European Organisation for the
Safety of Air Navigation

EUROCONTROL

EUROCONTROL Experimental
Centre

École nationale supérieure des
télécommunications

ISA project
(IPv6, Satellite communication and ATMode for ATN)

An IPv6 architecture for the Aeronautical Telecommunication Network

Thomas QUINOT

Issued: September 30, 1998
Identification: `$Id: ATNv6.tex,v 1.20 1998/09/30 14:58:06 tel1 Exp $`

The information contained in this document is the property of the EUROCONTROL Agency and
no part should be reproduced in any form without the Agency's permission.
The views expressed herein do not necessarily reflect the official views or policy of the Agency.

Contents

List of Figures

Chapter 1

Introduction

EUROCONTROL is investigating the means of establishing a general data network for ground-based air traffic control stations as well as aircraft, that shall be able to support four classes of services:

Air Traffic Service (ATS) i. e. Air Traffic Control proper;

Aircraft Operational/Administrative Communication (AOC/AAC) i. e. internal airline communication related to aircraft operation, respectively to company management;

Aircraft Passenger Communication (APC) i. e. data communication as a value-added service for airline customers.

The current ICAO[1] standard for the air traffic data network is ATN/OSI; it is based on the OSI protocols, which adaptations specific to aeronautical applications. The goal of this study, as part of the ISA project at EUROCONTROL Experimental Centre, is to provide an assessment of the adequacy of IPv6, and of its ability to fulfil the requirements of aeronautical applications as set forth by ICAO documents such as the ATN/OSI SARPs (Standards And Recommended Practices) [14]. The interest of IPv6 for air applications is twofold:

- IPv6 is an industry-driven, emerging set of standards that arised and is still being developed based on the needs of users and the actual implementation of new protocols and methods. Conversely, the ATN/OSI standards were completed before any serious implementation work was started.
- IPv6 are free, open standards; IPv6 RFCs are readily available at no cost, whereas ATN/OSI documents are costly ISO standards, and ATN/OSI networking systems often involve a lot of proprietary equipment.

There are five main requirements for an Air Traffic Network (ATN):

1. *Security.* All transmissions shall be identified and authenticated; integrity of messages shall be assured; confidentiality of data shall also be provided.

[1]International Civil Aviation Organization

2. *Quality of Service.*[2] The system shall guarantee that ATS communication permanently gets a sufficient bandwidth; the remaining bandwidth shall be shared between AOC and AAC, and subsidiarily by APC: the protocols shall map the priority scale (ATS, then AOC, then AAC, then APC) in each layer.
3. *Routing.* The system shall provide an adequate address space, and a means to route data from any point in the address space to any other point. In particular, each ground station and aircraft shall be referable.
4. *Mobility.* During a flight, an aircraft shall remain reachable through the ATN, even though it may not be in constant, physical connection with a single infrastructure point: the airborne Intermediate system (IS) mobility shall be taken into account. If APC service is to be provided, passenger equipment shall behave as though it was directly connected to its home network: End system (ES) mobility shall also be provided.
5. *Multicast.* The system shall allow the delegation of some ATC functions to aircraft (as opposed to ground-based control centres). In order to perform embedded collision avoidance (with applications such as EUROCONTROL's FREER, Free-Route Experimental Encounter Resolution [11]) without an excessive data flow requirement, aircraft need to be able to dynamically constitute multicast groups.

The ISA project at EUROCONTROL Experimental Centre aims at assessing whether IPv6 can meet these requirements. The present study contributes to this assessment by providing an implementation scheme of the ATN requirements as documented by the mandatory parts of the ATN/OSI standards, using IPv6 as the underlying network protocol stack.

We will particularly focus on three of the ATN requirements: routing, mobility and multicast; they constitute the core features of the network, and we consider that IPv6 provides important improvements over the OSI protocols in these domains.

[2]This requirement makes IPv6 far easier to use than IPv4, because, unlike IPv4, IPv6 provides protocol elements that allow the upper layers to require a specific QoS.

Chapter 2

Routing and mobility

2.1 Introduction

Aeronautical applications require a data communication network for data transfer involving aircraft, ground based, air traffic related operations, and third party applications. This network, currently under development, is referred to as the ATN (Aeronautical Telecommunication Network) in the remainder of this paper. It is specified by ICAO in the ATN/OSI Standards and Recommended Practices documents [14].

It shall be able to provide reliable data communication for aeronautical applications. It shall enable connections between ground based and airborne applications, and between two ground applications.

Innovative applications such as FREER also require the capability to establish connections between two airborne applications. This requirement is set forth by the Air Traffic Requirements document of the ISA project [23].

The ATN shall also allow users to establish communication with airborne applications using an address space independent of the aircraft's current position. We thus have to take into account the movement of aircraft in the organisation of the network.

2.2 Current practice

2.2.1 Presentation

In the current OSI based ATN, all airborne applications are connected to an airborne router (airborne BIS — Boundary Intermediate System). The BIS has a connection to the AES (Airborne Earth Station), a telecommunication device that handles the set-up of an X.25-like connection to a ground-based GES, via a relay satellite (AMSS for example). Each GES is connected to one or more ground-based IS that is part of the ground OSI ATN. In the ground network, IDRP [16] is used as the external routing protocol, because it provides a mechanism for policy based routing.

IDRP is also used to simulate mobility support. During flight, the AES transparently handles handover from one satellite link to another. The AES has a hardcoded, static table that lists the set of available GESs. It can thus make a connection with any of these GESs when a data link can be established, and route data packets accordingly. In this current, simple topology, airborne routers perform very limited routing tasks, most of the routing being assured by ground based equipment: each GES is attached to a BIS that performs IDRP. Movement of aircraft across GES cover zone is thus propagated into backbone routing tables.

2.2.2 Route initiation mechanisms in ATN/OSI

The ATN/OSI SARPs specify the procedures according to which route information is exchanged upon establishment of an air/ground data link connection. SARPs distinguish between several mechanisms for route initiation, depending on which end of the link initiates communications.

In air initiated route initiation, the airborne router decides to establish a connection. This process is controlled by a subsystem called the IS-SME (Intermediate System — Systems Management Entity). For data links where the communication equipment does not provide connectivity information, the IS-SME continuously issues X.25 connection requests on the data link, with callee address chosen in a list of hard-coded ground station addresses. It thus polls ground stations until a valid connection is established. In the case of data links that provide connectivity information in the form of "Join Event", the IS-SME only makes an X.25 call request when the data communication equipment signals that a new node has become available. The generated Join Event contains a DTE address for the newly connected station.

In ground initiated route initiation (on data links where the *ground station* generates Join Events), the ground router similarly makes an X.25 call request on receipt of the join event.

As soon as the ground based and airborne router have established a point-to-point X.25 connection, and have validated each other's address with respect to communication policy, they perform ES-IS [15] in order to exchange configuration information: the ground based and airborne routers announce their addresses (NET — Network Entity Title) to one another using an ISH (Intermediate System Hello) PDU with a NET option. Upon reception and validation of the remote IS NET by each router, an IDRP activate action is invoked by the IS that had the *responder* role in the route initiation process. A BIS-BIS connection is thus established for IDRP routing information exchange.

2.3 Analysis of architecture

The ground network shall implement policy based routing; this is a requirement set forth by ATN users. For this reason, IDRP shall be used as the ATN/IPv6 external gateway protocol in the ground based network. Conversely, there is no need for policy based routing on the air/ground link. We therefore do not consider it a requirement that IDRP be used on this connection.

The air/ground network raises an additional issue with respect to routing: the maintenance of a continuous binding between an airborne application and a fixed address or set of addresses in an address space which is fixed beforehand. During flight, aircraft shall be able to connect to the underlying, ground-based network *via* any number of successive ground stations that they get to establish contact with.

The data link media used for these connections include, but are not limited to, satellite communication (AMSS, future LEO systems), Mode-S and VDL.

In this paper, we present several network architectures for the IPv6 ATN that address this issue.

2.3.1 External routing

2.3.1.1 Presentation

The first solution that we will examine consists in defining one IPv6 autonomous system for each aircraft. In this option, each aircraft shall embed a BIS that will exchange IDRP data with ground based BISs so that ATN backbone routers always know an up-to-date route to the aircraft's address space.

Aircraft movements are thus interpreted as network topology changes, because a fixed part of the address space (a specific aircraft's network prefix) changes its attachment point to the rest of the network, and theses changes are taken into account by means of an external routing protocol, which reflects network topology in ISs' forwarding information bases according to routing data they gather from their peers and policy rules set forth by administrators.

The aircraft shall always announce route information for its embedded subnets; ground BISs that communicate with airborne BISs shall announce routes for all destinations to their mobile peers, and routes for all their mobile peers to their ground peers, i. e. ultimately to backbone routers.

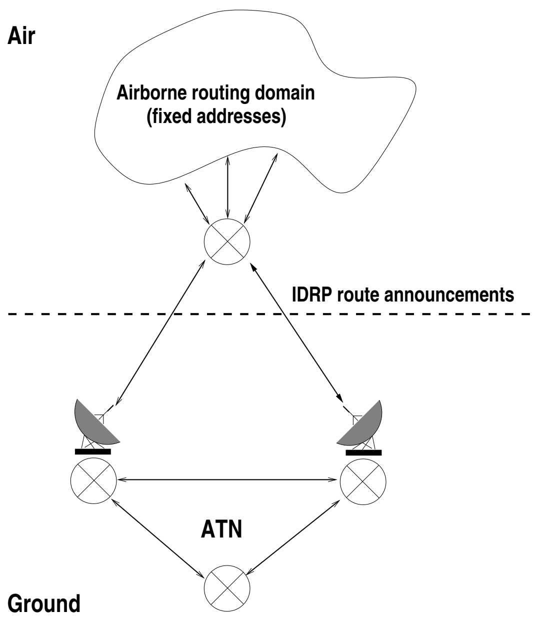

Figure 2.1: External routing

2.3.1.2 Operation

Operation is very similar to that of ATN/OSI. Whenever a point-to-point data link is established between a ground based and an airborne router, configuration information shall be exchanged. For this purpose, ATN routers shall send Router Advertisement ICMPv6 packets according to [9] and [21]. The Source link layer address may be used to notify the remote router of the link layer address of the local router.

The Neighbor Discovery draft [21] does not specify how routers shall process received router advertisement. It is up to the ATN designer to provide specific semantics for the purposes of the ATN. For the external routing approach, we specify that the IPv6 address contained in the received Router Advertisement shall be registered as the address of an IDRP neighbour BIS. The ATN router which assumed the responder role in data link initiation shall then invoke an IDRP activate action to start a BIS-BIS connection, as specified in [16]. Conversely, the ATN router which assumed the initiator role during data link initiation shall start listening for an IDRP incoming connection from the other router.

2.3.1.3 Performances and benefits

It is not currently known as a sure fact whether IDRP performed by an airborne BIS can react to data link handover within a delay short enough to be compatible with the contact continuity criteria set forth by the ATN user requirements. IDRP requires that routes to all aircraft be propagated to the ATN backbone routers, which can take a non-negligible amount of time depending on the number of simultaneously moving aircraft and the topology of the ATN.

There is no widely deployed ATN architecture that would permit a real-life-scale evaluation of IDRP's performance. Network infrastructure parameters such as maximum number of hops between a GES and the nearest backbone router, number of backbone routers and of routing domains (autonomous systems), or peering matrix of backbone routers. Future studies and simulations are necessary to establish the appropriateness of IDRP in the ATN.

Using IDRP to simulate mobile networking is the solution adopted in ATN/OSI, because the OSI protocols lack any real mobility support. We believe however that adopting this architecture for the IPv6 ATN brings no real benefits compared to ATN/OSI, because it does not take advantage of the advanced mobility features of IPv6. We recommend that a different approach be retained, using one of the two solutions described in 2.3.2 and 2.3.3.

2.3.2 Pure Mobile IP

2.3.2.1 Presentation

An alternative to having each aircraft perform IDRP consists in having no on board router, and considering every airborne applications as an IPv6 mobile entity. (Some kind of switching equipment is nonetheless required on board, in order to satisfy the quality of service requirements for the ATN.)

In this option, airborne systems change IPv6 addresses during flight each time they cross a cover zone border and establish communication with a new ground router. This shall be implemented by having all ground routers send IPv6 router advertisements announcing a specific network prefix for each mobile network in their scope (each such network corresponding to one aircraft).

On board devices thus gain knowledge of their new address using stateless address autoconfiguration, as specified by [24]; each IPv6 application acts as an IPv6 mobile entity and keep in touch with ground-based applications using the IPv6 mobility protocols.

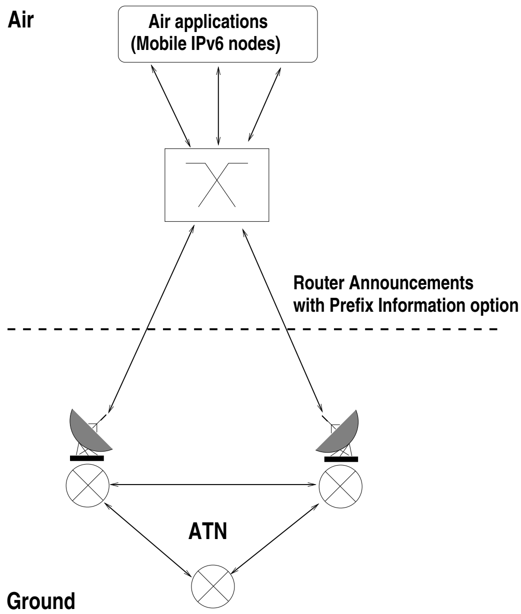

Figure 2.2: Pure mobile IP

2.3.2.2 Operation

Like in the external routing approach, we assume that the ground based and airborne telecommunication devices cooperate to establish a point-to-point link between a ground router and an aircraft's internal network. We do not require the underlying data link to provide connectivity information as used in 2.2.2.

The ground router shall assign a subnet prefix to each air/ground link that it may establish, and make each air/ground interface an advertising interface as defined by [21]. In other words, the ground router shall continuously send ICMPv6 Router Advertisements with a Prefix Information option on the air/ground links.

These advertisements shall be relayed by the airborne communication equipment and transmitted to all hosts on board the aircraft. Since Router Advertisements are multicast packets addressed to “all hosts on link”, no router shall stand between the ground router and the airborne applications.

The advertised router lifetime and the preferred lifetime for the advertised prefix shall be set to values sufficiently low to provide timely reaction to changes of network access point.

Upon receiving the Router Advertisements, all IPv6 hosts shall perform address autoconfiguration for the new prefix. The processing of Router Advertisement messages by IPv6 hosts is specified by the Neighbor Discovery and Stateless Autoconfiguration drafts [21, 24]. The newly configured address shall then be considered as a new care-of address for the host. When the current default router becomes unreachable, the IPv6 node shall select one of its new care-of addresses for use as its new care-of address, as specified in the IPv6 Mobility draft [20].

IPv6 hosts may also decide to switch to a new primary care-of address before losing contact with their current default router based on signal strength or link quality data provided by the underlying data link equipment. In this case, the switch operation is performed while the previous and new point data links are both active. This minimises the switch time and incurred packet loss; experiments have been conducted in wireless network contexts that show that this switching procedure can take as little as ten milliseconds in particularly good conditions [1]. In an ATN situation, where the round-trip time for transmission of the Binding update and Binding acknowledgment to and from a mobile host's home agent can be longer that in this experiment, we expect a few tenths of second to be a reasonable estimate of the switching time.

2.3.2.3 Performances and benefits

Each time an aircraft crosses a zone boundary, all on board nodes have to detect the movement and update their care-of address information. In particular, all nodes have to send binding updates to their respective home agents and correspondents. This process might consume a lot of network bandwidth.

Redundancy of the telecommunication equipment can be provided by having several, separate data link connections with one or more ground stations; on board hosts can simultaneously autoconfigure for multiple care-of addresses. They can dynamically switch to a new primary care-of address, should their current default router become unreachable.

The Mobile IPv6 Home agent currently constitutes a single point of failure. Solutions have been studied in the past to address this problem in Mobile IPv4 [4, 19]. An Internet draft was produced to specify a Home Agent Redundancy Protocol [4], which did not require specific support in the Mobile IP implementation. A similar protocol could be devised for HA redundancy in Mobile IPv6; as no support is required in the Mobile IPv6 specification, this task could take place in parallel with general Mobile IPv6 development with minimal interference. An Internet draft by David B. Johnson of Carnegie Mellon University is also expected soon that will describe a mechanism for avoiding the HA single point of failure.

2.3.3 Airborne IS *not* performing an EGP

2.3.3.1 Presentation

In this third approach, each aircraft embeds a router. The purposes of this router are multiple. First and foremost, it achieves isolation between airborne subnetworks that carry different classes of service; notably, it permits a complete separation between passenger communication and air traffic services. This is an important advantage compared to the previous solution in terms of systems security; proven IP firewalling and security techniques can be implemented on the on board router to protect the safety-critical ATS from eavesdropping as well as tampering from the passenger network.

The embedded router can also perform such tasks as efficient multicast routing, e. g. for air-air cluster communication, or protocol encapsulation, e. g. for support of legacy CLNP applications over the ATN/IPv6 infrastructure.

These functions do not require the airborne router to perform any external routing protocol such as IDRP. Instead, it uses an auto-configuration mechanism similar to [24] to change its addresses when crossing a cover zone boundary, and announces these changes to other nodes on board, using the Router renumbering protocol for other on board routers, and using Router Announcements with Prefix Information options for hosts. The Router renumbering protocol is defined in [8], and host configuration is done according to [24].

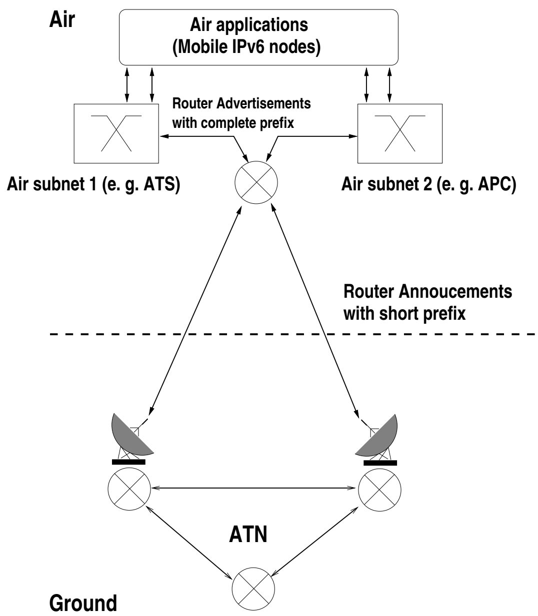

Figure 2.3: Embedded router

2.3.3.2 Operation

As in 2.3.2.2, we suppose that data link layer equipment establishes a point-to-point connection between ground based routers and the on board IS.

When a data link comes up, the ground IS shall start announcing a prefix on it using ICMPv6 Router Advertisement messages with Prefix Information options. The advertised prefix shall be shorter than 64 bits, to allow for several on board subnets. On reception of these messages, the onboard router shall form new prefixes for each of its links (air/ground radio links, on board subnets) by concatenating the advertised prefix with a link identifier. The on board IS thus performs an operation similar to stateless address autoconfiguration (which is normally defined only for IPv6 hosts).

Once the onboard IS has reconfigured all its interfaces, it can start advertising new (64-bit) prefixes on each on board subnet so that embedded end-systems can in turn perform traditional stateless autoconfiguration. Airborne hosts subsequently behave as IPv6 mobile nodes, and send Binding updates to their home agents and correspondent nodes to inform them of their new care-of addresses.

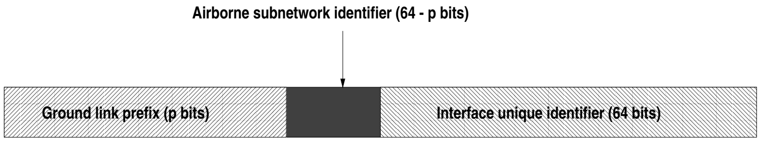

Figure 2.4: Address breakout for embedded router approach

2.3.3.3 Performances and benefits

The autoconfigured router approach suffers the same problem as the no-router approach: each embedded end-system must perform IPv6 mobility and send out binding updates each time the aircraft crosses a cover zone boundary, thus creating an important network load burst.

The airborne router approach provides enhanced security and flexibility. Traffic from the different on board subnetworks can be separated and filtered, for maximal security. The on board router can perform multicast routing and manage multicast groups by exchanging membership information with ground based or satellite-borne routers (see chapter 3 for a detailed discussion of multicast cluster operation in the IPv6 ATN). In addition, the embedded router can perform tunneling and proxying, in order to alleviate the air/ground subnetwork load during handover. Operational details on proxying and tunneling are given in 2.4.

In this approach, provisions can be take for complete equipment redundancy, in order for the on board router not to constitute a single point of failure. Two routers, attached to separate data communication equipment, can coexist on the mission-critical subnetworks (namely, the ATS subnetwork). Both routers shall announce prefixes to the attached end systems, and end systems shall autoconfigure for both addresses. One of the routers shall be considered primary at any given time.

Should the primary router fail, end systems shall notice the loss of the connection and elect the other router as primary. The procedure by which failure is noticed by end systems and switching to the new primary router is performed is identical to the procedure by which a new primary router and care-of address are chosen during hand over.

2.4 Proxying and tunneling

When an aircraft crosses a cover zone boundary, all IPv6 mobile nodes located on its embedded subnetworks have to send out Binding updates to their respective home agents, as well as all their correspondent nodes. When their number grows, this operation can create an important network load on the air/ground data link. Proxying and tunneling can help reduce the cost of this operation.

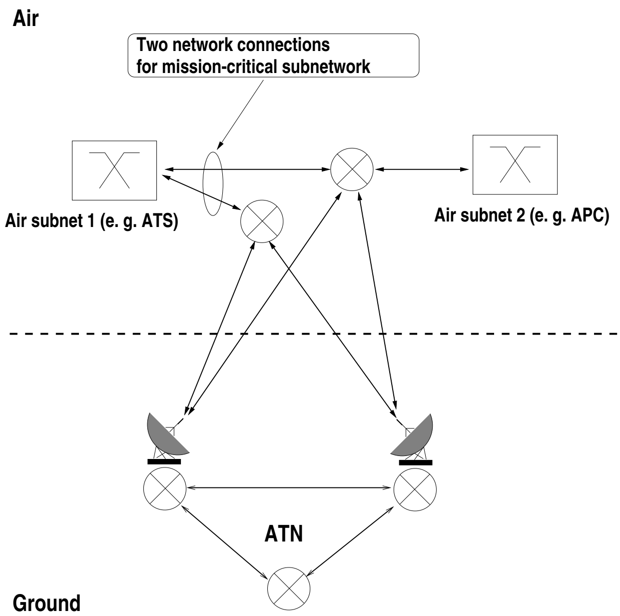

Figure 2.5: Redundant implementation of the on board router approach

2.4.1 Proxying

If a set of air applications send out all their requests and receive all their responses through an application proxy or a tunnel endpoint, then only the proxy (resp. tunneling entity) has to establish connections with ground based remote correspondents. As a consequence, the application nodes do not have to send out Mobile IPv6 Binding Updates, because they do not have any ground correspondents; only the proxy still has to perform updates. This can help reduce the network load if several on board applications establish connections with the same correspondent nodes.

Proxying is used daily in the Internet as a means of reducing traffic and limiting security hazards on corporate networks; it is a proven technique, with numerous existing implementations. In particular, a considerable, if not predominant, part of the WWW traffic on the Internet is carried through caching proxies. Multiple retransmissions of data is thus avoided, and user applications do not need to reside on an open subnetwork in order to use the WWW: they only need to be able to establish a connection to their local proxy.

Network load can be further reduced if the on board proxy is, in turn, connected to a ground based one (as is often done currently with hierarchies of WWW caching proxies). In this case, the on board proxy only has to notify *one* ground correspondent when it changes its care-of address. It may also choose to use a new ground proxy for future connections when it changes its care-of address. The means by which the on board proxy discovers the ground proxy which is closest to its point of attachment to the network is not presently defined. Ground routers may maintain a list of the nearest proxies for different applications; this information could be provided to passenger applications using the Service Location Protocol [13], or another similar system.

Proxying offers a method for cutting the cost of mobility updates for all applications where connections are initiated by the air extremity, because air applications can be kept informed of the proxy addresses they have to use. Ground applications, on the other hand, have no means of determining which proxy serves a given mobile node. If a ground node wants to initiate a proxy-supported connection with an airborne application, the application shall define a "call-back" mechanism whereby the ground application shall send a request directed to the mobile node (using its home address), to which the mobile node shall respond by establishing a connection to the requesting ground station through the appropriate proxy.

2.4.2 Tunneling

Tunneling is functionally equivalent to the above two-proxy technique; if a tunnel is established between an airborne endpoint and a ground system, then when the airborne tunneling entity changes its primary care-of address, it only has to notify the update to the remote tunnel endpoint.

Tunneling also provides a possibility to support legacy OSI applications during a transition from ATN/OSI to ATN/IPv6: IPv6 provides several facilities for transport of OSI datagrams in IPv6 packets. A mapping of NSAP addresses as used in the OSI ATN is readily available: OSI ATN uses ICD (International Code Designator) format addresses as defined by the OSI Network Service Definition [17], and RFC1888 [3] provides an algorithm for translating ICD NSAP addresses into IPv6 addresses. The addressing architecture in ATN/OSI is very sparse; it makes important use of ASCII-encoded textual designators as part of the addresses. It should therefore be easy to design an addressing scheme for ATN/IPv6 that will permit a simple mapping of OSI NSAP addresses.

CLNP tunneling has the disadvantage that all CLNP data must detour through a fixed ground-based tunnel endpoint before reaching the aircraft (tunneling prevents optimal routing); furthermore, this single tunnel endpoint might constitute a single point of failure, which may not be appropriate with respect to the safety requirements of ATN. For these reasons, porting of all ATN applications to IPv6 may constitute a more appropriate solution.

2.5 Conclusion

We have proposed three solutions for routing and mobility in the IPv6 ATN. We recommend that aircraft take a router aboard, performing IPv6 mobility as described in 2.3.3. We also recommend that proxying and tunneling be used in order to reduce the network bandwidth requirement on the air/ground data link, as described in 2.4.

Chapter 3

Multicast

3.1 Introduction

ATN/OSI has no multicast facilities. One of the objectives of the ATN/IPv6 project is to provide multicast service for FREER-1 [11], which requires the possibility to distribute information in real-time over a group of aircraft. Aircraft running the FREER-1 application have to send out position reports and trajectory intentions to potentially conflicting aircraft, including at least all aircraft that are flying within a minimal radius of 90 nautical miles [10]. In the rest of this document, the 90 nautical miles zone around an aircraft that shall be covered by position reports is referred to as a FREER-1 reporting sphere.

In order to cut down bandwidth demand, the ATN/IPv6 requirements document specify that position reports shall be propagated to a cluster consisting of up to 30 neighbouring mobile entities. Each such cluster corresponds to a set of potentially conflicting aircraft. Smaller clusters shall then be determined using the position reports exchanged through top-level clusters. Conflict clusters shall be constituted of up to four aircraft with conflicting routes. Inside conflict clusters, aircraft do not only exchange position reports, but also trajectory intentions. The requirements for two-level clustering and for limitation of cluster population to 30 mobile entities are subject to re-evaluation according to the conclusions of this report.

Our objectives are to design an architecture for the exchange of position reports and trajectory intention information using IPv6 multicast facilities. The underlying data communication system shall be either geostationnary satellites or low earth orbit satellites (hereafter GEOs, respectively LEOs).

3.2 Current practice

Current FREER-1 uses NAVSTAR GPS and/or GLONASS for precise position determination, and ADS-B (Automatic Dependent Surveillance Broadcast) radio messages for position reports. The propagation of position reports is bounded *de facto* by the limited range of ADS-B broadcasts. The typical range is from 120 to 150 nautical miles; the minimum required range is fixed at 90 nautical miles. Aircraft are thus aware of the position of all mobile entities in a sphere with a radius at least equal to 90 nautical miles around them. Trajectory intentions are also transmitted in ADS-B messages, as a set of Trajectory Change Points.

3.3 Analysis of architecture

3.3.1 Geostationnary orbit satellites — One-hop network

3.3.1.1 Presentation

We will first discuss how the ATN/IPv6 multicast requirements can be met using a set of geostationnary satellites. Due to the large coverage of GEOs, we assume that position reports exchange is only required within the cover zones of each satellite: an aircraft shall report its position to the GEO it currently has in view, and the satellite shall transmit it in a single hop to all neighbouring aircraft. The GEO shall behave as a fixed station regarding routing and mobility issues, as discussed in chapter 2; it shall embed a router that is capable of managing multicast groups, in order to forward position reports only to those mobile entities that may be interested in them.

Aircraft shall embed a GPS or GNSS device for precise determination of their position. They shall request to join and leave position report multicast group according to their movement. We will present three different methods for group management:

Centralised grouping. Aircraft report their position to an ad hoc application on board the satellite. This entity is globally aware of aircraft position information in its cover zone, and uses this global knowledge to correctly relay position reports at the application level.

Sector based grouping. Air space is statically divided into a number of adequately sized sectors, each of which is assigned one multicast group address. Each aircraft reports its position to the address of its current sectors and receives reports from its current sector and all adjacent sectors.

Cluster based grouping. Multicast groups are formed in a dynamic and distributed fashion using an election algorithm.

3.3.1.2 Centralised grouping

The first option for limiting position report bandwidth is not actually based on multicast, but on multiple unicast transmissions at application level. In this approach, a GEO embeds a user application that collects position reports from all aircraft in its cover zone, and relays them to all interested aircraft.

The embedded application has to track the position of all aircraft in its cover zone, and for each position report received, has to go through this set of aircraft and forward the position report if the distance between the position of the just-received report is within FREER-1 specified range of the recorded aircraft position.

In this first solution, we do not actually make use of IPv6 multicast, since we transfer the responsibility of choosing recipients for a given data packet to a central entity running on the satellite, in the application layer. This approach has the disadvantage of not being distributed at all, and exhibits a *single point of failure* in the satellite embedded application. We do not recommend this solution as a first choice for implementation of FREER-1 over IPv6.

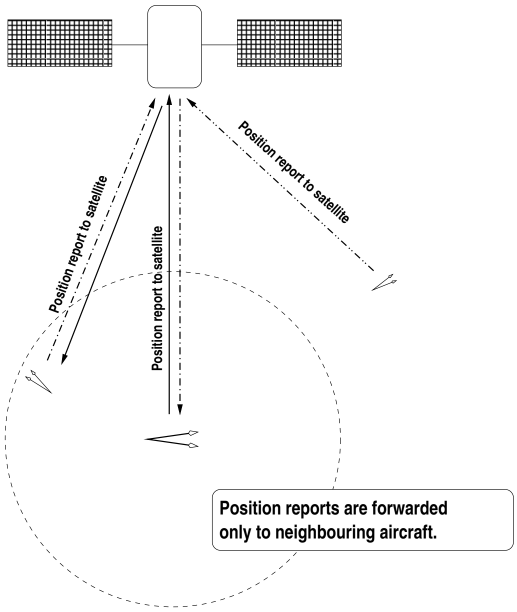

Figure 3.1: Grouping of aircraft by centralised application

3.3.1.3 Sector based grouping

In this second approach, air space shall be divided in a number of adequately sized sectors. Namely, the portion of space determined by a sector and all the immediately adjacent sectors shall be sufficiently large to enclose a sphere centred on any point in the first sector, with a radius equal to the minimal FREER-1 required range of 90 nautical miles. In the situation illustrated on figure 3.2, the aircraft shall report its position to the sector it currently inhabits, and shall receive position reports from other aircraft in all hatched sectors.

The propagation of position reports is thus effectively limited by the cell size. All aircraft receive position reports from all other aircraft within a minimal radius of 90 nautical miles, and receive no position reports from aircraft farther than 325 nautical miles

3.3.1.4 Cluster based grouping, according to distance

It is possible to devise a distributed algorithm by which aircraft shall constitute clusters and elect one clusterhead for each cluster; the clusterhead's identifier shall be used to craft a multicast group address that will be used for position report exchange among cluster members. In order to perform clustering and clusterhead election, all aircraft shall periodically report their position to all others, using a reserved multicast group for "All aircraft on same satellite".[1]

[1] If one wants to decide whether to transmit one's position reports to a given aircraft based on correspondent's distance, one has to be aware of correspondent's position. Position reports for out-of-range aircraft thus *have* to be periodically transmitted.

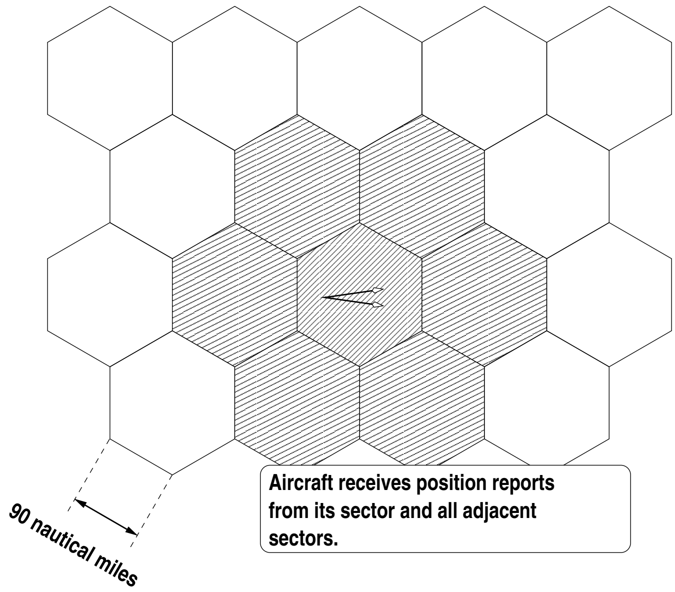

Figure 3.2: Grouping of aircraft using static air space sectors.

Together with these periodical reports, aircraft shall announce their clustering status (current cluster ID and status: clusterhead, cluster member or undecided). They shall then perform the following algorithm, which is a variation on LCC (Least Clusterhead Change) [7]:

1. Consider only announcements within a radius of 90 nautical miles. In this approach, the clustering is fully based on actual aircraft distance.

   Initially constitute clusters according to Lowest ID [12]: consider self as cluster head if no neighbours are clusterheads and all undecided neighbours have higher ID.

2. When a non-clusterhead node in cluster i moves into a cluster j, no clusterhead in clusters i and j shall be changed (only cluster members are changed; non-clusterheads cannot challenge clusterhead status).

3. When a non-clusterhead node moves out of all clusters and does not enter any existing cluster, it becomes a new clusterhead and forms a new cluster containing only itself.

4. when clusterhead $C(i)$ from cluster i moves into cluster j, it challenges the corresponding clusterhead $C(j)$. Either $C(i)$ or $C(j)$ shall give up clusterhead position according to lowest ID.

5. Nodes which become separated from a cluster will recompute the clustering according to lowest ID.

(Lowest ID can be replaced by any agreed, well-defined priority policy such as highest connectivity.)

Position reports shall be sent to the multicast group address corresponding to all clusters where the aircraft is clusterhead or cluster member. Every aircraft shall join the multicast groups for all clusters that contain potentially conflicting aircraft. This ensures sufficient propagation of position reports.

3.3.1.5 Performances and benefits

We presented three node grouping strategies in GEO context:

Centralised grouping does not make use of IPv6 multicast, and does not provide a distributed, fault-tolerant method for node grouping. It requires specific application-layer support on the satellite.

Sector based grouping is easy to implement, reactive, and very resistant to failure. Its efficience shall be assessed by simulation; it is felt that in the low-density air space covered by FREER-1 objective, this strategy will lead to efficient use of bandwidth resources.

Distance-controlled cluster based grouping also has the advantages of a distributed algorithm, but is much more complicated than sector based grouping. Its efficience is questionable, because it requires periodic global diffusion of position reports, and involves generation of duplicate position reports, defeating in this respect the purposes of multicast communication.

For these reasons, we recommend that, should geostationnary satellites be used as the communication infrastructure for FREER-1 over IPv6, air space sectorization be considered as a preferred approach for multicast group constitution.

3.3.2 Low earth orbit — Multi-hop network

3.3.2.1 Presentation

We are now considering the use of a constellation of low earth orbit satellites (LEOs) as the data link medium for air/air multicast communication. LEOs require different routing and packet delivery mechanisms than GEOs, because they are in constant motion relative to other satellites in the constellation and to aircraft in their coverage zone; besides, one cannot assume that two aircraft that need to exchange position reports will always share at least one common-view satellite, because the coverage area for a LEO is significantly smaller than that of a GEO, and is not sufficient to permanently encompass the typical radius of a FREER-1 reporting sphere. It is therefore necessary to provide multi-hop routing of position reports across the satellite constellation.

In this situation (*ad hoc* network of mobile nodes), unicast routing shall be implemented using a clustering algorithm. Cluster based unicast routing is documented in [7, 5, 18]. Multi-hop multicast routing in such a network, where all end systems and routers are constantly roaming and where no permanent infrastructure exists, is still an open research topic. It is technically possible to make use of unicast-independent multicast routing protocols, but this approach is very inefficient and inadequate in a context where routes are constantly changing due to router mobility.

Several protocols have been proposed to provide multicast service in an *ad hoc* network. CBT [2], suggested for use in *ad hoc* networks in [5], and Shared Tree

Wireless Network Multicast [6] both require routers to keep state information about multicast group membership, which is inappropriate in a network where hosts and routers are moving at high speed. The authors of [6] do mention that their protocol exhibits good performances only in networks of slowly moving nodes.

A novel approach of multicast routing in *ad hoc* networks is given in [22]. It makes use of flooding to afford an efficient, mobility tolerant, stateless multicast routing. That paper summarizes the state of the art on routing and multicast in *ad hoc* networks, and highlights the important research directions on this subject. For the purposes of this study, we will assume that a protocol for multicast routing has been chosen for LEO-supported applications.

We then propose two solutions for node grouping:

Sector based grouping. As in 3.3.1.3.

Cluster based grouping. A single multicast group is formed for all FREER-1 entities, but its scope is limited to a small number of routing clusters, thus effectively restraining the propagation of position reports only to potentially conflicting aircraft.

Since a LEO network does not provide a single contact point, We consider that centralised grouping as described in 3.3.1.2 is not an option.

3.3.2.2 Sector based grouping

The grouping method described in 3.3.1.3 can be applied in a LEO context without modification.

3.3.2.3 Cluster based grouping, according to hop count

The proposed routing algorithms for a LEO-supported air/air communication network group nodes in a number of clusters, each with a clusterhead node, which has a direct, bidirectional connection with every node in the cluster. Clusters are interconnected by gateway nodes, which are non-clusterhead nodes that have simultaneous connections to two or more clusterheads (and are therefore members of more than one cluster).

Knowing the configuration of the LEO constellation, it is possible to determine a hop count n such that if a packet is relayed n times, it is guaranteed to be delivered at least 90 nautical miles away from its origination point. A single multicast group address shall be assigned to "All FREER-1 entities"; aircraft shall send out position reports to this address with a TTL of n, and shall join this multicast group.

3.3.2.4 Performances and benefits

We presented two node grouping strategies in LEO context:

Sector based grouping has already been described earlier; it has the advantage of being applicable without changes in both GEO and LEO contexts.

Hop count-controlled cluster based grouping has the qualities of a distributed algorithm. Since it relies on the underlying routing infrastructure to perform the grouping, it does not involve any additional cost in the on board application to perform grouping. This solution also makes a very efficient use of address space, because it only requires one multicast group address for all FREER-1 entities. Contrary to distance-controlled clustering, it does not involve multiple multicast groups nor message duplication.

For these reasons, we recommend:

- that sector based grouping be considered the preferred solution for node grouping if a system compatible with both GEO and LEO satellites is to be developed;
- that hop count-controlled, cluster based grouping be considered and assessed, if the system to be developed will only use LEO satellites as the telecommunication infrastructure for FREER-1 over IPv6.

3.4 Conclusion

In this chapter, we studied the formation of multicast groups for the propagation of position reports between FREER-1 autonomous entities using satellite data links for air/air communication. If these satellites are GEOs, we recommend that a grouping scheme based on a static airspace sectorization be used, as described in 3.3.1.3. If they are LEOs, we recommend that sector based grouping or hop count-controlled, cluster based grouping be considered.

The formation of restricted multicast groups for the exchange of trajectory intentions has not be discussed here and is left for further study.

Bibliography

[1] Mary G. Baker, Xinhua Zhao, Stuart Cheshire, and Jonathan Stone. Supporting Mobility in MosquitoNet. In *Proceedings of the 1996 USENIX Technical Conference*, Jan. 1996.

[2] A. Ballardie. Core Based Trees (CBT) Multicast Routing Architecture. Request for comments 2201, Sep. 1997. IESG experimental protocol.

[3] J. Bound, D. Carpenter, *et al.* OSI NSAPs and IPv6. Request for comments 1888, Aug. 1996. IESG experimental protocol.

[4] Bjorn Chambless and Jim Binkley. HARP - Home Agent Redundancy Protocol. Internet Draft draft-chambless-mobileip-harp-00.txt, Oct. 1997. *Work in progress.* Expired April, 1998.

[5] Ching-Chuan Chiang and Mario Gerla. Routing and Multicast in Multihop, Mobile Wireless Networks. In *Proceedings of the Sixth IEEE International Conference on Universal Personal Communication (ICUPC'97)*, Oct. 1997.

[6] Ching-Chuan Chiang, Mario Gerla, and Lixia Zhang. Shared Tree Wireless Network Multicast. In *Proceedings of the Sixth IEEE International Conference on Computer Communications and Networks (ICCCN'97)*, Apr. 1997.

[7] Ching-Chuan Chiang, Hsiao-Kuang Wu, Winston Liu, and Mario Gerla. Routing in Clustered Multihop, Mobile Wireless Networks with Fading Channel. In *IEEE Singapore Conference on Networks (SICON'97)*, Apr. 1997.

[8] Matt Crawford and Bob Hinden. Router Renumbering for IPv6. Internet Draft draft-ietf-ipngwg-router-renum-03.txt, Mar. 1998. *Work in progress.* Expires September 12, 1998.

[9] Stephen E. Deering and Robert M. Hinden. Internet Protocol, Version 6 (IPv6) Specification. Internet Draft draft-ietf-ipngwg-ipv6-spec-v2-02.txt, Aug. 1998. *Work in progress.* Expires February 6, 1999.

[10] Vu N. Duong, Eric Hoffman, and Jean-Pierre Nicolaon. Initial Results of Investigation into Autonomous Aircraft Concept (FREER-1). Technical report, EUROCONTROL Experimental Centre, Jun. 1997. Presented at the first USA/Europe ATM R&D Seminar.

[11] EUROCONTROL Experimental Centre. FREER project home page. `http://www.eurocontrol.fr/projects/freer/`.

[12] Mario Gerla and Jack Tzu-Chief Tsai. Multicluster, mobile, multimedia radio network. *ACM-Baltzer Journal of Wireless Networks*, 1(3):255–265, Jul. 1995.

[13] Erik Guttman, Charles Perkins, John Veizades, and Michael Day. Service Location Protocol, Version 2. Internet Draft draft-ietf-svrloc-protocol-v2-08.txt, Jul. 1998. *Work in progress*. Expires January 30, 1999.

[14] International Civil Aviation Organization. *Aeronautical Telecommunication Network (ATN) Standards and Recommended Practices (SARPs)*, 2.2 edition, Jan. 1998.

[15] International Organization for Standardization. Information processing systems — Telecommunications and information exchange between systems — End system to Intermediate system routeing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service (ISO 8473). International standard ISO 9542:1988, 1988.

[16] International Organization for Standardization and International Electrotechnical Commission. Information technology — Telecommunications and information exchange between systems — Protocol for exchange of inter-domain routeing information among intermediate systems to support forwarding of ISO 8473 PDUs. International standard ISO/IEC 10747:1994, 1994.

[17] International Organization for Standardization and International Electrotechnical Commission. Information technology — Open Systems Interconnection — Network Service Definition. International standard ISO/IEC 8348:1996, 1996.

[18] Mingliang Jiang, Jinyang Li, and Yong Chaing Tay. Cluster Based Routing Protocol (CBRP) Functional Specification. Internet Draft draft-ietf-manet-cbrp-spec-00.txt, Aug. 1998. *Work in progress*. Expires February 1999. Implementation expected for 43rd IETF, December 7–11, 1998.

[19] David B. Johnson. Mobile IP Virtual Home Agents. In *Proceedings of the Thirty-Fourth IETF*, Dallas, Dec. 1995.

[20] David B. Johnson and Charles Perkins. Mobility Support in IPv6. Internet Draft draft-ietf-mobileip-ipv6-05.txt, Mar. 1998. *Work in progress*. Expires September 13, 1998.

[21] Thomas Narten, Erik Nordmark, and W. A. Simpson. Neighbor Discovery for IP Version 6 (IPv6). Internet Draft draft-ietf-ipngwg-discovery-v2-03.txt, Aug. 1998. *Work in progress*. Expires February 3, 1999.

[22] Katia Obraczka and Gene Tsudik. Multicast Routing Issues in Ad Hoc Networks. In *Proceedings of the Seventh IEEE International Conference on Universal Personal Communication (ICUPC'98)*, Oct. 1998.

[23] ISA Project. Air Traffic Requirements. EEC Note XX/98, Telecom Centre of Expertise — EUROCONTROL Experimental Centre, 1998.

[24] Susan Thompson and Thomas Narten. IPv6 Stateless Address Autoconfiguration. Internet Draft draft-ietf-ipngwg-addrconf-v2-02.txt, Feb. 1998. *Approved by IESG for publication as Draft Standard RFC.*

Internet Drafts are draft documents valid for a maximum of six months. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "work in progress."

To learn the current status of any Internet Draft, please check the listing, `1id-abstracts.txt`, contained in the Internet Drafts Shadow Directories:

- `ftp.is.co.za` (Africa),
- `ftp.nordu.net` (North Europe),
- `ftp.nis.garr.it` (South Europe),
- `ds.internic.net` (US East Coast),
- `ftp.isi.edu` (US West Coast),
- `munnari.oz.au` (Pacific Rim).

ISO standards can be obtained through ISO national member bodies (listing available from `http://www.iso.ch/addresse/membodies.html`). General ISO information can be found at `http://www.iso.ch/`.