SIMULATION STUDY OF THE HOLD DOWN TIMER IN THE AERONAUTICAL TELECOMMUNICATIONS NETWORK (ATN)

Introduction

MITRE's Center for Advanced Aviation System Development (CAASD) has been tasked to participate in validation of the ATN design. To investigate some aspects of the ATN, CAASD is using a simulation of the ATN using MIL3 Incorporated's OPNET. This simulation has recently been used to investigate the effect on mobile routing of the hold down timer, or minimum route advertisement interval, in the Interdomain Routing Protocol (IDRP, ISO/IEC 10747 [1]).

The simulation assumes that IDRP can be configured to support routing policies proposed in the ATN Manual [2] and refined in "The ATN Routing Concept," 10 October 1994, by Tony Whyman [3]. These routing policies are designed to minimize the amount of routing information by eliminating the need to distribute routing updates to all routers. As of the writing of this paper, these proposed routing policies have not been applied to real IDRP implementations. CAASD, among others, is working to do so.

The results presented below do not depend on whether IDRP is on the aircraft; therefore, the results apply to both Package 1 and the system as currently specified in the ATN Manual. The simulation used the Federal Aviation Administration's (FAA's) portion of the ATN; however, the general conclusions should be applicable to other parts of the ATN.

Purpose of the Hold Down Timer

Before describing the simulation experiments and results, it is important to clarify the purpose and function of the hold down timer in IDRP. As defined in ISO/IEC 10747,

> minRouteAdvertisementInterval determines the amount of time that must elapse between advertisements of routes to a particular destination from a single BIS....Two UPDATE PDUs sent from a single BIS that advertise feasible routes to some common set of destinations received from BISs in other routeing domains must be separated in time by at least minRouteAdvertisementInterval.

ISO/IEC 10747 allows the timer to be set between 5 and 1800 seconds. In the Internet, this type of timer is generally set between 1 and 6 minutes. The timer's intent is to prevent temporary surges in routing updates. The timer limits the maximum rate at which a router can advertise a route to a destination. So if a route to a destination changes rapidly, say from a failure condition like a link between Boundary Intermediate Systems (BISs) that toggles rapidly in and out of service, the number of updates gets damped. Also, if multiple changes to the route arrive during the hold down interval, only the latest information should be forwarded when the timer expires. Thus, the timer restricts not only the *rate* at which routes can be forwarded, but also may reduce the *amount* of routing information.

Hold Down Timer and the ATN

Since one of the concerns within the ATN is reducing the amount of routing updates generated by mobiles, some view the hold down timer as one of the key benefits of using IDRP for mobile routing. For example, [3] argues that

> ...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.

The paper goes on to give an example where the hold down timer reduces the frequency of routing updates without adverse impact on mobile routing.

In many cases, the hold down timer does have the positive effect described in [3]; however, CAASD's simulation has discovered situations in which the hold down timer will make it temporarily impossible to send data to a mobile. This result is described below.

"Implementation" of the Hold Down Timer in the ATN Simulation

The hold down timer in the simulation works as follows: When a router has a feasible route to advertise, it first checks if it has advertised a route to this destination within the last minRouteAdvertisementInterval seconds. If not, the router sends an "Add" Update immediately and records the time. If the router has recently advertised to the same destination, it holds the Add Update until the timer expires. If another Add Update to the destination is already pending, the new one replaces it. If a route becomes unavailable and there is no replacement route, the router forwards a Withdrawal Update immediately. If, when the Withdrawal Update is sent, an Add Update to the same destination is pending, the Add Update is discarded. Sending either an Add Update or a Withdrawal Update resets the hold down timer.

Description of Experiment

For these experiments, the simulation included a Federal Aviation Administration (FAA) ground topology of 22 routing domains, one for each Air Route Traffic Control Center (ARTCC). Within each routing domain was a single router, acting as a BIS. The ground routers were sparsely interconnected by point-to-point links, as shown in Figure 1. (An ARTCC is an en route control facility, sometimes referred to as a "center.") The simulation could have assumed that all routers shared a common subnetwork; however, using a mesh topology increased the potential for hold down effects and gave more insight into performance of the global ATN, where not all ATN routers will be on a common ground subnetwork.

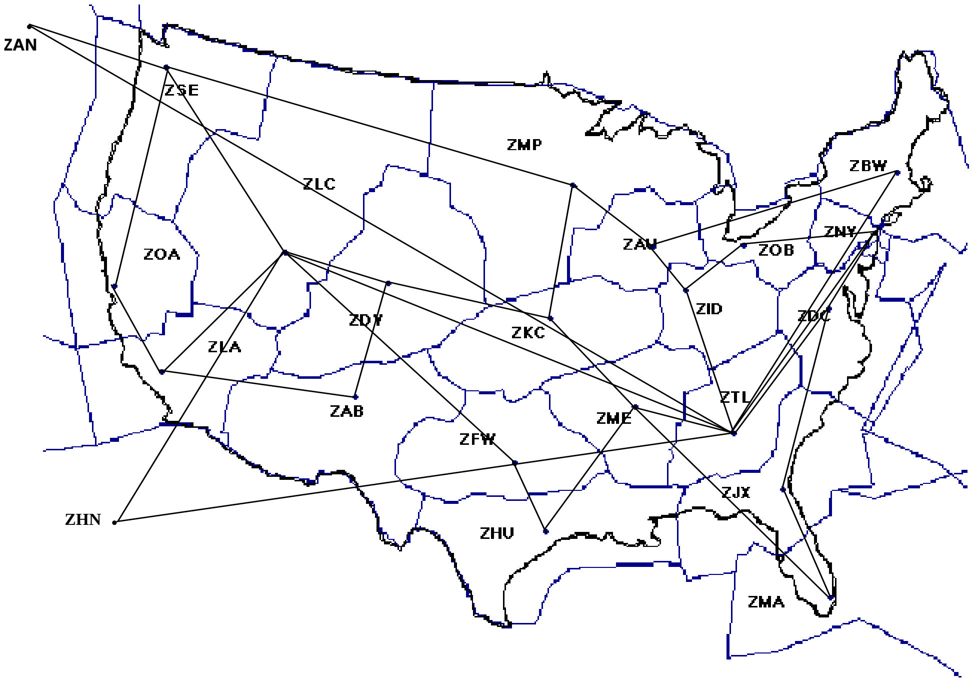

Figure 1. Network Topology

Two of the routers, Atlanta (ZTL) and Memphis (ZME), acted as "backbone" routers for this ATN "island." These routers were chosen to minimize the amount of routing updates in the overall network. [4] explains how this determination was made.

IDRP in the simulation was configured to follow the policies required in the ATN Manual and "The ATN Routing Concept." Following these policies, when an aircraft connects to a ground routing domain, the ground router sends an IDRP Add Update toward the nearest backbone router. (For robustness, a real network might forward updates to more than one backbone router via diverse paths.) In the simulation, this is accomplished by having each ground router's IDRP configured to know the next hop router on a single default route to the nearest backbone router. The two backbone routers exchange all routes with each other. If a ground router receives a CLNP packet whose address is not in its Routing Information Base (RIB), it forwards the packet by default toward the backbone, which is advertising a route to all aircraft. If a backbone router does not have an explicit route to the aircraft, the router should forward the packet to a Home domain outside of the FAA island. Since all aircraft in the simulation are over the FAA island, it is assumed that if the FAA's backbone does not have a route to the aircraft, neither will the home; therefore, in our simulation, the backbone discards the packet when it does not have a route.

To keep run times reasonable, the simulation included only aircraft with connections to Kansas City (ZKC), Denver (ZDV), Albuquerque (ZAB), Fort Worth (ZFW), and Houston (ZHU). Figure 2 shows how these five ground routers were connected and indicates airspace adjacencies. The arrows indicate where IDRP updates were sent following the shortest path to the backbone.

An average of about 30 aircraft were in each domain's airspace. Aircraft flight times in an ARTCC's airspace were exponentially distributed with a mean of 27 minutes. The minimum time in an ARTCC's airspace was 60 seconds. When an aircraft crossed a routing domain boundary, it established an IDRP connection with the next BIS. Approximately 60 seconds after establishing the new IDRP connection, the IDRP connection to the previous domain was disconnected, simulating 110 to 120 seconds of overlapping coverage (including 50 to 60 seconds to establish the new IDRP connection). (Overlapping coverage is discussed further later in the paper.) Mode S was the only air/ground ATN subnetwork in the simulation.

Each aircraft ran several applications. For one application, the aircraft sent a short request to an end system in one of the 22 ground routing domains chosen at random. The ground end system sent a response back to the aircraft. Getting the response to the aircraft required mobile routing. If the ground end system's local router sending the response did not have an explicit route to the aircraft, it forwarded the packet to the nearest backbone node by default. Packets were dropped if routing information has not propagated to the backbone. Thus, even though aircraft flew over only five routing domains, all 22 ground routing domains participated in the simulation.

Results

Figures 3 shows the time from when an update was created until it was received at its destination. Some updates are processed at one or more intermediate routers before reaching the backbone. The hold down timer was set to 275 seconds. Statistics were recorded from 20,000 to 30,000 seconds into each simulation run. Notice that the maximum delay was longer than the hold down timer value (275 seconds). The bottom chart in Figure 3 is the cumulative distribution function (cdf). The cdf is the probability that the time to distribute an Add Update was less than a given length of time. Figure 3 shows that in most cases (>95% for this simulation), updates were not delayed, but there was a very long tail for those that are delayed.

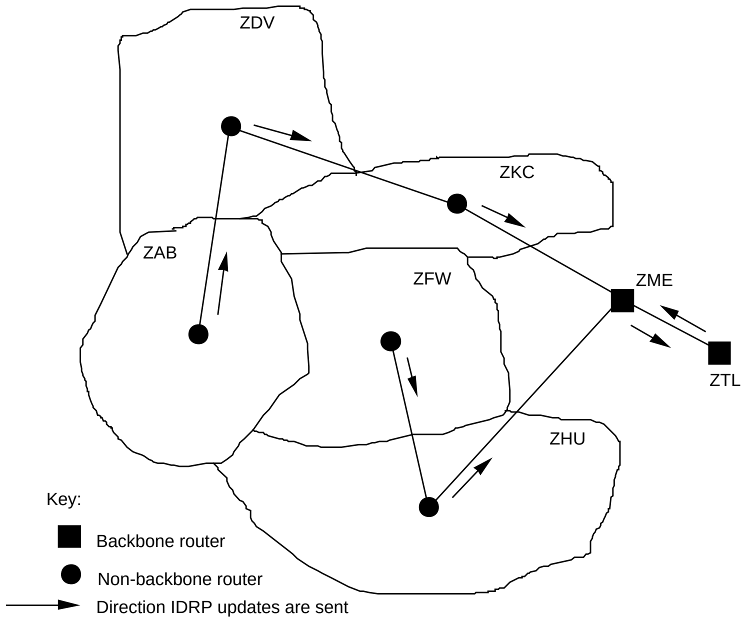

Figure 2. IDRP Update Forwarding to Backbone

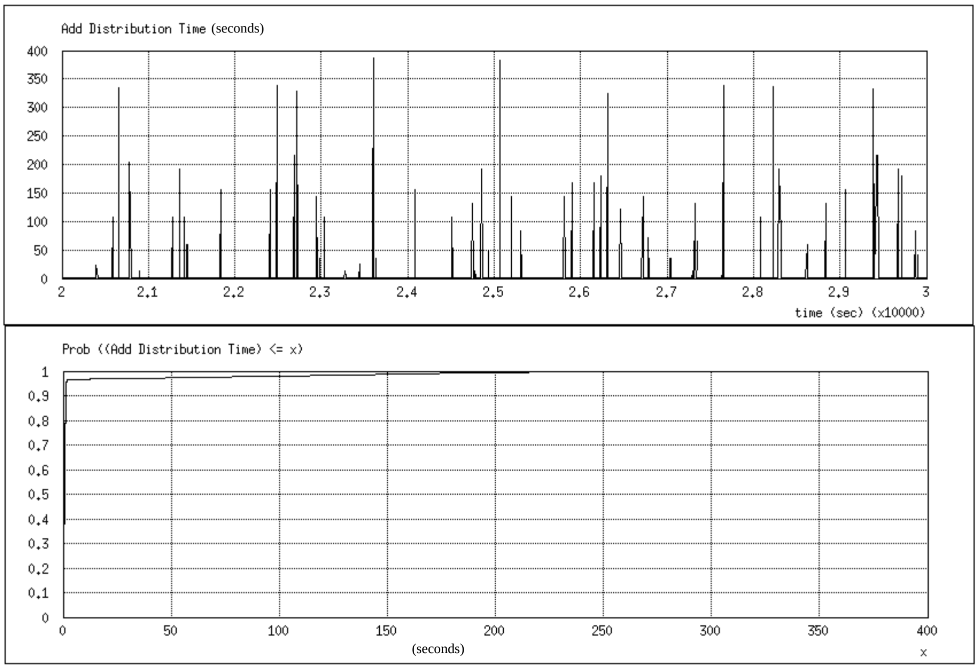

Figure 3. Time to Distribute Add Updates

Packets to several aircraft were dropped because of the hold down timer. For example, the simulation included a flight illustrated in Figure 4. This flight illustrates what can happen when an aircraft flies close to the border of coverage areas or when a holding pattern crosses subnetwork coverage areas. The figure shows the time that certain events occurred and highlights the portion of the flight affected by the hold down timer. All times in the following discussion are referenced from when the aircraft entered coverage of the Houston (ZHU) routing domain.

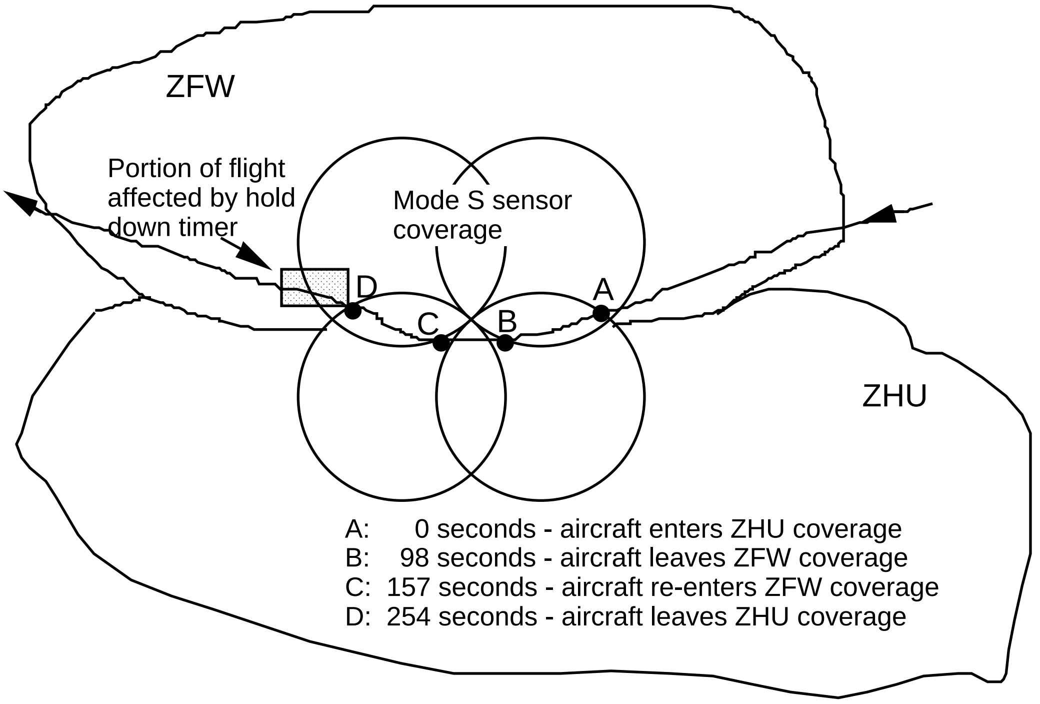

Figure 4. Sample Flight Path

The sequence of IDRP updates and ZME’s route to this aircraft are shown in Table 1. In the Updates Sent column, (A) indicates the update is an Add; (W) indicates a Withdrawal.

Because the aircraft left ZFW and returned shortly thereafter, ZFW was forced to hold its advertisement when the aircraft returned. This allowed enough time for the aircraft to leave overlapping coverage with ZHU, so ZHU, having no replacement route, sent ZME a withdrawal. Once ZHU did find out about the route via ZFW, it could not immediately forward the route. Thus, from 254 seconds until 529 seconds, the backbone routers had no route to the aircraft. During this 4.5 minute interval, the two backbone routers discarded 20 packets to this aircraft.

Table 1. Sequence of Events and IDRP Updates

Time (seconds into simulation run)	Events and Updates Sent	Update Distribution Time (sec)	ZME RIB Contents
A 0	Aircraft enters ZHU coverage		
50	IDRP connection with ZHU established		
50	ZHU -> ZME (A)	0	ZME - ZHU - A/C
50	ZME -> ZTL (A)	0	no change
B 98	Aircraft leaves ZFW coverage		
98	ZFW -> ZHU (W)	0	no change
C 157	Aircraft re-enters ZFW coverage		
206	IDRP connection with ZFW re-established		
D 254	Aircraft leaves ZHU coverage		
254	ZHU -> ZME (W)	0	**No route to this aircraft**
254	ZME -> ZTL (W)	0	**no change**
373	ZFW -> ZHU (A)	167	**no change**
529	ZHU -> ZME (A)	323	ZME - ZHU - ZFW - A/C
529	ZME -> ZTL (A)	323	no change

Figure 5 shows the length of time that the backbone routers did not have any route to the aircraft as a function of the hold down timer for this particular flight scenario. Figure 5 also shows the number of packets to the aircraft that were dropped because of the lack of a route at the backbone. If the hold down timer was less than 156 seconds (the time between events B and D), ZFW was able to advertise that it again had a connection to the aircraft before ZHU lost its connection to the aircraft and sent a withdrawal. If the hold down timer was longer than 156 seconds, then ZHU first sent the withdrawal and had to wait the entire length of the hold timer before advertising the route via ZFW. So beyond the threshold (156 seconds in this case) the length of loss of communication is equal to the hold timer interval.

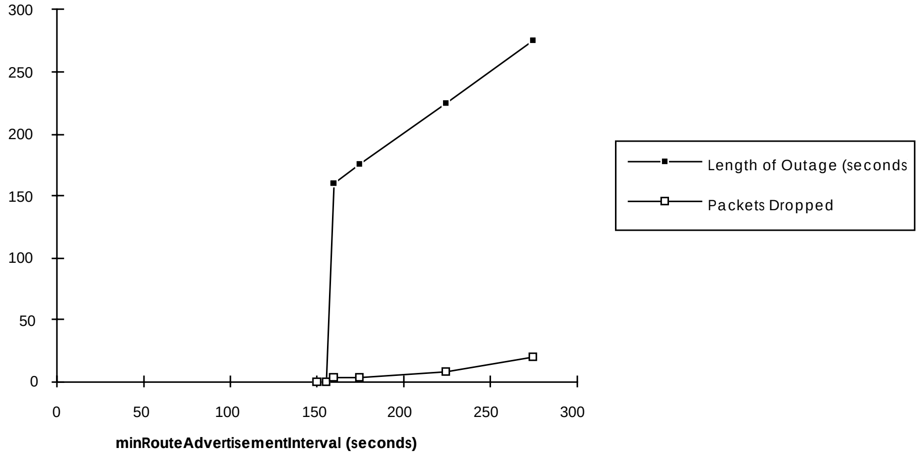

Figure 5. Communications Outage vs. Hold Down Timer

The discontinuity of the length of time that the backbone had no route is a surprising effect. A similar effect could occur with a fixed hold down timer and slightly different flight times across a center boundary. In this situation, different aircraft following slightly different flight paths could experience surprisingly different service from the ATN.

It was noted earlier that a positive aspect of the hold down timer is that it will reduce the number of routing updates. The simulation did show a reduction in the number of routing updates received at the backbone, but only by a few percent.

Hold Down Timer and Intradomain Routing

In this simulation, the loss of communication only affected end systems which were not local to ZFW and had to use the backbone to reach the aircraft. The simulation assumes there is a single BIS within each routing domain and all CLNP packets originating in a routing domain sent to a destination outside the routing domain go through that BIS. For robustness, it may be desirable to have more than one BIS in a routing domain. In this case, packets originating in the routing domain with a direct connection to an aircraft could be routed to a BIS that does not have the direct connection. In this case, the packets are forwarded to the backbone and the hold down timer could impede communication between a controller and an aircraft in local airspace.

Subnetwork Coverage and Frequency of Hold Down Effects

Knowing that the hold down timer can have negative consequences for mobile routing, it would be interesting to know how often such effects will occur. The frequency of hold down effects will be strongly influenced by how flight paths intersect subnetwork coverage areas and how much subnetwork coverage areas overlap. Because the simulation did not include a detailed model of coverage areas, including cones of silence and terrain masking, the simulation was not able to provide a good answer to this question. These coverage issues are discussed below.

Amount of Overlapping Coverage

The simulation has shown two regions of operation with respect to the length of overlapping coverage. In the first case,

time in overlapping coverage < time to establish IDRP connection.

In this case, the previous router will advertise a withdrawal before the new router can advertise a feasible route. If the withdrawal is forwarded through the new router, it will have to hold the advertisement of its direct route for nearly the full hold timer interval. Alternatively, if the new route is advertised to the backbone via the previous ground router, that router will have to hold the new route until its hold down timer expires since it sent the withdrawal. In either case, the backbone will not have a route to the aircraft for nearly the length of the hold down timer. The worst case of this condition is a “break-before-make” air/ground subnetwork, where there is no overlapping coverage.

In the second case,

time in overlapping coverage < time to establish IDRP connection + hold down timer.

When this condition is true, the hold down timer may prevent communication with an aircraft that leaves and re-enters a center.

The simulation results described above are for the second region, where

time in overlapping coverage > time to establish IDRP connection

but

time in overlapping coverage < time to establish IDRP connection + hold down timer.

Cone of Silence

Other aspects of air/ground subnetworks may cause the hold down timer to take effect. Mode S has a “cone of silence” or “zenith cone” directly above the sensor as illustrated in Figure 7. The cone exists because radar antenna patterns are designed to provide coverage at a distance rather than directly overhead. The diameter, d, of the cone at an altitude, a, is given by

$$d = 2 * a * \tan 54^{o} = 2.75 * a$$

For an aircraft 30,000 feet above the antenna, the cone is approximately 13.6 nautical miles (nmi) in diameter. At 400 knots, an aircraft would take 122 seconds to fly directly across the cone. If the cone is not covered by an adjacent sensor, the router will receive a leave event after the aircraft has been in the cone for approximately 5 scans (24 seconds for terminal radars and 60 seconds for en route radars). Whether there is overlap depends on altitude and sensor range and location. If a leave event is generated, the hold down timer may delay advertisement of a route to the aircraft after it re-enters the coverage area. Current plans propose 118 terminal sites with 100 nmi range and 25 en route sites with 200 nmi range. Projected coverage for this configuration does leave some cones of silence uncovered by adjacent sensors for some altitudes (e.g., 30,000 feet).

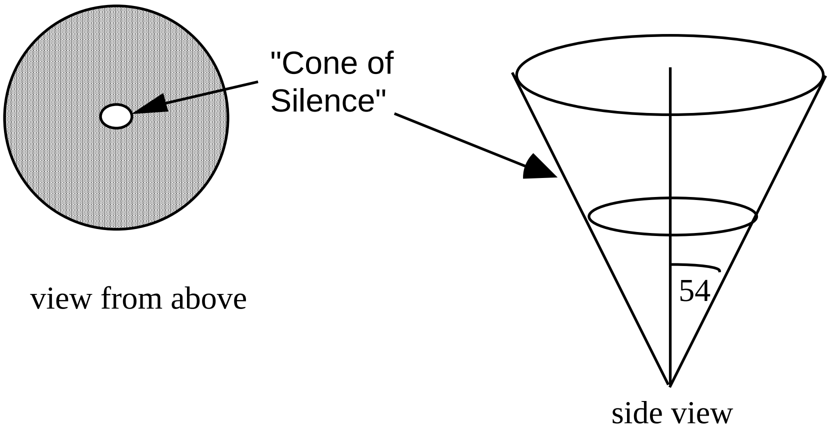

Figure 7. Mode S Cone of Silence

Terrain Masking

Terrain masking is another aspect of subnetworks which may cause the hold down timer to take effect. Terrain masking creates jagged edges on coverage areas as is shown in Figure 8. Aircraft that fly across these jagged edges will be in and out of coverage quickly. The hold down timer may damp and thus delay readvertisement of aircraft reachability.

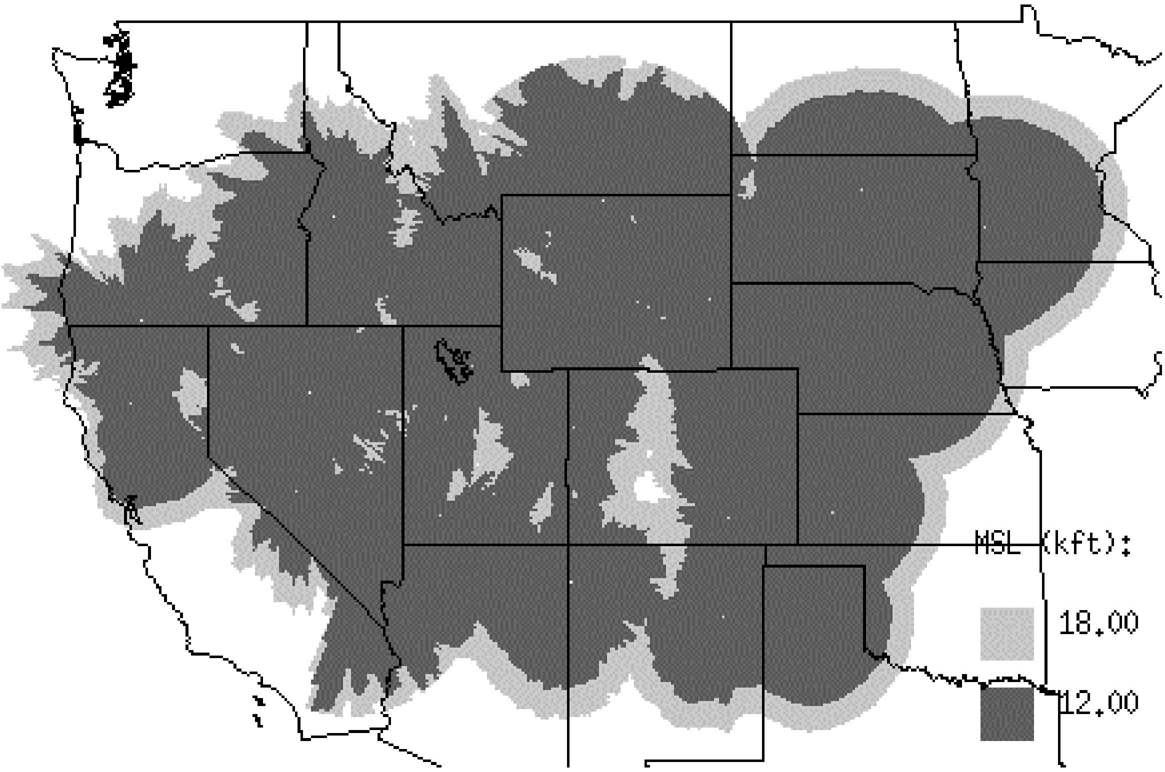

Figure 8. Mode S Coverage at 12000 and 18000 Feet

Because Mode S will be a means of surveillance, it is likely that gaps in coverage from cones of silence and terrain masking will be minimized for controlled aircraft. Any gaps in coverage that are left are likely to be in locations where air/ground communication is not highly critical. Even so, the potential effect on data link communication should be kept in mind as the overall system is designed and deployed.

Hold Down Timer Implementations

Another factor that can influence how often update packets are delayed is the way the hold down timer is implemented. As was quoted early in this paper, the ISO/IEC 10747 description of the hold down timer is as follows:

> minRouteAdvertisementInterval determines the amount of time that must elapse between advertisements of routes to a particular destination from a single BIS....Two UPDATE PDUs sent from a single BIS that advertise feasible routes to some common set of destinations received from BISs in other routeing domains must be separated in time by at least minRouteAdvertisementInterval.

ISO/IEC 10747 goes on to say that “any technique that ensures that the separation will be between one and two times the value minRouteAdvertisementInterval is acceptable.” A simple way to meet this requirement would be to subject all add updates to the hold down timer and send updates only at fixed intervals, with the interval equal to the hold down timer. The simplicity of this technique might make it attractive to vendors since it would have no noticeable consequences for fixed routing; however, such an implementation would have serious consequences for the ATN.

The hold down implementation in the CAASD simulation is more complex in order to guarantee the fastest possible distribution of routing information. Standards-compliant, real world implementations may not be as responsive, in which case hold down effects in the ATN will be significantly more frequent than in our simulation.

Strategies to Minimize Hold Down Effects

If IDRP continues to be the mobile routing solution, then steps should be taken in the overall ATN system design to minimize the negative impact of the hold down timer. A temptation will be to set the hold down timer very low; however, it would be risky to set it lower than experience in the Internet has proven necessary. Potential steps are listed below. The cost of each must be weighed.

- Use make-before-break air/ground subnetworks
- Design air/ground subnetworks to ensure several minutes of overlapping coverage
- Ensure that cones of silence are covered by adjacent sensors
- Use IDRP implementations that distribute routing information as soon as the hold down timer expires
- Interconnect all BISs on an ATN island with a common subnetwork so that routing information does not cross multiple internetwork hops before reaching the backbone
- Allow a service provider to hide most subnetwork transitions
- Limit the ground network design to one very high availability BIS per routing domain to prevent the hold down timer from affecting distribution of routing information within a routing domain

Conclusions

The simulation has shown that the hold down timer will affect distribution of routing information, which in some cases will temporarily prevent communication with aircraft. These outages may seem unpredictable to those using the network. The simulation is not currently able to predict how often such outages will occur because of the difficulty in accurately modelling subnetwork coverage.

The aeronautical community must consider whether they are willing to accept a mobile routing strategy which will occasionally prevent data link communication with an aircraft for up to several minutes. Other mobile routing strategies, such as Mobile IP [5], Cellular Digital Packet Data (CDPD) [6], and the Dynamic Extensible Routing Protocol (DERP) [7], do not suffer the same loss of communication as IDRP experiences because of the hold down timer. Comparisons of IDRP based mobility with other strategies should be made. If IDRP continues to be the basis of the aeronautical community's mobile routing, then the hold down timer should not be set artificially high to reduce the amount of routing information, and subnetworks and the ground network topology must be designed to minimize hold down timer effects.

Acknowledgments

The author would like to acknowledge that Ted Signore, CAASD, and Dallas Scott, CAASD, suggested this line of investigation and provided significant technical assistance. Leslie Crane, CAASD, helped the author understand Mode S coverage and provided the coverage map in Figure 7. Many others offered helpful comments.

References

1. "Information Processing Systems - Telecommunications and Information Exchange between Systems - Protocol for Exchange of Inter-domain Routeing Information among Intermediate Systems to Support Forwarding of ISO 8473 PDUs," ISO/IEC 10747, 18 October 1993

2. ATN Manual, Second Edition, 19 November 1993.

3. "The ATN Routing Concept," Tony Whyman, 10 October 1994.

4. CAASD, "The ATN Simulation Model, Status Report," Enclosure to F088-L-139, 30 July 1993.

5. "IP Mobility Support," Internet Draft, C. Perkins, ed., 4 January 1995.

6. "Cellular Digital Packet Data System Specification" , Release 1.0, 19 July 1993.

7. "An Architecture for Quasi-Distributed Internetwork Level Mobile Routing," Dallas Scott, March 1992