Network and Architectural Considerations in the Design of an Autonomous Air Traffic Control System for Commercial Airplanes

By Lanny Fields

Senior Engineer
Honeywell International, Phoenix, AZ

For SAE 574:
Net-Centric Systems Architecting and Engineering
Prof. Ken Cureton
Fall 2008
Section #32344

Table of Contents

Table of Figures

Table of Tables

1. Abstract

The current trend toward the use of digital systems in airplane and air traffic control has been made possible with the exponential increase in the computational capabilities of processor-based systems. Another well-known trend in the aerospace industry is the ever-increasing amount of air and ground traffic resulting from airlines and airports attempting to accommodate the needs of the flying public – a growing population in and of itself. The intersection of these suggests that one possible solution to alleviate air travel congestion could be the automation of air traffic control and allowing it to have direct control over airplane flight paths. Such a system would, in theory, reduce the workload of the flight crew and the air traffic controllers, as well as increase traffic flow. The system would also be very complex even by modern air traffic and data processing standards; great care would need to be taken in developing its architecture in order to properly design the system.

This paper presents several analyses of such a conceptual system from a “net-centric” perspective. The system’s operation is described from the context of a flight along with the network organization. The analyses include the architecture of the system and issues related to growth, scalability, robustness, and flexibility. Other analyses include some of the “-ilities" such as availability, testability and maintainability. An examination of the system from the perspective of failures is presented: failure modes, effects, criticality and mitigation, fault tolerance and redundancy management. Collaboration among ground stations and airplanes will also be discussed along with issues related to information assurance: security, data integrity, authentication and trust between collaborators.

The concept of an automated air traffic control system which controls airplanes requires a high degree of operational integrity and availability. The overall benefit to the flying public is to allow more airplanes to be in the sky and on the ground – an obvious necessity as air traffic continues to increase. However, with increased complexity comes increased risk and it is important to fully understand the issues in order to mitigate those risks. By applying techniques for analyzing networked systems, the problem can be reduced to something more manageable, thereby lessening the risk associated with developing the system.

About the author: Lanny Fields has worked as a systems engineer in aerospace for 15 years on numerous fly-by-wire flight control systems. His master’s degree coursework at USC in the area of systems architecting and engineering has engendered a keen interest in architecting, modeling and analysis of complex systems, including net-centric architectures and systems. He is currently working on the Boeing 787 Dreamliner program.

2. Background

Air traffic congestion is rapidly becoming one of the major commercial transportation challenges at the start of the 21st century as more people take to the skies for their travel needs. "Forecasts indicate a significant increase in demand, ranging from a factor of two to three by 2025…. In short, U.S. competitiveness depends upon an air transportation system that can significantly expand capacity and flexibility, in the presence of weather and other uncertainties, while maintaining safety and protecting the environment. "[1] The FAA's current plan to comprehensively upgrade the existing air traffic control system to meet this projected demand is collectively called the Next Generation Air Transportation System, or NextGen for short.

This paper will describe the concept of an Automated Air Traffic Control System (AATCS) which is designed to enhance the existing capabilities of NextGen in and around airports. Airplanes normally receive much of their real-time in-flight data from air traffic control over voice communication with the pilot, particularly during take-off and landing. Pilots then act upon the instructions that they received verbally, however, the system does not efficiently allow pilots to use all of the available airspace and it is prone to error. The AATCS seeks to improve upon this by allowing the flight path of the airplane to be controlled by commands from stations on the ground, with the pilot still able to assume control during degraded, reversionary and emergency modes of operation. One aspect of NextGen called "Free Flight" moves away from point-to-point travel in unrestricted airspace and allows pilots more discretion in flight planning to avoid zig-zagging between air traffic control stations, however, this freedom is curtailed while in restricted airspace.[2] This renders the AATCS less suitable for use during an airplane's cruise mode, but more desirable in the vicinity of an airport.

The potential benefits from using a system like the AATCS stem from its net-centric system characteristics, which are distinguished from network-centric characteristics by a key concept: an "enhanced ability to operate and use a system that has been enabled by network technology" by the primary enabling characteristics of a net-centric system: time, location independence and collaboration.[3] The combination of the enabling characteristics allows for rapid access and comprehension of time-critical information, rapid decision-making and action, the ability for nodes to communicate without knowledge of physical location, and permitting information exchange between nodes to facilitate the decision-making process.[4] The net-centric nature of the AATCS would effectively permit airplanes to fly more closely spaced apart while maintaining safe navigation throughout the departure and arrival patterns. This type of system requires a high level of security and a robust architecture as well as bandwidth and computational power in order to manage the complexity of the data being processed, transmitted and acted upon in real-time, such as a net-centric system is capable of providing.

The AATCS system's architecture will be presented and then evaluated from several perspectives that are considered to be net-centric in nature. Examples or excerpts of some of these analyses will be provided; exhaustive analyses, such as complete failure modes, effects and criticality analyses for every element in the AATCS, are outside the scope of this paper.

3. AATCS System Description

3.1. AATCS System Overview

The Automated Air Traffic Control System consists of two primary system element types: ground control stations and airplanes. The two types are connected via an air-to-ground wireless network and are in constant communication with the other nodes in the network. Each system element type also communicates with other network members of its own type: ground stations within the vicinity of an airport are linked to each other and airplanes communicate with other airplanes within range. Ground stations have additional interfaces with secondary system elements such as external data sources. Airplanes possess their own internal networks which connect on-board subsystems to flight control computers. Each element and its architecture and interfaces are described in further detail in this section.

A top-level diagram of the system is shown in Figure 1 below.

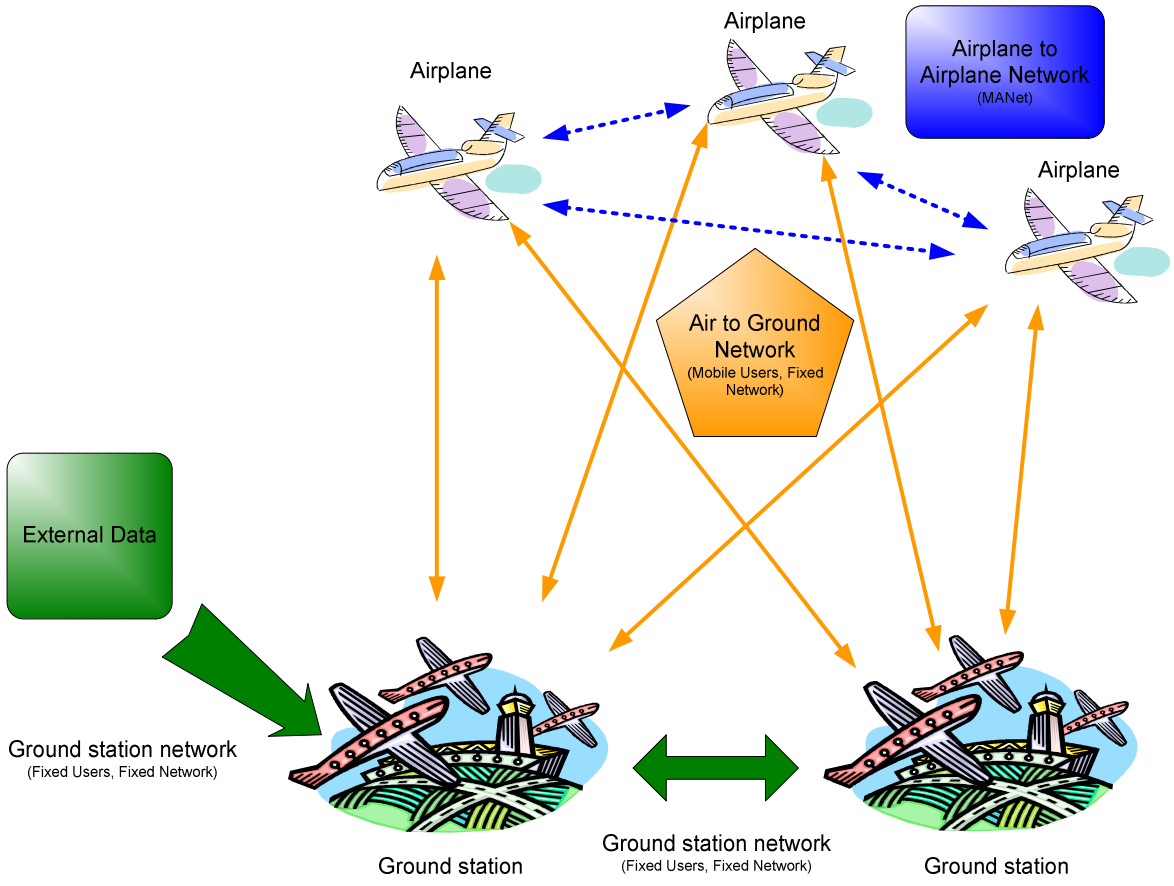

Figure 1 – Automated Air Traffic Control System

3.2. AATCS System Operation

The operation of the system is best described in the context of a flight.

When an airplane has taxied onto the runway and is ready to depart, it connects to the air-to-ground network and, after authentication, begins processing the flight commands it receives from the ground stations. At this point, the flight commands are nothing more than instructions to hold for take-off. The airplane also connects to and similarly authenticates with the air-to-air network. After the airplane verifies that the data received is valid, the pilot engages the automatic control system and allows the flight commands from the ground stations to have full authority. The airplane then accelerates through take off and rotation into the air. As the airplane follows the computed trajectory during climb to its cruise altitude, it eventually loses contact with the ground stations at the point of departure. If the commands become invalid or communication is lost at any point during these maneuvers, the system automatically disengages and the pilot takes over.

The pilot is assumed to take control at this point for Free Flight during cruise for the reasons previously mentioned in the Background section. However, the pilot could decide to allow the automatic mode to continue computing the flight vector and fly the plane based on the last valid commands received and its current position, with updates provided by any "waypoint" ground stations it connects to and authenticates with en route. The airplane does not attempt to connect with another airplane in an ad-hoc air-to-air network until it reaches its destination.

As the airplane enters the airspace of the destination airport, it once again connects to and authenticates with the local air-to-ground and air-to-air networks. It performs the same actions as during take-off, though in reverse. The pilot, if in command, relinquishes control of the airplane after the data from the local networks has been validated. The airplane then automatically slots itself for approach and landing, in accordance with the ground station's instructions. After landing, the airplane taxies off the runway and transitions back to pilot control before reaching the gate. One possible improvement might be to include automated maneuvers to guide the airplane all the way back to the gate, however, the discussion and analysis of that part of the system is beyond the scope of this paper.

3.3. AATCS System Elements

3.3.1. Ground Station Network

The ground station network is comprised of ground control stations which continually receive and process data from external sources, which are described in further detail below. The ground control stations also communicate with each other and verify their results against the results received from other ground stations in the airport's network. The processed results, which are the flight commands for airplanes in the network, are broadcast wirelessly while the stations and the external data sources are all connected via a fiber optic backbone.

A notional diagram of the regional ground control station network is shown in Figure 2.

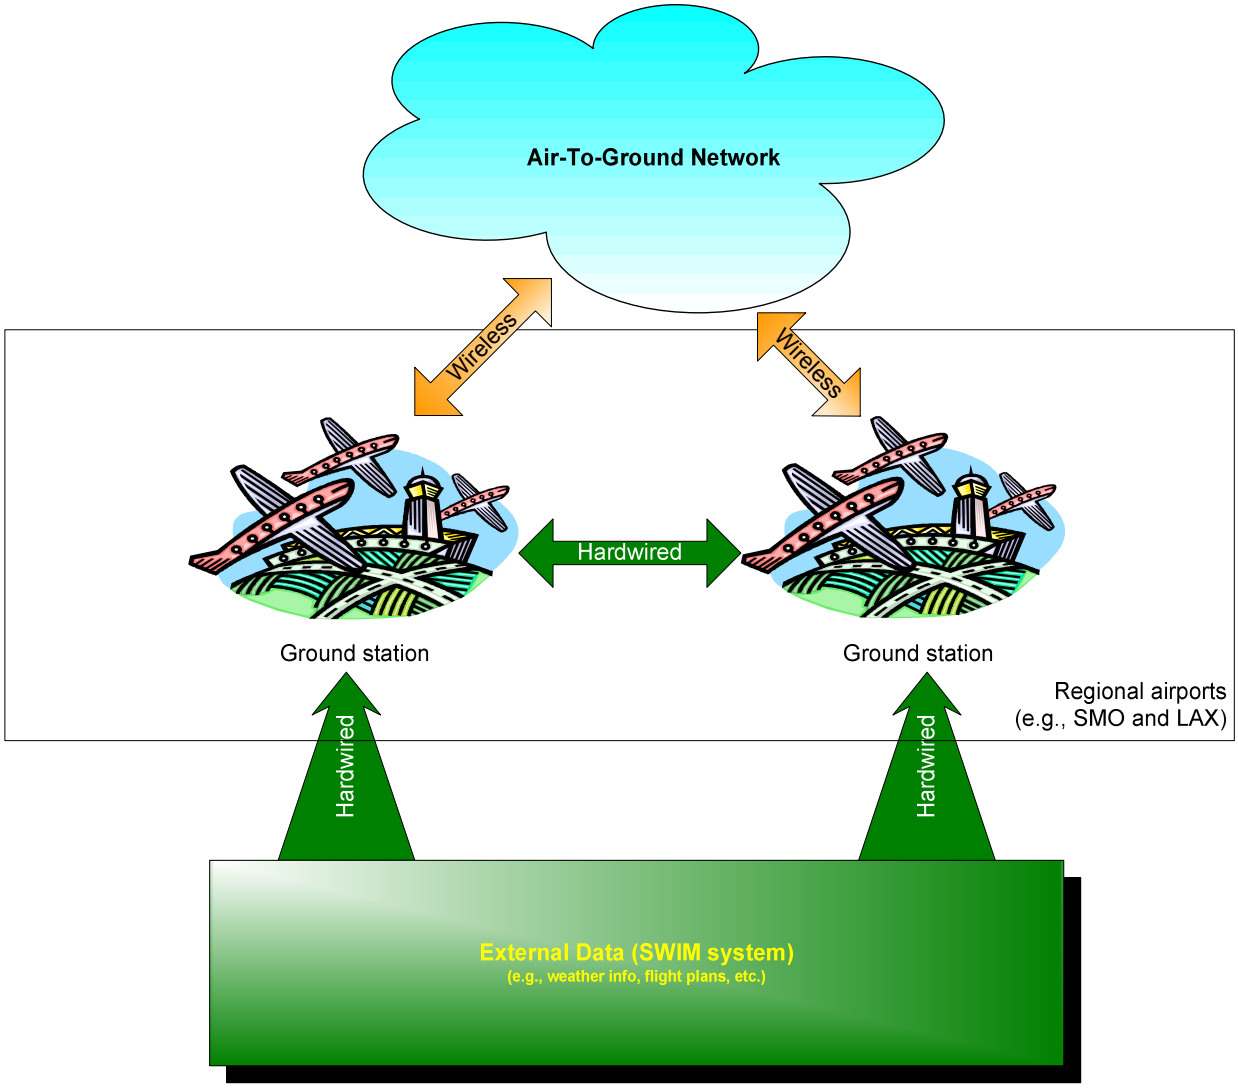

Figure 2 – Regional Ground Control Station Network

3.3.2. Ground Control Stations

The ground control stations are located in close proximity to airports and can be co-located with air traffic control. Each station receives flight-related data on the ground station network from the new NextGen System-Wide Information Management (SWIM) system, which is a service of the National Airspace System (NAS) that provides "surveillance, weather, and flight data, aeronautical and NAS status information"[5]

The ground control station processes the data continually and transmits flight path corrections on a real-time basis (once per second) to the airplanes within its airspace. These corrections are commands which are received and interpreted by the flight control computers on an airplane. The ground control stations also receive flight vector and status information from airplanes.

The stations' processors and data storage for non-repudiation are located on-site in separate server rooms. Station personnel monitor the data in real-time on displays and consoles, advising pilots and remote ground station personnel as needed during normal operations and continuously during an emergency situation.

3.3.3. External Data Sources (SWIM System)

External data from a variety of sources is required in order to analyze and understand the impact of effects in real-time which could affect air travel:

- Local weather data from the immediate airfield and surrounding airports
- A new national weather data "enterprise service dissemination of common weather observations and forecasts to enable collaborative and dynamic NAS decision making" known as the NextGen Network Enabled Weather (NNEW)[6]
- Electronically filed flight plan information
- Coordinates and vectors of the local airspace
- Reports of turbulence and flight path deviations from other airplanes
- Redundant data from stations within the regional ground control station network

This data is concentrated using the NextGen SWIM system service and is broadcast to ground control stations over the ground station network.

3.3.4. Air-to-ground network

The air-to-ground network is the wireless communication link between the ground stations that are within the range of airplanes in the local airspace. The ground stations provide flight commands to the airplanes based on the data received from the external sources and from flight vector information from the airplanes. Due to the critical nature of the data transmitted on the network, two transceiver channels are required to satisfy the corresponding FMECA events. Airplanes connect to the network and authenticate before receiving and accepting airplane-specific flight commands from the ground stations. An airplane will disconnect from the network when out of range from the ground stations for departures, or when taxiing from the runway for arrivals.

3.3.5. Airplanes and the Flight Control System (FCS) Network

Airplanes in the automated air traffic control system process the flight path commands received from the ground control stations using electronic units which will be referred to as flight control computers (FCCs) for the purposes of this paper. These FCCs are connected to the airplane's dual-redundant transceiver units over similarly redundant high-speed data bus links. The FCCs

receive the flight path commands and data from other on-board sensors which relay inertial and local air data in order to provide commands to fly the airplane. The gust suppression system in the Boeing 777, which senses aerodynamic impulses and provides commands to actuators to counteract the gust and smooth out the ride for passengers, is one example of an on-board sensor suite[7].

Control laws resident in the FCC software then compute commands to drive actuators, which in turn move the flight surfaces and allow the airplane to maneuver as commanded. The FCCs also annunciate status messages on the displays and back-drive the cockpit controls in order to enhance the pilot's situational awareness. They provide status and actual flight path data to the ground stations, and provide certain redundant SWIM system data to other airplanes. The flight control system receives the data and uses it only as a monitor to indicate potentially corrupted ground station data; in this case, an alert is provided to the pilot, who can then decide whether or not to disconnect from the AATCS and pilot the airplane manually. The level of redundancy of the flight control system's elements serves to mitigate many of the FMECA events associated with the flight control system, due to the flight critical nature of the FCS.

A top-level diagram of the airplane flight control system is shown in Figure 3 below.

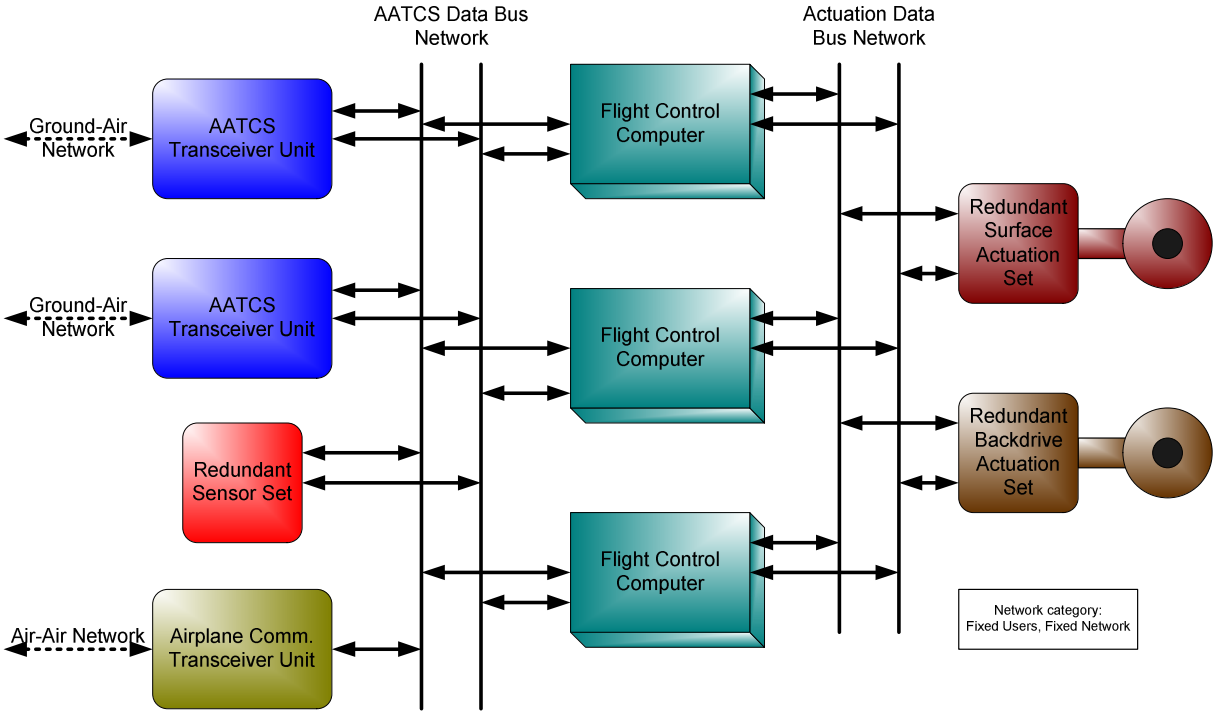

Figure 3 – Airplane Flight Control System Network (simplified)

3.3.6. Airplane to Airplane Network

The airplane-to-airplane network is a wireless network that connects the various airplanes within the vicinity of the airport. The data communicated on this network consists of redundant SWIM system data received from the ground stations that is rebroadcast to other airplanes for the purpose of monitoring data received from ground stations. Because this data is non-critical (i.e., not directly used in calculating commands to fly the airplane), only one transceiver channel is required to satisfy the corresponding FMECA events.

3.4. Network Organization and Architecture

The organization of the network must be defined before the commencement of any net-centric analysis to be performed on the system. Fundamental questions should be asked: what are the users, what constitutes the network, and in what manner do they connect to the network? How do they collaborate? What is the architecture of the nodes in the network?

3.4.1. AATCS Network Categorizations

There are five network categories[8] into which the AATCS could be classified, as shown in Table 1 below:

Table 1 – Description Of Network Types

Network Type	**Examples**
Fixed users of a fixed network	On-campus public computing center with dedicated workstations and server room
Fixed users of a mobile network	DirectTV subscribers
Mobile users of a fixed network	Cellular phone networks, wireless networks in hotels
Mobile users of a mobile network	Satellite phone networks
MANet: mobile ad-hoc network	Future Combat Systems, the Borg collective consciousness[9]

In order to determine the fundamental organization of the AATCS network, the concept of the users of the network must first be established. The users of a network are defined as those that use the services provided by the network, the most basic of which is communication with another user. Access of shared data and sharing of computational resources are other characteristics of a user on a network[10]. The airplanes are considered to be the users in the air-to-ground network because they receive flight path correction data from the ground control stations (a type of network service) as well as data from other airplanes – other users on the network. Airplanes are considered users in the air-to-air network as well; they use data received from other airplanes, even though in a limited capacity. Ground stations are users in the ground station network because of their usage of distributed services and shared data such

as the SWIM system data. Similarly, flight control computers are the users in the flight control system network for their usage of sensor data and flight commands.

The entire automated air traffic control system can thus be categorized as a hybrid of three network types. The ground control station network, which includes the links to external data sources, is considered to be a fixed network of fixed users due to the unmoving nature of the ground stations and the data sources. The air-to-ground network connecting the airplanes to the ground stations consists of mobile users of a fixed network; the airplanes are mobile and connect to and use the data provided by the ground stations. The air-to-air network is defined by the number of airplanes in the network at any given time, which can vary as airplanes connect and disconnect from the network. The air-to-air network is thus considered to be a mobile ad-hoc network or MANet. Finally, the airplane flight control system network is also characterized as a fixed network of fixed users – the components and subsystems resident on the airplane are not mobile.

Notional diagrams of the network organization of the AATCS have been previously shown in Figure 1, which depicts the ground to ground, ground to air, and airplane to airplane networks in a top-level diagram, Figure 2, which notionally depicts the ground control station network, and Figure 3, which shows a representative airplane flight control system network.

3.4.2. User Collaboration

User collaboration in a net-centric system demonstrates the concept that "the collaborative effort is *much more* than the sum of capabilities of individuals."[11] Wikipedia is one example of this, where individuals collectively update a shared knowledge database to produce a vast online encyclopedia that anyone can access. In the AATCS, the results of collaboration are not quite as dramatic as this example; however, the same effect is still achieved.

In the air-to-air network, the users (airplanes) collaborate by sharing SWIM system data received from the ground stations and using them as a monitoring check against the data directly received from the ground station. The airplanes cooperatively act as a group self-checking mechanism against the potential event of corrupted data being transmitted to one airplane. The FCCs can also be viewed as collaborators to a limited extent. Although not described in detail in this paper, each FCC provides certain data cross-channel to its redundant counterparts which serve to equalize commands transmitted to the airplane surface actuation subsystems. This equalization data is summed with the outgoing commands to reduce force fights and associated airplane structure fatigue by reducing the difference between position commands to different actuators driving the same surface.

In both AATCS examples, the result of the collaboration provides additional insight into what is transpiring at a higher 'level' of system operation and allows one or more parts of the subsystem to act collectively toward a more desirable outcome than they would have been able to attain as individuals.

3.4.3. Network Node Architecture

In order to properly analyze the system from a net-centric perspective, the architecture of the network protocol first must be established. One way to reduce the complexity of the network protocol is to model it in a layered architecture, where each layer is independent from the rest, performs specific services, and has flexible, well-defined interfaces with neighboring layers to allow updates to the layer while minimizing impact on those neighboring layers.[12] This paper will focus on the air-to-ground network as an example.

The network protocol for the air-to-ground network is designed to be as simple as practical because of the large bandwidth needed to transmit the entire data stream to each individual airplane from the ground stations. It is a four-layer protocol based on the TCP/IP 4-layer model. Each layer has its own services which, in aggregate, implement the functionality of the layer. Security services form the backbone of the protocol and are used in each layer. The security services are briefly mentioned here and are described in more detail later in the paper.

The “top” layer is considered to be the furthest abstracted from the physical characteristics of the network and is called the application layer. The services associated with this layer deal primarily with the data used by the airplane or the ground station. In the message that the ground stations transmit to the airplane, this would include translation algorithms on the airplane which disassembles the SWIM system data into constituent components – weather, flight commands, etc. – and formats them to be correctly used by the flight director and navigation algorithms in the flight control computers. The message that the airplanes transmit back to the ground station would include the flight vector data and airplane-level mode and status information. Data encryption is also performed in this layer.

The transport layer is immediately beneath the application layer and is responsible for the higher-level connection and transmission protocols in the AATCS, which is essentially same as the transmission control protocol (TCP). The AATCS TCP consists of the same three-way handshake to establish a connection and the receiver also provides acknowledgement that the message was received for non-repudiation. However, there is no re-transmission scheme due to the need for minimal latency in the flight commands. A missed or corrupted TCP packet means that the entire message is discarded, the event is logged (again, for non-repudiation), and the next transmission simply occurs with the freshest data available, as though the missed message had never happened. Authentication and the tunneling VPN function are also part of this layer.

The packet layer resides below the transport layer and performs many of the functions associated with further fragmentation of the packets, packet routing, packet re-assembly from fragments, and delivery. Each airplane that connects to the network is assigned a temporary address that is geographically (airport) dependent and valid only for the duration of the session. Packet-level CRC encoding and decoding also occurs at this layer.

The network layer is the “lowest” layer and encompasses many of the physical and hardware functions of the protocol. Data packet encapsulation is performed at this layer. The temporary network address assigned in the packet layer is decoded into a unique physical address to ensure that the packet is received by the correct airplane. Packets which do not have the correct physical address are ignored. Physical signaling characteristics, including start and stop frames, packet-level checksums, bit times, transmission voltage levels and signal recovery methodologies in the receiver, also occur at this layer.

Between each layer is an inter-partition boundary which serves to logically isolate and provide interfaces between neighboring layers. A header is pre-pended to outgoing data at each layer boundary which provides information needed by the peer layer on the receiving side such as port addressing, security and status information.

A functional representation of the network node architecture is shown in Figure 4.

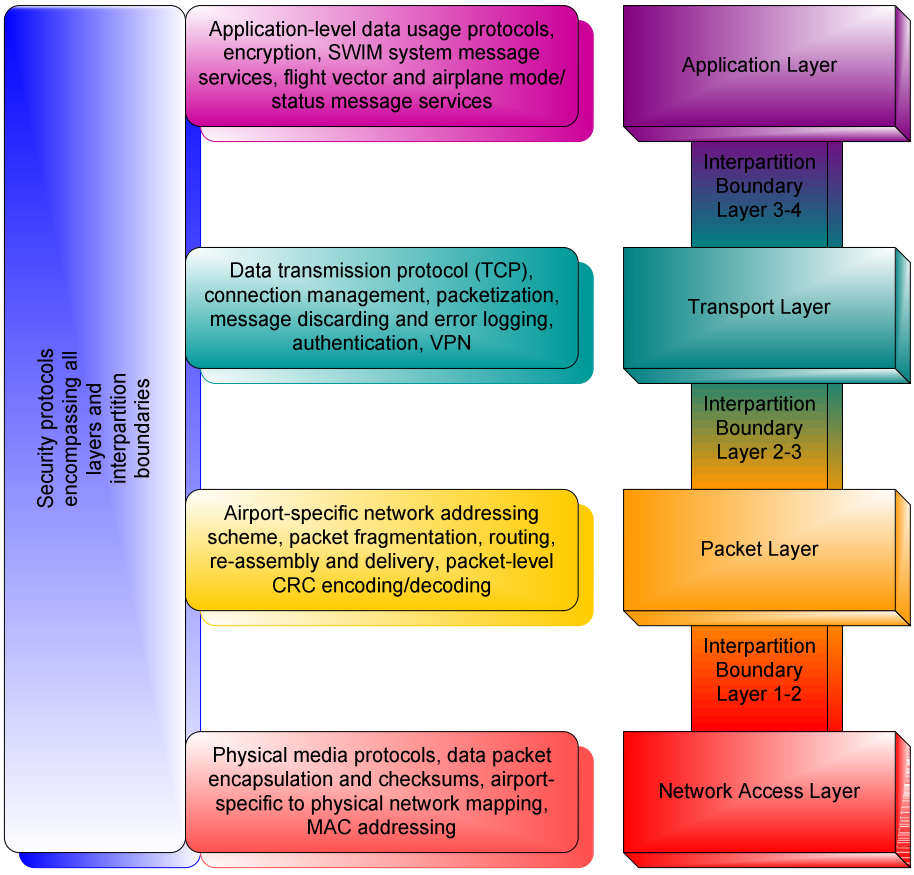

Figure 4 – Network Architecture Model of Air-To-Ground Network Node

4. AATCS Net-Centric System Analysis

The AATCS can now be analyzed from a net-centric perspective based on the description of the system's operation and architecture, which were defined in the previous section. System characteristics such as availability, failures, testability, integrity, security, robustness and the capability for future growth are examined in detail in this section.

4.1. Availability

The availability of a system is a "measure of the effectiveness of the fault tolerance of the system" by which "you can provide without a reasonable doubt that rigorous methods are used to provide measurable confidence that information systems are protected against attack or failure."[13] It is also defined as "the measure of the degree to which an item is in an operable and committable state when called for at an unknown (random) time."[14]

Any system may be simply defined as being in two states: available, when the system is performing its intended functions, and unavailable, when the system is not performing its intended functions. Although this may seem trivial, it is the availability and fault tolerance requirements of the system that can drive the system architecture's redundancy, testability and maintainability characteristics.

The availability requirements of the AATCS can be split into two areas, critical and non-critical functional operation of the system, as further described below.

4.1.1. Availability of Critical System Functions

The critical functions of the AATCS encompass the ground control station network, the air-to-ground network and the airplane flight control system network. These system elements are required for the AATCS to operate safely and successfully: the ground control stations receive data from the external data services and from the other ground stations, then transmit flight commands to the airplanes, which receive the data and process them and their sensor data into commands to actuators which fly the airplane. The reversionary mode of operation, direct pilot control, is also considered a critical function which takes control in the event that all communication with ground stations fail or when flight commands are detected to be untrustworthy.

As indicated above, availability also means that functions or operations are not functional or operational when undesired. For example, a maintenance test which initiates erroneously during flight has the potential to reduce system availability in a catastrophic manner, as would the loss of operation of all three FCCs.

The AATCS is required to operate with the chance that a hazard or failure which renders the system unavailable is extremely improbable. The FAA has determined this to be "so unlikely, not expected" in the system, equivalent to a 1e-9 probability of occurrence.[15] This is the number that modern airplanes are being designed to – for example, "before the 777 could carry

passengers, Boeing had to demonstrate that the likelihood of a power failure was less than one in a billion. That means it should never happen for as long as 777s are flying."[16] Although this number may also seem extremely improbable that it would ever be achieved, it can be demonstrated analytically that the availability requirement is met by system redundancy, component failure rates, and maintenance actions such as tests which run periodically. This analysis is sometimes done using a fault tree analysis, as shown in the simplified example in Figure 5 below. The fault tree consists of AND and OR gates which represent combinatorial probabilities of events – AND gates specify that all events must occur for the next higher-level event to occur, while OR gates specify that one or more events can occur for the next higher-level event to occur.

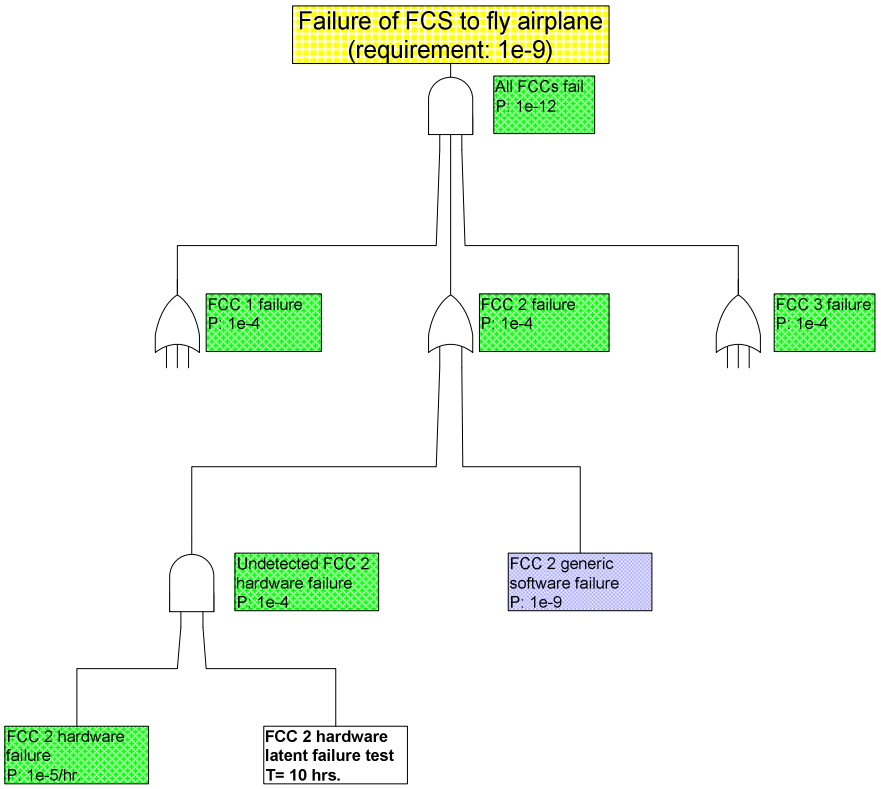

Figure 5 – FCS Fault Tree Analysis Example

The top-level event, shown in yellow, shows that the availability requirement of the FCS to operate without a catastrophic failure is one minus the failure rate, or 1 – (1e-9), which is 99.9999999% of the time, or "9 9's". The fault tree demonstrates this by showing that the probability of all three FCCs failing is 1e-12; the AND combination of the three FCCs' failure rates means that all three must fail in order to lose operation of the FCS. This is the first point in the fault tree that shows how the level of redundancy allows the availability requirement to be met.

The fault tree then shows that the FCC 2 failure can be caused by either an undetected hardware fault or a generic software fault. This is true for the other FCCs as well – only one FCC branch is shown for clarity. The OR condition means that the probabilities are summed and the hardware branch, which is five orders of magnitude greater than the software branch, dominates. The reason for the very low probability assigned to a generic software fault is due to adherence to design assurance processes such as DO-178B level A. This is described later in further detail.

The failure rate for the FCC hardware is 1e-5 per hour and because there is a maintenance test to check for latent failures in the hardware, the exposure time of the failure is limited to the test interval – 10 hours, in this case. Thus the combined probability of an undetected latent hardware failure is the multiplication of the two, or 1e-4. This shows how maintenance and test functionality can be driven by availability requirements.

4.1.2. Availability of Non-Critical System Functions

The non-critical functions of the AATCS include the airplane-to-airplane network, the data received from the network and the data transmitted to the network. The use of airplane data in the flight command processing algorithms is used only as a "sanity check" to monitor the commands received from the ground stations. Because the data is supplemental and is not directly used in the computation of actuation commands, the non-critical functions are not required to be available 100% of the time. The redundancy can thus be lower for non-critical systems and a fault tree analysis can again demonstrate that the availability requirement for the non-critical system function is met.

4.2. Failure Modes and Effects Criticality Analysis (FMECA)

"The principle of FMECA is to consider each mode of failure of every component of a system and to ascertain the effects on system operation of each failure mode in turn."[17]

The FMECA presented in Table 2 below is representative of the analysis performed on each unique type of unit or element within the AATCS. The unit chosen for this example is the flight control computer. Several types of failures, including hardware, software and network failures, are shown. The table lists the failure modes – each manner in which the FCC can fail, one per row. The effect of the unmitigated failure upon the system is described next. There are three columns which describe the probability of the failure occurring (low, medium or high); the severity of the unmitigated failure's effect upon the system (low, moderate or high); and the associated risk (low, moderate or high), which is the qualitative multiplication of the probability and the severity. Finally, the mitigation or containment of the failure and/or the action taken to prevent the failure is described. This serves to drive the design of the FCC and, in some cases, the design of system components which interface to the FCC. Common mode failures drive the usage of dissimilarity in the system. Redundancy and designing failure mitigation into the system increases the complexity of the unit and the system, which reduces the mean time before failure (MTBF), but increases the overall availability of the system. The MTBF is the inverse of the failure rate, which is used in the fault tree analysis example shown previously.

Italicized failure modes in the table below the dashed line indicate modes which are of low consequence or are otherwise highly improbable. The effects of such failures are not considered to be significant enough to warrant further analysis or mitigation.

Table 2 – Flight Control Computer (FCC) FMECA Table					
Failure Mode	**Effect(s)**	**Probability**	**Severity**	**Risk**	**Mitigation/Action**
Intermittent failure causes rapid changes in actuator command engagement state	Oscillatory command to actuator(s), possible flutter condition created	High	High	High	Monitoring, latching and complete disengagement after certain number of cycles, redundant FCCs
Infiltration of network	System is no longer trustworthy, commands overridden, airplanes could be remotely hijacked and ordered to fly off-course into buildings, mountains, etc.	High	High	High	Data encryption, authentication and related security measures, disengagement of automatic control by pilot
Power input short circuit	No power to FCC, possible damage	Medium	High	Moderate to High	Redundant FCCs, install fuses and transorbs
Wrong software loaded in FCC	FCC configuration and operation is different from other redundant units	Medium	High	Moderate to High	Procedural mitigation, System configuration check by FCCs, annunciation of problem to pilot, non-dispatch condition
Generic software or firmware failure occurs in FCC	One or more actuator commands from FCC do not match expected commands	Low	High	Moderate to High	Two dissimilar lanes in FCC architecture for self-checking

Table 2 – Flight Control Computer (FCC) FMECA Table					
Failure Mode	**Effect(s)**	**Probability**	**Severity**	**Risk**	**Mitigation/Action**
Command link to actuator open or shorted	No commands to actuator	Medium	High	Moderate to High	Actuator redundancy, redundant FCCs
Loss of all communication	Unable to receive flight control data	Low	High	Moderate to High	Pilot assumes control of airplane
Databus receiver fails open	Unable to receive data from ground control station	Low	High	Moderate	Redundant receiver
Data packet CRC failures (persistent)	Data in all packets have been corrupted and are untrustworthy	Low	High	Moderate	Redundant transmission channel and receiver
Power input open circuit	No power to FCC	Medium	Moderate	Moderate	Redundant FCCs
Data packet CRC failure (low rate of intermittence)	Data in packet has been corrupted and is untrustworthy	Low	Moderate	Low to Moderate	Redundant transmission channel and receiver
Physical damage to external connector during transportation of unit	*Cannot install flight control computer in airplane*	*Medium*	*Low*	*Low to Moderate*	*Replace damaged FCC with new FCC*
No data from airplane to airplane network	*Supplemental data for automatic control modes (monitoring of ground station data) not available*	*Low*	*Low*	*Low*	*Use ground station data for automatic control modes, direct pilot control for reversionary mode*

4.3. Fault Tolerance and Redundancy Management

Fault tolerance and redundancy requirements are driven primarily from availability requirements and failure modes, effects and criticality analyses. One definition of fault-tolerant airplane design is that it allows the system "to withstand single or multiple failures which results in either no loss of functionality or a known loss of functionality or reduced level of redundancy while maintaining the required level of safety."[18] Redundant elements in the system is one means of accomplishing this, typically in conjunction with the use of reversionary or degraded modes in the software which are entered when certain classes or types of failure occur. Figure 6 shows a combination of reliability block diagrams and mode transition diagrams, simplified for the purpose of explanation.

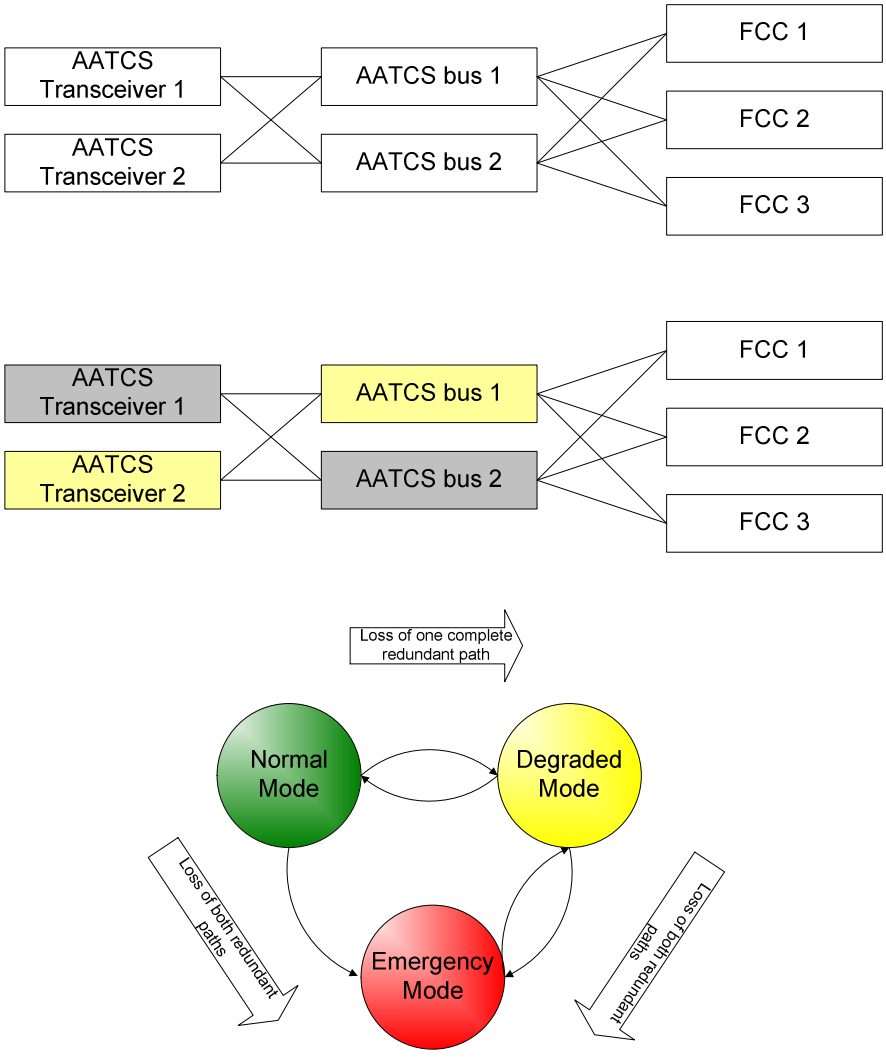

Figure 6 – Reliability Block Diagram and Mode Transition Diagram Examples

In this FCS-specific example, the reliability block diagrams show the dual-redundant natures of the AATCS transceivers and the AATCS data buses, which route the data to the triplex-redundant FCCs. The mode transition diagram depicts three operational modes: normal mode, degraded mode and emergency mode. The first reliability block diagram at the top shows the system in a fault-free state. This is analogous to normal mode operation. If a failure occurs in either transceiver or either data bus, the level of redundancy for the input data path to the FCCs is reduced from 2 to 1. At this point, the system transitions to the degraded mode of operation and, though not included in the diagram, a cockpit annunciation would result in order to warn the pilot of the loss of redundancy. This transition is a form of graceful degradation.

The second reliability block diagram shows a scenario where two elements have failed, shown with grey coloring. Should there be a failure of either redundant component, shown in yellow, complete loss of the function would result and the mode would transition from degraded to emergency. The backdrive controls would disengage from the cockpit controls and the pilot would take control manually.

It should be mentioned that while the FCS can degrade in any direction between normal mode and emergency mode, there is no capability of “up-moding” from an emergency mode directly back to a normal mode of operation. This is done as a safety precaution, to not allow the AATCS from regaining immediate control without passing through the degraded mode first. Although not explicitly stated, there would likely be some kind of pilot acknowledgement of fault recovery required in order to upmode back into the normal mode of operation.

Ground stations have similar levels of redundancy in their hardware and software to allow a path of graceful degradation to reversionary or emergency modes when failures occur. As with the airplane element in the AATCS, the ground station personnel assume direct control when the emergency mode becomes active.

4.4. Testability, Diagnostics and Maintainability

In addition to fault tolerance and redundancy, certain aspects of diagnostics, testing and maintenance have also been added to the system architecture. The maintainability of the system of the AATCS is driven by availability requirements of the system in a deployment and overall in-service context and not only while the system is active during flight. Built-in test capabilities are designed into the hardware and the software to enable testing of functions when the system is operating in a maintenance mode. The requirements for what additional circuitry and code result from the mitigation/action column in the FMECA. The usage of the maintenance functionality can then be applied to the fault tree analysis.

Invasive tests of the system are not initiated during operation in order to preserve the availability of the AATCS. This means that the airplane can run its tests only when parked at the gate or when removed from service. The ground stations coordinate with other ground stations in the area and take one set of redundant components off the networks before executing their tests, similar to server maintenance performed in a commercial computing environment.

Ground station personnel and pilots are similarly 'removed from service' periodically to re-certify their knowledge of AATCS emergency mode operation and procedures.

4.5. Integrity and Security

An automated system that takes care of the transportation of people on each of the approximately 87,000 flights per day in the U.S.[19] must have high integrity and security built into the architecture to guarantee the public's safety. This means not only system integrity but also data integrity, both of which are further described below.

4.5.1. System Integrity

System integrity is a measure of the level of trust that the system's behavior will function as expected during operation. For flight equipment, this assurance is broadly recognized as the outcome of successful certification by airworthiness authorities following the demonstrating compliance to validated requirements at every level of system, software and hardware development.

In the AATCS, all airplane hardware and software are designed to the highest level of rigor by providing compliance to well-known standards and processes. DO-178B ("Software Considerations in Airborne Systems and Equipment Certification") provides guidelines for the development of software in flight-critical systems[20]. DO-254 ("Design Assurance Guidance for Airborne Electronic Hardware") similarly provides guidelines for complex hardware development in flight-critical systems, with "complex" being defined as FPGAs, PLDs, ASICs and other programmable devices which exhibit software-like functionality[21]. Because the loss of functionality or erroneous operation within the AATCS is considered catastrophic, AATCS equipment is required to follow level A development for DO-178B and DO-254, which is the most stringent level of development. This requires a high degree of independence between the development and verification phases, to ensure that the design process in and of itself does not introduce errors which affect system operation.

In addition to following these processes, AATCS equipment is required to be compliant to the systems development process outlined in ARP 4754 ("Certification Considerations for Highly-Integrated or Complex Aircraft Systems"), which is an Aerospace Recommended Practice that provides guidance related to certification aspects from the overall system perspective, with a focus toward "ensuring that safety is adequately assured through the development of process and substantiating the safety of the implemented system"[22].

Other established internal and external standards and processes within the company or entity responsible for equipment development also provide system assurance, usually within the context of the aforementioned standards. Examples of these include the ISO-9000 series and CMMI (Capability Maturity Model, Integrated) types of certifications.

Ground stations also must provide systems which have a high degree of system integrity, to match the level of integrity of the airplane systems. As the metaphor states, a chain is only as strong as its weakest link, and airplane system integrity is of little use if the ground stations were to endure a generic fault which corrupts the flight commands. Therefore, ground station equipment is required to be developed and certified to the same standards as airborne equipment.

Similarly, there is a requirement for personnel to be 'certified', to ensure that each role in using the system will be performed properly. Ground station personnel and airplane pilots are required to go through background checks that are part of the normal employment application process for air traffic controllers and commercial pilots, respectively. Additional time must be spent training in standard and emergency procedures that are specific to the AATCS.

4.5.2. Data Integrity

Data integrity means ensuring the trustworthiness of the data on the network – "ensuring that data is unchanged from its source and has not been accidentally or maliciously modified, altered or destroyed."[23] The AATCS is a high-integrity system, due to its level of criticality in flight safety, and as such, data integrity is required at multiple levels in the network node architecture shown in Figure 4.

The method of analyzing data integrity involves examining it from one basic perspective: did the data arrive at the destination unmodified in any way? The source of a potential modification could be internal to the system – an intermittently stuck bit in the receiver, for example, where a certain bit position is always a 0 or a 1 for an indeterminate period of time. The source could also be external, such as electromagnetic interference. The end result for both is the same – corrupted data.

In the AATCS, checksums and cyclic redundancy checks (CRCs) are used by the network architecture to help guarantee that the data arrived unmodified. The computation of a CRC "resembles a long division operation in which the quotient is discarded and the remainder becomes the result…. An n-bit CRC…will detect any single error burst not longer than n bits, and will detect a fraction $1-2^{-n}$ of all longer error bursts."[24] Checksums are 16 bits in length and are calculated by adding all of the bits in the data packet and discarding the overflow or carry bits when the number "rolls over" from 65,535 (hexadecimal number 0xFFFF) back to 0. The checksum is appended to the packet and is recalculated on the receiving side. Errors introduced into the message would then show up as a mismatch between the receiver-calculated checksum and the transmitted checksum. In order for this method of checking to be defeated, the checksum would have to be modified along with the data such that the changed to the modified data would also be summed into the current checksum, which is an unlikely occurrence for random errors. Another way is for the total number of "bit flips" to be exactly 2^{16}, exemplifying the modulo-65,536 nature of the protection. Again, for random errors, it is highly unlikely that the number of bits affected will be exactly 2^{16} every time.

An example of how checksums are applied in the AATCS and used for error checking is shown in Figure 7.

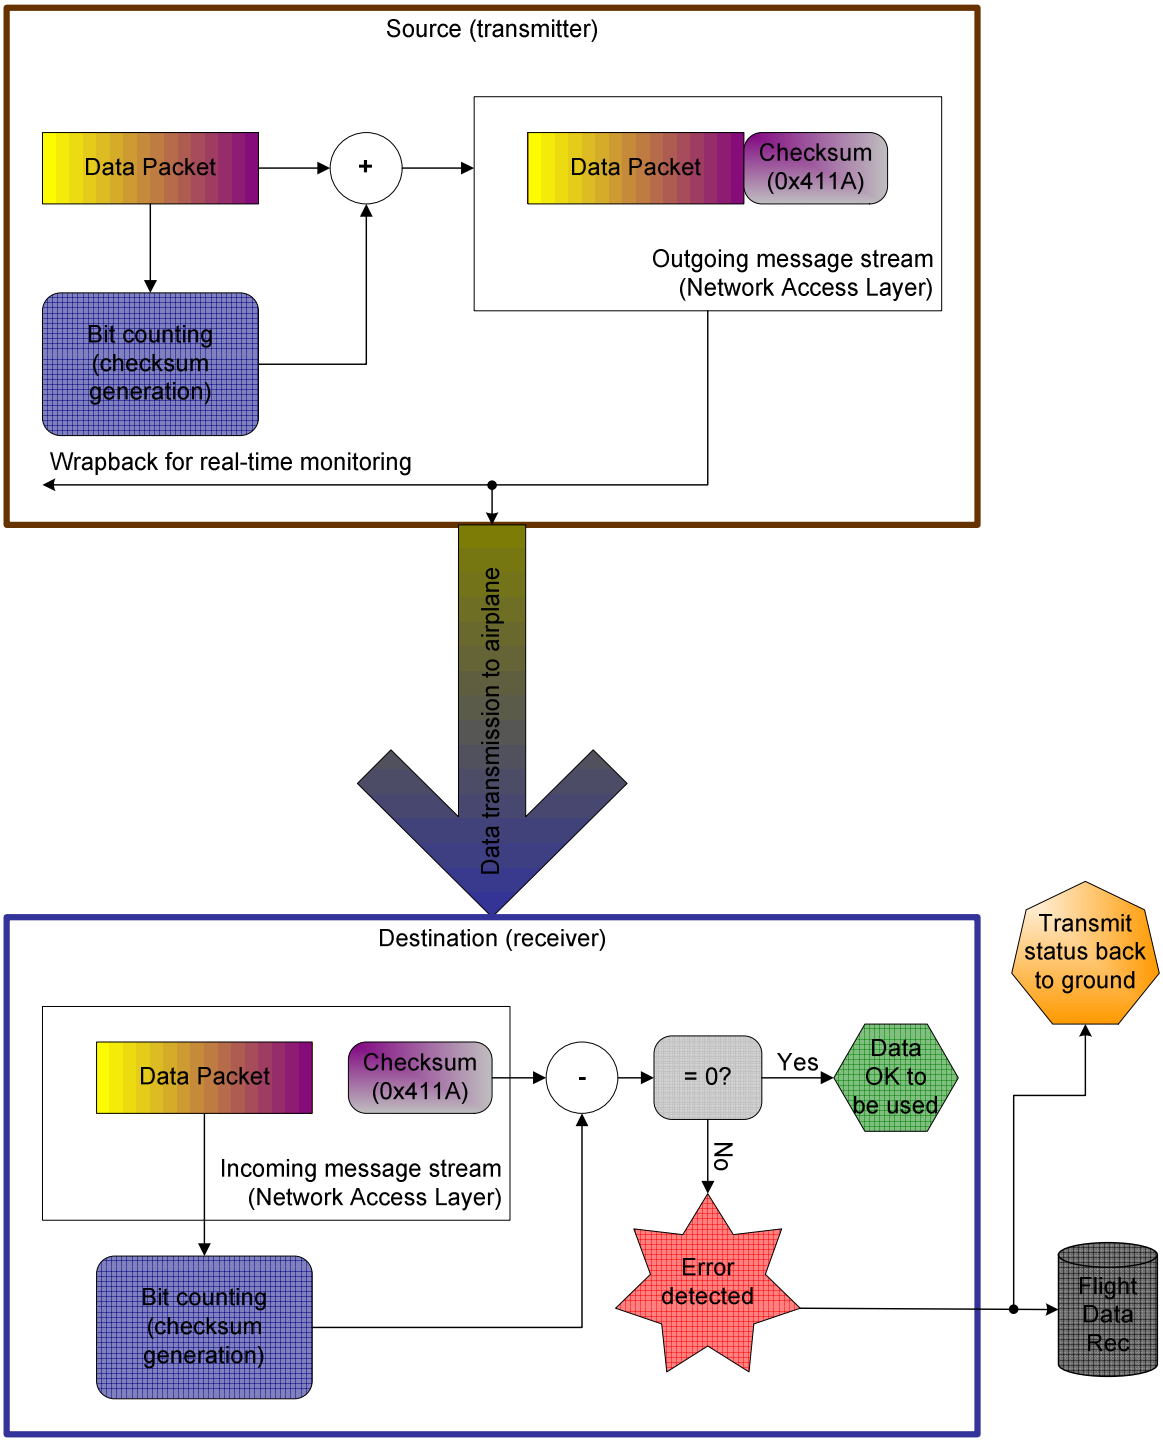

Figure 7 – Checksum Applied To AATCS Data Packets

In the example, the ground station calculates the checksum and appends it to the outgoing data packet. In addition to transmission through the air-to-ground network, the data packet is also wrapped back in order to verify that the data sent out was the data provided to the transmitter. After the data packet is received, the data's checksum is again calculated and verified against the transmitted checksum. If the two match, the data is usable and forwarded to the FCCs. If they do not match, an error is declared and the packet is discarded. Additionally, the error is time-stamped and the error count is incremented before both parameters are stored to the flight data recorders for non-repudiation. The status of the error is transmitted back to the ground station, where it is similarly stored.

Data integrity is part of the overall suite of functions which comprise information assurance. Authentication, confidentiality, non-repudiation and security management work in concert with data integrity to ensure that the information transmitted over the AATCS networks are safe and secure.

4.5.3. Authentication

Authentication is another part of information assurance and is defined as collective "security measures designed to establish the validity of a transmission, message, or originator, or a means of verifying an individual's authorization to receive specific categories of information."[25] In essence, how can the system guarantee that the user who is transmitting data is really the user and not someone else?

The security of the AATCS system is highly dependent on the trustworthiness of the users who connect and disconnect from the air to ground network. Authentication measures are required for airplanes to securely gain access to the network and communicate with ground stations and other airplanes upon entry to the ground station's local airspace. Identity of the airplane is proven in three stages. Before the airplane attempts to connect to the local airport's air to ground network, a list of airplanes expected to arrive or depart is generated based on filed flight plans and corrections en route through the SWIM system. Airplanes attempting to connect to the network who are not on the list cannot connect. The next stage involves the use of a pre-determined authentication code that is based on a pseudorandom sequence generator, similar to what is used by RSA's SecurID system[26], where the pilot chooses a password and files it with the flight plan, as well as entering it into a display panel on the airplane, which is transmitted to the FCCs, before departing. The pseudo-randomly generated code (PRC) is automatically transmitted by the FCCs rather than manually entered. Finally, individual packets are encrypted and transmitted by establishing a tunneling virtual private network (VPN) between the airplane and the ground station. Encryption also plays a role in confidentiality, as described below.

Figure 8 depicts stage 1 of the authentication measures in the AATCS air to ground network. Stages 2 and 3 are shown in the section describing confidentiality.

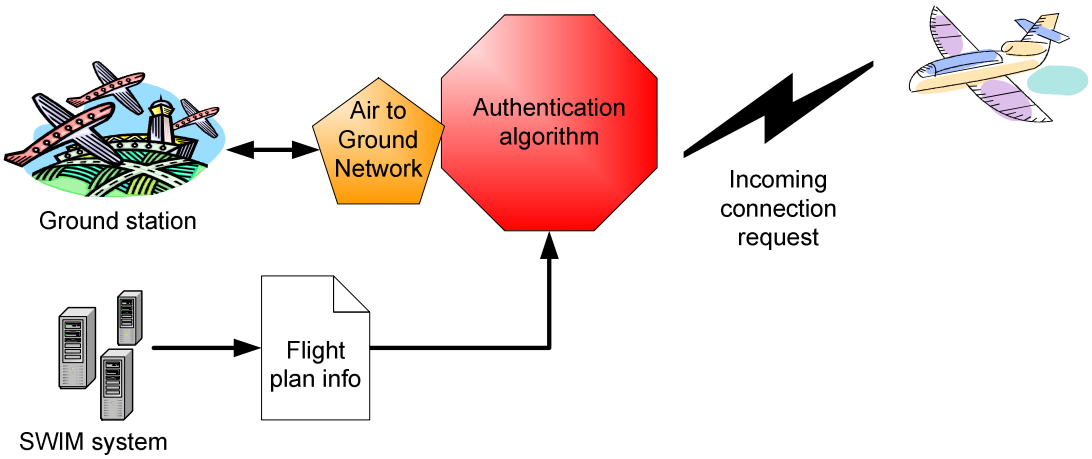

Figure 8 – AATCS Flight Plan Authentication Example

Ground stations also require authentication for personnel to access the monitoring capability of the AATCS data; current authentication measures used by air traffic controller is sufficient. Pilots are not required to provide authentication beyond what is normally provided to access to airplanes in an airport.

4.5.4. Confidentiality

Confidentiality provides "assurance that information is not disclosed to unauthorized persons, processes or devices (entities)" and "includes methods of keeping things secret."[27] Encryption plays a significant role in confidentiality (and information assurance in general) and is introduced in the previous section that describes authentication. Encryption and the use of the tunneling VPN secures the data. Also, the short lifespan of both the data in the packet (one packet is transmitted per second) and the encryption algorithm (the pseudorandom code changes every minute) makes it difficult and extremely unlikely for an outside entity to break the encryption and access the data in real-time.

Figure 9 depicts a combined example of the use of a virtual private network and encryption in the transmission of data from the ground station to the airplane. This example shows the pilot entering a password into the SWIM system, which is then used with the PRC by the security function to encrypt the packet and transmit the message over the VPN. On board the airplane, the AATCS transceiver unit decrypts the received message using the password that was entered by the pilot via the flight display before flight and the PRC. The decrypted message is transmitted to the FCC over the AATCS data bus. If the decryption fails, the packet is treated as never having been sent in the first place; i.e., the data is dropped. If enough packets fail, the system drops into either the degraded mode or the emergency mode, depending on the level of data path redundancy remaining in the system.

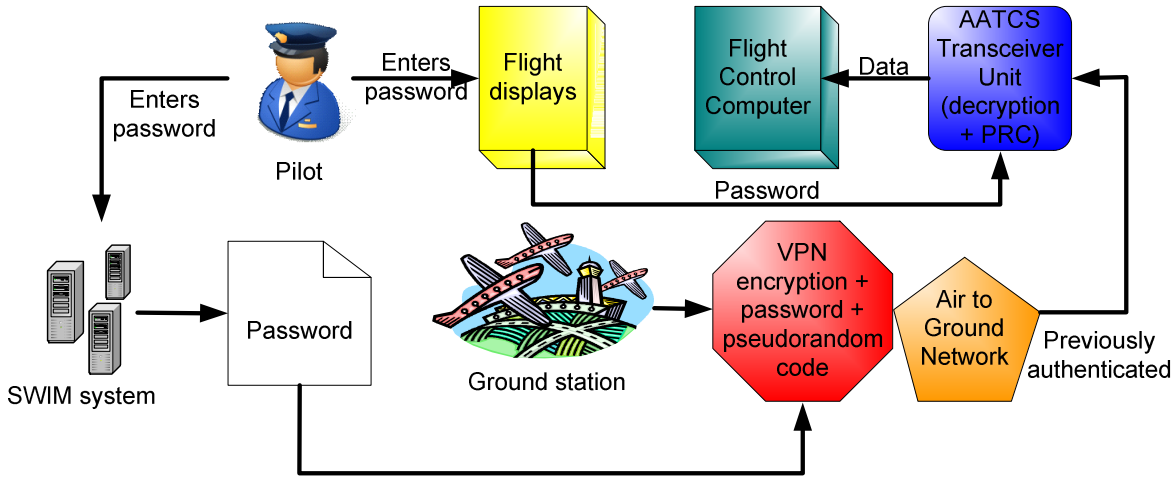

Figure 9 – AATCS Confidentiality Encryption Example

Access to the AATCS by humans is provided in a role-based control method where neither the pilot nor the ground station personnel can alter the data; they can only disconnect the system and assume control in an emergency situation.

4.5.5. Non-Repudiation Measures

Non-repudiation measures provide "assurance that the sender of data is provided with proof of delivery and the recipient is provided with proof of the sender's identity, so neither can later deny having processed the data."[28] Although this concept is more frequently used in commercial computer networks, there is a similar application in airplanes which is required by the FAA as part of Federal Aviation Regulation (FARs), §135.152 Flight Data Recorders. This section lists the parametric flight data to be stored in the flight data recorders "in the event of an accident or occurrence that requires the immediate notification of the National Transportation Safety Board…."[29] In the AATCS, similar requirements to store certain parameters in the data traffic received and transmitted by the airplane, as well as authenticated user access to the network and airplane modes would be levied by the FAA in order to be able to reconstruct events leading to accident or event under investigation and thus prove what had transpired. The data would be stored in the airplane as part of the information in the flight data recorders. The checksum errors logged is shown in Figure 7 as an example.

Dedicated data servers in the ground stations similarly collect parametric data received from airplanes in the network, SWIM system messages, and the authentication information from any user who attempts to access the network – successfully or not. Each entry is time-stamped to provide a history of which events occurred when. Data packets transmitted from the ground station to an airplane are wrapped back to the ground station computers and verified for integrity of the transmission, as shown in Figure 7. This data is also stored on the servers. Although not shown in the figure, which illustrates the ground to air path of the data, the transceivers also wrap back-transmitted data in order to verify what was sent.

Due to the enormity of data that is stored for non-repudiation, a daily backup to external memory (e.g., RAID arrays) is performed. Data is required to be kept on the server for a minimum of one week and the oldest data begins to be purged after every two weeks. Audits of stored data are performed on a daily basis by ground station personnel examining metric reports which summarize the data. Further analysis is possible by accessing server or backup records.

4.5.6. Security Management Method

The security management method of the AATCS requires a Defense-In-Depth Security strategy in order to ensure the highest level of protection in the system. While the Enclave Security strategy provides boundaries that separate a local computing environment from external access and provide means for controlled access, such as logging into a secured home wireless network, the ability of viruses, Trojan horses and other similar insider attacks to breach the enclave and affect the local computing environment makes it less suitable for the AATCS. "To fully protect the information during its lifetime, each component of the information processing system must have its own protection mechanisms. The building up, layering on and overlapping of security measures is called defense in depth."[30]

The AATCS uses several layers of security to protect the system extensively throughout rather than at just one boundary. Checksums and CRCs are used to protect the integrity of the data being transmitted. Authentication barriers prevent access to the air-to-ground network except for those airplanes whose filed flight plans or flight corrections en route show that the airplane attempting to connect is where it is supposed to be. Encryption prevents unauthorized access to the transmitted data as well as maintaining confidentiality of the data. Message wraparound checking and message and error store-downs form an audit trail of "who did what when" for non-repudiation.

While it may be relatively easier to defeat one security method, the combination of all security methods makes it very difficult and extremely unlikely that an outside entity would be able to intrude upon the network and the transmitted data. In fact, security must be of paramount concern in the design of a system such as this – no outside entity must ever possess the capability to breach the network and insert flight commands which could lead to the destruction of the airplane, the deaths of the passengers, and cause 'collateral damage' to others in the air or on the ground.

4.6. Future Considerations

Air travel is expected to grow significantly in the coming years; despite increasing fuel costs, the recent economic downturn and the demonstrated need for airlines to cut corners by offering "pay-as-you-go" services, in many cases. The enormous popularity of the 787 Dreamliner airplane, for example, is an indication of airlines' desire to continue to serve the flying public while realizing the fuel and maintenance savings that the airplane provides. There is still a need

for people to travel by air to distant locations – whether for business or for pleasure – and the AATCS architecture must be flexible to accommodate this and other (things) such as unintended uses of the system.

4.6.1. Scalability and Growth

As previously stated, the need to accommodate more airplanes connecting to the local AATCS network will be driven by the increase in air traffic. While this will be limited at major airports, this will not necessarily be the case at smaller airports, where smaller regional jets such as the Embraer 170 airplane and VLJ (very light jet) airplanes like the Eclipse 500 could be utilized. With the increase in passenger travel from larger and smaller airports, the ground stations which serve primarily as waypoints along flight routes will see the largest increase in network traffic. This will, in turn, create higher demands on the processing power and transmission bandwidth of the AATCS ground stations. Technological capabilities have increased exponentially over the last thirty years and similar increases in the future should mitigate many of the needed increases in scalability. However, an increase in the number of servers and hardwired connections could be a limiting factor. This suggests that, to plan for growth, the architecture of the AATCS ground stations should be distributed and allow for the addition of new equipment.

Airplanes will also see an increase in the number of other airplanes that they connect to, as well as other ground stations which will need to be built to accommodate air travel. It is anticipated that this will not require airplane systems to be as scalable as the ground control stations. The AATCS network architecture and protocols should be scalable and have spare bits, bandwidth, etc. designed into the system to allow for a reasonable maximum number of airplanes and ground stations within the networks. 30% - 50% growth or spares is typical.

In addition to the number of airplanes, the speed of faster airplanes could become a growth issue should supersonic flight become permissible in the future. Transmission protocols and processing power would need to accommodate flight command updates with sufficiently small lag time or propagation delay in the flight control system on the airplane and out to the actuators. The flexibility in the architecture to incorporate upgrades and technological improvements will be key in keeping the propagation delay low and is discussed further in the next section.

4.6.2. Robustness and Flexibility of Architecture

Robustness is a measurement of how well a system can operate despite abnormal or unexpected inputs or stimuli. Because of the high integrity of the AATCS system's design, a single fault or series of faults leading to catastrophic events is extremely remote. If something does happen to corrupt the system to the point where it is impossible to continue to be used, the reversionary mode of direct pilot control is always available. This presupposes the detectability of such faults in the system, which in turn defines the stringent monitoring and security requirements of the system.

The network architecture needs to be flexible enough to incorporate major changes as the system evolves from the perspective of user services and unexpected or unanticipated usage. While it is difficult to plan for 'unknown unknowns', some reasonable growth in air traffic can be anticipated and designed into the system, as described previously. Similarly, attempts can be made to envision changes to the overall architecture which do not involve scalability or growth.

The security infrastructure is perhaps the most likely area to change due to the ever-evolving methods of attack and attempts to access the network without proper authorization. Designing the network to accommodate the removal of an older protocol, for example, and update to a newer or different protocol is highly desirable.

The usage of the network is unlikely to change with vehicle type; as long as new commercial air transport vehicles carry AATCS equipment on board, there should be no difference to the ground station.

5. Summary/Conclusions

This paper has presented an overview of the automated air traffic control system (AATCS) and several net-centric analyses which illustrate various aspects of the architecture, from a variety of perspectives. The system is designed to supplement and utilize NextGen capabilities such as the System Wide Information Management (SWIM) system. By examining the system from a net-centric perspective, which entails an understanding of the enhanced ability to provide rapid access, processing and comprehension of data in a network environment, the complexity of the system can be characterized in a manner that can be better understood.

The analyses highlighted primary areas of interest: the system and network architectures; the availability of the system and how the system recovers from or otherwise mitigates failures; system and data integrity; and system security concepts such as authentication, confidentiality and non-repudiation. Although the design of the system was shown that, if properly designed, it can meet the stringent availability requirements typical of modern airplane design, there are still risks in attempting to certify and use such a system.

Perhaps the largest single design challenge involves the security of the data throughout the system. Information assurance is of utmost importance; the commands issued by the ground control stations as well as the SWIM data that is used in the calculation of flight commands must not be tampered with or otherwise corrupted by an outside entity. Although a pilot can manually disconnect the AATCS and fly the airplane in an emergency, it would be far worse for corrupted data to appear as valid data and cause a catastrophic loss of the airplane, the people on board, and others. Data security is certainly a high-profile challenge given the events of Sept. 11, 2001. The thought of such another scenario using the AATCS should be convincing enough of the need for the highest security levels in the system.

There are also technological requirements that the system design must ensure meeting, such as the bandwidth to handle the additional overhead of security functions on the data. Redundancy by definition introduces complexity and processors in the airplane flight control system must also be able to handle the processing of all of the data in real-time with minimal propagation delay. Future issues regarding growth and the capability to handle more airplanes demonstrate that architectural flexibility, robustness and scalability must be designed into the system to ensure that upgrades can be performed and done so cheaply with minimal impact, to allow continued deployment for decades.

Although the concept appears viable, the complexity and risk associated with the automatic control of airplanes, no matter how safe the system is, makes it difficult to convince the airworthiness authorities and the flying public that it is safe enough to put into operation. However, as airports become more crowded and the upper limit on the number of airplanes

that can arrive and depart is reached, it might become necessary to “think outside the box” and examine concepts once thought to be too unsafe.

6. Acknowledgements

The concept that this paper is based on originated from a discussion with Brad Betters of Tandel Systems, Bruce Vacey of Honeywell International and the author. The participants of the discussion have given the author verbal permission to further expand upon the original idea of an automated air traffic control system. The analyses and all unattributed material in this paper are solely the creation of the author.

7. References

[1] *"NASA & The Next Generation Air Transportation System (NextGen)"* (n.d.), retrieved October 5, 2008, from www.aeronautics.nasa.gov/docs/nextgen_whitepaper_06_26_07.pdf.

[2] "Free flight (air traffic control)" (July 18, 2008), retrieved October 21, 2008 from Wikipedia (Free Flight): http://en.wikipedia.org/wiki/Free_flight_(air_traffic_control)

[3] Cureton, K. (Aug. 26, 2008). *SAE 574 Lecture #1: Net-Centric Systems Architecting and Engineering, Fall 2008*. University of Southern California.

[4] Cureton, K. (Aug. 26, 2008). *SAE 574 Lecture #1: Net-Centric Systems Architecting and Engineering, Fall 2008*. University of Southern California.

[5] "Fact Sheet – System-Wide Information Management (SWIM)", (May 2, 2006), retrieved October 23, 2008 from Google (System-Wide Information Management): http://www.faa.gov/news/fact_sheets/news_story.cfm?newsId=7129

[6] "OEP Plan Reference Sheet NNEW" (June 19, 2007), retrieved October 23, 2008 from Google (NNEW): http://www.faa.gov/about/office_org/headquarters_offices/ato/publications/oep/version1/reference/nnew/

[7] Spitzer, C., (Ed.). (2001). *The Avionics Handbook*. CRC Press LLC.

[8] Cureton, K. (Sept. 2, 2008). *SAE 574 Lecture #2: Networked System Categories, Fall 2008*. University of Southern California.

[9] "Borg (Star Trek)" (Sept. 27, 2008), retrieved October 5, 2008 from Google (The Borg): http://en.wikipedia.org/wiki/Borg_(Star_Trek).

[10] Tanenbaum, A. (1988), *Computer Networks*, New Jersey: Prentice-Hall.

[11] Cureton, K. (Sept. 2, 2008). *SAE 574 Lecture #2: Networked System Categories, Fall 2008*. University of Southern California.

[12] Cureton, K. (Sept. 9, 2008). *SAE 574 Lecture #3: Essential Concepts Of Network Nodes, Fall 2008*. University of Southern California.

[13] Cureton, K. (Sept. 30, 2008). *SAE 574 Lecture #5a: Information Assurance For Net-Centric Systems, Fall 2008*. University of Southern California.

[14] Hines, J. (July 14, 2008). *Systems Engineering Theory and Practice, SAE 541, Summer 2008, Session 7*. University of Southern California.

[15] "FAA System Safety Handbook, Appendix B: Comparative Risk Assessment (CRA) Form" (Dec. 30, 2000), retrieved October 26, 2008 from Google (extremely improbable 1e-9): http://www.faa.gov/library/manuals/aviation/risk_management/ss_handbook/media/app_b_1200.PDF

[16] Levin, A., "777's power loss concerns aviation officials," (Feb. 26, 2008), retrieved Oct. 26, 2008 from Google (Boeing 777 failure one in a billion): http://www.usatoday.com/travel/flights/2008-02-26-777_N.htm

[17] O'Connor, P. (2002), *Practical Reliability Engineering*, West Sussex: John Wiley & Sons, Ltd.

[18] Spitzer, C., (Ed.). (2001). *The Avionics Handbook*. CRC Press LLC.

[19] "Air Traffic Control: By The Numbers" (n.d.), retrieved Oct. 18, 2008 from Google (number airline passengers per day): http://www.natca.org/mediacenter/bythenumbers.msp.

[20] "DO-178B", (Oct. 15, 2008), retrieved Oct. 19, 2008 from Google (DO-178B): http://en.wikipedia.org/wiki/DO-178B

[21] "DO-254", (Oct. 10, 2008), retrieved Oct. 19, 2008 from Wikipedia (DO-254): http://en.wikipedia.org/wiki/DO-254

[22] "SAE ARP 4754 "Certification Considerations for Highly-Integrated or Complex Aircraft Systems"" (Apr. 10, 1996), retrieved Oct. 19, 2008 from Google (ARP 4754): http://www.skybrary.aero/bookshelf/books/334.pdf

[23] Cureton, K. (Oct. 7, 2008). *SAE 574 Lecture #5b: Information Assurance For Net-Centric Systems, Fall 2008*. University of Southern California.

[24] "Cyclic redundancy check" (Oct. 22, 2008), retrieved October 27, 2008 from Wikipedia (CRC): http://en.wikipedia.org/wiki/Cyclic_redundancy_check

[25] Cureton, K. (Oct. 7, 2008). *SAE 574 Lecture #5b: Information Assurance For Net-Centric Systems, Fall 2008*. University of Southern California.

[26] "RSA SecurID" (2008), retrieved October 27, 2008 from Google (RSA): http://www.rsa.com/node.aspx?id=1156

[27] Cureton, K. (Oct. 7, 2008). *SAE 574 Lecture #5b: Information Assurance For Net-Centric Systems, Fall 2008*. University of Southern California.

[28] Cureton, K. (Oct. 7, 2008). *SAE 574 Lecture #5b: Information Assurance For Net-Centric Systems, Fall 2008*. University of Southern California.

[29] *FAR/AIM 2009* (2008). Newcastle, WA: Aviation Supplies & Academics, Inc.

[30] "Information Security" (Oct. 14, 2008), retrieved Oct. 20, 2008 from Wikipedia (defense-in-depth): http://en.wikipedia.org/wiki/Information_security#Defense_in_depth