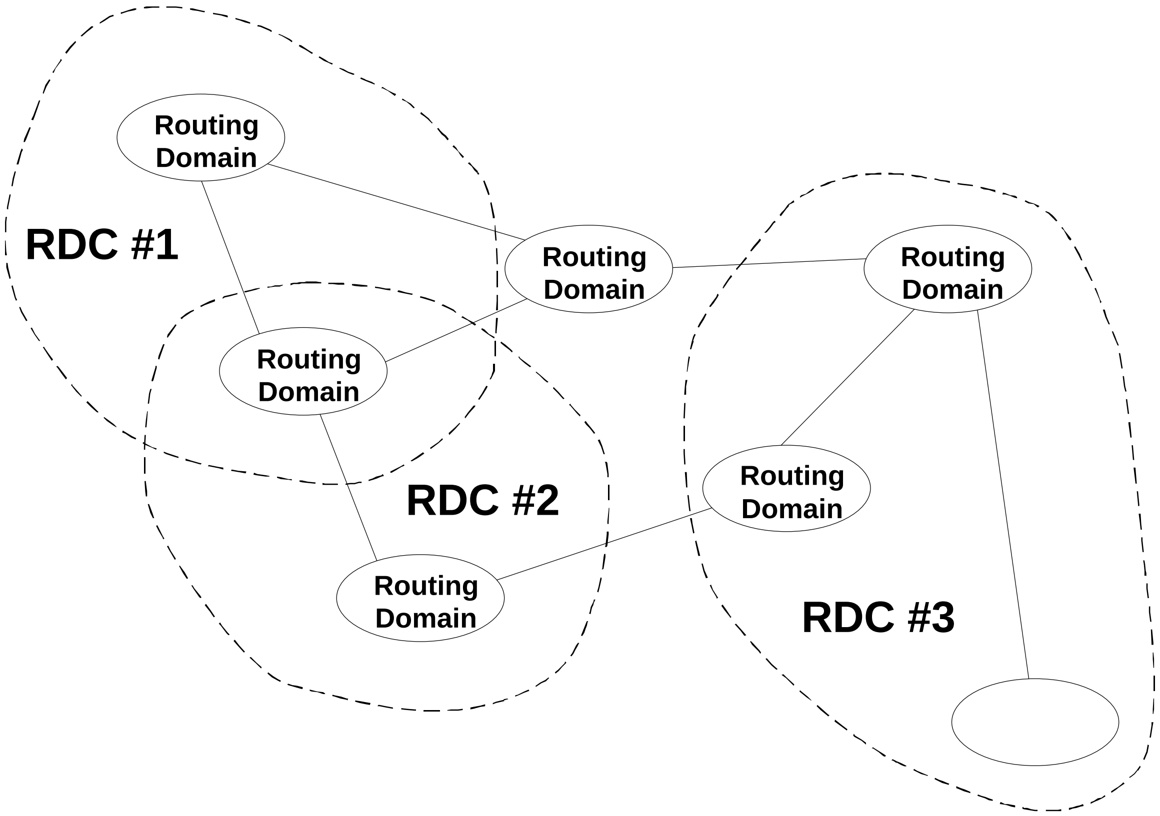

Figure 1-8 Routing Domain Confederations

When a route is advertised to a Routing Domain outside of the RDC, the Routing Domain Identifiers in the route's path information that identify Routing Domains in the RDC, are replaced by a single identifier - that of the RDC itself. Through such a mechanism, path information may be reduced, and a scaleable routing architecture achieved. The proper deployment of RDCs ensures that the diameter of the ATN Internet is not limited by the capability of routers to process path information.

The entire ATN is itself specified as an RDC. This enables the interconnection of the ATN with other Internets without limiting interconnection scenarios due to the total network diameter becoming greater than the maximum permitted by routers' path information handling capabilities.

1.6.2 Route Aggregation and RDCs

Route Aggregation is the process by which two or more routes are combined into a single route that replaces the original route. Route Aggregation is complemented by Route Information Reduction, which is the process by which the set of address prefixes that identifies the destination(s) of a route is replaced by fewer shorter address prefixes.

Typically, when two or more routes are aggregated, the destination address prefixes are combined together as the destination of the aggregated route, and Route Information Reduction is then applied to this combined set of address prefixes. The result of these two processes is a overall reduction in the number of routes without increasing the detail associated with the route's destination.

Route Aggregation may occur at any at point in an Internet. However, RDC boundaries can provide an ideal point at which to perform aggregation. This is not only because the path information can also be simultaneously reduced, but also because RDCs help simplify the management of Route Information Reduction.

For example, if an RDC is formed from all Routing Domains with a common six octet address prefix, then whenever a route exits to that RDC, it is possible to aggregate all routes to destinations inside the RDC and for the destination of that route to always comprise that six octet address prefix only. This being irrespective of whether all Routing Domains within the RDC are currently online, or whether all address combinations are even allocated. No ambiguity exists because of the fact that no Routing Domains with addresses deriving from that six octet address prefix can exist outside of the RDC.

Although it is not essential to define an RDC for this purpose, RDCs simplify the management of the reduction of addressing information.

RDCs thus perform an essential and key role in the implementation of a scaleable Internet. By reducing path information at RDC boundaries and providing a straightforward approach to controlling route aggregation, RDCs can be used to ensure that routing is not a limit on the size of the ATN Internet.

1.7 Routing in the ATN Ground Environment

1.7.1 A General Model for ATN Routing

A general model of ATN Routing is shown in Figure 1-9. 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. ATM avionics on board aircraft are then linked to the rest of the network through~~,~~ satellite, VHF and Mode S datalinks, as appropriate, and may have more than one air/ground datalink in use simultaneously.

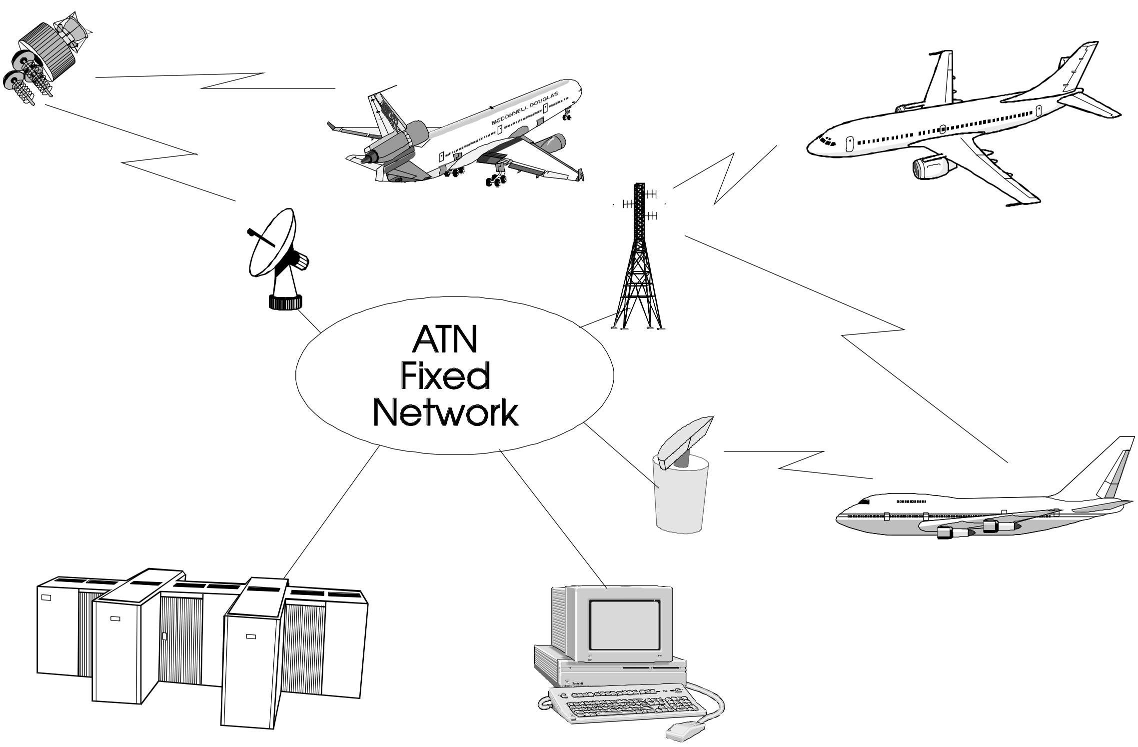

Figure 1-9 General Model of ATN Routing

1.7.2 Routing in the Ground Environment

The fixed network is not a single entity but itself consists of many different networks all linked together, as illustrated in Figure 1-10. 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~~; h~~. Host computers are directly connected to a nearby data network, typically a LAN.

User data is switched by the routers as discrete packets formatted according the ISO 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.

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.

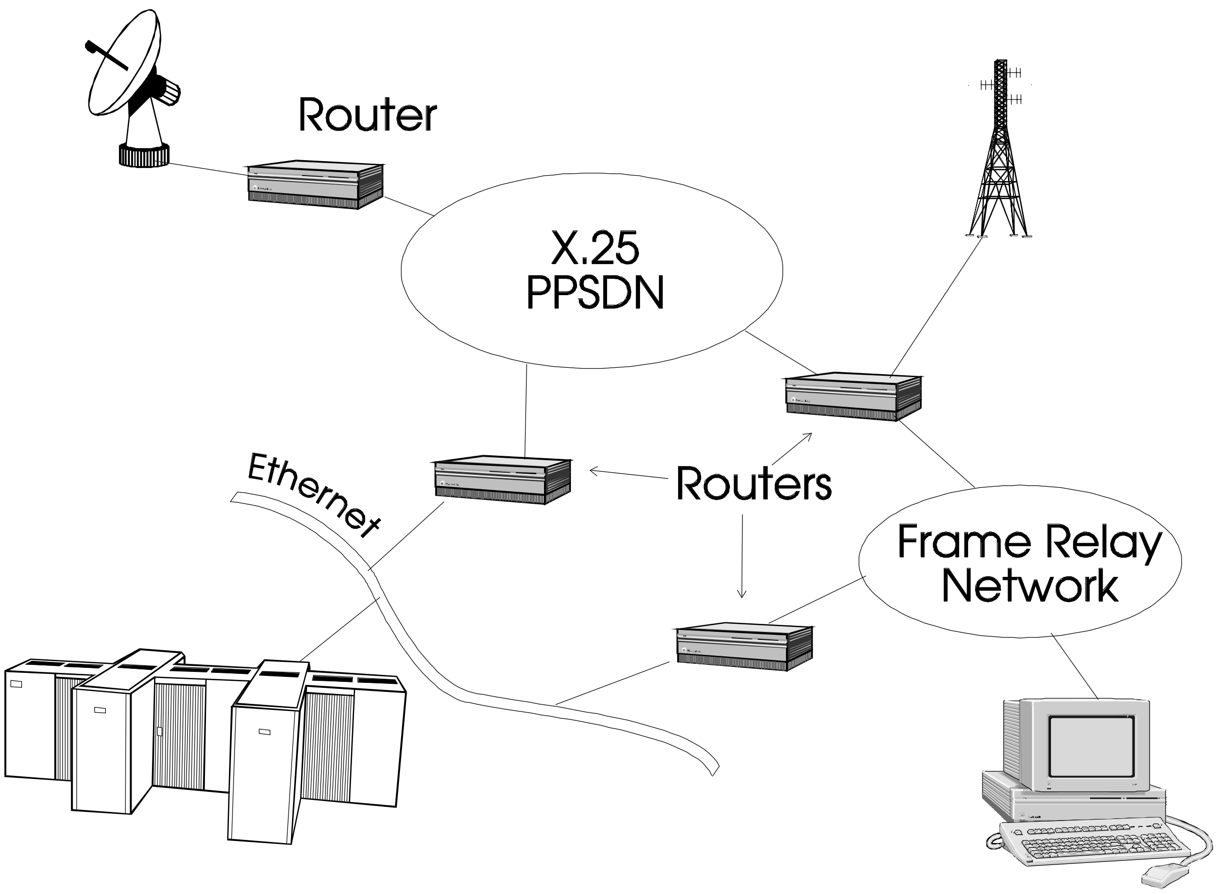

Figure 1-10 The ATN Ground Environment

1.7.3 The ATN Ground Environment

The ATN Ground Environment will comprise an Administrative Domain for each organisation participating in the ATN, and each such organisation will implement one or more Routing Domains, with IDRP used to exchange routing information.

In the ATN, in addition to the ATN wide RDC, it is anticipated that Administrations and ATN regions will, wherever possible, organise their addressing plans and form RDCs, such that route aggregation can keep the amount of routing information passed between organisations and regions to an absolute minimum.

The ATN Addressing Plan apportions a separate part of the address space to each ICAO Administration and to each IATA airline and other organisations. This allows for great flexibility in use, however, participating organisations are strongly recommended to co-ordinate the allocation of addresses to maximise the possibilities of route aggregation. 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.

1.8 The Mobile Routing Concept

1.8.1 Mobility and Routing Domains

While the scaleability of an Internet ~~demands that~~is enhanced when Routing Domains near to each other are characterised by similar address prefixes, this is not an absolute requirement. Routing Domains can be adjacent, have totally dissimilar address prefixes and still interconnect successfully. Furthermore, with a dynamic routing protocol, such as IDRP, two Routing Domains need only to interconnect when they need to, and ~~are~~can both be active on the same network. The onward re-advertisement of routes can inform the rest of the ATN Internet about such a temporary connectivity while it exists, and the loss of connectivity when it occurs. A Routing Domain can thus temporarily join an Internet at one point of attachment, then disconnect and join the Internet at some other point, the only impact being in the efficiency of routing information distribution, and eventually on scaleability.

This property of the routing architecture and of IDRP, is exploited by the ATN to support Mobile Routing.

In the ATN, the systems onboard an aircraft form a Routing Domain unique to that aircraft and characterised by one address prefix for ATSC systems, and another for AISC systems. As an aircraft proceeds on its route, it interconnects with ground based Routing Domains over the various air/ground networks, the actual network used and Routing Domain interconnected with dependent on the aircraft's actual position, and the airline's routing policy. Routing Information is then exchanged between ground Routing Domains, using IDRP, so that all ground Routing Domains are aware of the current route to that aircraft. This is illustrated in Figure 1-11.

In this example, there are four ground based Routing Domains RD1 through to RD4. RD1, RD2 and RD3 all support air/ground datalinks, while RD4 depends on the other three for air/ground communications. The aircraft currently has communications over air/ground datalinks with both RD2 and RD3.

Using IDRP, both RD2 and RD3 advertise a route to the aircraft's systems, to RD4. RD4 chooses between these two available routes using its own Routing Policy, which might, for example, favour the route through RD3. Similarly, the aircraft's router must choose between the routes to RD4 offered by RD2 and RD3. It need not make the same choice as RD4.

As the aircraft continues on its journey, it may lose communication with RD3. For example, it goes out of range of the VHF datalink it was using to communicate with RD3. RD3 informs RD4 of this situation by issuing the appropriate IDRP protocol action to withdraw the route, and RD4 now changes to using the route offered by RD2, as it is now the only route to the aircraft. The aircraft's router also recognises the loss of communication with RD3 and must now route all traffic via RD2.

Further on the journey, the aircraft comes into contact with an air/ground datalink offering communication with RD1. A datalink is established and routing information exchanged. RD1 now advertises the new route to the aircraft, to RD4. RD4 now once again has two routes to the aircraft and must make a choice between them using its local routing policy rules. It might, for example, now prefer the route through RD1, in which case all data to the aircraft is now routed via RD1. The router in the aircraft also goes through a similar decision process.

While the topology of the ATN ground environment is much more complex than the above example, this is essentially how mobile communications is implemented by the ATN.

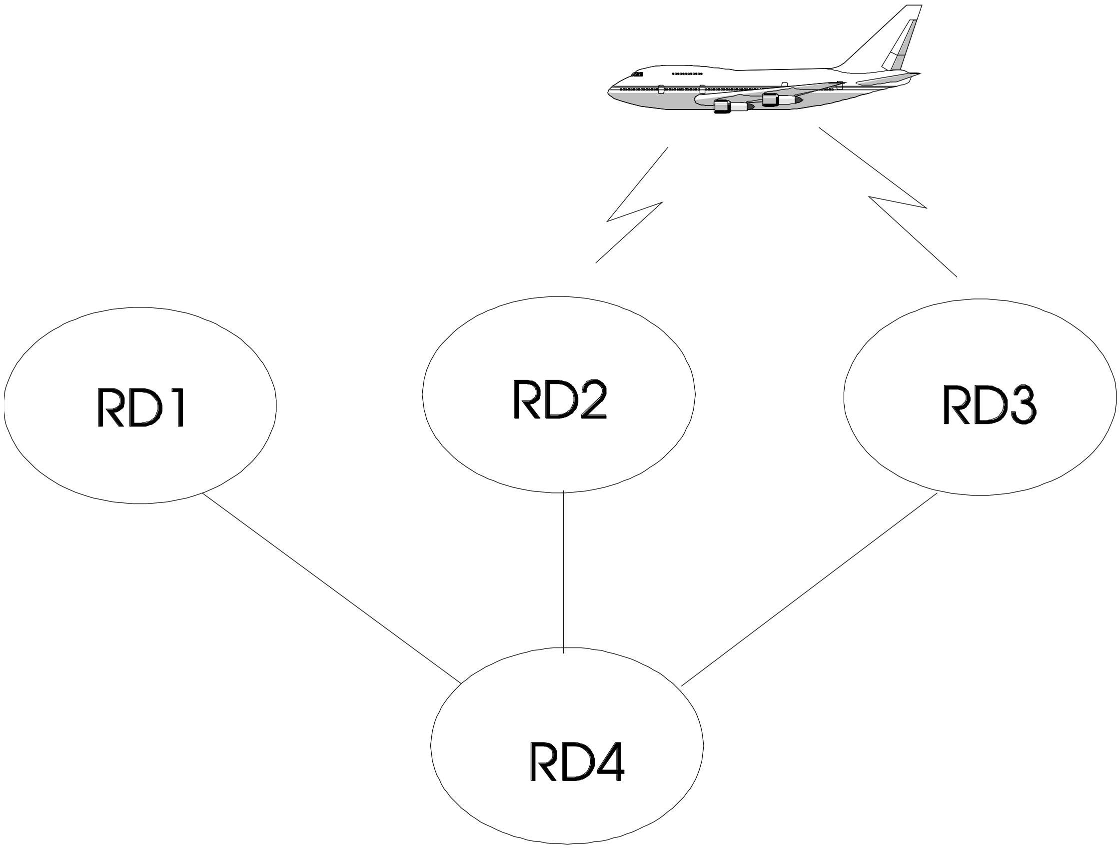

Figure 1-11 Mobile Routing Example

1.8.2 Containing the Impact of Mobility

While the principles of mobile routing outlined ~~above~~in the previous section are straightforward they are not scaleable using the existing IDRP mechanisms associated with Route Aggregation and RDCs. The problem is that even if an aircraft is given an address prefix similar to the address prefixes that characterise the ground Routing Domains at the start of its journey, such a similarity is unlikely to be maintained for the duration of the flight. Route Aggregation possibilities are thus very limited.

Instead, an alternative mechanism has been developed to permit mobility within a scaleable Internet architecture, building on two concepts: the ATN Island, and the “Home” domain (see 1.8.4 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.

1.8.3 Routing to Mobiles within an ATN Island

The ATN island exists for the exclusive purpose of supporting routing to mobiles. An ATN Island is simply an ATN region comprising a number of Routing Domains, some of which support air/ground datalinks. These Routing Domains form an RDC, as illustrated in Figure 1-12, and an ATN Island is essentially an RDC in which certain Routing Policy rules are followed. All ATN Routing Domains that have air/ground datalink are members of an ATN Island and, although most ATN Routing Domains which do not have air/ground datalink capability will also be members of ATN Islands, they do not have to be and can still have access to routes to aircraft if they are not a member of an ATN Island RDC. Routes to destinations in ground based Routing Domains will be exchanged by ATN Routing Domains, both within an Island and between Islands. However, this is outside of the context of the ATN Island. The ATN island exists to support routing to mobiles and only applies to this case.

Within each ATN Island, at least one Routing Domain forms the Island's *backbone*. This is anothermay be only one RD or may actually be an RDC comprising all backbone Routing Domains in the same ATN Island.

Within the ATN Island, the Backbone RDC provides a default route to *all aircraft*, as illustrated in Figure 1-12, this is advertised to all other Routing Domains within the Island as a route to the common address prefix for all aircraft.

Routing Domains with routes to aircraft then have a simple routing policy rule to determine to which adjacent Routing Domain they must advertise such a route[1]. This is the Routing Domain currently advertising the preferred route to *all aircraft*. This will be a backbone Routing Domain if such a Routing Domain is adjacent, otherwise it will be(or a Routing Domain that provides a route to the backbone.). Either way the impact of such a policy rule is that the Backbone RDC is always informed about routes to all aircraft currently reachable via datalinks available to the Island's Routing Domains, and can thus act as default route providers for packets addressed to airborne systems.

[1] A route to an aircraft is readily identifiable from the destination address prefix, as all address prefixes that characterise an aircraft Routing Domain descend from a unique address prefix.

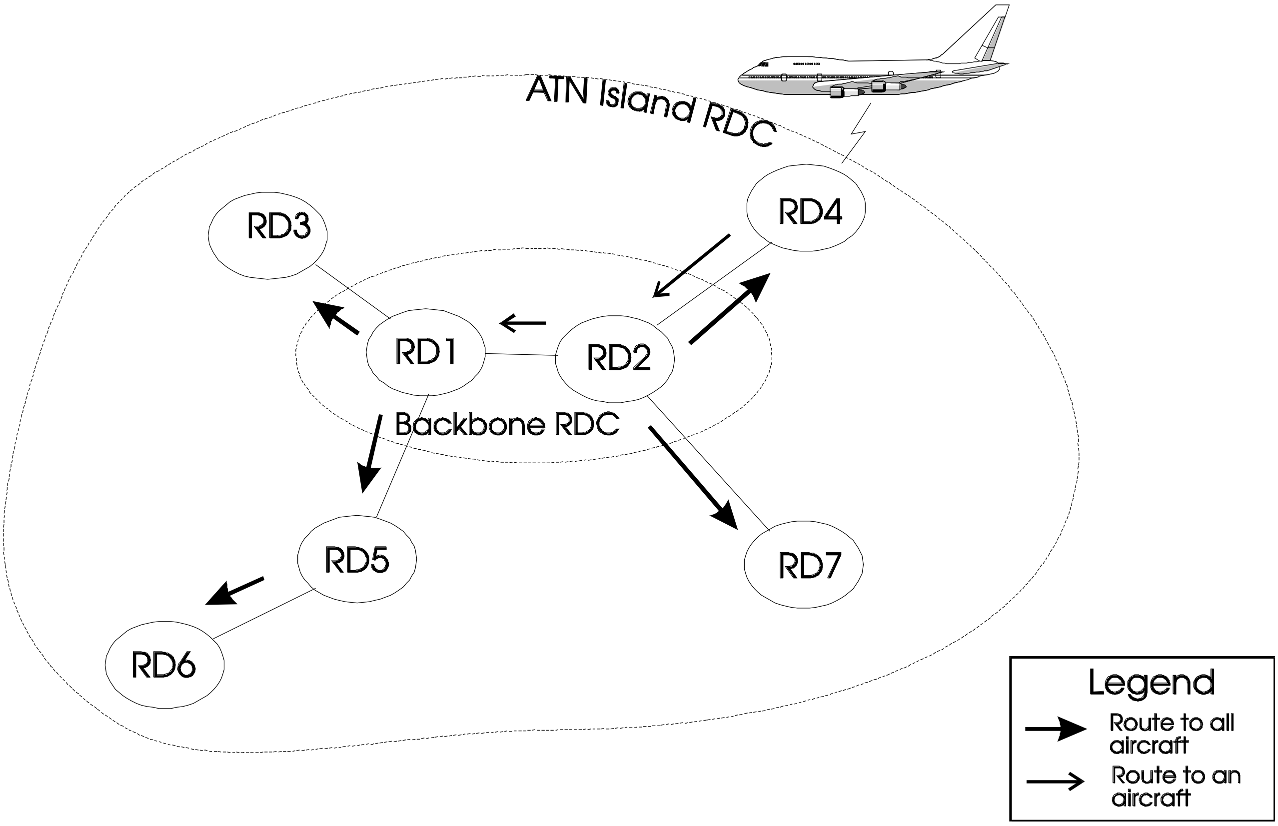

Figure 1-12 Mobile Routing Within an ATN Island

Routing Domains off the backbone also have a simple routing decision to make when they need to route a packet to a given aircraft. It is routed along the explicit route to the aircraft if it is known by them, or on the default route to all aircraft via the backbone~~, otherwise.~~. Routing with IDRP always prefers routes with the longest matching address prefix. ~~Therefore,~~Since the default route to all aircraft is always a shorter prefix of that for an explicit route to an aircraft, ~~and t~~the explicit route to an aircraft will be preferred (since it will always have a longer matching address prefix). This routing strategy happens automatically without any special provisions.

The example above is not the only policy rule that can apply to routes to aircraft. Routes to aircraft can be advertised to any other Routing Domain within the Island, provided that a policy rule is set up to allow this. This may be because there is a known communication requirement which makes bypassing the backbone desirable, or because it is desirable to provide a second (hot standby) route to aircraft from the backbone. The architecture accommodates these requirements. The only limitation on this is that imposed by the overhead of supporting routes to mobiles (see 1.8.7 below).

Within the Backbone RDC, all Routing Domains must exchange all routes to aircraft, which are advertised to them, they are then able to act as default routers to any aircraft currently in communication with the ATN Island. However, because the backbone routers need to know routes to all such aircraft, their capacity places a limit on the number of aircraft that can be handled by an ATN Island and hence on the effective size of the Island.

The ATN Island is only the first part of achieving a scaleable routing architecture for mobile routing. Its true benefit is to focus the overhead of handling the potentially large number of routes to aircraft on a few specialised routers in the backbone. Off the backbone, a Routing

Domain with an air/ground datalink needs only the capacity to handle the aircraft supported by its datalink, and there is a similar impact on Routing Domains that are Transit Routing Domains providing a route between the backbone and an air/ground datalink equipped Routing Domain. For all other Routing Domains on the Island, there is no impact on routing overhead due to aircraft.

In the absence of a backbone, all routers within the Island would need to be explicitly informed with a separate route to each aircraft, if they were to be able to route to any aircraft currently in contact with the Island. This is because there is very little probability of route aggregation with routes to aircraft.

1.8.4 Routing to Mobiles between ATN Islands

ATN Islands can be set up such that their geographical spread matches Air Traffic Control communication requirements and, for ATC purposes, there may not be a requirement to provide inter-Island communications in respect of aircraft. However, airline operational requirements are perceived to require this, and hence the mobile routing concept is developed to provide a greater level of scaleability.

The mechanism used to achieve this derives from the concept of the "Home" domain.

Aircraft for which inter-Island communications are required must have a "Home" domain, which is a Routing Domain in an ATN Island's backbone. This "home" need not be in either the ATN Island through which the aircraft is currently reachable, or in the ATN Island with which communication is required. The role of the "Home" domain is to advertise a default route to all the aircraft belonging to an airline, or the General Aviation aircraft of a given country of registration. This default route is advertised to all other ATN Island's backbone routers.

The operation of the "Home" domain is illustrated in Figure 1-13. In this example, ATN1 is the ATN Island acting as the "Home" for all aircraft belonging to the same ~~as~~ airline as the aircraft illustrated as currently reachable via ATN4. ATN1 advertises the default route to all such aircraft to all Islands in which it is in contact and, depending on local policy this route may be re-advertised to other Islands. In the figure, ATN3 re-advertises the default route on to ATN4.

The backbone routers of an ATN Island have a simple policy rule to implement for each explicit route to an aircraft that they have available. If a default route to all the aircraft in the aircraft's airline or country of registration exists[2] then the actual route to the aircraft is advertised to the Routing Domain advertising that default route. Otherwise, the explicit route is not advertised outside of the Island. In Figure 1-13, 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.

[2] Such a route is generated by the "Home" Domain , and is readily identifiable from the destination address prefix, as all address prefixes that characterise an aircraft belonging to the same airline descend from a unique address prefix.

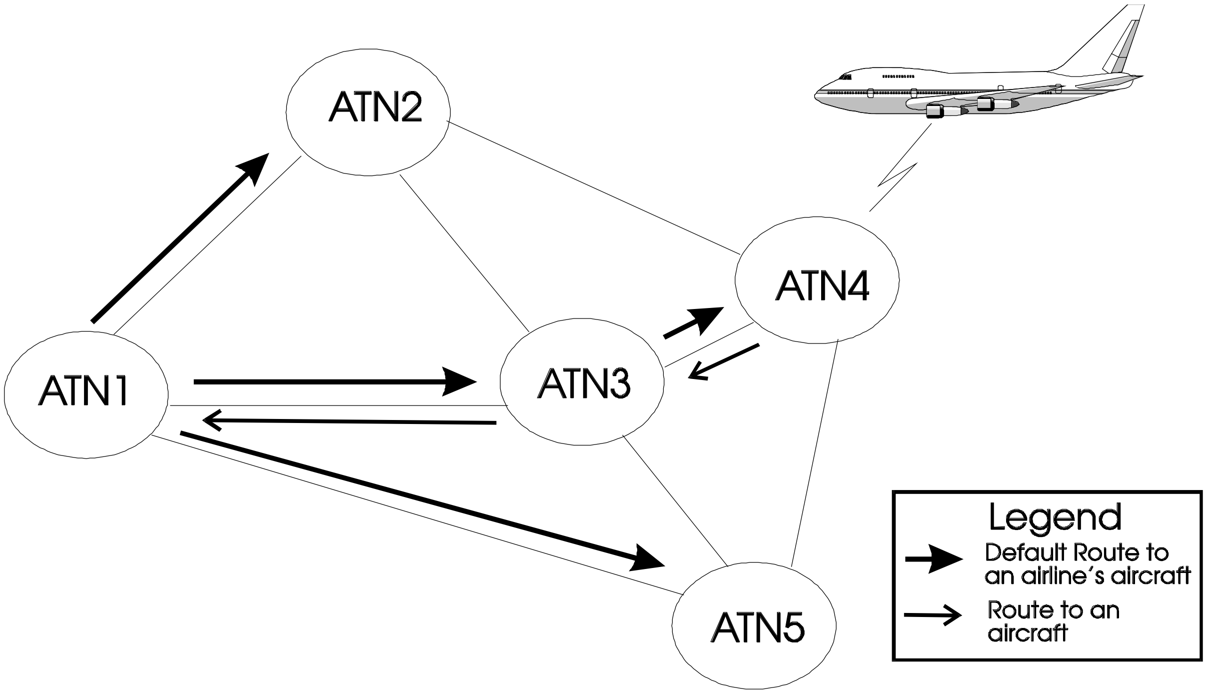

Figure 1-13 Inter-Island Routing

The impact of this rule is that the “Home” is always kept aware of routes to all of “its” aircraft. As it is also providing the default route to such aircraft, routers on other ATN Islands (e.g. ATN2) that have packets to route to one of that “Home’s” aircraft will by default send those packets to the “Home” Routing Domain (ATN1), where the actual route to the aircraft is known, and thus the packet can be successfully routed to the destination aircraft (via ATN3 and ATN4).

In the above example, this is clearly non-optimal as ATN4 can be reached directly from ATN2. However, the loss of optimal routing is acceptable as, otherwise a scaleable architecture could not have been developed.

The impact of this strategy on routing overhead, is that an ATN Island backbone has to be capable of handling routes to all aircraft currently in contact with the Island, and all aircraft for which it is the “Home”. ~~Thus, and assuming that all ATN Islands are fully interconnected, if there are at most ‘n’ aircraft in contact with the Island, and the Island is~~

~~“Home” to ‘m’ aircraft then:~~

~~n + m < “maximum number of routes to mobiles that can be handled by a backbone router”~~

~~has to be true.~~

However, this ~~limi~~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 th~~is limit.~~e 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.

1.8.5 ATN External Interfaces and Mobiles

As ~~discussed above~~presented in previous sections, the ATN is itself an RDC and this will prove to be very useful should it ever prove necessary to provide access to ATN mobiles to other Internets. This is because at an RDC boundary, such as at the ATN boundary, Route Aggregation and reduction of path information can readily take place. In this case, it is possible to aggregate all routes to aircraft into a single route with a destination given by the address prefix for all aircraft. As the path information for such an aggregated route is also collapsed to a single ATN RDC identifier, the complexity of routing information exported at the ATN boundary can be kept to a simple single route that is independent of the number of aircraft and ATN internal complexity.

1.8.6 Impact on Air/Ground Datalinks

A final limiting factor on the ATN is the capacity of the air/ground datalinks. At present, these are low bandwidth communications channels and only the minimum routing information can be transferred over them.

IDRP is potentially an ideal protocol for this environment. Techniques such as RDCs and Route Aggregation can be used to minimise the information contained in each route. Furthermore, two or more routes to the same destination that differ only in security parameters, or service quality metrics, can be ~~multiplex~~combined together into a single message keeping the actual information exchanged to a bare minimum.

In addition, IDRP is a connection mode protocol and, as such, once a route has been advertised between a pair of Boundary Intermediate Systems it does not have to be retransmitted during the lifetime of the connection. A BIS-BIS connection is kept alive by the regular exchange of small "keepalive" packets, and once routing information has been exchanged it remains valid for the lifetime of the connection without having to be retrans~~i~~mitted.

The ATN uses these properties of IDRP to keep the transfer of routing information over an air/ground datalink to a minimum. When the datalink is first established, the airborne router will advertise a route to internal destinations for each combination of traffic (security) type and QoS metric supported. These routes will be combined into a single protocol message and downlinked for onward distribution through the ground ATN.

The ground router will also uplink routes to the aircraft and to keep the information down to a minimum, a further RDC is defined, comprising all ground ATN Routing Domains. This RDC, the "ATN Fixed RDC" ensures that for each uplinked route, the path information is collapsed to a single identifier, that for the ATN Fixed RDC.

The actual routes uplinked are subject to the policy of the ground router's Routing Domain. However, it is anticipated that routes will be provided to at least:

- the local Routing Domain (typically that providing Air Traffic Services), and
- the ATN as a whole,

in addition to other routes as determined by local policy.

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

1.8.7 The Impact of Routing Updates

~~The above discussion has illustrated how~~As indicated in the previous section, a scaleable routing architecture can be developed in support of mobile routing. It is now necessary to consider the factors that limit the number of routes to aircraft that an ATN Router can handle.

Each route known to a router occupies a certain amount of data storage and, while data store can be a limiting factor on the total number of routes handled, it is unlikely to be so in this case. The number of route updates that a router can handle is more than likely to be the limiting factor.

In the ground environment, route updates will usually only occur when changes occur in the local region of the Internet (changes further away are hidden by route aggregation). Typically the introduction of a new Routing Domain or interconnection, or the removal or loss of one of these will cause a change. However, the frequency of update is unlikely to be high.

However, with mobiles, such as aircraft, the situation is very different. Aircraft are constantly on the move, changing their point of attachment to the ATN, and hence generating routing updates. The impact of these updates needs to be minimised if the number of aircraft that can be handled by an ATN Island is to be maximised, and an important and useful feature of IDRP can be exploited in order to help meet this objective.

1.8.7.1 “Hold Down” Timer Use

Vector distant routing protocols, such as IDRP, typically implement a “hold down” timer, which introduces a minimum delay between the receipt of a route and its re-advertisement. This timer is used to avoid instability due to frequent route changes, and the actual value of the timer is then usually a trade-off between a short timeout to give rapid response and a long timer to keep down routing overhead and minimise instability.

However, under IDRP, routing events that indicate a major change (i.e. new route or loss of a route) are not subject to a hold down timer, only those that report a minor change to an existing route are subject to a hold down timer. This means that IDRP is very responsive to connectivity changes while avoiding instability due to minor changes. For example, consider a simple extension to the previous example, illustrated in Figure 1-14.

In this example, RD4 provides a route to the aircraft, to RD5. When the aircraft loses contact with RD3, RD4 is immediately informed, as there is an effective zero length hold down timer for withdrawn routes. However, while RD4 recognises this event and switches to the route provided by RD2, it does not necessarily inform RD5 of this now minor change to the route immediately (the route still exists, only the detail of the path is different), and anyway, the update must be sent not less than the period **minRouteAdvertisementInterval** since any previous update. In this example, it should be noted that the minor change will not affect RD5’s routing decision, as it has no alternatives available.

Sometime later, the aircraft comes into contact with RD1. RD4 is immediately informed as this is a new route. However, even if RD4 switches to this new route, it does not inform RD5 of the change until the **minRouteAdvertisementInterval** has again expired.

This has important implications for the design of an ATN Island. If an Island’s air/ground datalinks are all connected to Routing Domains which are themselves adjacent to the Backbone RDC, all connectivity changes will be immediately reported to the Backbone giving a high route update rate. On the other hand, if there are intermediate Routing Domains between the backbone and the Routing Domains connected to air/ground datalinks, then the update

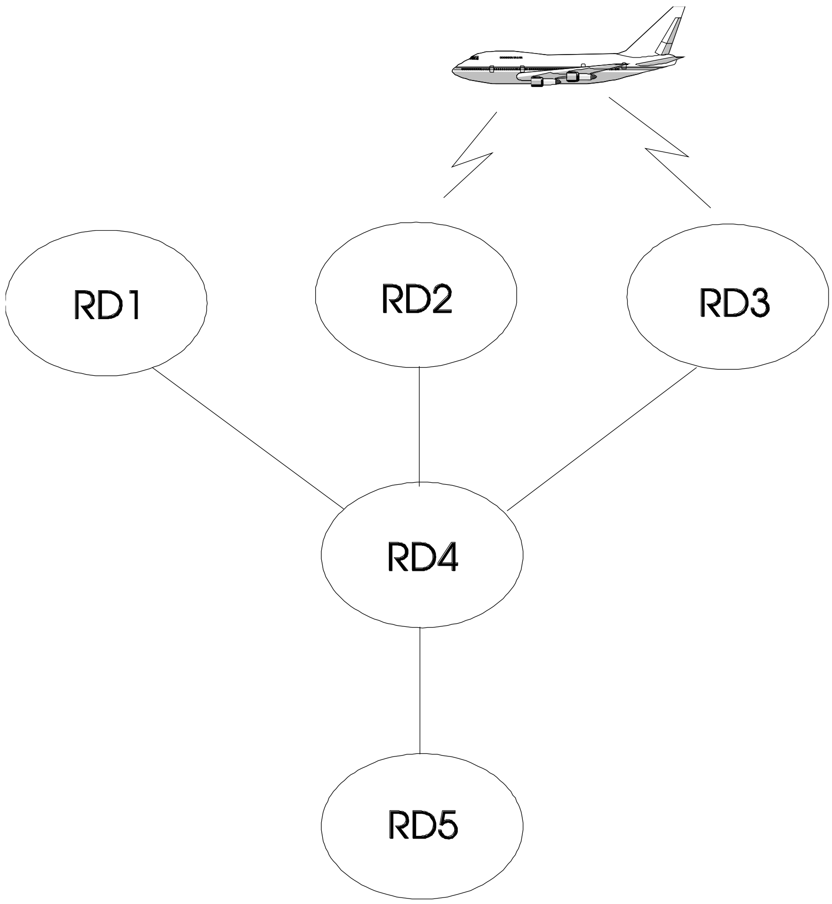

Figure 1-14 Impact of a Hold Down Timer

frequency can be significantly reduced, without affecting the responsiveness to real connectivity changes.

This is an important benefit derived from using IDRP to support mobile routing compared with, for example, a directory based approach to mobile routing. Under a directory based approach, there would be a central directory server on each ATN Island (c.f. the Backbone), updates on the position of aircraft would be sent direct to the directory, and other routers would consult the directory in order to determine the current location of a specific aircraft. In terms of overhead, this situation is analogous to an ATN Backbone Routing Domain directly connected to each Island Routing Domain with air/ground datalink capability, and the directory has to be able to take the full update rate. IDRP can, however, distribute the update load throughout the ATN Island.

Routes advertised to an aircraft's "Home" are also affected by the hold down timer and, in this case, RDCs and the Hold Timer work together to keep the routing overhead to an absolute minimum.

As an ATN Island is an RDC, routes advertised to other Islands have their path information for the transit through the RDC replaced by a single RDC identifier, and therefore, in many cases, changes in the route will not even be visible to another ATN Island. When changes are visible (e.g. a change in hop count or QoS metric), and such changes can be kept to a minimum by careful network design, then the Hold Timer limits the rate at which such changes can be advertised and prevents minor changes which are also short lived, being exported outside of the Island.

1.8.8 Failure Modes

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

1.8.8.1 Loss of the "Home"

Loss of the "Home" may come about from either the loss of the Routing Domain advertising a route to the "Home" for a given set of aircraft, or the loss of the communications path to it. The consequence of either failure is clear: the affected aircraft are now only reachable from systems on the ATN Island to which they are currently adjacent.

In practice, there should not be a single point of failure related to the "Home" Routing Domain. A Routing Domain may comprise many Boundary Routers, each of which may advertise the route to the "Home". Only loss of all of these Boundary Routers 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 Boundary Routers, or between different pairs. Only if all such links are lost, will total loss of communications occur.

Therefore, it will always be possible to design a network topology that will avoid the loss of the "Home" being due to any single failure, and which can ensure that the probability of loss of the "Home" is kept within acceptable limits. Where inter-Island communications are required in support of air safety, then the design of the Inter-Island ATN topology must be supported by an appropriate failure mode analysis to ensure that safety limits are maintained.

1.8.8.2 Failure of an ATN Island Backbone

Failure of an ATN Island may also result from the failure of the Routing Domain(s) that comprise an Island's Backbone, or of communications paths with an Island's backbone. The consequence of such a failure is that the aircraft currently adjacent to the Island are only reachable from the Routing Domains supporting air/ground datalinks with those aircraft, and any other Routing Domains on the Island to which routing information to those aircraft is advertised according to explicit policy rules.

For similar reasons to those already d~~iscussed above~~etailed in 1.8.8.1, there is no need for loss of an Island Backbone to be due to a single point of failure, and an appropriate network design should be developed for each ATN Island to ensure that the probability of the loss of the backbone is within acceptable limits.

1.8.9 Optional non-Use of IDRP

Simple networks can often avoid dynamic routing mechanisms in favour of statically defined routing tables, initialised by a System Manager. However, even in the early ATN, the existence of Mobile Systems does not permit the general use of static routing techniques. Aircraft may join and leave the air/ground subnetwork(s) at any time and this dynamic behaviour must be recognised by the routers and reflected in the routing tables. Some dynamic adaptive routing protocol is needed to support this requirement. IDRP is specified for this purpose. However, implementing IDRP functionality on an airborne router may not be practicable in ~~early~~the early stages of ATN implementation.

An alternative approach is possible using provisions in the ISO 9542 ES-IS protocol. An exchange of Intermediate System Hello (ISH) PDUs is already required as part of the route initiation process, and, in a limited topology, an exchange of ISH PDUs can be sufficient to provide the exchange of dynamic routing information necessary to support mobile routing. Furthermore, a regular exchange of ISH PDUs (part of the normal operation of ISO 9542) can be used to keep the link between ground and airborne routes "live" in the absence of IDRP.

Such a use of the ISH PDUs depends upon an assumed relationship between the Network Entity Title (NET) of each router - which is essentially the router's address - and the NSAP Addresses in the ground and airborne End Systems. The NET is exchanged as part of the ISH PDU. When the Air/Ground router receives an ISH PDU from an airborne router, it may infer from the ATN Addressing Plan the common NSAP address prefix of all NSAPs onboard that aircraft. This being the first eleven octets of the NET. This NSAP Address Prefix may then be used as the destination of a route to the NSAPs onboard that aircraft and the route entered into the ground router's Forwarding Information Base. It is then possible for the ground End Systems to send data to airborne End Systems on that aircraft.

The same process may also take place on the Airborne Router, on the receipt of an ISH PDU from the Air/Ground router, enabling airborne End Systems to send data to ground End Systems. The routing information remains current until either a regular exchange of ISH PDUs ceases, or the subnetwork connection is cleared, when the ground and airborne routers remove the associated routes from their forwarding information bases.

The architecture of a ground router implementing this functionality is illustrated in Figure 1-15. The architecture is straightforward enough with the ES-IS protocol active on both subnetworks. Both protocol entities update the Forwarding Information Base (FIB) which is, in turn, used by the Forwarding process to route packets.

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

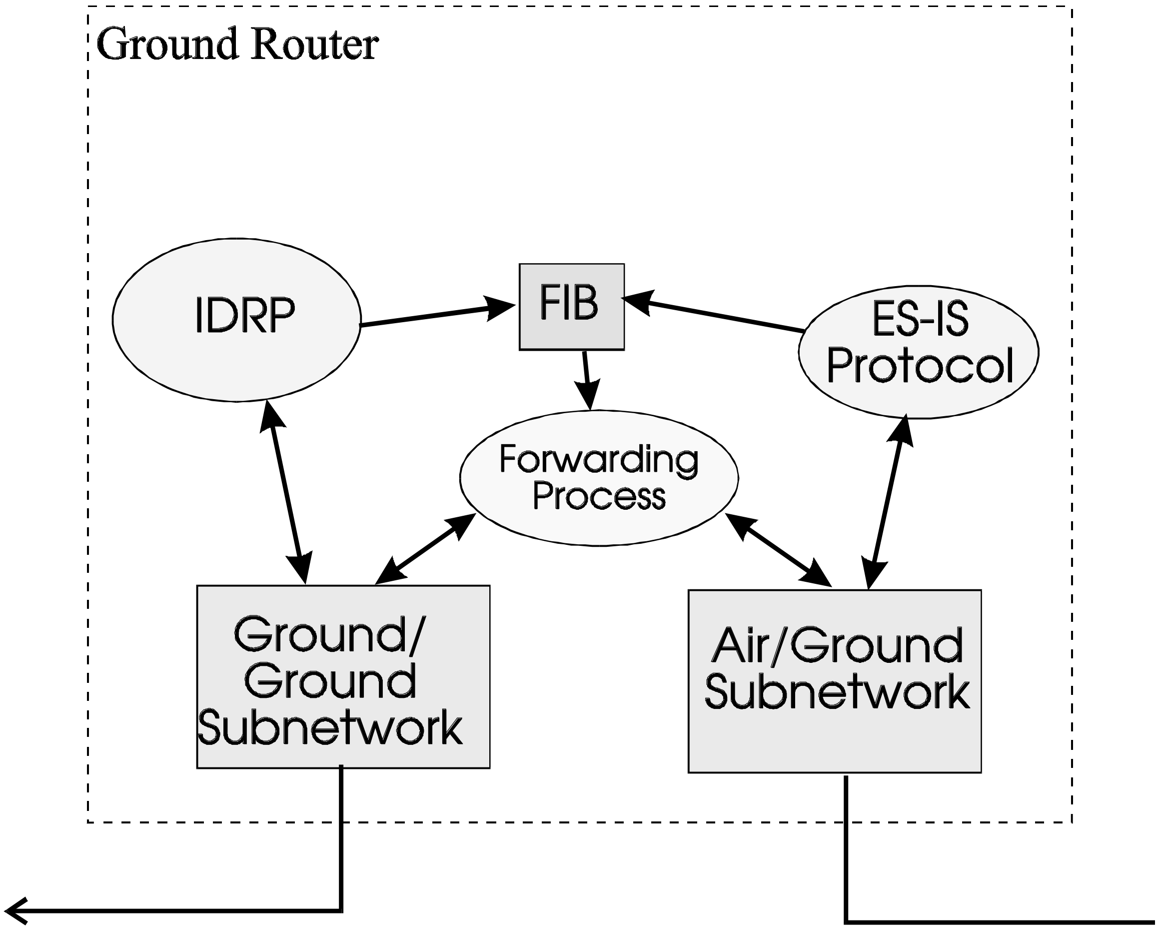

Figure 1-15 Architecture of an Initial Ground Network Router

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

1.9 Route Initiation

1.9.1 The Purpose of Route Initiation

ICAO has adopted the use of Policy Based Routing procedures for routing between ATN Routing Domains (RDs), including the support of routing to mobile systems. Dynamic Routing Information is exchanged using the procedures specified in ISO 10747 and used and disseminated according to local routing policies specified in accordance with the ATN SARPs. However, before routing information can be exchanged between any two Routing Domains, it is first necessary to establish a communications path between Boundary Routers[3] in each of those RDs. The establishment of such a communications path is known as “Route Initiation”.

[3] The term Boundary Router may be read in most cases, as synonymous with the architectural entity “Boundary Intermediate System” (BIS). In practice, a Boundary Router includes a BIS along with Management Entities and and End Systems component to support such entities.

Route Initiation procedures are required whenever two ATN RDs need to be interconnected. Since the ATN SARPs specify that, on board an aircraft, the communications systems and the applications processors that they serve comprises a Routing Domain, Route Initiation procedures also apply to the establishment of air/ground communications.

Route Initiation commences when the decision is made to establish a communications path between two ATN RDs. Route Initiation finishes upon the initial exchange of routing information between the Boundary Routers, or the unsuccessful termination of the Route Initiation procedure.

Note: Boundary Routers within the same RD also exchange dynamic routing information using ISO 10747. The Route Initiation procedures are the same as for inter-domain connections except that both Routers will be under the control of the same administrator.

1.9.2 Ground-Ground Route Initiation

1.9.2.1 The Communications Environment

Ground-Ground communications typically use long lasting physical or logical communications paths. Route Initiation can normally be regarded as a rare event and will often be only semi-automated.

The communications networks in the ATN ground environment are outside the scope of the ATN SARPs, but can be assumed to include:

1. X.25 Public and Private Data Networks
2. Leased Lines
3. Integrated Services Digital Networks (ISDNs)
4. Frame Relay Services
5. The Public Switched Telephone Network (PSTN).

The actual choice of communications network is a matter for bilateral agreement between the organisations and states that wish to interconnect their RDs, and will depend on local availability, tariffs and policies. In many cases, high speed (e.g. V.32bis or V.34) Modems and the PSTN will be used as a backup for a dedicated data network.

The communications protocols used to provide the data link will also depend upon the communications network used and bilateral agreement. In the case of X.25 data networks, Frame Relay and communications services provided via the ISDN D-Channel, then the communications protocols are mandated by the data network provider. In the case of Leased Lines and the ISDN B-channel, then HDLC LAPB (ISO 7776) is the likely choice. For the PSTN, the asynchronous communications provided by V.32bis and V.34 Modems makes the Point-to-Point Protocol (PPP) as specified in RFC 1548, the likely choice.

Note: Route Initiation is not necessarily synonymous with the establishment of an uninterrupted communications link between two Boundary Routers. For example, the speed at which an ISDN B-Channel is established is such that it may be practicable to break the communication circuit during idle periods and re-establish it when there is data to send, whilst still maintaining a logical communications path between the two Boundary Routers. Route Initiation is concerned with the establishment of the logical communications path.

1.9.2.2 Summary of Procedures

The sequence of procedures for a typical ground-ground Routing Initiation is illustrated in Figure 1-16, and summarised below. They are ~~discuss~~escribed in greater depth in the following sections. This illustrates the co-ordination of two Systems ("A" and "B") interconnecting over a common network. The procedures are:

1) Adjacent BIS MOs are established in both Systems. In each case, an MO is established to identify the other system and contains the parameters necessary to create and maintain a BIS-BIS connection with that system. Both systems will also have been configured with appropriate SNDCFs associated with each attached subnetwork.

2) A communications path is established over the subnetwork; typically one system is initiator and the other responder.

3) Establishment of the communications path is notified to the Systems Manager.

4) In response, the Systems Manager for each system adds a route to the local FIB and to the remote System, and

5) invokes the IDRP "Start Event" action, or re-run the decision process if a BIS-BIS connection already exists with the remote system.

6) On successful establishment of the BIS-BIS connection, Route Initiation completes.

Note: while the Systems Manager may be a real person explicitly issuing commands, the "Systems Manager" in the above description may alternatively be a procedural script carrying out an automatic action in respon~~d~~se to a Systems Management Notification.

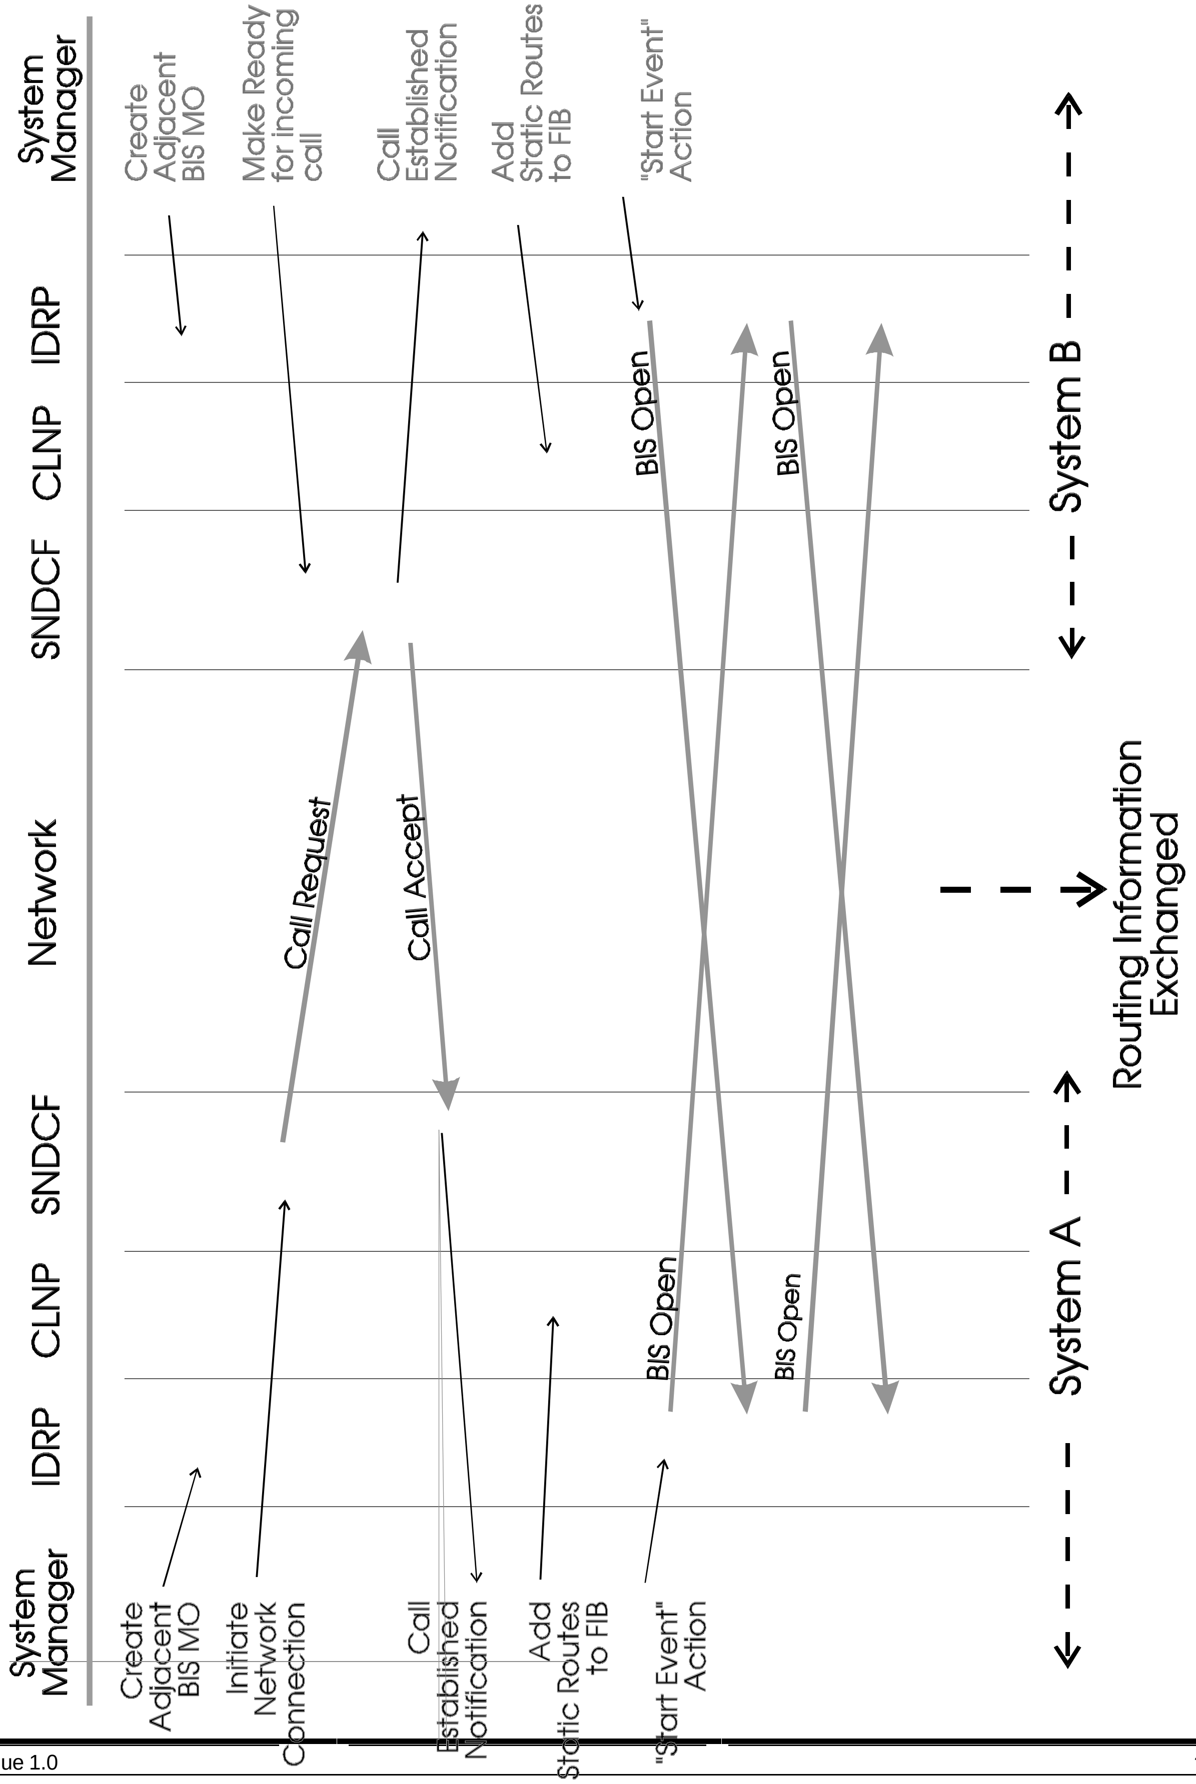
System Manager
IDRP
CLNP
SNDCF
Network
SNDCF
CLNP
IDRP
System Manager
Create Adjacent BIS MO
Initiate Network Connection
Call Established Notification
Add Static Routes to FIB
"Start Event" Action
Make Ready for incoming call
Call Request
Call Accept
BIS Open
BIS Open
BIS Open
BIS Open
System A
System B
Routing Information Exchanged

Figure 1-16 Ground-Ground Route Initiation Sequence

1.9.2.3 Initial Route Initiation

Route Initiation begins with the decision to establish a communications path between a pair of Boundary Routers, including the decision on which communications networks to use. The first procedure is to establish the underlying communications circuit between the Boundary Routers and hence to establish the logical communications path.

These procedures will be data network dependent and will require some sort of interaction between the respective Systems Managers. Typically, one Boundary Router will need to be in a passive state awaiting an incoming event (e.g. an X.25 call indication or a PSTN Ring Indication), while the other takes an active role and initiates circuit establishment (e.g. by generating an X.25 call request, or “dialling” the telephone call).

When appropriate to the type of data network used, the QoS, Security and Priority requested on any such call request, should be satisfactory for the exchange of routing information.

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.

1.9.2.4 Route Initiation in CLNP

The ATN SARPs specify the use of the Connectionless Network Protocol (CLNP) specified in ISO 8473 for ATN subnetwork independent communications. Establishing a data link (e.g. an X.25 virtual circuit) is a necessary condition for data to be exchanged between two Boundary Routers using CLNP, but not a sufficient condition. In order for the data link to be used by the CLNP Network Entity, and hence as a communications path for the forwarding of data packets, it is necessary to:

1. Assign an appropriate Subnetwork Dependent Convergence Function (SNDCF) to interface the data link to the Network Entity;

2. Update the Forwarding Information Base (FIB) to record statically known routes available over the data link and via the remote Boundary Router.

The former is necessary in order to match the characteristics of the actual network and communications protocol used over that network to the characteristics assumed by the CLNP Network Entity. The second is necessary in order to permit the exchange of dynamic routing information.

The SNDCF is typically specified for a network type and associated at system configuration time with a physical communication port. In most cases, the assignment of the SNDCF is implicit in the network over which communications is established, and no explicit action will need to be carried out to assign the SNDCF. Indeed, most implementations will required assignment of the SNDCF prior to establishment of the data link. However, for some network types there may be alternatives chosen at connection establishment time.

The FIB may be updated with any statically known routes that are known *a priori* to exist via the newly established data link, where a route consists of an NSAP Address prefix paired with an identifier for a data link. When forwarding data packets, the CLNP network entity locates the longest matching NSAP Address Prefix in the FIB, when matched against the packet’s destination NSAP Address, and then queues the packet for transmission over the associated data link. Multiple FIBs may also exist, matching different QoS and security requirements. So that Routing Information may be exchanged, the FIB associated with the QoS level used for the exchange of Routing Information, must be updated to include, as a minimum, a route to the network entity located on each Boundary Router. to which a data link has been established.

Therefore, once a data link has been established to a remote Boundary Router, 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 Boundary Router at the other end of the newly established data link. As a minimum, this prefix may be the complete NET. A; and,

b) the data link to that remote Boundary Router.

Note 1: the reverse must also take place when the data link is terminated i.e. the above route must be removed from the FIB.

Note 2: alternatively, such routes may be entered into the FIB at system initialisation. However, this strategy only gives satisfactory results if there is one and only ononly if there is a single possible data path to the remote Boundary Router.

1.9.2.5 Route Initiation in IDRP

The ATN SARPs specify the use of the Inter-Domain Routing Protocol (IDRP) specified in ISO 10747, detailed in ISO 10747, for the exchange of dynamic routing information between Boundary Routers. Once a communications path has been established between two Boundary Routers and sufficient static routing information has been entered into the local FIB in order to enable the forwarding of data packets to the remote Boundary Router itself, IDRP may be used to exchange dynamic routing information.

IDRP may only exchange dynamic routing information when a BIS-BIS connection has been established. This is a logical connection established by using the IDR PP 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 Boundary Routers.

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 Boundary Router. The information held includes the NET of the remote Boundary Router, authentication data, the specific IDRP procedures used to establish the BIS-BIS connection and timer values. One such MO exists for each remote Boundary Router with which IDRP may exchange routes. Typically, this MO is setup in advance of the underlying communications path, and will usually be created once agreement to interconnect has been reached.

Once the FIB has been updated with a route to the remote Boundary Router, the "start event" action is requested of the Adjacent BIS MO associated with that Remote Boundary Router. This initiates the procedures for creating the BIS-BIS connection and is followed by the exchange of dynamic routing information. It is the final action of the Route Initiation procedure.

During establishment of the BIS-BIS connection either or both IDRP entities will take an active role in connection establishment, or one will be active and the other passive. The role, active or passive, is determined by information configured into the Adjacent BIS MO. If one IDRP entity is to be passive, then Systems Managers must ensure that the other is configured in the active role. If both IDRP entities are configured in the active role, then the BIS-BIS connection establishment procedures are less efficient, than if one is in the passive role. However, given that the loss of efficiency is small and typically of no consequence given that ground-ground BIS-BIS connections are usually long lived, Organisations and States are recommended by the SARPs to always configure the Adjacent BIS MOs for BIS-BIS connections between ground ATN Boundary Routers for BIS-BIS connection establishment in the active role. This is to avoid to risk of both being configured in the passive role by mistake.

However, there is one exception to the above. That is when the newly established communications path is to a remote Boundary Router with which a BIS-BIS connection already exists. This is possible when multiple networks are available between the same pair of Boundary Routers. Multiple concurrent connections may be desirable in order to give high availability through redundancy and to provide additional data transfer capacity.

IDRP permits only a single BIS-BIS connection between a given pair of Boundary Routers, 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 Boundary Router and only invoke the Start Event Action if one does not already exist. This action will in any case, be ignored if issued when a connection does already exist.

However, other action may be appropriate if there is a need to recognise the different QoS that may be available when a new communications path is opened up (or lost), or a change occurs in the Security Types that may be supported by alternative communications paths to the same remote Boundary Router. In such cases, the SARPs require that the IDRP Decision Process ~~is~~be aware of the aggregate QoS and Security Restrictions over the communications paths to a given remote Boundary Router (Adjacent BIS). The SARPs require the Decision Process to update the QoS on received routes (when processing the adj-RIB-in) to reflect the QoS of the communications path and to use this updated QoS when determining the degree of preference of the route and when re-advertising it.

The~~y~~ SARPs also require that the Decision Process does not place in the IDRP adj-RIB-out, any routes with Security Types incompatible with any restrictions that exist on the aggregate communications path. For example, if none of the available communications paths to a given remote Boundary Router permits the transfer of “Administrative” data, then a route with a Security Type reflecting administrative data may not be placed in the Adj-Rib-out for that Router (and hence advertised to it).

Therefore, whenever an additional communications path to a given remote Boundary Router becomes available (or is lost), the Systems Manager must cause the IDRP Decision Process to be re-run, instead of invoking the Start Event.

1.9.3 Air-Ground Route Initiation

Air-Ground Route Initiation is similar to ground-ground Route Initiation, but differs for the following reasons:

I. ICAO specified subnetworks are used for air-ground communications with their procedures for use mandated by SARPs rather than subject to bilateral negotiation.

II. Route Initiation typically starts as soon as communication is possible e.g. an aircraft coming into range of a Mode S Interrogator, and, in consequence Route Initiation starts as soon as the Systems Manager is notified of the possibility of communications (e.g. capture by a Mode S Interrogator).

III. 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.

IV. 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.

V. 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.

1.9.3.1 Communications Environment

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

1. The Aeronautical Mobile Satellite Service (AMSS)
2. The VHF Data Link (VDL)
3. The Mode S Data Network

In each case, ITU recommendation X.25 provides the data network access procedures, and the responsible ICAO Panel's have required that:

a) AMSS communications are "air initiated", that is the aircraft is responsible for initiating communication with the ground

b) VDL communications are similarly air initiated.

c) Mode S communications are "ground initiated" that is a ground ATN Router attached to a Mode S data network is responsible for initiating communications with an aircraft.

1.9.3.2 Summary of Procedures

The Air-Ground Route Initiation procedures are illustrated in Figure 1-17, and summarised below. They are d~~iscuss~~escribed in greater depth in the following sections. This figure illustrates the case where a Join Event is generated by the air-ground subnetwork. If the subnetwork cannot generate a Join Event then the procedures start with the Call Request, as part of a polling procedure. System "A" is the initiator and System "B" is the responder. If the air-ground subnetwork is air-initiated then System "A" represents the Airborne Router, and System "B" the Ground Router. If the air-ground subnetwork is ground-initiated, then System "A" represents the Ground Router, and System "B" the Airborne Router.

The Route Initiation Procedures are:

1) When an aircraft attaches to an air-ground subnetwork, a Join Event is generated, potentially to both Airborne and Ground Routers. If received by System "B", the Join Event is ignored; System "B is ready to receive incoming calls as soon as it attaches to the Mobile Subnetwork.

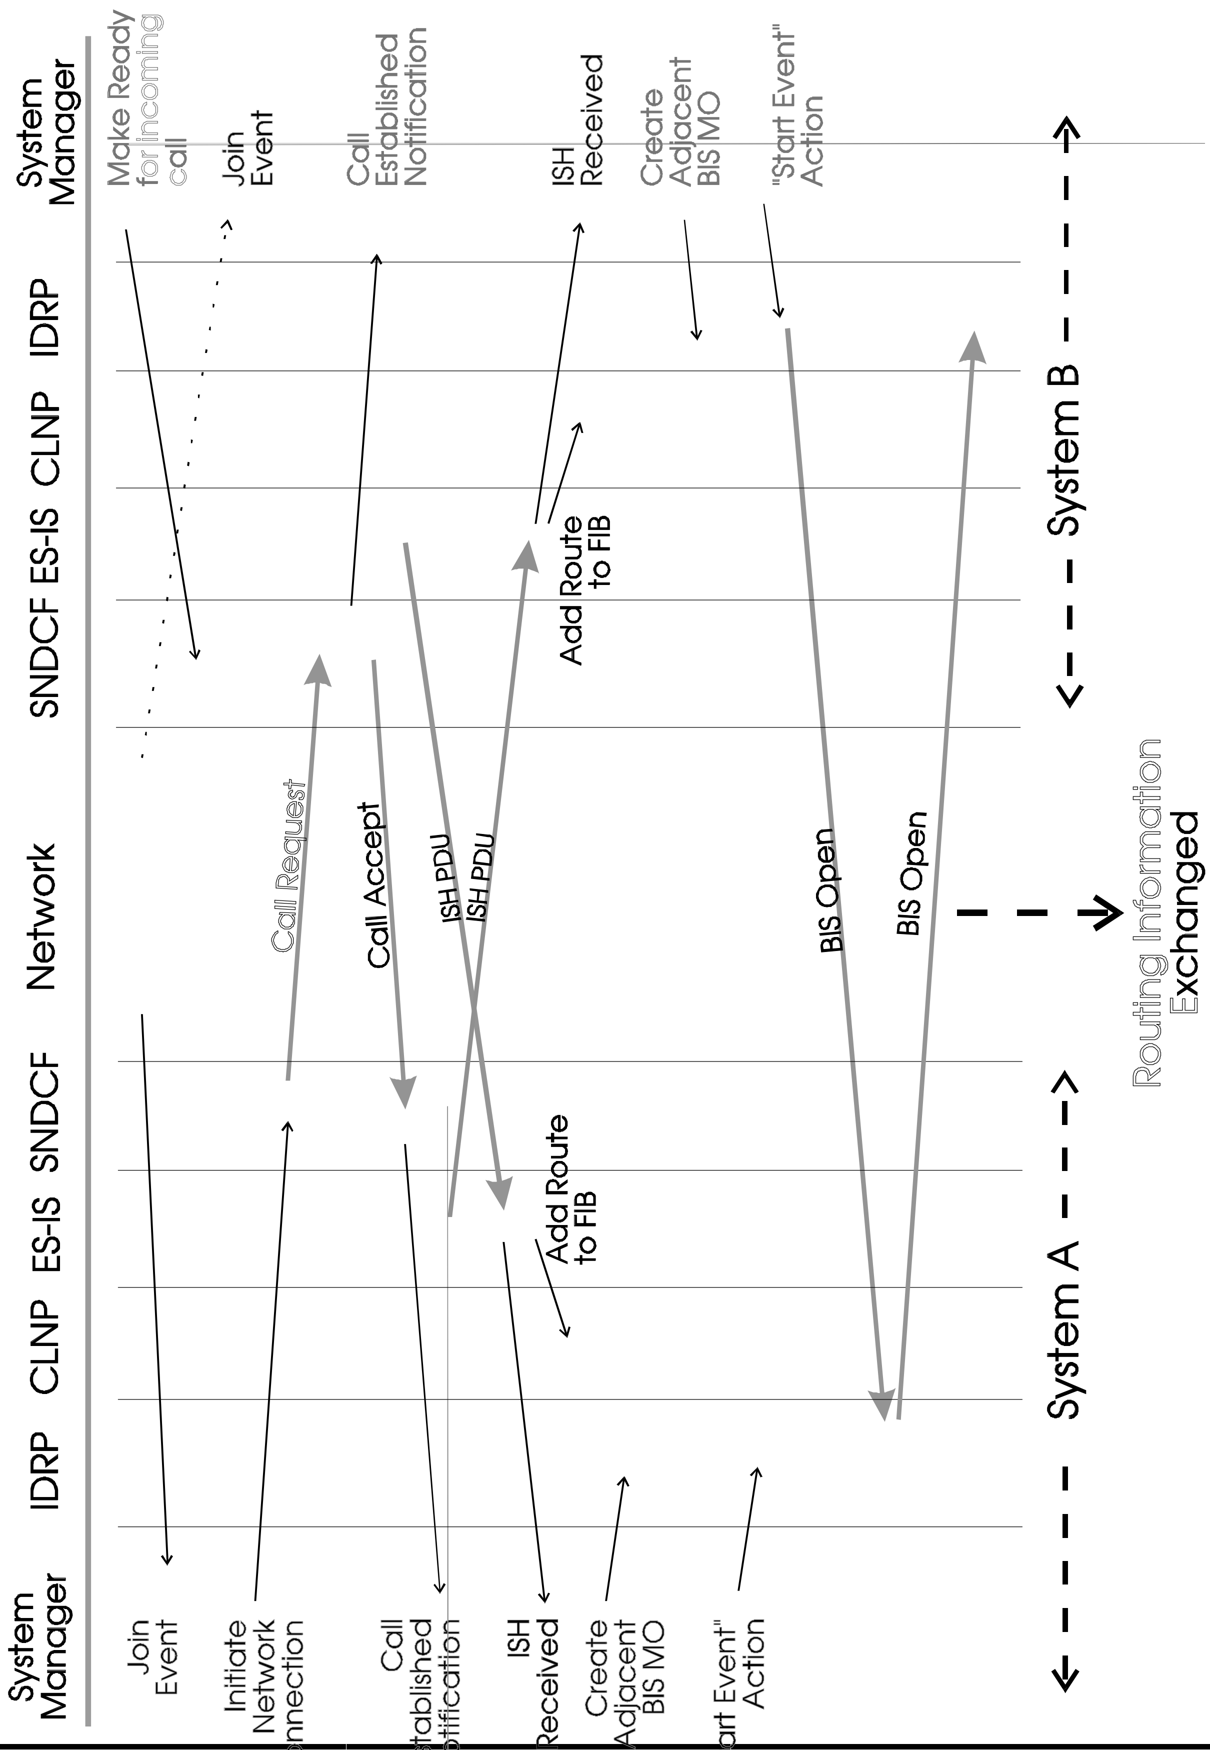
System Manager
IDRP
CLNP
ES-IS
SNDCF
Network
SNDCF
ES-IS
CLNP
IDRP
System Manager
Make Ready for Incoming call
Join Event
Call Established Notification
ISH Received
Create Adjacent BIS MO
"Start Event" Action
Call Request
Call Accept
ISH PDU
ISH PDU
Add Route to FIB
Add Route to FIB
BIS Open
BIS Open
Join Event
Initiate Network Connection
Call Established Notification
ISH Received
Create Adjacent BIS MO
"Start Event" Action
System A
System B
Routing Information Exchanged

Figure 1-17 Air-Ground Route Initiation Procedures

2) System “A” acts on a Join Event by initiating the establishment of a virtual circuit to the address given by the Join Event, provided such a connection is permitted by local policy, or

3) if polling, System “A” issues a Call Request to the next address on its poll list.

4) When an incoming call is received by System “B”, it accepts the call if permitted to do so by local policy, and generates and sends an ISH PDU to System “A” over the newly established virtual circuit. This ISH PDU includes the NET of the System “B” Network Entity.

5) When System “A” receives a Call Accept, it too generates an ISH PDU, and sends it to System “B” over the newly established virtual circuit. This ISH PDU includes the NET of the System “A” Network Entity.

6) On receipt of the ISH PDU, both systems update their local FIB to include the routing information received on the PDU, and

7) if one does not already exist, the local IS-SME creates an Adjacent BIS MO for the remote system identified by the ISH PDU, and issues a “Start Event” action to that MO. The Adjacent BIS MO created in System “A” identifies the system as being in the passive role, while the System “B” MO identifies the system as being in the active role. Hence on receiving the start event, System “A” simply listens for an incoming BIS OPEN PDU, while System “B” generates one and sends it to System “A”. System “A” responds to the OPEN PDU, with its own OPEN PDU.

8) Alternatively, if a BIS-BIS connection already exists with the remote system, then the IDRP Decision Process is re-run.

9) Once the BIS Open PDUs have been exchanged, the Route Initiation procedures have been completed.

1.9.3.3 Initial Route Initiation

In the air-ground environment, Route Initiation starts with the notification that an aircraft has come into contact with an air-ground subnetwork, and that a BIS-BIS connection should be established, so that dynamic routing information may be exchanged. In order to ensure the automatic and timely execution of these procedures, a management entity is required by the ATN SARPs to be implemented in each airborne Router and each ground Router with air-ground connectivity. This known as the “Intermediate System - Systems Management Entity” (IS-SME).

Note: The IS-SME is part of the Systems Management Agent for that Router and may also implement other functions outside of the scope of Routing Initiation.

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

1) Air-Ground subnetworks that can recognise when an aircraft has come into contact with the subnetwork (e.g. logged on to a satellite, or captured by a Mode S Interrogator) and hence that a communications path may be established with that aircraft, and which report this event.

2) Air-Ground Subnetworks which have no mechanism for recognising the above event and/or reporting it.

In the former case, Route Initiation procedures commence when the air-ground subnetwork reports this event - known as the “join” event. In the latter case, Route Initiation additionally includes procedures to allow support for Route Initiation in the absence of such an indication.

Note: Only when air-ground communications are air-initiated is it possible to establish communications without a join event.

1.9.3.3.1 The Join Event

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

1) A subnetwork identifier allowing the Boundary Router to associate the event with an air-ground subnetwork to which the Router is connected.

2) The address on that subnetwork of the remote airborne or ground Router.

3) The expected lifetime of the adjacency i.e. how long a communications path is expected to be available.

A Ground Router will typically receive a join event for each aircraft that joins each air-ground subnetwork to which the ground Router is attached. The receipt of such join events will therefore be a regular activity. An airborne Router will typically receive a join event for each ground Router on an air-ground network at the time it comes into contact with that air-ground subnetwork.

On receipt of a Join Event, an ATN Ground Router will, if communication is ground initiated, issue a call request to the subnetwork Address reported by the Join Event and thence establish a virtual circuit with the corresponding Airborne Router. An ATN Ground Router will ignore any Join Events received from air-initiated Air-Ground subnetworks.

Likewise, on receipt of a Join Event, an ATN Airborne Router will, if communication is air initiated, issue a call request to the subnetwork Address reported by the Join Event and thence establish a virtual circuit with the corresponding Ground Router. An ATN Airborne Router will ignore any Join Events received from ground-initiated Air-Ground subnetworks.

In each case, the QoS, Security and Priority requested on the call request should be satisfactory for the exchange of routing information. A local policy decision may also be taken to ignore a Join Event from certain sources.

1.9.3.3.2 The Join Event for Subnetworks that do not support ATN Systems Management

It is anticipated that not all ICAO air-ground subnetworks will support the ATN Systems Management protocols. In order to provide the equivalent of the join event, this Guidance Material provides the following guidance describing an alternative procedure for passing a join event to an air-ground Router. Future ICAO SARPs for air-ground subnetworks which do not specify support of ATN Systems Management should specify the following procedures or an equivalent procedure.

1) A communications path (e.g. a virtual circuit) is established between the ATN Router and a subnetwork processor (e.g. Mode S GDLP) by a Systems Manager and kept open as long as both Router and subnetwork are active.

2) Join events are passed from subnetwork processor to Router over this subnetwork connection and as discrete items of data (e.g. as a single packet), and passed to the IS-SME.

3) The Join Event packet is formatted as a sequence of fields according to Table 1.

1.9.3.3.3 Procedures for Air-Ground Subnetworks that do not Provide a Join Event

With this class of subnetwork, it is necessary to adopt a polling strategy in order to establish air/ground communications, and an Airborne Router must "poll" a list of Ground Routers that has been configured by the System Manager.

A suitable "poll" is a periodically repeated Call Request packet addressed to the DTE Address of a Ground Router. Such call requests are regularly repeated until they are answered with a Call Accept from the addressed Ground Router, and an Airborne Router may cycle through a list of Ground Router DTE Addresses until a connection is established. The QoS, Security and Priority requested on this Call Request should be satisfactory for the exchange of routing information.

Once a virtual circuit has been established, the Router may cease to cycle through its poll list, until the connection terminates (e.g, because the aircraft goes out of range of the mobile subnetwork), when it must resume polling for another connection. However, this may lead to unnecessary gaps in communications availability. Furthermore, not all ground Routers will support all security types required by the aircraft. The airborne Router is thus recommended to continue to cycle through its poll list, even when subnetwork connections exist, and to poll the remaining DTE Addresses on the poll list. Polling need only stop when the Router has made sufficient air/ground connections to satisfy its requirements for each supported traffic type, QoS and availability. Polling may resume when these requirements cease to be met

Note: Typically, there will be many more Airborne Routers on a mobile subnetwork than there are Ground Routers, regardless of the subnetwork's coverage area. Hence, while an Airborne Router can be expected to be configured with a complete list of Ground Router DTE Addresses, it is unlikely to be practicable for a Ground Router to be configured with a complete list of Airborne Router DTE Addresses. This is why subnetworks which do not provide information to DTEs on the connectivity status of other DTEs are only considered suitable for air-initiated BIS-BIS connections.

1.9.3.4 Route Initiation in CLNP

As a result of the handling of the Join Event or the “polling” procedure described above, a virtual circuit will have been established between Airborne and Ground Routers. The Mobile SNDCF specified in the ATN SARPs should also have been assigned to support the use of this virtual circuit by CLNP. As with ground-ground Route Initiation, it is now necessary for the IS-SME to add to each Router's FIB, a route to the NET of the remote Router's Network Entity, using the newly established virtual Circuit.

Field	Size, octets	Format	Status	Contents
Message ID	1	binary	required	’1’
Length	1	binary	required	Total message length, in octets
Version	1	binary	required	’1’
Lifetime	2	binary	required	Lifetime of link, in seconds
SNPA	var	type/len/value	optional	Remote ATN Router DTE address(es) now available

Notes:

1. The length field defines the length of the entire message, including the message identifier field

2. The value of the lifetime field is determined by the subnetwork processor. This value should be set to the expected time (in seconds) that connectivity over the mobile subnetwork is expected. A typical value would be on the order of 600 - 1200 seconds (10 - 20 minutes). Note that if air/ground connectivity is still possible shortly before expiration of the lifetime, the SP should re-issue the routing initiation event.

3. The SNPA field contains the subnetwork address of the remote Router. For example, the routing initiation event delivered to the aircraft Router contains the SNPA of the ground Router(s). The actual SNPA may have a different format or length for each subnetwork (for an 8208 subnetwork, the SNPA is the equivalent to the DTE address). The three subfields, type, length, and value are set as follows:

    a) a one-octet type field is set to ’1’, indicating the field as type "SNPA"

    b) a one-octet length is set to the length of the remote Router SNPA address

4. the variable-length value contains the actual DTE address of the remote Router

5. Multiple SNPA fields may be included within a single routing initiation event to report the reachability of several Routers simultaneously .

6. The VER field should be set to ’1’.

7. The value of the type field identifying the following data to be of type "SNPA" should be set to ’1’

Table 1 Join Event Format

However, all each Router knows at this point is the DTE Address of the other Router. In order to avoid the maintenance problem inherent in managing lookup tables that would enable a correspondence to be made between a DTE Address and a NET, a dynamic procedure has been specified by the ATN SARPs.

An ISO 9542 IS Hello (ISH) PDU is used for this purpose. This is sent either as data, once the connection has been established, or as part of the Call Request/Call Confirm dialogue when “Fast Select” is supported by the air-ground subnetwork. Both Airborne and Ground Routers generate an ISH PDU that reports their NET to the other Router. On receipt of an ISH PDU, each Router updates its FIB with a route to the remote Router, using the NET supplied by the ISH PDU and associating this NET with the subnetwork connection over which the ISH was received, as the forwarding path.

Note: this procedure is also used to negotiate the interim procedures used when IDRP is not supported by the Airborne Router.

1.9.3.5 Route Initiation in IDRP

Route Initiation in IDRP in the air-ground case is then almost identical to the ground-ground case, except that the SARPs require that one Router is in the passive mode and the other in the active mode. This is because the efficiency improvement gained by this approach is worthwhile in the air-ground environment, and the active and passive roles can be unambiguously identified when ICAO air-ground data networks are used.

The SARPs specify that for air-initiated air-ground subnetworks (i.e. AMSS and VDL), that the Ground Router takes on the active role and the Airborne Router takes on the passive role. For ground-initiated air-ground subnetworks (i.e. Mode S), the SARPs specify that the Airborne Router takes on the active role and that the Ground Router takes on the passive role. This approach will permit the exchange of route initiation data to take place in the shortest timeframe.

The Adjacent BIS MO, if it does not already exist, must be created in response to a notification that an ISH PDU has been received over a new subnetwork connection. It is necessary to create this MO in response to receipt of the ISH PDU, because it is not realistic to pre-configure an Adjacent BIS MO for every ~~aircraft or Ground Router that~~Airborne or Ground Router to ~~could be connected to~~which it could be connected.

An IDRP “Start Event” is then invoked by the IS-SME, provided that a BIS-BIS connection does not already exists with the remote system. If a BIS-BIS connection does already exist then, as in the ground-ground case, and for the same reasons, the IS-SME must cause the IDRP Decision Process to be re-run.

1.9.4 Air-Ground Route Initiation without IDRP

Due to avionics limitations, the ATN SARPs permit, as an interim measure, the existence of ATN Airborne Routers which do not support IDRP. Modified Route Initiation procedures are specified to identify such Airborne Routers and thence to infer the routes that would have been distributed had IDRP been implemented.

Note 1: The identification of routes by inference is only possible because aircraft are required by the ATN SARPs to be End Routing Domains. That is they do not relay data between ground stations or to other aircraft, and hence only provide routes to their local Routing Domain.

Note 2: The consequence of this procedure is that aircraft cannot be dynamically informed about ground route availability. Therefore, until this interim measure has been withdrawn, the ground ATN environment must be constructed to ensure a higher level of availability than would have been necessary had dynamic information been available to all aircraft. This is because, when aircraft make assumptions about ground route availability, those ground routes must exist within the margins of tolerance necessary for air safety.

1.9.4.1 Summary of Procedures

The procedures for Air-Ground Route Initiation without IDRP are illustrated in Figure 1-18, and summarised below. They are d~~iscuss~~ecribed in greater depth in the following sections. The figure illustrates the case where Air-Ground Routing is ground-initiated. The Route Initiation Procedures are:

1) When an aircraft attaches to an air-ground subnetwork, a Join Event is generated, potentially to both Airborne and Ground Routers. If received by System "B" (the Airborne Router), the Join Event is ignored. System "B is ready to receive incoming calls as soon as it attaches to the Mobile Subnetwork.

2) System "A" (the Ground Router) acts on a Join Event by initiating the establishment of a virtual circuit to the address given by the Join Event, provided such a connection is permitted by local policy, or

3) if polling, System "A" issues a Call Request to the next address on its poll list.

4) When an incoming call is received by System "B", it accepts the call if permitted to do so by local policy, and generates and sends an ISH PDU to System "A" over the newly established virtual circuit. This ISH PDU includes the NET of the System "B" Network Entity, with the NSEL set to the conventional value of hexadecimal **FE**.

5) When System "A" receives a Call Accept, it too generates an ISH PDU, and sends it to System "B" over the newly established virtual circuit. This ISH PDU includes the NET of the System "A" Network Entity.

6) On receipt of the ISH PDU, both systems update their local FIB to include the routing information received on the PDU, and

7) System "A" generates the derived routes using the NET of System "B", inserts them into the IDRP RIB, and invokes the IDRP Decision Process.

8) System "B", generates the derived routes from its local "look up" table and inserts them into its local FIB. If for any derived route, an alternative route exists via a different Ground Router to the same destination then only that with the highest degree of preference as indicated by the look up table is inserted in the FIB.

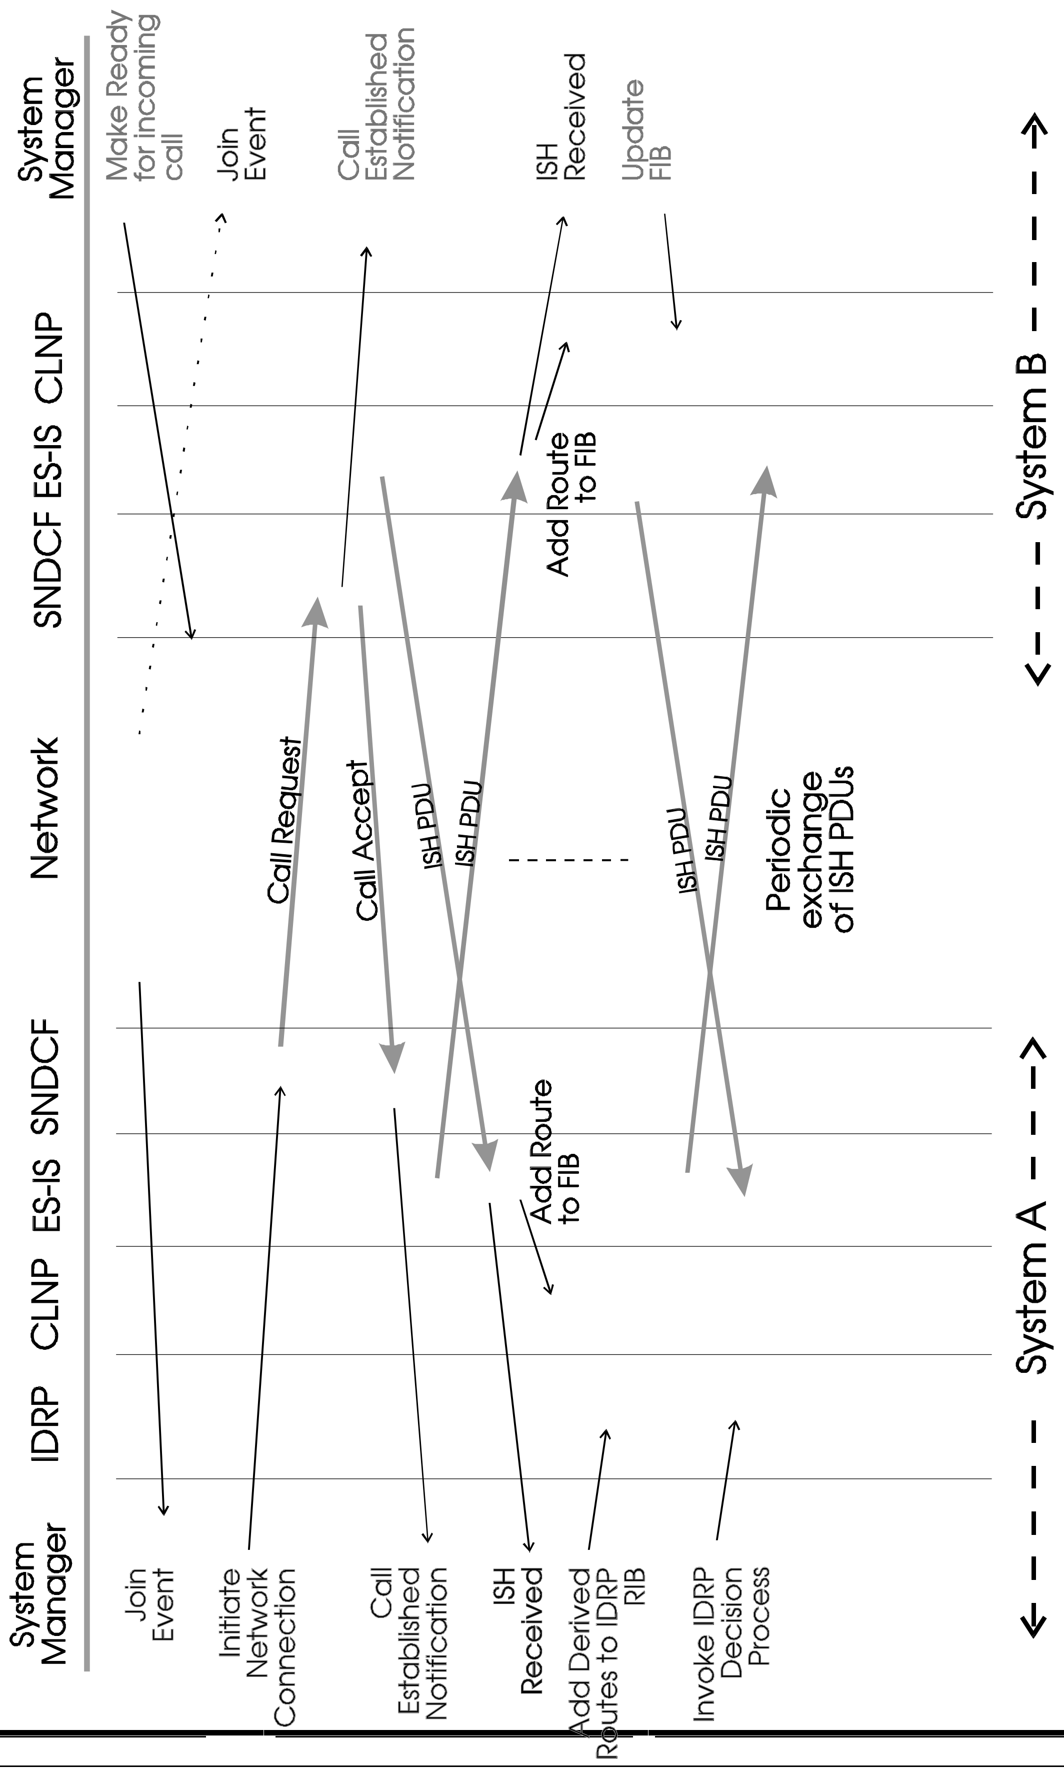
System Manager
IDRP
CLNP
ES-IS
SNDCF
Network
SNDCF
ES-IS
CLNP
System Manager
Join Event
Initiate Network Connection
Call Established Notification
ISH Received
Add Derived Routes to IDRP RIB
Invoke IDRP Decision Process
Make Ready for incoming call
Join Event
Call Established Notification
ISH Received
Update FIB
Call Request
Call Accept
ISH PDU
ISH PDU
ISH PDU
ISH PDU
Periodic exchange of ISH PDUs
Add Route to FIB
Add Route to FIB
System A
System B

Figure 1-18 Air-Ground Route Initiation without IDRP

1.9.4.2 Initial Route Initiation

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

1.9.4.3 Route Initiation in CLNP

The ATN SARPs require that the NET of an ATN Router's Network Entity has a Network Selector (NSEL) of zero. This is in accordance with ISO 10589. The SARPs further specify that Airborne Router's that do not support IDRP over the air-ground data link, have an alias NET with an NSEL value of hexadecimal '**FE**', and that this NET is used in the ISH PDU passed over the air-ground data link.

Note: that support of ~~an~~ NET with an NSEL of zero is necessary in such Airborne Routers when, for example, they also support ISO 10589 within the aircraft.

Receipt of an ISH PDU with a~~n~~ 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 d~~iscussed below~~etailed in the following section.

1.9.4.4 IS-SME Procedures without the use of IDRP

1.9.4.4.1 In the Ground Router

When the IS-SME receives a notification that an ISH PDU has been received from an Airborne Router that does not support IDRP, it must derive the routes that are available via the Airborne Router and add these routes to the local IDRP Routing Information Base (RIB). IDRP may then update the FIB and distribute these routes in the normal fashion.

The derivation of routes is possible because the aircraft is known to comprise an End Routing Domain, and from knowledge of the ATN Addressing Plan it is possible to determine an NSAP Address Prefix common to all systems in the aircraft from the NET of the Airborne Router. Further, from *a priori* knowledge of ITU restrictions that may apply to each air-ground data network and the Quality of Service offered by each such data network, the distinguishing path attributes appropriate to the routes may also be determined.

The number of routes derived by the Ground Router in respect of a specific Airborne Router will be determined by the number of different Application Security Types permitted by ITU restrictions to pass over the air-ground subnetwork multiplied by the number of QoS metrics appropriate to the network. Each such route will have as its Network Layer Reachability Information (NLRI), an NSAP Address Prefix constructed from the first eleven octets of the received NET. That is because the ATN Addressing Plan results in a common eleven octet prefix for all NSAP Addresses and NETs in one aircraft's Routing Domain, which may therefore be determined by inspection of any NSAP Address or NET from any system in that Routing Domain.

The IS-SME must then add those routes to the IDRP RIB and run the IDRP Decision Process, which then disseminates those routes and adds them to the FIB in line with the existing Routing Policy, and provided that they are a preferred route to the Airborne Router.

The actual strategy for doing this is implementation specific. However, a likely strategy is for the IDRP implementation to allocate special "adj-RIB-ins" (one per RIB-ATT) for holding routes received by mechanisms outside of the scope of IDRP. The Decision Process will then consider such routes along with those in “normal” adj-RIB-ins. The only distinguishing aspect of such routes is that they will include the “EXT_INFO” path attribute. This is a flag that enables Routing Policy to differentiate between routes that have been advertised by IDRP throughout and those which have been learnted through some other mechanism, perhaps less reliable. As in the general case, the Decision Process must be able to associate this special Adj-RIB-in with the connections to the Airborne Router, and the QoS provided by these connections . This is so that when computing the degree of preference for each such route, or when copying them to the loc-RIB, the Decision Process can update their QoS to reflect the current communications paths that exist to the Airborne Router.

If additional subnetwork connections are opened up (or lost) to an Airborne Router then, instead of generating the derived routes, as before, the IS-SME must cause the IDRP Decision Process to be re-run.

Finally, in this interim role, the IS-SME must also determine when the assumed routes are no longer valid. This event occurs when either the air-ground subnetwork connection is lost or when the periodic exchange of ISH PDUs ceases. On the occurrence of either such event, the routes generated above must be withdrawn.

Note: that in contrast with the use of IDRP over an air-ground data link, when the ATN SARPs recommend that for reasons of efficient bandwidth utilisation, ISH PDUs are not periodically transmitted, in this case they must be periodically transmitted in order to maintain the “liveness” of the routes.

1.9.4.4.2 In the Airborne Router

The IS-SME procedures are in this case, similar to the ground case, except that:

a) the NLRI of the generated routes cannot be simply derived from the Ground Router’s NET. This is because the Ground Router is typically part of a Transit Routing Domain, and the destinations of the onward routes that it offers will not have any known relationship to its NET.

b) The generated routes must be directly added to the FIB as IDRP is not present to do this on behalf of the IS-SME, or

c) if ISO 10589 is implemented, the generated Routes are used to generate Reachable Address MOs and the ISO 10589 entity is used to update the FIB.

In order to determine the NSAP Address Prefixes for the generated routes, lookup tables will have to be provided so that given the NET of a Ground Router, the Airborne Router can identify the NSAP Address Prefixes for destinations reachable via that Ground Router. Furthermore, such look up tables will have to provide:

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

ii) The Capacity, Hop Count and QoS information for such destinations in a manner sufficient to enable alternative routes to be discriminated between. i.e. an indication of relative preference for each supported metric.

Operationally, there will be a need to ensure that such tables are up-to-date with information appropriate to the Flight Region(s) through which the aircraft will fly, prior to each flight. The actual implementation of this procedure is dependent on the systems involved.

The IS-SME will have to keep dynamic information on which routes are available via each Ground Router with which it is in contact. This information is derived from the look up table and *a priori* information for each Air-Ground Subnetwork supported. When multiple subnetwork connections exists to a given Ground Router then the routing information will be determined taking into account the characteristics of each such subnetwork.

When routes to the same destination are available via different Ground Routers, then the IS-SME will have to choose between them based on the degree of preference given by the look up tables.

The IS-SME is also responsible for maintaining the FIB with an up-to-date set of available preferred routes determined as above. It must add such routes to the FIB when they become available, and remove them when the reverse is true, Alternatively, if ISO 10589 is implemented, then the IS-SME may make such routes available to 10589 by creating a Reachable Address MO for each such route, and removing the MO when the route ceases to be available. The ISO 10589 implementation may be relied upon to maintain the FIB with this routing information.

1.10 Quality of Service Maintenance

In response to User Requirements from Air Traffic Services Communications (ATSC) users, the ATN provides a classification mechanism for ATN Routes. The classifications reflect the Quality of Service available over the route taking into account availability, capacity and transit delay. Class “A” is the best while Class “H” is the worst.

These classifications do not reflect dynamic conditions but are assigned statically by network managers and reflect the result of capacity planning work.

In the CLNP header of each user data packet, the sender may then identify the minimum route class that the packet should follow; this reflects the application requirements. The ATN Routers will then choose the lowest classification route available that meets the user requirement and, if one cannot be found, then the route with the highest classification albeit lower than that required by the application.

1.11 **Priority**

1.12 **Security**

As an operational network, ATN Security has to be taken seriously. There are three aspects to ATN Security:

1. Maintaining Regulatory and National Restrictions on the use Air/Ground Data links.

2. Maintaining Restrictive Routing Policy Requirements

3. Protecting the ATN Against mis-use.

The first two aspects are dealt with procedurally. Information is included with each use on the type of Air/Ground data link over which it is available and any restrictions that apply to each such data link. When packets are then forwarded a route, if any, is chosen that is permissible for the packet’s application data to pass over and which is in line with any routing controls that the sender has specified in the packet header.

The final aspect of ATN Security requires specific security mechanisms to be effective. Security Mechanisms in this area are currently *tba*.

The IDRP protocol supports a range of authentication mechanisms (referred to as authentication type 1, 2 and 3). Authentication type 1 provides an unencrypted checksum, 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 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.