Modular, Cost-Effective, Extensible Avionics Architecture for Secure, Mobile Communications

William D. Ivancic, NASA/GRC
NASA Glenn Research Center
Cleveland, Ohio 44135
(216) 433-4000
wivancic@grc.nasa.gov

Abstract—Current onboard communication architectures are based upon an all-in-one communications management unit. This unit and associated radio systems has regularly been designed as a one-off, proprietary system. As such, it lacks flexibility and cannot adapt easily to new technology, new communication protocols, and new communication links. This paper describes the current avionics communication architecture and provides a historical perspective of the evolution of this system. A new onboard architecture is proposed that allows full use of commercial-off-the-shelf technologies to be integrated in a modular approach thereby enabling a flexible, cost-effective and fully deployable design that can take advantage of ongoing advances in the computer, cryptography, and telecommunications industries.

TABLE OF CONTENTS

.....	
1. INTRODUCTION.....	1
2. CURRENT ARCHITECTURES.....	1
3. FUTURE IP-BASED ARCHITECTURE	5
4. POLICY-BASE ROUTING	6
5. LAYER-2 TRIGGERS	6
6. VOLUME	7
7. SUMMARY.....	7
REFERENCES.....	8
BIOGRAPHY	9

1. INTRODUCTION

In October, 2003, NASA embarked on the ACAST project (Advanced CNS Architectures and System Technologies) to perform research and development on selected communications, navigation and surveillance (CNS) technologies to enhance the performance of the National Airspace System (NAS). The Networking Research Group of NASA's ACAST project, in order to ensure global interoperability and deployment, formulated their own salient list of requirements. Many of these are not necessarily of concern to the FAA, but are a concern to

those who have to deploy, operate and pay for these systems. These requirements were submitted to the world's industries, governments, and academic institutions for comments. Overall responses showed a consensus agreement on six major points [1,2]:

- (1) It is critical that any new technologies being deployed provide a positive return on investment (ROI).
- (2) Network Centric Operations (NCO) will be a major technology in future airspace systems and the next generation Internet Protocol, IPv6 will be the protocol of choice.
- (3) Links should be shared, and the system should be provider-independent. This makes QoS a requirement.
- (4) A common global security structure must be developed and IPsec is probably the best choice. Some work still needs to be done regarding IPsec multicast, envisioning a certificate-based security architecture, and figuring out how exactly to do QoS with respect to wireless links and encryption.
- (5) The system must be able to share network infrastructure.
- (6) The system must be extensible to meet future needs.

2. CURRENT ARCHITECTURES

Aircraft Communications Addressing and Reporting System

The current avionic communication architecture is shown in figure 1 [3]. This architecture has evolved since the early 70's and is based on the Aircraft Communications Addressing and Reporting System (ACARS). The ACARS network is modeled after the point-to-point telex network where all messages come to a central processing location. The datalink service provider (DSP) routes the message to the appropriate end system using its network of land lines and ground stations. The two DSPs available to the airlines are Aeronautical Radio, Incorporated (ARINC) and Societe

¹ U.S. Government work not protected by U.S. copyright

Airlines originally operated the ACARS network for their maintenance, flight and cabin operations. Since the late 80's, equipment such as the Communication Management Units (CMU) or Air Traffic Service Unit (ATSU) has been used to deliver additional information for example pre-departure clearance, oceanic clearance and digital-airport terminal information service over the airline data link. During the early 1990s, a datalink interface between the CMU and flight management system (FMS) was created to enable flight plans and weather information to be sent from the ground to the CMU. Soon after, an interface between the CMU and the flight data acquisition and management system / aircraft condition monitoring system (FDAMS/ACMS)³ was created. This system is capable of identifying abnormal flight conditions and automatically sending real-time messages to an airline thereby enabling airlines to better monitor their engine performance and to identify and plan repair and maintenance activities.

The ACARS system is comprised of an airborne subsystem and a ground system. The airborne subsystem consists of the CMU, radio systems, a display screen and printer. The ground system consists of all the ACARS remote transmitting/receiving stations, and the datalink service providers computer and switching systems. The Air Carrier Command and Control and Management Subsystem is part of the ground infrastructure and provides all the ground-based airline operations including: operations control, maintenance and crew scheduling.

ACARS messages may be of three types: Air Traffic Control (ATC), Airline Operational Control (AOC), and Airline Administrative Control (AAC). ATC messages are used by aircraft crew to request clearances, and by ground controllers to provide those clearances. AOC and AAC messages are used to communicate between the aircraft and its base. Various types of messages are possible. These include fuel consumption, engine performance data, and aircraft position as well as free text data.

Typical ACARS messages are only 100 to 200 characters in length. Such messages are made up of a one-block transmission from (or to) the aircraft constrained to be no more than 220 characters within the body of the message. For downlink messages which are longer than 220 characters, the ACARS unit will split the message into multiple blocks with an additional constraint that no message may be made up of more than 16 blocks. Thus, ACARS utility is limited to relatively short messages.

Initially, ACARS used only very high frequency (VHF) Data Link (VDL) communication. ACARS has since

Figure 1 - Typical ACARS Onboard Network

expanded to other communication links such as satellite communication (SATCOM) and high frequency (HF). ACARS has also been updated for use over aviation VHF link control using VDL Mode 2. VHF is the most commonly used and least expensive link. Since transmission is line-of-sight, VHF is not available over the oceans. Current satellite communications is via the INMARSAT's satellite network which provides nearly global coverage – except over the poles [4]. The SATCOM links are extremely expensive services relative to other options. Thus, other solutions are desirable. The most recently established subnetwork is HF and is used to provide coverage in the Polar Regions.

Data link messages can be sent either via a VHF, HF, or satellite network. The current software within the CMU automatically decides the most efficient (and cheapest) path for delivery of the message, depending on the location of the aircraft. *It is important to note that today, critical ATC messages are delivered by the same path as other messages, such as AOC flight plan uplinks. There is currently no priority assigned to ATC messages.*

Future Air Navigation System

In the early 80's, the airlines were increasingly using digital links between their aircraft and their ground services for logistical flight management via ACARS. They were also looking for a means to offer telecommunication services to the passengers. In addition, information exchanges between collocated and/or remotely located ground systems required modernization of the existing links with deployment of ground networks.

³ These systems analyze engine, aircraft, and operational performance conditions

Figure 2 - ATN Data Communication Environment

The airlines also realized that standardization of a single system which utilized various digital communications technologies and ensured interoperability would improve performance, maintenance and administrative management while simultaneously reducing risk and cost [5].

In 1983, ICAO chartered the Special Committee on the Future Air Navigation System (FANS) to study the current air traffic infrastructure and recommend changes to support the anticipated growth in air traffic over the next 25 years⁴. The FANS committee identified these needs:

- Replacement of the current analog radios with digital air/ground communications;
- Use of satellite and HF communication systems to provide communication where deployment of line-of-sight systems is not practical such as in the oceanic domain;
- Global Interoperability;
- Network-enabled systems to support automation in the airplanes and on the ground;

- Transition to a Global Positioning System (GPS)-based navigation and landing systems; and,
- Installation of flight service automation to enable pilots to plan and file flight plans without reliance on flight service specialists.

Aeronautical Telecommunication Network

As a result of the FANS studies, ICAO decided to standardize the aeronautical network and modernization of the information exchanges between the ground and the aircraft, for air traffic control purposes. This is done by adopting common interface services and protocols based on the Open System Interconnection (OSI) model of the International Standardization Office (ISO). This model distributed the telematics functions in 7 layers of functional responsibility thereby providing a mechanism for standardization of the different data links, and their complementary use. These standards are known as the Aeronautical Telecommunication Network (ATN) standards.

ATN could be used for all digital communications by the aeronautical community. The ATN consist of application entities such as Controller and Pilot Data Link Communication (CPDLC), Automatic Dependant Surveillance (ADS) and communication services (routing)

⁴ It is now 2005 – 22 years later, and only an extremely small portion of FANS has been deployed – particularly regarding the Aeronautical Telecommunication Network

which allow avionics, air-to-ground and ground networks to interoperate. The ATN has been designed to provide data communications services to Air Traffic Service provider organizations and Aircraft Operating agencies. Communication traffic that was envisioned to reside on the ATN included: Air traffic services communication (ATSC); aeronautical operational control (AOC); aeronautical administrative communication (AAC); and aeronautical passenger communication (APC). As a result of global deployment of the Internet Protocol suite, TCP/IP, passenger service is currently running over Internet Protocol networks. Furthermore, these networks are likely to replace ATN as IP technologies and protocols can now meet the salient requirements of ATN and are globally deployed.

While waiting for Aeronautical Telecommunication Network (ATN) to develop and become available, Boeing built a FANS application to run on the existing ACARS system. The application includes CPDLC and ADS. This avionics package became known as FANS-1. The Airbus Industry equivalent system is known FANS-A. These systems are known collectively as FANS-1/A. FANS -1/A uses the network configuration shown in figure 1. On today's aircraft, FANS-1/A and ATN ATS Data link applications cannot be used simultaneously on the same aircraft as one is ACARS-based and the other ATN-based. Furthermore, simultaneous use of ATN and FANS applications is not expected to be retrofit to current generation aircraft because of legacy equipment and architectures [6].

Figure 2 shows the basic structure of the ATN. The main components of the ATN are the end systems (ES), the ATN routers and subnetworks. The subnetwork is part of the communication network, but is not part of the ATN. It is defined as an independent communication network based on a particular communication technology (e.g. X.25 Packet-Switched Network) which is used as the physical means of transferring information between ATN systems. The ATN routers are responsible for connecting various types of subnetworks together. ATN end systems host the application services as well as the upper layer protocol stack in order to communicate with peer end systems. Note that the air-ground subnetworks consist of the HF, VHF, Satellite, and Mode-S links. These are the same links and infrastructure as ACARS. The only difference is that ATN routers are connected to the end of the radio systems. The ACARS onboard architecture shown in figure 1 closely corresponds to the ATN onboard architecture with the CMU taking on additional ATN router functionality.

The ATN supports communication between ground-to-ground systems and air-ground systems. The ground-to-ground systems include: airline systems and ATS systems, ATS to ATS systems and airline systems. The air-ground systems include: airline and aircraft systems; and ATS and aircraft systems. IP routers are currently performing many of the connectivity between ground systems as commercial

Figure 3 - ATN Island Routing Domain Confederation

off the shelf (COTS) equipment provides the necessary functionality at a fraction of the cost of an ATN router. Furthermore, much of ATN was based on X.25 packet switching, a technology that is no longer supported by the commercial community.

The key differences between an ATN inter-domain router and a standard OSI router are: possibility of applying a specific set of routing policies in support of mobile communication (e.g. which link is least expensive or best quality); support provided for (currently undefined) ATN security functions; and use of compression for air-ground routers to increase the efficiency when using bandwidth-limited air-ground data links.

Today, the major features that ATN routers are left to perform are the air-ground and ground-ground communication with regard to mobility. In ATN, keeping track of the location and routing to aircraft – mobility – is performed using the Inter-Domain Routing Protocol (IDRP) and by confining the domains to a relatively small group, sometime described as islands. A distributed IDRP directory using Boundary Intermediate Systems (BISs) is implemented along with a two level directory approach using an ATN Island concept consisting of backbone BISs and a home BISs concept [Figure 3]. This is done to limit the convergence time of route updates. If the routing structure were to become too large, convergence times would become unacceptable.

It is extremely important to note that using a routing protocol to handle mobility effectively requires one to own the entire infrastructure. This is necessary simply because one generally is not allowed to inject routes into another's infrastructure. Furthermore, using ATN protocols requires deployment of ATN ground infrastructure at the end of each radio system. Thus, when new technologies become available, it is very difficult to utilize them without building out an ATN specific implementation.

3. FUTURE IP-BASED ARCHITECTURE

By using an IP-based network rather than an ATN-based network, we can meet the salient requirements highlighted in the Introduction section. Aircraft mobility can be handled by mobile-IP. The onboard mobile router takes care of all mobility such that none of the end systems on the mobile networks need to be mobile-aware. Mobile routing is link independent and does not inject routes into the infrastructure [7-10]. This allows for use of shared infrastructure. One does not have to own the infrastructure and can easily insert new link technologies as they mature.

A number of architectural variations should be investigated regarding mobile networking. These include: placement of the home agents, geographically distributed home agents, security considerations, policy issues for deployment of multiple mobile networks within the same mobile router, and quality-of-service over open networks. For example, should the mobile network be located in the airlines domain, the civil air authority's (CAA) domain, or both? The latter implies that the airline and CAA networks are interconnected at the mobile router. This has security issues that need to be addressed but could possibly be resolved through good architecture design and use of IP security [Figure 6].

Figures 4, 5 and 6 show IP-based onboard architectures that enable low-cost system deployment. The entire system is COTS based. No special requirements are placed upon the equipment. This architecture uses encryption devices and firewalls to securely isolate critical subnetworks. With ATN, such security is not easily implemented as the equipment is not readily available. Because of this, ATN requires a very high trust relationship between the aircraft networks and the ground networks.

Figure 4 shows an IP-based transitional architecture. One cannot expect the onboard avionics to change for a long period of time. Current avionics systems are designed for safety and redundancy well before the Internet technologies came into being. Thus, the avionic backplane and bus are not IP compliant and will require a specialized controller/gateway to provide and interface between the IP network and the avionic equipment and sensors. Notice that the avionic subnetwork and the passenger subnetwork are securely isolated but that both can use a common mobile router.

Figures 5 and 6 are similar to 4 with the exception that the CMU is no longer needed to provide and interface to the onboard avionics as the avionic are now fully IP compliant. In figure 5, the aircraft operations subnetwork is securely isolated from the air traffic control network. In figure 6 the ATC and AOC subnetworks are separate, but protected by the same firewall/encryption unit. In effect, one can trade some security for simplification of the network and

reduction in equipage. For these architectures, all links carry IP packets. Current avionics links are designed for very small messaging. For a fully IP-based network, the avionics radios' media access and data link layers would need to be developed to better handle IP packet-based communications.

Figure 4 - IP-Based Transitional Architecture

Figure 5 - IP-Based Architecture with ATC and AOC Combined

Figure 6 - IP-Based Architecture with AOC and ATC Separation

4. POLICY-BASE ROUTING

In ATN, the inter-domain routing protocol is used to propagate routing policy. Each routing domain contains its own routing policy. Routing policy is advertised outside the domain by the boundary intermediate system router (BIS). The ATN routing policy is used to determine the “best route” to take when more than one link is available to and from the aircraft. Although this requirement has existed within the specification from the beginning, its use has been limited to date and operationally untested for the following reasons: there currently are not enough ATN users to tax the system; system deployment is minimal; and, the airlines generally only have one link active. For cost reasons, SATCOM is not turned on unless needed. Furthermore, two simultaneous VHF radios are not active simultaneously.

Initial commercial implementations of mobile networking for IPv4 only allowed for one link to be used at any given time, even if two or more links were available [11]. Work within the industry and in the Internet Engineering Task Force (IETF) network mobility (NEMO) and Mobile Nodes and Multiple Interfaces in IPv6 (monami6) working groups are addressing this issue [12-14].

Figures 7 through 9 illustrate the advantages of policy-based routing in a mobile network. Consider the mobile network having three links available. One link has been classified as highly reliable but relatively low rate. This link is reserved for command and control. The second link is a low latency, low bandwidth link. The third link is high-rate for passenger services⁵. Assume policy is set with the following rules:

- (1) Only ATC traffic is allowed to use the reliable link.
- (2) Data precedence is set such that ATC is highest priority, AOC is next highest and passenger traffic has lowest priority.
- (3) ATC and AOC traffic are allowed to use the low-latency link
- (4) ATC, AOC and passenger traffic are allowed to use the high-rate link.
- (5) Link preference for ATC is reliable link – highest, low-latency link – middle, high-rate – last.
- (6) Link preference for AOC is low-latency followed by high-rate.

Figure 7 shows all links active. Figure 8 shows that ATC traffic can be delivered even if all other links are unavailable.

⁵ The passenger link may be classified as secondary, but being a money generating link with the potential for real-time, directed advertising riding on this link, the availability will likely be as good or better than other links.

Figure 9 shows that ATC and AOC traffic have precedence over passenger traffic and could use the high-rate link if their preferred links are unavailable. Figure 9 is of greatest interest because one could conceivably make this the preferred link for all traffic if safety-of-flight QoS requirements could be met. Doing so would release

Figure 9 - Policy-Based Routing, All Links Active

Figure 8 - Policy-Based Routing, Critical Link Active

Figure 7 - Policy-Based Routing, Passenger Link Active

spectrum to ATC and AOC as many users could be using the high-rate links when available.

5. LAYER-2 TRIGGERS

Current avionic links provide for some minimal quality-of-service and message prioritization. This is performed within the radio or between the CMU and the radio with prioritization being preconfigured. Since the messages are small and the link capacity is low there is little need to have a feedback mechanism between the radio and the router to enhance QoS. Current and future high-rate links would benefit greatly by having a standardized feedback mechanism between the radio systems and the router. Such mechanism could indicate if a link is available and the quality and bandwidth of the link. The former is important for fast handovers between links. The latter is of particular importance for bandwidth-on-demand systems. For

instance, the Boeing Connexion outbound radio link can operate from approximately 16 kbps up to 1 or 2 Mbps. This rate is continually varying depending on outbound traffic demands and satellite network congestion. Assuming the interface between the router and Connexion radio is an Ethernet connection, some type of layer-2 trigger or feedback to the router is necessary to determine the available data rate. If the interface is serial, having the radio provide the clock may solve the data rate problem.

Air traffic control and management applications are very short messages. Therefore, it is not necessary for the air traffic control and management applications to know what link is being used or what bandwidth is available. These applications have already been developed to operate over extremely bandwidth limited systems. For future air safety applications such as transmission of secure video, the application would have to be link-aware or be developed in a manner that enables the application to figure out the type of link it is transitioning and operate accordingly.

6. VOLUME

In order to obtain a positive return on investment (ROI), the overall system costs must be affordable. The system costs include equipment, installation, deployment, down-time losses during installation, and infrastructure. One of the most likely ways to achieve positive ROI is by volume production and reuse of existing technologies. In the US alone, it is estimated that commercial airlines make up only 4% of the active civil aircraft – approximately 15,000 out of a total of 215,000 aircraft [15]. “Airbus forecasts that of this total, 16,600 new passenger aircraft of more than 100 seats will be needed in the coming 20-year period, creating an average 830 deliveries per year [16].” In contrast, today, 700 million cars are globally deployed. This is for a human population of 6 billion. Toyota expects to produce 9.2 million vehicles in 2006. General Motors produce approximately 9.1 million vehicles in 2005 [17]. Tens of thousands of aircraft over a 20 year period is not large volume. Millions of units of anything per year is.

Internet technology and mobile networking is a technology that will be integrated into automobiles. The car-to-car consortium is dedicated to the objective of further increasing road traffic safety and efficiency by means of inter-vehicle communications [18]. The Internet Car (iCar) project in Japan is working to make automobiles nodes on the Internet. iCar is researching how to connect automobiles to the Internet, how to obtain drive-by data from automobiles via the Internet, and how to design the mechanisms to share information between automobiles effectively [19]. The Internet ITS (Intelligent Transport System) Consortium is an organization in Japan exploring the possibility of ITS and other related information services. Several member organizations are jointly developing various applications and trying them out now. Applications

being developed for cars trucks and busses are numerous and include:

- Car-to-car communication
- Driver assistance information where the location and other information about each vehicle was exchanged by car-to-car communication.
- ITS taxi service where the taxi company runs a system to distribute the best taxi based on the locations, idle/operation information and customer preference/location.
- Probe servers where a probe server shares information gathered by various probes from different vendors and distributes it in a uniform manner. The server can collect car inspection information and maintenance log, as well as recall information and tell when a given part needs to be exchanged, based on mileage meter and used period of time.
- Probe data analysis and synthesis where time/location data among various probe data can be integrated to create traffic information. The system allows prediction of traffic jams for user-specified day of week and time, as well as telling the best route to the destination.
- Vending machine networks where vending machines⁶ can become wireless LAN access points, to offer broadband wireless communication infrastructure.
- Large volume content distribution service where encrypted data contents can be downloaded onto car-equipped devices and decryption key can be sent later to enable a new type of distribution, which lowers communication cost and makes download operation transparent.
- Next-generation road service where computer-assisted road service automates the process of locating and failure of a broken-down car and towing it to a desired destination.

These types of technologies and applications are appropriate for deployment considerations in general, business class, military and commercial aircraft.

7. SUMMARY

Current avionics communication architectures are based upon an all-in-one communications management unit. The origin of these systems can be traced back to global teleprinter network, telex, established in the 1920s! Today ACARS is widely deployed in commercial airlines. The ATN network is an attempt to modernize ACARS, using most of the existing radio technologies with limited modifications. These systems are designed to be deployed in a closed, aeronautics-only network. In addition the systems lack flexibility and cannot adapt easily to new technologies, new communication protocols, and new

⁶ Vending machines are widely deployed in Japan and will be network to provide information on content and restocking. Thus, they can be used to provide connectivity to the wired Internet for other systems such as cars.

communication links. Use of the same Internet technology as being developed for other mobile vehicles – in particular automobiles – will enable low-cost, highly reliable systems that can provide a positive return on investment, share network infrastructure and be extensible to meet future needs.

REFERENCES

- [1] “NASA Request for Comments on Global Air Space System Requirements,”
<http://roland.grc.nasa.gov/~ivancic/RFI/rfi.html>,
November 2005
- [2] W. Ivancic, “NASA’s Proposed Requirements for the Global Aeronautical Network and A Summary of Responses,” 2005 NASA ICNS Conference & Workshop May 2 - 5, 2005 Fairfax, VA,
http://roland.grc.nasa.gov/~ivancic/papers_presentations/2005/ICNS2005_NASA_RFC_Paper_Final.pdf,
November 2005
- [3] Collins CMU-900, 523-0810056-00111J 2M-9/02, 2002.
http://www.rockwellcollins.com/content/pdf/pdf_2251.pdf, November 2005
- [4] <http://aero.inmarsat.com/services/>
- [5] “Aeronautical Telecommunication Network (ATN) Comprehensive ATN Manual (CAMAL) Part I - Introduction and Overview,” prepared for the ATNP Working Groups by FANS Information Services Ltd, January 1999
- [6] “FANS-1/A Technical Capabilities,” ICAO Data Link Steering Group, DLS G/2 – WP, July 29, 2005
- [7] V. Devarapalli, R. Wakikawa, A. Petrescu, P. Thubert, “RFC 3963 - Network Mobility (NEMO) Basic Support Protocol,” January 2005
- [8] C. Perkins, “RFC 3344 - IP Mobility Support for IPv4,” August 2002,
- [9] W.Ivancic, D. Stewart, T. Bell, P. Paulsen, D. Shell: “Securing Mobile Networks in an Operational Setting,” IEEE Computer Communications Workshop 2003, October 2003
- [10] W. Ivancic, P. Paulsen, D. Stewart, D. Shell, L. Wood, C. Jackson, D. Hodgson, J. Northam, N. Bean, E. Miller, M.Graves and L. Kurisaki: “Secure, Network-Centric Operations of a Space-Based Asset: Cisco Router in Low-Earth Orbit (CLEO) and Virtual Mission Operations Center (VMOC),” NASA/TM-2005-213556, May, 2005

- [11] "Cisco 3200 Series Mobile Access Router Software Configuration Guide," October 14, 2004
- [12] S. Gundavelli, "IP Mobility – Motivation and Protocols," presentation at the Native6Inc Advanced Mobility Workshop, July 2005, <http://www.native6.com/assets/PDF/workshops/amw-sgundavelli-071805.pdf>
- [13] <http://www.ietf.org/html.charters/nemo-charter.html>
- [14] <http://www.ietf.org/html.charters/monami6-charter.html>
- [15] United States Department of Transportation, Bureau of Transportation Standards http://www.bts.gov/publications/national_transportation_statistics/2005/html/table_01_11.html, November 2005
- [16] Airbus Global Market Forecast 2004 – 2023 http://www.airbus.com/en/myairbus/global_market_forecast.html, November 2006
- [17] Lewis, L., "Toyota to overtake GM by end of next year," The Times, October 27, 2005, <http://business.timesonline.co.uk/article/0,,9067-1844828,00.html>, October 2005
- [18] <http://www.car-to-car.org/>, November 2006
- [19] <http://www.wide.ad.jp/project/wg/iCAR.html>, November 2005

BIOGRAPHY

Will Ivancic is a senior research engineer at NASA's Glenn Research Center working in the networking and advanced communication technology development. Mr. Ivancic's work includes: advanced digital and RF design, communications networks, satellite onboard processing, and system integration and testing. Mr.

Ivancic's recent work has concentrated on research and deployment of secure mobile networks for aerospace and DoD networks