

AERONAUTICAL TELECOMMUNICATIONS NETWORK PANEL

Working Group Two

San Diego 17.10.94-28.10.94

The ATN Routing Concept

Presented By Eike Meyenberg and Henk Hof

Prepared by Tony Whyman

SUMMARY

This paper attempts to explain the ATN Routing Concept, both how it applies to the ATN Ground Environment and to Mobiles. The text is offered as supplementary to the ATN Manual 2nd edition. The guidance material in the ATN Manual explains the ATN Routing Concept at a level of great technical detail, while this paper aims to take a broader view providing an explanation of how the ATN works.

TABLE OF CONTENTS

1. Introduction	1
1.1 Scope	1
1.2 Purpose of Document	1
1.3 References	1
2. ATN Routing	1
2.1 ATN Routing Models	2
2.1.1 A General Model for ATN Routing	2
2.1.2 Routing in the Ground Environment	2
2.1.3 Routing in the Air/Ground Environment	3
2.1.4 Administrative Domains and Routing Domains	4
2.1.5 Addressing	4
2.1.6 The Inter-Domain Routing Protocol (IDRP)	6
2.1.7 Policy Based Routing	7
2.1.8 Routing Domain Confederations	7
2.1.9 The ATN Ground Environment	9
2.2 Routing to Mobile Hosts	10
2.2.1 Mobility and Routing Domains	10
2.2.2 Containing the Impact of Mobility	11
2.2.3 Routing to Mobiles within an ATN Island	11
2.2.4 Routing to Mobiles between ATN Islands	13
2.2.5 ATN External Interfaces and Mobiles	15
2.2.6 Impact on Air/Ground Datalinks	15
2.2.7 The Impact of Routing Updates	16
2.2.8 Failure Modes	17
3. Recommendation	18

1. Introduction

1.1 Scope

This paper provides background and guidance material on the ATN Routing Concept specified by the ATN Manual 2nd Edition. In writing the paper, it has been assumed that a known defect in the specification of routing policy in an ATN Island has been resolved as outlined in a companion paper [2].

1.2 Purpose of Document

This paper has been written in response to concerns expressed about the existing guidance material in the ATN Manual, which it is believed does not properly explain the ATN Routing Concept.

1.3 References

- | | | |
|----|-------------------------------|--|
| 1 | | ATN Manual 2nd Edition |
| 2. | DED1/EAS3/STA/DOC/
AW94_04 | Proposed Resolution for ATN Island Routing Policy Defect |
| 3. | ISO/IEC 10747 | Protocol for the Exchange of Inter-domain Routing Information among Intermediate Systems to support Forwarding of ISO 8473 PDUs. |
| 4. | ISO 8473 | Protocol for providing the Connectionless-Mode Network Service |
| 5. | ISO TR 9575 | OSI Routing Framework |

2. ATN Routing

The ATN has two major objectives:

- to support air/ground communications between, for example, Air Traffic Control Centres and Aircraft, and
- to support ground-ground communications between and amongst, for example, Air Traffic Control Centres and Airlines.

In themselves, these are straightforward objectives and can be easily expressed. However, from a technical point of view, realising them is much less straightforward.

In the air/ground case, the basic requirement is to develop an air/ground data link to support the ATN Concept, and several technologies are available for this. Communications Satellites can provide an almost worldwide air/ground datalink capability; Mode S SSR can support air/ground data link wherever suitably equipped Mode S SSR is in place; and VHF Data Link can be deployed over most of the world's land masses. In addition, when at rest at an airport terminal, aircraft can also be connected to a high capacity Local Area Network to supplement the air/ground datalink.

However, none of these technologies can be implemented as a single system, no matter how reliable that system can be designed. Several geostationary satellites and ground stations are required to provide a worldwide datalink service based on communications satellites, and a very large number of ground stations will always be required for any land based air/ground datalink. In each case, some mechanism will be necessary to determine through which ground station a given aircraft can be reached, at any one time.

This is a routing problem and will have to be solved regardless of the air/ground communications technology. The ATN ground environment will consist of a number of ground stations and a fixed data communications network linking the ground stations to ATC Centres and other ground based ATN users. Some mechanism is required to enable data originated by a ground based user and addressed to systems onboard an aircraft, to be sent through the fixed data network to the ground station through which that aircraft is currently reachable, and where it may then be uplinked.

Each of the technologies identified above has its advantages and disadvantages and, in many ways, they can be regarded as complementary. There is also the requirement to be able to take advantage of high bandwidth Local Area Networks when aircraft are at rest and attached to airport terminals, and it is also highly desirable that there should be no unnecessary restrictions on either incorporating or migrating to any new air/ground data communications technologies that may be developed in the future. Therefore the ATN has to be independent of any one air/ground data communications technology.

In the ground-ground case, there also needs to be a fixed data communications network linking the ATC Centres, etc., and the same ground network should be useable for both ground-ground communications, and in support of air/ground communications.

As with the air/ground environment, it is unlikely that a single communications network technology will be preferred in the ground-ground environment, or that it will even be desirable to mandate a single technology.

The ATN is therefore conceived of as an internetwork - a federation of many different fixed and air/ground network technologies, generally referred to as an Internet. The ATN can thus bring together and capitalise on the many different networks already in existence, and include new networks using state of the art technologies. These networks are interconnected through *routers* - the ATN's data switches - and it is the routers that navigate a way through the ATN Internet switching and routing user data as discrete packets.

The ATN Routing Concept is concerned with how the ATN routers perform this task.

2.1 ATN Routing Models

In order to understand the ATN Routing Concept, it is first necessary to discuss general routing principles, and how they apply to the ATN.

2.1.1 A General Model for ATN Routing

A general model of ATN Routing is shown in Figure 2-1. In this model, the ATN consists of a fixed ground network which links satellite, VHF and Mode S ground stations together with ground based Host computers, including both large scale data processing engines and workstations. 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 2-1 General Model of ATN Routing

2.1.2 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 2-2. The ATN ground environment will consist of multiple networks, owned by different administrations and organisations, and implemented using many different technologies. In some cases, these will be existing networks with spare capacity made available to the ATN. Others will be new networks implemented specifically to support ATN use. There will be X.25 Private Packet Switched Data Networks (PPSDNs), Frame Relay Data Networks, Integrated Services Digital Networks (ISDNs), Local Area Networks (LANs) e.g. Ethernet, and others. These networks are then linked together through routers which provide the connectivity between the different types of data network, and to the air/ground networks; host computers are directly connected to a nearby data network, typically a LAN.

User data is switched by the routers as discrete packets formatted according to 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 2-2 The ATN Ground Environment

2.1.3 Routing in the Air/Ground Environment

As illustrated in Figure 2-3, the data communications equipment on board an aircraft also includes a router. In this case, the router links the communications equipment specific to each air/ground datalink to an internal avionics LAN, to which avionics equipment with datalink capability is attached.

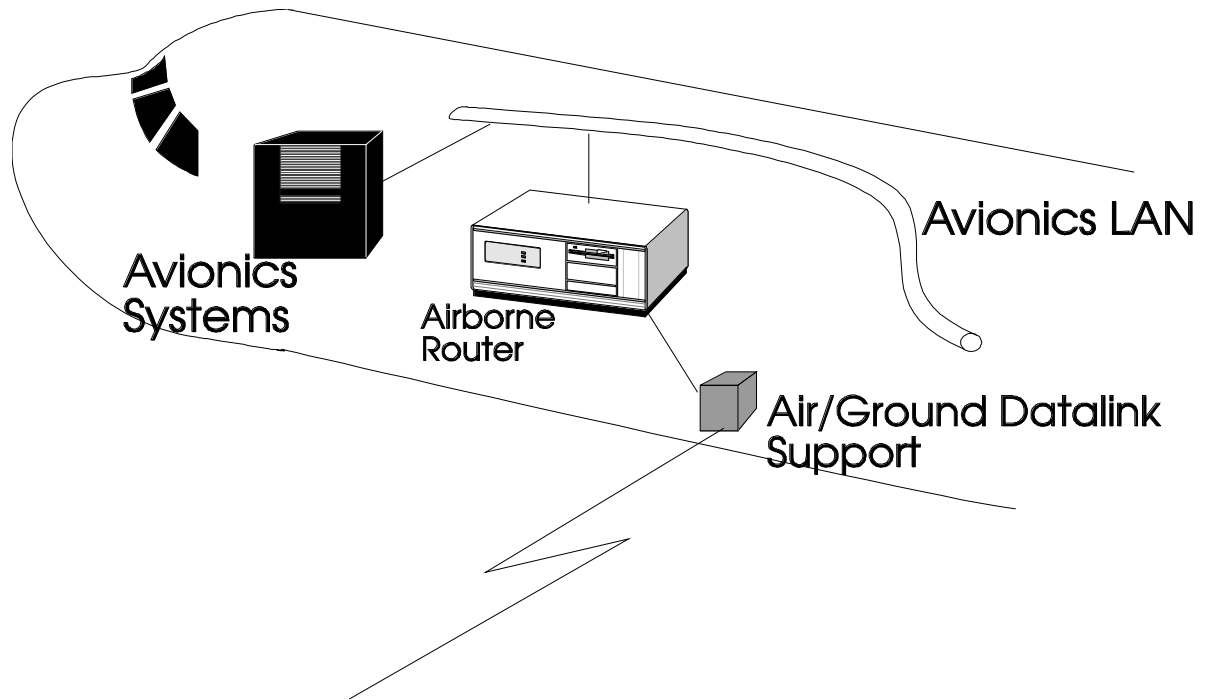


Figure 2-3 The Airborne ATN

While the airborne case is generally as simple as illustrated, the ground fixed network is much more complex than the simplified diagram suggests, and it is therefore necessary to impose a structure on it that enables a scalable architecture to be developed.

2.1.4 Administrative Domains and Routing Domains

In order to develop a scalable routing architecture, the ATN has adopted the ISO Routing Framework, presented in ISO TR 9575 and illustrated in Figure 2-4. 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).

2.1.5 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.

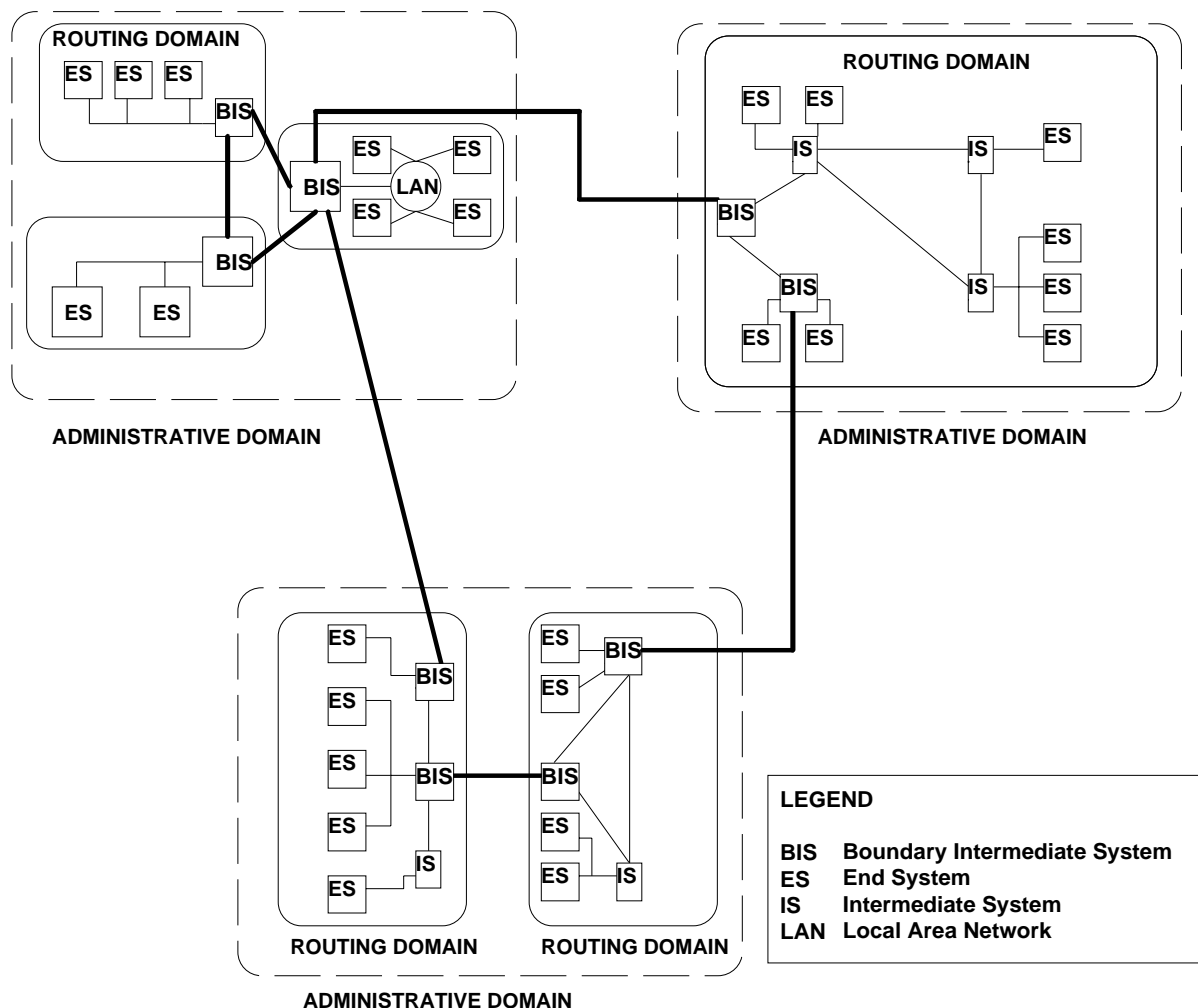


Figure 2-4 ISO Routing Framework

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 2.1.8 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.

2.1.6 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.

2.1.7 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 known 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.

2.1.8 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.

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 2-5 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 2.1.8.2 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 the Internet.

2.1.8.1 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.

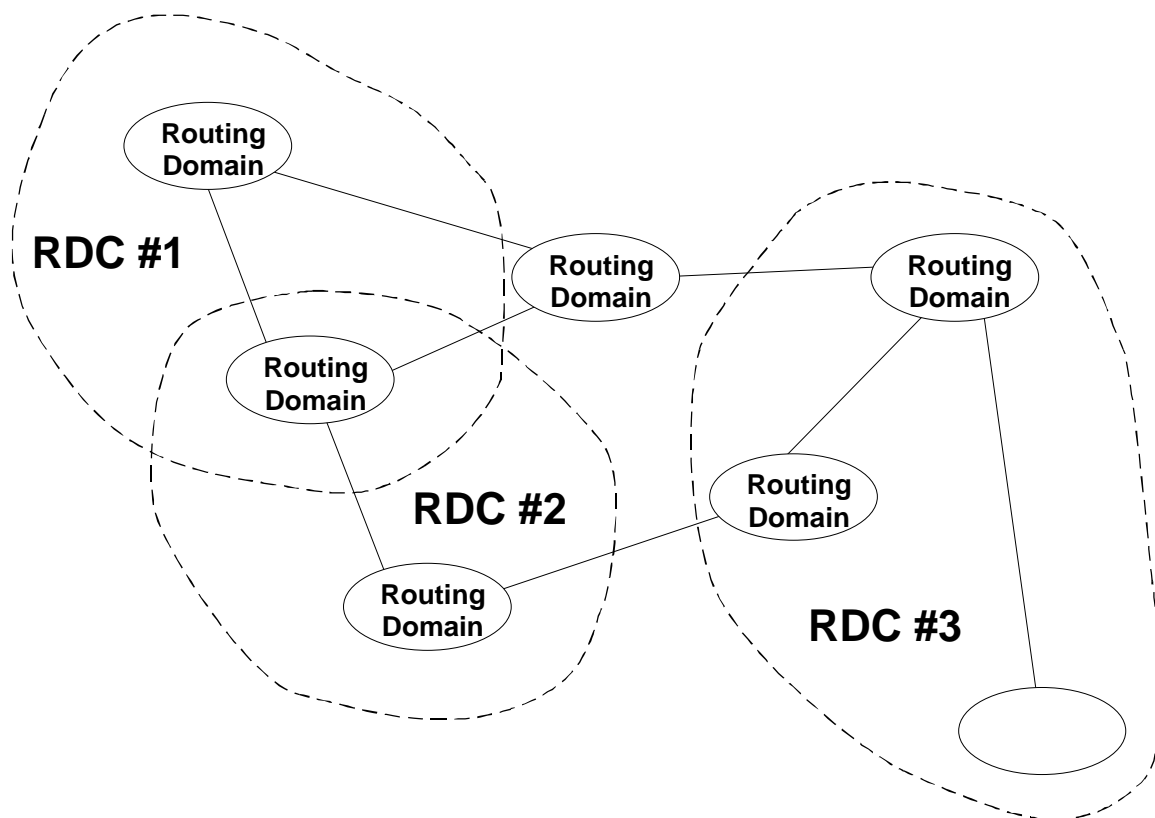


Figure 2-5 Routing Domain Confederations

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.

2.1.8.2 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 an overall reduction in the number of routes without increasing the detail associated with the route's destination.

Route Aggregation may occur at any 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 exists 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.

2.1.9 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 address 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.

2.2 Routing to Mobile Hosts

2.2.1 Mobility and Routing Domains

While the scalability of an Internet demands that 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 both active on the same network. The onward re-advertisement of routes can inform the rest of the ATN Internet about such a temporary connectivity while it exists, and the loss of connectivity when it occurs. A Routing Domain can thus temporarily join an Internet at one point of attachment, then disconnect and join the Internet at some other point, the only impact being in the efficiency of routing information distribution, and eventually on scalability.

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 2-6.

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 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.

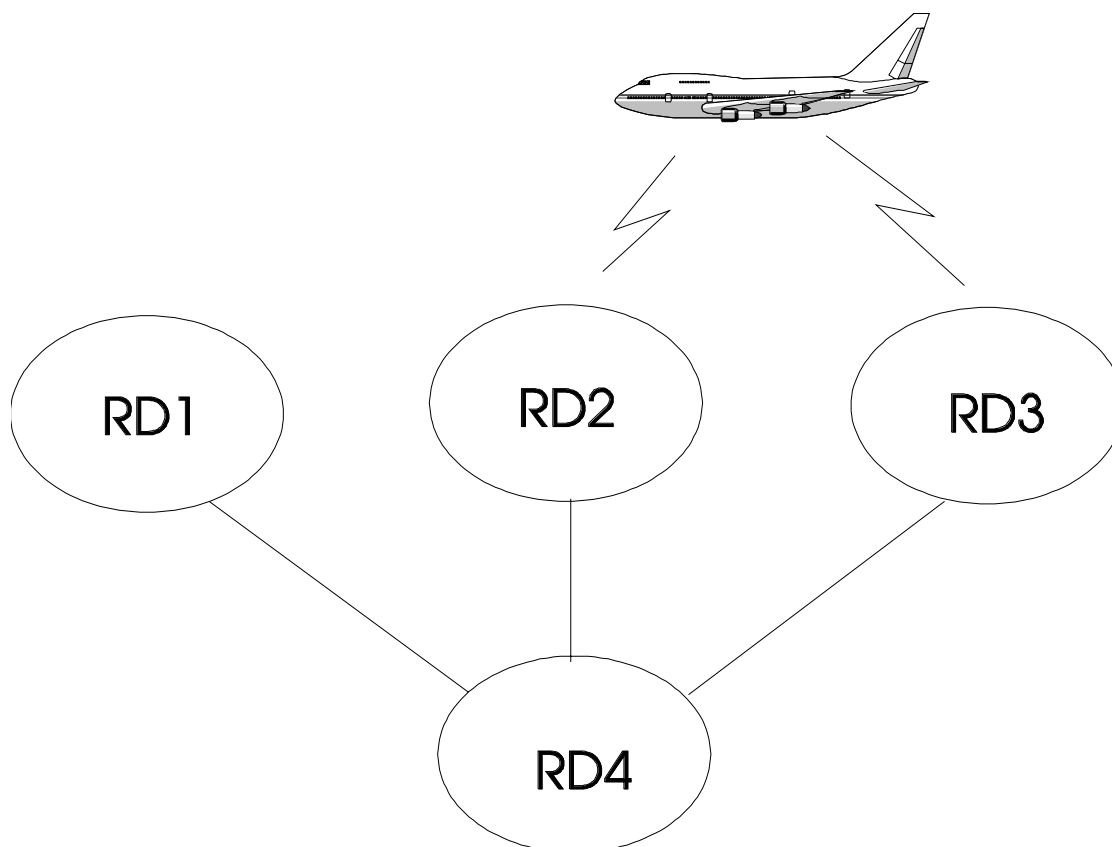


Figure 2-6 Mobile Routing Example

2.2.2 Containing the Impact of Mobility

While the principles of mobile routing outlined above are straightforward they are not scaleable using the existing IDRPs 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 2.2.4 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.

2.2.3 Routing to Mobiles within an ATN Island

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 2-7, 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 another 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 2-7, 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 route¹. 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 a Routing Domain that provides a route to the backbone. Either way the impact of such a policy rule is that the Backbone RDC is always informed about routes to all aircraft currently reachable via datalinks available to the Island's Routing Domains, and can thus act as default route providers for packets addressed to airborne systems.

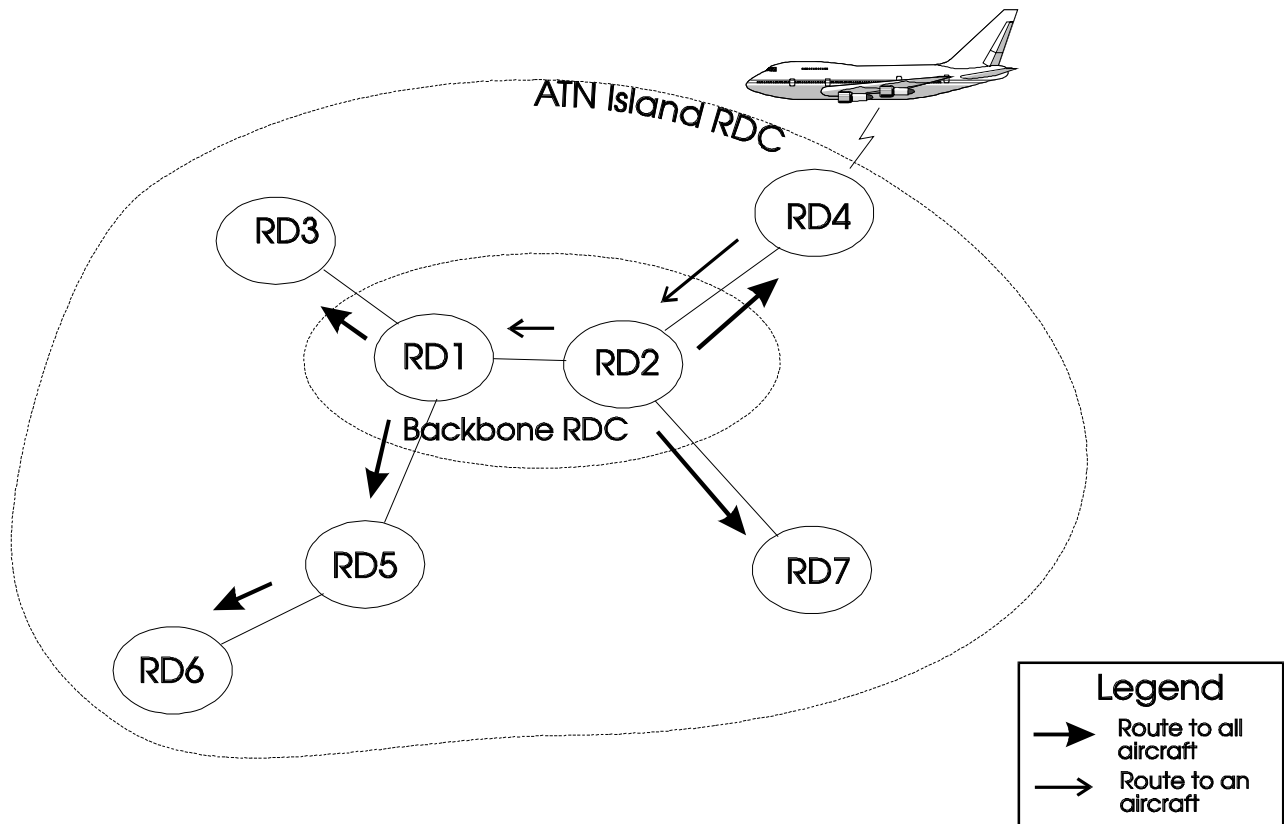


Figure 2-7 Mobile Routing Within an ATN Island

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 IDRPs always prefers routes with the longest matching address prefix. Therefore, the default route to all aircraft is always a shorter prefix of that for an explicit route to an aircraft, and this routing strategy happens automatically without any special provisions.

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.

The 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 2.2.7 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.

2.2.4 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 2-8. In this example, ATN1 is the ATN Island acting as the “Home” for all aircraft belonging 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 exists² 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 2-8, 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.

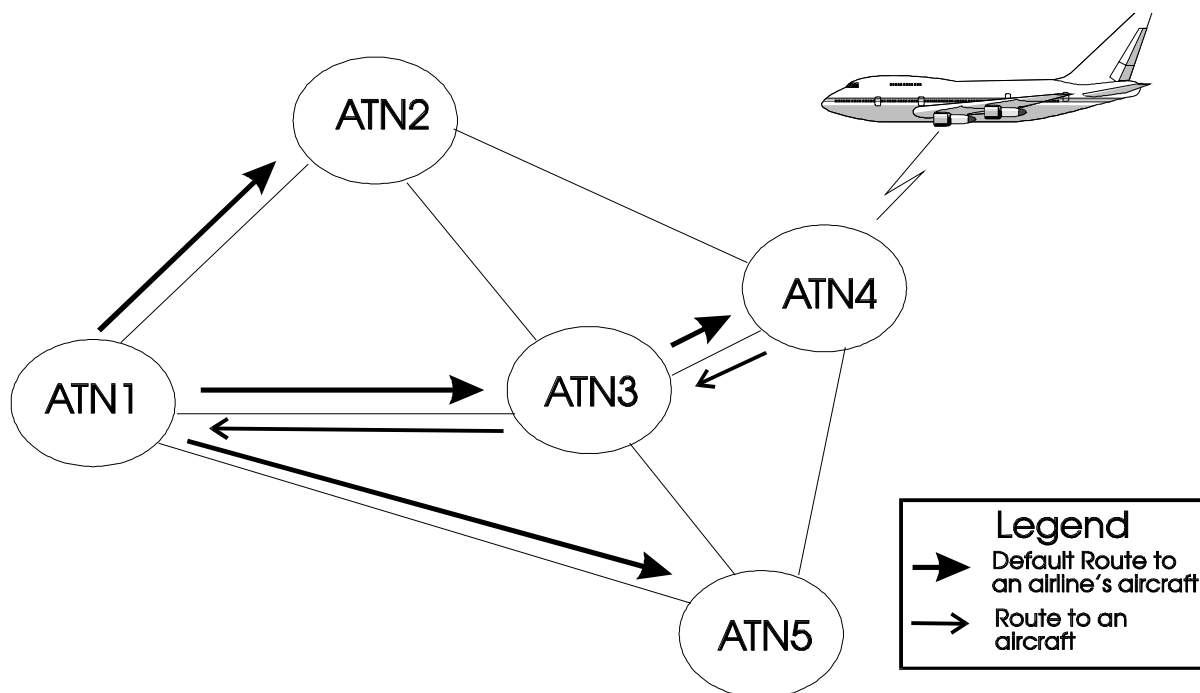


Figure 2-8 Inter-Island Routing

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:

$$n + m < \text{“maximum number of routes to mobiles that can be handled by a backbone router”}$$

has to be true.

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.

However, this limit 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 this limit. The routing architecture thus allows for a much larger number of mobile systems than that permitted by a single ATN Island.

2.2.5 ATN External Interfaces and Mobiles

As discussed above, 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.

2.2.6 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 multiplexed together into a single message keeping the actual information exchanged to a bare minimum.

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

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

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

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

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

in addition to other routes as determined by local policy.

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

2.2.7 The Impact of Routing Updates

The above discussion has illustrated how 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.

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 2-9.

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

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

This has important implications for the design of an ATN Island. If an Island’s air/ground datalinks are all connected to Routing Domains which are themselves adjacent to the Backbone RDC, all connectivity changes will be immediately reported to the Backbone giving a high route update rate. On the other hand, if there are intermediate Routing Domains between the backbone and the Routing Domains connected to air/ground datalinks, then the update frequency can be significantly reduced, without affecting the responsiveness to real connectivity changes.

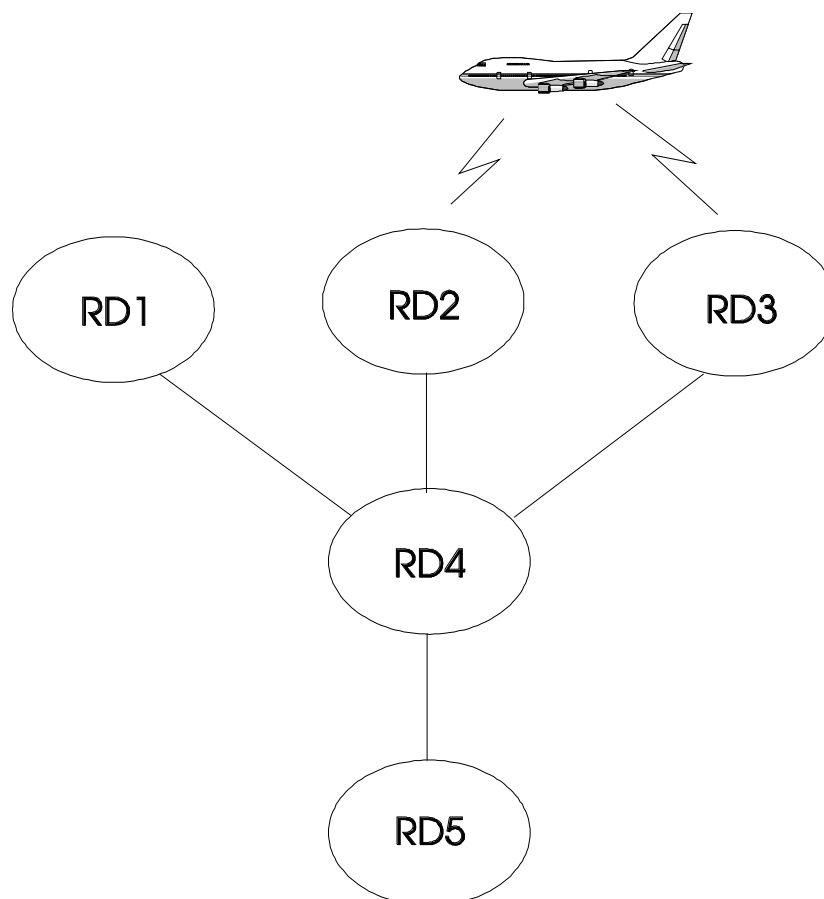


Figure 2-9 Impact of a Hold Down Timer

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

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

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

2.2.8 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.

2.2.8.1 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.

2.2.8.2 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 discussed above in 2.2.8.1, 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.

3. Recommendation

The panel is invited to consider the contents of section two of this paper as additional ATN Guidance Material.