Data Collection Systems

By

Don Burtis

July 2012

(Revision 2)

AeroVoodoo LLC ✈ www.aerovoodoo.com ✈ +1 (505) 715-2610

Table of Contents

List of Figures

List of Tables

Introduction

Even before the early days of the last millennium, data collection and analysis has been a principal means of discovering and proving advances in science. "No one in the history of civilization has shaped our understanding of science and natural philosophy more than the great Greek philosopher and scientist Aristotle (384-322 BC), who exerted a profound and pervasive influence for more than two thousand years"[1]. And about 1300 years before Aristotle, Egyptian doctors were using data collected from patients to determine the best treatment for their diseases[2]. While the scientific method wasn't well-documented until the mid-1200's[3], its informal use (and more importantly the use of collected data) was clearly used many centuries before.

In aviation, data collection has been a cornerstone of validating the design of an aircraft; the impact of adding a nifty accessory; maintaining airworthiness; and supporting recent SMS[4] initiatives that formalize processes for safe operation of aircraft. The cost of modern data collection has decreased dramatically as manual analog measuring devices have given way to automated digital data systems. New technology and consumer-driven initiatives such as cloud storage and on-demand availability of immense compute power have also contributed to reduced end-to-end data system costs.

As more aircraft data is collected, new opportunities for monetizing analyses of that data are emerging. The monetization takes two forms: the traditional reduction in maintenance costs as data supports lengthening service intervals (for example), and new possibilities for revenue from providing analyses of individual aircraft operation and performance compared to norms across multiple aircraft (just one example of how data may be monetized).

In this monograph, data collection systems from an aviation viewpoint will be explored. Sensors that convert a physical property into an electrical signal are briefly discussed, along with the methods used to transport the digital representations of the measured physical property. The storage, retrieval, and analysis of the digital data are discussed in the context of how they relate to the data system architecture.

Although much of the discussion is focused on aviation data that is created in-flight, the same principles apply to data that is collected during testing of aviation components on the ground.

1. Gary B. Ferngren in his 2002 book "Science and Religion: A Historical Introduction".
2. Approx. 1600 BC — An Egyptian medical textbook, the Edwin Smith papyrus, (circa 1600 BC), applies the following components: examination, diagnosis, treatment and prognosis, to the treatment of disease, which display parallels to basic empirical methodology. (Lloyd, G. E. R. "The development of empirical research", in his Magic, Reason and Experience: Studies in the Origin and Development of Greek Science).
3. In 1265, Roger Bacon described a scientific method, which he based on a repeating cycle of observation, hypothesis, experimentation, and the need for independent verification. He recorded the manner in which he conducted his experiments in precise detail so that others could reproduce and independently test his results.
4. SMS is "Safety Management System", a formalization and extension of existing safety-related initiatives and requirements, backed by the International Civil Aviation Organization (ICAO).

Collection of data always raises the specter[5] of trampling on individual's rights to privacy: it is easier to focus on the engineering decisions that have to be made concerning data collection and retrieval rather than some of the business and legal issues.

Who owns the data? Some will argue it belongs to whoever collects and retrieves it. But should that be the case? If I buy a car and use my built-in GPS navigation system while driving around, does the GPS data belong to me – or the car company? Same car, now I use a portable GPS navigator while driving around. Does the GPS data still belong to the car company? Does the answer change if the data is related only to engine performance? Issues like this must be considered when creating a system that some would argue allows others to spy on your actions.

Some car manufacturers finesse this issue by admitting the data belongs to the owner, not the maker of the car. But, the car makers assert that signing up for OnStar or its equivalent automatically grants them a license to use that data as they see fit.

What about liability? Suppose data collected during flights prove that an aircraft misbehaves under certain conditions? Does the aircraft company have an obligation to disclose that? Most would likely fix the issue, but without admitting any fault on their part.[6]

What about government-ordered release of that data? Can someone claim you violated their privacy by releasing their data to a government agency?

How long should the data safely stored in your data center be kept? Maximizing maintenance savings or revenue from monetization of data argues to keep the data forever. But does that present an opportunity for discovery and use of the data during future law suits?

If you share some of the data with a third party (engine vendor, for example), how much can and should you share? If you share location data, have you violated any privacy of the aircraft owner?

Think of a fleet of aircraft being managed by a fractional share operator who bills the individual aircraft owners for time they used the jointly owned aircraft based on your shared flight data. Who is responsible if you send incorrect data (inadvertently, of course) to the fractional share operator? Do you need a larger liability insurance policy to cover potential risks such as this?

Suppose you discover a new, unique way of analyzing or using the collected data that allows you to make millions of dollars selling the results to others. Do you or should you pay part of the profits back to the owners of that data? Or does your license to the data entitle you to keep all the profits?

5. Indeed: "specter" is defined in a popular on-line dictionary as "Something widely feared as a possible unpleasant or dangerous occurrence."
6. Remember the advice your auto insurance company gives for when you are in an accident? Never admit any liability or fault.

The automobile industry has started to sort out these issues as court cases start to hinge on data collected by the air bag controllers. Was that person really just going 40 miles an hour like he claimed? Some court cases have been won and some lost because of the answer supplied by the collected data.

AeroVoodoo cannot answer these legal and business questions for you. These examples just show a few of the issues that require any data collection and retrieval system to be subjected to your own careful legal and business review as it is being defined.

Data System Architecture

Aircraft data sensors exist at many locations within the airframe. Engine temperature sensors are of course located on the engines. Cabin temperature sensors are located where the crew and passengers are, and so on and so on. Each sensor must be examined to determine where its analog to digital conversion will occur, and where the digital data will be stored. In some cases, some of the sensors are part of a control system – such as the sensors that are used to control a jet engine. All modern jet engines are controlled by digital electronic engine controllers that assist the pilot to control the fuel flow based on current ambient conditions such as air pressure, air speed, air temperature, engine temperature, fuel temperature, and other parameters. These engine control devices necessarily have analog to digital converters for all of these measured parameters.

Similarly, the auto pilot is controlled in part[7] by a combination of air pressure, air speed, attitude, and location (from GPS). Since these sensors and the engine control sensors have analog to digital conversion in an existing box or boxes, there is an appealing argument to use the already converted data as digital inputs to the data storage device.

This architecture strategy works well for a data collection system that is being designed as part of a new aircraft; but becomes less attractive if the data collection system is being designed for after-market installation. In this latter case, it may be less expensive and require less certification justification if duplicate sensors are used.

After the data sensor locations and types are defined, each new sensor will require a tactical decision about where the analog to digital conversion will take place: individually at the sensor, or at a common intermediate point (where many sensors will share a single or multiple analog to digital converters) outside of the data collection box, or inside the data collection box. The tradeoff is the cost and weight of the intervening wires versus the cost and weight of multiple conversion boxes. There is no single answer because many variables must be considered. For temperature sensors, running each sensor to a central data collection box requires consideration of how the cold junction of the thermocouples are handled and the problem of using intermediate connectors with the same dissimilar metals as the thermocouple. Ouch. On the other hand, having an analog to

7. Flight Plan and Navigation Database are other ingredients for controlling the aircraft. Since these are known ahead of time, there is no need to record their current values.

digital converter at each temperature sensing point becomes expensive and may not work because of the thermal environment surrounding some temperature sensing points[8].

Finally, the location of the data collector itself must be defined. The method for data retrieval must be considered. If data retrieval is exclusively wireless, then the box may be located almost anywhere (again considering how to minimize wiring cost and weight). If the data is collected by removal of a memory storage element such as an SD Card, then the box must be located where a maintenance person or crew member can easily access the storage element. Another consideration for data collector location is crash survivability. This is discussed in the "Data Collection" section later in this monograph.

Certification requirements for most airborne electronics create an interesting perspective on how much number crunching should be done in the air versus on the ground. Changes to airborne software requires an appropriate level of oversight and proof to the regulatory agencies that the change will not reduce safety in the air in addition to proving no impact on the validity of the collected data itself. On the other hand, changes to software on the ground only need to prove no impact on the data validity, and the oversight is one's own company rather than a regulatory body. Based on this, one should strive to make the software on the aircraft as simple as possible and plan for all data analysis to occur in the ground-based data center.

Another important point of doing little to the data while still in the air is that once the data is on the ground, all data analysis may take place on a copy of the collected data, leaving the original collected data intact. That way, if any errors are discovered that affect the validity of the data, the ground-based software can be corrected and run against a new copy of the original data[9].

Should the collected data be encrypted while in the data collector? Should all data retrieved be encrypted while in transit between the aircraft and a ground-based data center? Those questions can only be answered after thinking about what advantage is gained in each case. In general AeroVoodoo believes encryption adds unnecessary complexity to the airborne software and should not be used. Of course your business goals may dictate a different answer[10].

A data collection system may be segmented into five activities, which are discussed in the following sections:

- Data Creation
- Data Collection
- Data Retrieval
- Data Storage
- Data Analysis

8. This of course is especially true for measurement of the engine temperature.
9. Even if the software bug is discovered years later!
10. Military data collection, for example, might require encryption protection to prevent location information to be used by a foreign entity.

Data Creation

Today's data is digital. Decades ago we gave up our motorized rolls of paper and moving-vane galvanometers with the little ink wells as our primary recording device for data. Now we convert all analog data into digital samples with contemporaneous time stamps so we know when the value measured existed at the sensor. Some scoff that we introduce an error by the very act of digitization. True enough, but advances in the analog to digital converters has decreased the error to infinitesimal amounts. And oh, by the way, those moving-vane galvanometers weren't error-free either!

Sensor Technology Discussion

A few kinds of sensors are digital (metal particle detectors, proximity sensors, rotation (as in flow meters), and simple micro switches come to mind). But temperature, pressure (stress), shock, accelerometers (attitude), magnetometer, and position sensors are analog. Each of these latter types of sensors requires a conversion circuit to sample their analog signal and convert it to a digital reading.

Temperature

Temperature measurements are generally done by allowing a thermal gradient to be applied to two dissimilar metals. One end of the metals are welded together and placed at the point to be measured[11]. The other ends of both metals are located in a common area (the "cold" junction) with a known temperature. The cold junction temperature may be known because of active regulation of its temperature or from an accurate measurement of its temperature. Most modern systems have sufficient computational power to use this latter scheme. Note that any intermediate connectors require pins and sockets to match the two dissimilar metals or the connector location becomes the cold junction.

Selection of the metals to be used depends upon the temperature range being measured.

Table 1 – Temperature Ranges for Common Thermocouple Types

Type	Temperature range °C (continuous)	Temperature range °C (short term)	Error Range One (°C)	Error Range Two (°C)
K (chromel/alumel)	0 to +1100	−180 to +1300	±1.5 from −40 °C to 375 °C ±0.004×T from 375 °C to 1000 °C	±2.5 from −40 °C to 333 °C ±0.0075×T from 333 °C to 1200 °C
J (iron–constantan)	0 to +750	−180 to +800	±1.5 from −40 °C to 375 °C ±0.004×T from 375 °C to 750 °C	±2.5 from −40 °C to 333 °C ±0.0075×T from 333 °C to 750 °C

11. This is called the "hot" junction.

Type	Temperature range °C (continuous)	Temperature range °C (short term)	Error Range One (°C)	Error Range Two (°C)
N (Nicrosil–Nisil)	0 to +1100	−270 to +1300	±1.5 from −40 °C to 375 °C ±0.004×T from 375 °C to 1000 °C	±2.5 from −40 °C to 333 °C ±0.0075×T from 333 °C to 1200 °C
R (platinum or a platinum–rhodium alloy for each conductor)	0 to +1600	−50 to +1700	±1.0 from 0 °C to 1100 °C ±[1 + 0.003×(T − 1100)] from 1100 °C to 1600 °C	±1.5 from 0 °C to 600 °C ±0.0025×T from 600 °C to 1600 °C
S (platinum or a platinum–rhodium alloy for each conductor)	0 to 1600	−50 to +1750	±1.0 from 0 °C to 1100 °C ±[1 + 0.003×(T − 1100)] from 1100 °C to 1600 °C	±1.5 from 0 °C to 600 °C ±0.0025×T from 600 °C to 1600 °C
B (platinum or a platinum–rhodium alloy for each conductor)	+200 to +1700	0 to +1820	n/a	±0.0025×T from 600 °C to 1700 °C
T (copper–constantan)	−185 to +300	−250 to +400	±0.5 from −40 °C to 125 °C ±0.004×T from 125 °C to 350 °C	±1.0 from −40 °C to 133 °C ±0.0075×T from 133 °C to 350 °C
E (chromel–constantan)	0 to +800	−40 to +900	±1.5 from −40 °C to 375 °C ±0.004×T from 375 °C to 800 °C	±2.5 from −40 °C to 333 °C ±0.0075×T from 333 °C to 900 °C
Chromel/ AuFe	−272 to +300	n/a	Reproducibility 0.2% of the voltage; each sensor needs individual calibration	

Pressure (Stress)

Simple strain gauges are resistors configured to change resistance when the substrate they are mounted on is deformed. The substrate is bonded to the material being measured. Thermal compensation may be done by simultaneously measuring the temperature of the strain gauge, or by configuring a second, co-located strain gauge that is not subjected to deformation of its substrate.

Shock

Shock is usually measured by an accelerometer. Both piezoelectric and piezoresistive (MEMS[12]) accelerometers can be used. Since the output signal levels are so low, it is common to have processing electrics as part of the

12. “MEMS” stands for Micro Electronic Mechanical System.

sensor. This is especially true of the MEMS type. Note that if a shock in all directions needs to be measured, then three accelerometers (x, y, z directions) must be used.

Accelerometer

Acceleration is also measured with accelerometers – in this case three are used (one each for x, y, z directions). This kind of sensor is now contained in virtually every smart phone sold, so the prices have come down rapidly and significantly during just the past few years. These sensors are almost always the MEMS type. Unless the architecture strategy is to use the data from the aircraft AHRS (Attitude and Heading Reference System), this sensor is usually located in the data collection box and its output data refers to the aircraft as a whole. (Of course, if the measurement is meant to apply to some part moving in reference to the main airframe, the tri-axis sensor must be located on that part.)

Magnetometer

Another common sensor sued in modern smart phones is the magnetic sensor that provides a "North" reference (subject to the usual declination errors). Assuming the duplicate sensor architecture strategy is being used, the only reason for remote mounting of the magnetometer (away from the data collection box) is to move it away from interfering electrical fields or metal shields that interfere with its proper operation.

Rotation

Typical uses for measuring rotation are engine rotation speed and liquid flow (such as fuel flow). Both measurements usually are done with a Hall sensor that pulses each time a predetermined part of the rotating item passes a fixed point. The elapsed time between pulses is measured and can be converted to rotational speed. Other kinds of sensors can also be used, but the concept is the same.

Position (Linear or Rotational)

Measuring the position of an object is a lot harder than one would think. Mechanical linkages with the object being measured must account for backlash (hysteresis) and the effect of sensor friction on the object being measured. One sensing technique uses a variable reluctance device known as a linear variable differential transformer (LVDT). The only moving part is an iron armature that changes the coupling of multiple windings as the armature is partially withdrawn from the windings. Local electronics is often provided to both excite the windings and to measure the minute changes in reluctance. However, if a hostile thermal environment exists where the measurement must be made, the electronics can be located away from the LVDT sensor.

Metal Particle Sensors

This is an example of a special purpose sensor that is engineered to indicate when metal particles are present in an engine's oil. Engine oil flows through the switch which relies on the presence of metal particles in the oil to compete the switch's circuit. The contacts are often magnetized to attract any existing particles, so that the particles are trapped in the body of the switch.

Proximity Sensors

Proximity switches are used to sense a binary position of an object, such as a door being closed or not closed, landing gear down or not down, parking brake set or not set, etc. Proximity switches on aircraft typically sense the nearness of a metal "target" by changes in its magnetic field. Proximity switches require power and have their own internal circuitry to determine the switch's state (open or closed).

Note for some applications, such as landing gear position, two proximity switches would be used for each landing gear. One would be used to tell when the gear is "down and locked" and a second one would tell when the same landing gear is "up and stowed".

Simple Microswitches

Some temporary test setups will use microswitches to indicate the state of an object ("off" or "on"). When using microswitches a parameter sometimes overlooked is the "bounce" time of the switch. Even with over-center constructions that force the switch to close (or open) rapidly once a mechanical threshold is passed, the switch will open and close several times ("bounce") when it is changing state. The duration of the bounce is proportional to the mass of the moving contact of the switch, and typically is several milliseconds in duration. Therefore circuitry monitoring microswitches (or any mechanical switch) must allow for the bounce phenomenon.

Another issue to understand with mechanical switches (especially microswitches) is contact contamination when used in "low energy" switching applications. Proper cleaning of the switch contacts occurs when a brief arc is created as the switch closes or opens. The primary concern for low energy switching is contact contamination. Since low energy loads do not arc or burn the contact surface, they also do not arc or burn off the contaminants that may reside on the contact surface. These contaminants may cause erratic switch resistance and can stop current flow. The lack of sufficient contact arching can be mitigated by using special contact plating with a noble metal such as gold, platinum, rhodium, etc. This special contact plating is sometimes labeled "dry circuit" contacts, or "low energy" contacts. Even with "dry circuit" contacts, the voltage switched should be at least 0.5 volts[13] to ensure the contacts remain free of surface contaminants.

13. Microswitch (Honeywell) General Technical Bulletin #13 - Low Energy Switching (001008-3-EN), July 1997, page 7.

Data Collection

Data Collector

Data generated by sensors or read from data busses occurs at a faster rate than needed for later accurate reconstruction of events. Figure 1 shows the incoming data on the left and the data retrieved for later analysis on the right. The data collector reads all of the data readings into a transient memory. Then when it is time to save a reading for later use, the most recent data reading is "sampled" (it is selected) for placement into the permanent memory. Data readings not used (not sampled) are discarded. Later, the saved data is retrieved.

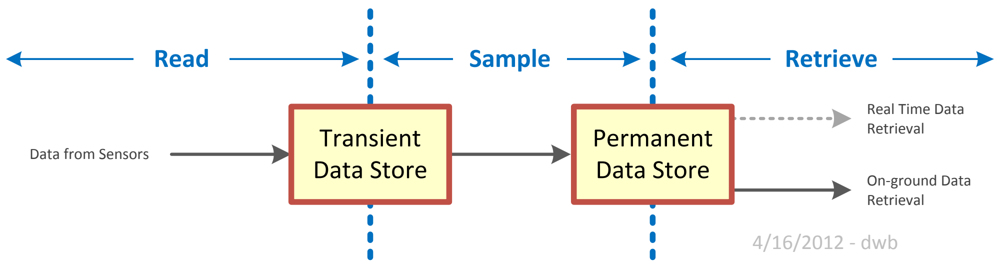

Figure 1: Definition of Data Collection Nomenclature

Breaking these activities down to their simplest form, Figure 2 shows where these activities occur.

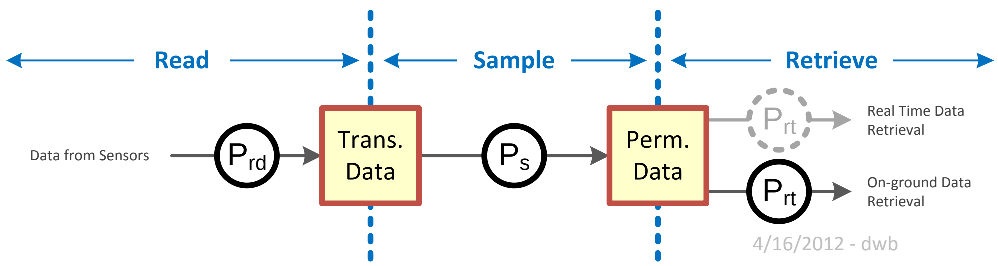

Figure 2: Data Collection Processes

The data collector may be thought of as a device that performs three functions:

- Read Data (P_{rd}).
- Sample and Store Data (P_s).
- Retrieve Data (P_{rt}). Note that an optional real time data retrieval process (for immediate ground communication) may also exist.

These processes run asynchronous to each other, so (for example), the sampling process never has to wait for the read process to finish processing a serial data input. Each of these processes is discussed later in this monograph.

If the data system also has real time data retrieval for all or part of the stored data, the data collector selects the data to be sent via SatCom (or other real time channel) for immediate use. This topic is covered in a later section "Data Retrieval".

Read Data Process

Every data reading is read into a transient data storage area. This memory holds one data reading for each sensor on the aircraft. When a new data reading has been received, it is time-stamped and overwrites the previous data reading. In this manner, the transient data storage always has a complete version of the most recent data reading for each sensor.

Sample and Store Data Process

The data collector regulates[14] how often the data from each sensor is saved for later use (that is, placed into "permanent" data storage). This process is called "sampling" the data of each sensor. The transient data storage contains the most recent data reading for each sensor. The sample and store process waits until it is time for a data reading of a sensor, then reads the most recent data reading from the transient data store and writes it to the permanent data store.

Retrieve Data Process(es)

The data collector may have one or more data retrieval processes. One process is dedicated to supplying all of the sampled data after the flight has completed. This may be performed by a ground maintenance person or a flight crew member. The process may involve removal and replacement of a flash memory card such as a compact flash card or a USB memory stick. Alternately, there may be wired or wireless access to the data collector that may be used as a data port for retrieving the data.

An optional retrieval port may be supplied to transmit important data immediately to the ground. An example of the kind of data that might be important enough is fleet dispatch data that might be used to reconfigure flight plans for an air taxi service due to cancellations or undesirable weather conditions. This dispatch data might include the current location of the aircraft and the amount of fuel on board. These facts could be used to provide the most efficient reconfiguration of the fleet's flight plans to accommodate unexpected conditions.

14. By means of its configuration which may be a writable data file or hard-coded in its program software.

As the cost of in-flight wireless data transmission is reduced by technology advances and mass use for other data purposes, a seductive argument emerges to send 100% of the collected data to the ground in real time during the flight. When this becomes practical, all the muss and fuss of attempts to preserve data during a crash go away.

Memory Types

Flash semiconductor memories have been getting smaller in size, larger in capacity, and lower in cost ever since they were invented around 1980. These devices have followed Moore's Law[15] for more than four decades. This historical behavior is projected to continue far into the future as well.

Moore's Law predicts that the number of transistors that can be placed on an inexpensive integrated circuit will double every two years. Stated another way, compute power will double every two years and memory prices will drop to half price every two years.

One unintended consequence is that semiconductor devices affected by Moore's Law do not stay in the marketplace very long. Memory chips that sold in the multi-millions 10 years ago have been obsolete and in some cases unobtainable for perhaps half of that time. Conversely, the life span of a typical aircraft today is measured in decades – so it is almost a guarantee that some of the semiconductor devices will no longer be available for repair of aviation grade compute devices. An important part of any avionics supplier's product support is a plan for obsolescence of components in their equipment.

On the other hand, Moore's Law has accurately predicted the price behavior of memory – especially flash memory which has been driven down in price as its use in mobile phones, cameras, and more recently personal computers has mushroomed. What was very expensive a few years ago is now available at a fraction of its former price. As an example, 16 Gigabyte USB flash memory devices are now selling for less than $12. Two years ago the same product sold for about $35. A few years before that, 4 Gigabyte USB flash memory devices sold for about $400. Since most retail sales channels have 100% or more markup from parts list to sales price, the actual flash memory chips for that USB device probably cost less than $4.

Flash memory cards are very hardy and withstand enormous G forces. The US Air Force has fished out flash memories from crashed aircraft and retrieved data from them[16]. However, they don't do so well in fires.

While the melting point of the silicon substrate is quite high at 2570 °F[17], the aluminum used to interconnect the active structures on the chip melts at a still-high temperature of ~1200 °F, the silicon dioxide chemical vapor

15. Forty five years ago Gordon Moore of Intel (while he was still with Fairchild Semiconductor) made this prediction in his paper "Cramming more components onto integrated circuits", Electronics, Volume 38, Number 8, April 19, 1965. When the historical data since 1970 is plotted, it closely follows the curve predicted by Moore's Law.
16. The crashes did not have associated fires. From a private conversation in 2006 with Chris Solon, team member of Air Force group who read the data.
17. IC manufacturing temperatures are from http://www.madehow.com/Volume-2/Integrated-Circuit.html

deposition is done at 752 °F which is still somewhat above the temperature of burning Jet A or Jet A-1 fuel[18]. The real culprit may be the solder used to attach the chips to the circuit board. Typical solders melt at ~365 °F[19].

Crash Survivability

FAA and EASA certified airborne flight data recorders must meet certain standards for crash survivability. These regulations all rely on the EUROCAE specification ED-112 "Minimum Operational Performance Specification for Crash Protected Airborne Recorder Systems" (March 2003).

This specification is aimed at preserving the complete functionality of any (and all) recording instruments on aircraft that require them. This includes all recording functions (Flight Data, Cockpit Voice, Images & CNS/ATM[20]).

The specification recognizes that only the memory element needs to survive the crash, and thus other electronic circuitry is excused from the testing.

Two fire tests are specified in ED-122 for the memory device (the rest of the electronics may be excused):

- **High Temperature Fire** (Section 2-4.2.4 on page 34). This test is exposure to 1100 °C (2012 °F) for a period of 60 minutes (1 hour)[21].
- **Low Temperature Fire** (Section 2-4.2.5 on page 36). This test is exposure to 260 °C (500 °F) for a period of 600 minutes (10 hours)[22].

Both of these tests are well above the melting temperature of solder. The only way for a flash memory device to survive these tests is by using special packaging that reflects and absorbs huge amounts of heat to delay its transmission to the sensitive memory element buried within. Some packaging tricks may use phase-changing materials such as paraffin wax that absorb a large amount of heat when melted.

The other tests in ED-112 are easily met with good, water-proof packaging of the memory element.

ED-112 specifies in great detail both the minimum data elements that must be recorded and the format they are recorded in. Many data collectors will not comply with ED-112 for this one reason alone. Therefore, identifying the data collector as a "Flight Data Recorder" or "FDR" must be avoided unless the data collector system is

18. Jet A (USA) and Jet A-1 (rest of world) have an open air burning temperature of 500 – 599 °F. (http://en.wikipedia.org/wiki/Jet_fuel).
19. Melting temperature of solder is from http://en.wikipedia.org/wiki/Solder.
20. CNS stands for Communications, Navigations, and Surveillance. ATM stands for Air Traffic Management.
21. This test is typically performed using a calibrated set of oil or gas burners. The memory is non-operational during the test, but has been warmed up to operational temperature prior to the test.
22. This test is typically performed using a calibrated thermal chamber. The memory is non-operational during the test, but has been warmed up to operational temperature prior to the test.

designed to meet all of the parts of ED-112 (and the associated FAA TSO-C124b "Flight Data Recorder Systems")[23].

Data Endianness and Bit Order

Much debate and misunderstandings have resulted from the two primary data storage methodologies used in computers. These methodologies are known as "Endianness" from an article by Danny Cohen in 1980. In an Intel White Paper on Endianness[24], "Little-Endian" is described as storing the least significant byte at the lowest memory address regardless if the data element width is 1, 2, or 4 bytes. "Big-Endian" stores the most significant byte at the lowest memory address. By this definition, Intel 80x86 processors are little-endian, and Motorola 68000 processors are big-endian.

Historically, asynchronous serial protocols were developed based on what became the RS-232 standard. They transmitted the least significant bit first. Because of the standardization of hardware to convert from byte to serial streams, virtually all communications protocols transmit the least significant bit first[25].

Data that is transmitted using Internet protocols has the most significant byte transmitted first, and is therefore considered big-endian. Since the data that will be stored by the data collector will likely be sent to a data center using data transmission techniques, the data outputs of the data collector should be big-endian.

However, the ARINC 429 protocol which was initially created before the concept of endianness was articulated, uses little-endian for both byte order and bit order – with the exception of the label (arguably the most significant byte in the 32 byte data packet) which is sent first, and with its most significant bit sent first[26].

Bit Numbering

Today bits within a byte (or word, etc.) are numbered from 0 to 7 (or 15, etc.) from right to left with the least significant bit numbered "0". However, since the genesis of the ARINC 429 specification predates the widespread use of computers in aviation, the bits were numbered from 1 to 32 from right to left, with the most significant bits of BNR data appearing on the left.

Wired Data Bus Technology Discussion

Historically, aviation data buses have been hard-wired and included parallel (one wire per bit) and serial (all bits share a wire) versions depending upon the transfer rates required. Technology advancements have made it

23. But to be useful, these non-FDRs still need to preserve the collected data through an aircraft crash and fire.
24. Intel White Paper "Endianness", November 15, 2004
25. One notable exception is the I2C protocol which transmits the most significant bit first.
26. Most hardware implementations of the ARINC 429 interface automatically reverse the bits of the label so the host sees the octal label number as if it had been sent LSB first.

practical to use serial data buses in nearly every instance of inter-box communications[27]. The media used for transmission is usually copper wires, but now fiber optics is being used for some data transmission, especially as the bandwidth requirements are increased to accommodate video data and high quality audio data now commonly used by in-flight entertainment equipment (IFE).

Some of the newer technologies have been borrowed from other industries. Engineers faced with $50 per node implementation costs for aviation-only data busses look longingly at their personal computer that uses a $5 chip to implement a bus that has 1000 times the speed when compared to the aviation-only data bus. We can thank those engineers for bringing CAN, Ethernet, TTP, and even the lowly RS-232 and RS-485 to aviation. Sometimes these buses need "enhancements" (native Ethernet is not deterministic), and sometimes the new technology doesn't catch on in aviation (Byteflight and FlexRay).

RS-232

Early data bus technology used distance and media type as differentiators. Distances up to tens or hundreds of feet were handled with relatively slow, large transition, single-ended binary signals[28]. Longer distances used variations of modulation schemes to create and decode a complex analog signal. A commonly used example of a simple short distance data bus is called RS-232[29], which defines transmit transitions of ±10 volts and receive thresholds of ±5 volts. At low enough bit rates, RS-232 can be reliably used to communicate over distances of hundreds of feet. No clock is transmitted between the transmitter and the receiver, but by phase-shifting the sampling clock (which is most often 16 times higher in frequency than the bit rate) at the beginning of the "start bit", the effects of wire capacitance is greatly reduced, although clock frequency skew between the transmitter and receiver may negate some of the reduction. The single-ended signaling method of RS-232 requires large voltage swings to account for sight differences that might occur in the local ground references between transmitter and receiver. RS-232 was intended to be used between two nodes. Multi-drop applications are possible, but not supported by all implementations of the interface circuits. RS-232 transmitters typically do not have any way to disable the transmitter, so multi-master networks are only possible with custom drivers.

Because of its extensive use in telephony and computers for the past 40+ years, the interface circuitry for RS-232 is available as integrated circuits at very low cost.

---

27. Intra-box communications still often uses parallel buses for speed reasons. A prime example is a computer main memory bus which is usually 64 bits wide in current personal computer designs.
28. Single-ended means the signal value at the receiver is created by changing the voltage on a single wire at the transmit end. The receiver compares the voltage between that wire and signal ground to pre-determined thresholds to identify which state the wire is in.
29. Originally defined in 1962 as RS-232 and most recently defined by *"TIA-232-F Interface Between Data Terminal Equipment and Data Circuit-Terminating Equipment Employing Serial Binary Data Interchange"*, issued in 1997. The original purpose of the specification was to connect mechanical teletype machines to modems for communications with main frame computers. While the official specification name has changed as the name of the responsible organization has changed, it is still most commonly known as RS-232. (The RS stood for "Recommended Specification").

RS-422/485

As data communications began to carry more information, a variant of RS-232 was created. It uses differential signaling with the data appearing on two wires. By responding to the voltage between the two wires instead of the voltage between the one wire used in RS-232, differences in the absolute voltage of the local grounds can be tolerated. No longer is the noise margin affected by voltage differences between the signal grounds of the two boxes. This new recommended standard, RS-422 also allows faster data rates – up to 10 Mbit/sec, or about 40 times faster than RS-232. Multiple receivers can be tied to a single transmitter, and RS-485 which is a variant of RS-422, allows party-line or multi-master operation by inclusion of a signal to disable each master's transmitter.

Neither RS-232 nor RS-422 specifies the content of the bits transmitted. They are merely specifications of the signaling method on the wire(s). However, the attractiveness of the low cost of implementation when compared to ARINC 429 has driven wider use of RS-232 and RS-422 in avionics in the past decade.

ARINC 429

At about the same time that RS-232 was being created, Aeronautical Radio, Inc. (ARINC) was tasked with creating a data bus to be used to communicate between avionics boxes supplied by different suppliers to the airlines. Initially released as one of the topologies defined in ARINC 419 released in 1966, the shielded, twisted pair interface was known as the DADS interface (Digital Air Data System). After release as ARINC 429 in 1978, it was known as the Mark 33 Digital Information Transfer System (DITS). A major improvement over RS-232 was the use of a differential signal where two wires are switched by the transmitter between +10 volts and -10 volts (measured between the wires). A third state where both lines are at the same voltage is also defined. The receiver identifies "high" as +6.5 to + 13 volts between the wire pair and "low" as -6.5 to -13 volts between the wire pair. If the voltage between the pair is from -2.5 to + 2.5 volts, the receiver identifies the state as "null". The signaling technique and the wide gaps between voltage thresholds provide excellent immunity from being affected by radiated electrical noise. ARINC 429 also defines the bit pattern, word length, simple error detection (word parity), and content meaning. The signaling speed is fixed at either 12.5 Kbits/sec or 100 Kbits/sec.

Special integrated circuits are available that contain the ARINC 429 interface circuitry. Because far fewer ARINC 429 circuits have been used compared to RS-232 and RS-422, the cost of the integrated circuits for ARINC 429 is relatively high – perhaps 100 times the cost of the integrated circuits for RS-232 or RS-422.

Even the ARINC 429 "high speed" option of 100 Kbits/sec is now too slow to handle all of the data traffic in modern avionics. As a result, ARINC 429 is gradually assuming a support role to connect to legacy equipment and low data rate peripherals. Although the data from a weather radar is handled by ARINC 453's 1 Mbit/sec bit rate, that too is insufficient for all the combined data handled by a modern multi-screen cockpit display system that includes the ability to display any data on any screen in case of individual screen failure.

ARINC 629

In the mid-1990's, Boeing decided to improve the data rate and robustness of ARINC 429. The result was ARINC 629, which runs at 2 Mbit/sec. It was used on Boeing's 777 aircraft, but failed to gain traction with the rest of the aviation industry because of (in part) its expensive custom interface hardware.

Ethernet

Ethernet was invented by Bob Metcalf when he worked at Xerox's Palo Alto Research Center (PARC) in the late 1960's. It has been enhanced over the years to include variations that transmit 1 Gigabit/sec data with low cost integrated circuits[30], using 4 unshielded twisted pairs. Slower variants (10 Mbit/sec) use a single twisted pair for each direction.

The avionics high speed data bus of choice that has emerged from major avionics companies such as Honeywell and Garmin is a variation of Ethernet[31].

While each of the major avionics companies will probably continue to use its proprietary version of Ethernet, there is a push to create a standard that will allow interconnection of avionics systems using the high bandwidth afforded by Ethernet and its variants.

AFDX

After the failure of ARINC 629 to gain traction with avionics suppliers because of the requirement for custom interface hardware, Airbus defined a new high speed data bus standard known as AFDX (Avionics Full-Duplex Switched Ethernet). AFDX uses commercial off-the-shelf (COTS) interface parts, thus the cost is kept low by the high volume of parts used by other industries. Now adopted by Boeing, AFDX is being officially defined by a subset of the ARINC 664 specification.

AFDX removes the transmission collision problem of half-duplex Ethernet by requiring separate transmit and receive paths. Other protocol changes provide guaranteed bandwidth and defined quality of service.

MIL-STD-1553

This bus was originally designed by the Society of Automotive Engineers (SAE) for military avionics in 1970[32]. The amount of data needed in military operations was quickly exceeding the capabilities of the existing data buses of the time.

30. Today's laptops for business usually have 1 Gigabit/sec Ethernet capability!
31. Each company has developed their own proprietary version of the software protocols that run on standard Ethernet interface circuits. These changes to the software protocols are required to meet the FAA regulations for simple analysis of data bus activity. Note that RS-422/RS-485 have been bypassed as choices by the major companies because of a lack of growth – these older serial interface protocols are near their peak bandwidth capability and are not the subject of on-going research and development to extend their usefulness.

The 1553 bus is specified to operate at 1 Mbit/sec, although at least one interface chip supplier talks about overclocking their chips to 2 Mbit/sec or even to 10 Mbit/sec[33].

The 1553 data bus uses balanced line differential signaling and may have up to 32 nodes on one bus. The specification details how to connect to the bus to reduce wavefront reflections[34] on the properly terminated bus.

Communications is half-duplex, and the signaling protocol is designed to eliminate any simultaneous transmission on the bus by two users.

The bus has been widely used in military and aerospace programs outside of the original airborne design target as well on aircraft. Some of the uses include satellites; space shuttle payloads; International Space Station; missiles; submarines; surface ships; tanks; howitzers; large transport aircraft including refuelers; bombers; tactical fighters; and helicopters. The Bay Area Rapid Transit (BART) in San Francisco uses the 1553 bus. And many foreign governments and international bodies have adopted the 1553 bus[35].

USB (Universal Serial Bus)

Compaq, Digital, IBM, Intel, Northern Telecom, and Microsoft created a new serial data bus standard in the mid 1990's to replace RS-232 and RS-422 on personal computers. About 10 years ago, the specification was updated and the speed increased from 12 Mbit/sec to 480 Mbit/sec. USB has started to show up in some avionics devices, notably the Entegra from Avidyne, where it is used for navigation data base updates (and general software upgrades as well). It was also used on the Eclipse 500 aircraft. The popularity of the iPod and iPhone has prompted the inclusion of USB ports on some automobiles and more recently in some entertainment products for aircraft.

CAN Bus (Controller Area Network)

The CAN bus was designed by Robert Bosch GmbH starting in 1983. Its design focus was the many systems in automobiles that needed to communicate with each other. The intent was (and still is) to allow many peripheral devices to communicate directly with each other without the need for a centralized host computer to monitor and regulate the flow of information. Because of the use in automobiles, major integrated circuit suppliers (including Intel and Philips) have created low cost interface circuits that implement the CAN protocol.

---

32. Work on the data bus design began two years earlier. The initial name for this data bus was "A2-K". In 1970, the final release of the specification occurred, and the document number was MIL-STD-1553 (USAF). The F-16 used the A revision (released in 1975) of the 1553 bus.
33. http://www.nationalhybrid.com:8080/index.php?option=com_content&task=view&id=36&Itemid=239
34. 1553 uses a shielded single twisted pair to transport the signals. The twisted pair uses termination resistors that match the impedance of the twisted pair. Any abrupt change in the impendence due to a local connection to the bus will create a reflection of the signal, much like occurs in an improperly terminated coax cable. The 1553 specification anticipates these disturbances and identifies requirements to reduce their effect on the signals.
35. NATO published STANAG 3838 AVS and the UK has published Def Stan 00-18 (Part 2). Both of these documents are close derivatives of MIL-STD-1553B.

The speed of the bus is 1 Mbit/sec up to 40 meters, but lower bit rates[36] can be used to extend the reach to more than 500 meters. A single pair of wires supports the data flow in both directions. The wire may be shielded twisted pair, unshielded twisted pair, or even two wires in a ribbon cable. Of course performance in noisy environments differs depending upon which wire regime is used. The signal encoding is done with the non-return to zero (NRZ) technique. NRZ provides the smallest number of voltage transitions for a given data rate. Since a clock is not transmitted between devices, a means must exist for maintaining bit synchronization when long sequences of the same bit value are transmitted. This means is called bit stuffing, where a single bit of the opposite state is inserted in the bit stream after 5 bits of the same value are sent. The receiver detects these stuffed bits and automatically deletes them from the data stream. Thus if six zeroes or six ones are seen, it is automatically an error.

CAN bus allows multiple masters to attempt to send a message at the same time. A clever prioritization scheme allows the higher priority message through, while signaling the lower priority device that its message must be repeated (usually a few bit times right after the higher priority message).

On the surface, the CAN bus appears to be quite similar to the 1553 bus. Both operate at 1 Mbit/sec. Both use differential signaling. Both are half-duplex and both support connection to many nodes.

There are four important differences. The CAN bus uses a different way of synchronizing and signaling which requires a technique called bit stuffing to compensate when too many bits of the same state are sent. 1553 uses Manchester encoding that guarantees a signal transition during every bit. The CAN bus implements a novel mechanism for arbitration that insures higher priority messages are not delayed. The CAN bus is also privately owned (by Robert Bosch GmbH) and royalties must be paid for its use[37]. And last, the high volume manufacture of automobiles justifies fierce pricing competition in CAN bus interfaces – which in turn reduces the price for implementing CAN bus.

ARINC 825 (Derived from CAN Bus)

Airbus and Boeing liked the CAN Bus, and decided to push for creation of a version that is aimed specifically towards use in aviation. The original CAN Bus leaves much to be determined by the individual system designer. That leads to subtle incompatibilities if equipment from multiple suppliers is connected using CAN Bus.

The Airlines Electronic Engineering Committee (AEEC) created the CAN Technical Working Group which included members from Airbus, Boeing, Rockwell Collins, and GE Aerospace. Their work resulted in the ARINC 825 specification ("General Standardization of Controller Area Network Bus Protocol for Airborne Use") which was published November 15, 2007. Other ARINC specifications using CAN Bus are emerging (ARINC 826 Data Load Standard for example) and are based on ARINC 825.

---

36. 10 Kbit/sec is the lowest bit rate. However all devices are required to support 20 Kbit/sec data rate.
37. Most times these royalties are collected from the maker of the chips that implement the protocol – such as microprocessors that have the CAN protocol built in.

The primary design goal of ARINC 825 was to integrate CAN Bus communications into the larger network architecture[38] used on the aircraft designed by Airbus and Boeing. The major enhancements of the CAN Bus spec is the addition of three more network protocol layers and a more flexible message priority scheme.

Data rates supported by ARINC 825 are 83.3 Kbit/sec, 125 Kbit/sec, 250 Kbit/sec, 500 Kbit/sec, and 1 Mbit/sec. The physical media (wires, connectors) is specified to follow airborne design guidelines for minimization of effects of HIRF and noise susceptibility (shielded twisted pair, metal connectors, etc.).

Byteflight and FlexRay

Byteflight was invented by BMW in the 1990's to create a high performance, deterministic data bus for automotive use. Semiconductor partners provided silicon that found its way into many BMWs over the next decade. But by 2005, BMW cars actually carried four other data busses in addition to a Byteflight bus, and one of them (Media Oriented Systems Transport, or "MOST") had more than four times the 10 Megabit per second capacity of Byteflight[39].

Although it is deterministic, has a high priority message structure, and inexpensive silicon was available, Byteflight was never embraced by the large commercial aviation companies. Its only aviation use was (and is) on smaller aircraft, in products produced by Avidyne.

During the first part of this century, BMW enlisted other automobile manufacturers (the FlexRay Consortium) to improve on Byteflight, especially in the area of fault tolerance and data capacity. The result was the introduction of FlexRay[40] in a production automobile in 2005. But by 2009, the Consortium disbanded and FlexRay is now being converted to an ISO standard. FlexRay has certain disadvantages such as lower operating voltage levels and asymmetry of the edges, which leads to problems in extending the network length. The full implementation of FlexRay is not used in aviation[41].

TTP (Time-Triggered Protocol)

The Time-Triggered Protocol data bus was developed in the early 1980's at the Vienna University of Technology. It has extensive health monitoring and fault isolation capabilities which caused it to be selected for Full Authority Digital Engine Control (FADEC) systems. It can accommodate both synchronous (25 Megabits per second) and asynchronous messages (5 Megabits per second)[42]. TTP has been used in military aircraft (F-16 fighter jet) and large commercial aircraft from both Airbus (A-380) and Boeing (787 Dreamliner).

38. Specific enhancements provide better support message transition from ARINC 825 to ARINC 664 and the reverse. This was accomplished by defining logical communications channels, individual station addressing capabilities, and one-to-many/peer-to-peer communications mechanisms.
39. MOST is now used in almost every major car brand in the world.
40. Byteflight is accommodated by the FlexRay protocols as a subset.
41. Avidyne still uses Byteflight in some of their products – it is not known if they have adopted FlexRay's fuller protocol.
42. The limitations are imposed by current silicon implementations of TTP.

One advantage TTP has over FlexRay is the ability to easily embed CAN bus messages within the TTP message format. This allows easier integration with multi-bus systems.

Wireless Data Transmission

Advances in wireless technologies, driven mostly by consumer telephony products such as the iPhone, have resulted in low-cost high-bandwidth data communication circuitry that is starting to be used as a data link between some sensors and the data collectors.

Although this section includes wireless technologies more suited for data retrieval than for data collection, they are all discussed here for easy comparison.

Wi-Fi

Wi-Fi ("Wireless Fidelity" – a term no longer used by the Wi-Fi Alliance, holder of the Wi-Fi trademark and logo rights) resulted from the 1985 release of the ISM Band for unlicensed use in the US. Subsequent derivative work by NCR and AT&T resulted in the precursor to the Wi-Fi specification (IEE 802.11). Some algorithms used in Wi-Fi were invented and patented by an Australian company CSIRO. The existence of these validated patents leads some observers to say that Wi-Fi was invented in Australia.

Wi-Fi development over the past decade has included multiple frequency bands (2.4 GHz, 5 GHz), higher transmission speeds (up to 600 Mbps), better security (WPA, WPA2), lower power (to compete with Bluetooth and ZigBee devices), and of course much lower cost.

The emergence of a low power Wi-Fi solution has the potential to severely limit the growth of ZigBee and the low power Bluetooth devices.

The combination of lower power and lower cost has allowed market entry of some sensor/data creation devices that convert and wirelessly transmit temperature readings from popular thermocouples. By using batteries to power the device, a truly wireless sensor is now practical and available.

Bluetooth

Bluetooth is a proprietary open wireless technology standard for exchanging data over short distances (using short-wavelength radio transmissions in the ISM band from 2.400–2.480 GHz) from fixed and mobile devices, creating personal area networks (PANs) with high levels of security. Created by telecoms vendor Ericsson in 1994, it was originally conceived as a wireless alternative to RS-232 data cables. It can connect up to seven devices, although the data is transmitted in an addressed mode (not broadcast).

The Bluetooth specification is managed by the Bluetooth Special Interest Group (SIG), which has more than 15,000 member companies in the areas of telecommunication, computing, networking, and consumer electronics. The SIG oversees the development of the specification, manages the qualification program, and protects the trademarks. To be marketed as a Bluetooth device, it must be qualified to standards defined by the SIG. A network of patents is required to implement the technology and the patents are only licensed to those qualifying devices; thus the protocol, while technically open, may be regarded as proprietary.

Several versions of Bluetooth exist. The original Bluetooth devices are capable of data rates of nearly 1 Mbps (0.724 Mbps). The later devices are capable of data rates up to 2.1 Mbps. Later versions of the Bluetooth specification (3.0 and up) refer to significantly higher data rates, but those are achieved by using a co-located, parallel Wi-Fi channel for the data transfers (and using the Bluetooth channel in a supervisory mode to control the data transfers).

ZigBee

ZigBee is a specification for a suite of high level communication protocols using small, low-power digital radios based on the IEEE 802.15.4 standard for personal area networks. Applications include wireless light switches, electrical meters with in-home-displays, and other consumer and industrial equipment that requires short-range wireless transfer of data at relatively low rates. The technology defined by the ZigBee specification is intended to be simpler and less expensive than other WPANs, such as Bluetooth. ZigBee is targeted at radio-frequency (RF) applications that require a low data rate, long battery life, and secure networking. ZigBee has a defined rate of 250 kbps best suited for periodic or intermittent data or a single signal transmission from a sensor or input device. ZigBee uses the 2.4 GHz frequency band.

As Wi-Fi enhancements reduce power consumption, ZigBee's advantages will disappear. As a potentially obsoleted technology, AeroVoodoo suggests ***not*** using ZigBee for data transfer.

Cellular Telephony Network Standards (LTE, EV-DO, WiMAX, etc.)

The data rates available from the US cellular carriers are increasing. As the carriers roll out their wireless upgrades (from "3G" to "4G") over the next few years, smart phones such as the iPhone will be able to view full-rate streaming video in high definition.

Cellular telephony networks already offer data transmission rates that favorably compete with other wireless networks such as Wi-Fi. The satellite services cannot compete on price or data transmission rate, but if you are out of cellular coverage, the satellite networks can provide coverage.

The US has 4 major cellular providers (Verizon, AT&T, Sprint, T-Mobile) and between them, probably a dozen different wireless technologies (some required to maintain backwards compatibility with their older phones). Since all of these technologies have been designed to transmit and deliver data packets compatible with the IP and UDP protocols used by the Internet, and competitive pressures keep the pricing very similar, the only interesting feature is their speed and coverage. Coverage changes often as the carriers add towers and additional antennae on existing towers. Check the coverage map each carrier has listed on their web site. The speeds

Iridium

Iridium[43] Communications Inc. (formerly Iridium Satellite LLC; NASDAQ: IRDM) is a company, based in McLean, VA, United States which operates the Iridium satellite constellation, a system of 66 active satellites used for worldwide voice and data communication from hand-held satellite phones and other transceiver units. The Iridium network is unique in that it covers the whole Earth, including poles, oceans and airways.

An important advantage of Iridium is the service is globally available (except at very high latitudes).

However, the service currently is very limited in data rate (only 2.4 Kbps) and pricey (about $2.50 for 10 Kbytes, compared to cellular data rates of about $3.00 for 200 Mbytes).

Iridium is planning a major system upgrade in the 2015-2017 time frames. It claims it will be offering data rates up to 1 Mbps after its new satellites are deployed.

Orbcomm

Orbcomm, Inc. is a company that offers global asset monitoring and messaging services from its constellation of 29 LEO (Low Earth Orbit) communications satellites orbiting at 481.5 miles (775 km). The protocol structure is oriented to very small (6 to 30 bytes) data payloads. The latency inherent in Orbcomm's network design prevents it from supporting certain safety-critical applications.

Orbcomm's spectrum of 137 to 150 MHz is at a much lower frequency than other commercial mobile satellite operators, which operate in the L-band (at 1.5-1.6GHz) and S-band (at 2.0-2.5GHz), and thus Orbcomm can operate with lower power, but also requires a much larger antenna than other networks such as Iridium or Globalstar. Most antennas are basic "whip" antennas that can be several feet long. Smaller, more compact designs are available but with performance trade-offs.

Orbcomm's most significant competitor is Iridium, which offers the lower latency and more capable Iridium SBD service which can send larger data packets with lower latency and a much smaller antenna.

The data transmission rate is double that of Iridium for the downlink, and the same as Iridium for the uplink. See Table 2 for details.

43. The company derives its name from the chemical element iridium. The number of satellites projected in the early stages of planning was 77, the atomic number of iridium, evoking the metaphor of 77 electrons orbiting the nucleus.

Globalstar

The Globalstar project was started in 1991 as a joint venture of Loral Corporation and Qualcomm. On March 24, 1994, the two sponsors announced formation of Globalstar LP, a limited partnership established in the U.S., with financial participation from eight other companies, including Alcatel, AirTouch, Deutsche Aerospace, Hyundai and Vodafone. At that time, the company predicted the system would launch in 1998, based on an investment of $1.8 billion. The Globalstar satellites are LEO types, and use S Band for downlink and L Band for uplink communications. The data transmission rate is 9.6 Kbps.

Currently Globalstar is experiencing degraded service (both voice and data) due to faster deterioration of its S Band satellite amplifiers than was forecast. They are launching a new series of satellites (three launches to date have placed 18 new satellites into orbit) to restore and enhance service, but details of when they will finally fix all of the current problems are not given on their web site.

Inmarsat

Inmarsat plc is a British satellite telecommunications company, offering global, mobile services. It provides telephony and data services to users worldwide, via portable or mobile terminals which communicate to ground stations through eleven geostationary telecommunications satellites. Inmarsat's network provides communications services to a range of governments, aid agencies, media outlets and businesses with a need to communicate in remote regions or where there is no reliable terrestrial network.

The company was originally founded in 1979 as the International Maritime Satellite Organization (Inmarsat), a not-for-profit international organization, set up at the behest of the International Maritime Organization (IMO), a UN body, for the purpose of establishing a satellite communications network for the maritime community. In 1999 it was converted to a for profit company and in 2005 listed on the London Stock Exchange.

Inmarsat offers a variety of data services – the Swift64 product is used extensively for aviation data. In addition to its commercial services, Inmarsat provides global maritime distress and safety services (GMDSS) to ships and aircraft at no charge, as a public service.

Swift 64 provides 64 Kbps service for each channel. Up to 4 channels may be grouped together to provide 256 Kbps service. Inmarsat operates in the L Band, so the antennae are larger than satellite services in the higher frequency bands (such as Iridium).

Wireless Data Transmission Rates

The data rates of the different wireless protocols are given in the table below. This information is current at the time of publication of this monograph, but is changing rapidly as the cost reduction of silicon circuits promotes ever more sophisticated information processing to be incorporated into the wireless nodes.

The technologies changing the fastest are the ones driven largely by consumer products, while the technologies changing the slowest are the satellite based one that require new satellites to be launched to effect major upgrades in their data transmission rates.

Table 2 – Data Transmission Rates of Various Wireless Technologies

Common Name	Governing Specification	Downstream (Mbit/s)	Upstream (Mbit/s)	Notes and Frequency Band
Wi-Fi (b)	IEEE 802.11b	11	11	2.4 GHz
Wi-Fi (g)	IEEE 802.11g	54	54	2.4 GHz
Wi-Fi (n)	IEEE 802.11n	600	600	2.4 GHz
Wi-Fi (a)	IEEE 802.11a	54	54	5 GHz
Bluetooth 1.1, Bluetooth 1.2	Bluetooth SIG	0.723	0.723	2.4 GHz
Bluetooth 2.0	Bluetooth SIG	2.1	2.1	2.4 GHz
ZigBee	IEEE 802.15.4 plus ZigBee Alliance Enhancements	0.25 (0.9 burst)	0.25 (0.9 burst)	2.4 GHz
Iridium	Iridium Communications, Inc.	0.0024	0.0024	1.616 – 1.6265 GHz (L Band)
Orbcomm, Inc.	Orbcomm, Inc.	0.0048	0.0024	137 – 150 MHz (VHF Band)
Globalstar	Globalstar	0.0096	0.0096	Downlink: 2.48 – 2.50 GHz (S Band) Uplink: 1.61 – 1.63 GHz (L Band)
Inmarsat plc (Swift 64)	Inmarsat plc	0.064 per chan	0.064 per chan	Downlink: 1.525 – 1.559 GHz (L Band) Uplink: 1.626 – 1.660 GHz (L Band) Up to 4 channels may be used to get faster data transmission rates up to 0.256 Mbps.
LTE	3rd Generation Partnership Project	100 Cat3 150 Cat4 300 Cat5	50 Cat3/4 75 Cat5	LTE-Advanced update expected to offer peak rates up to 1 Gbit/s fixed speeds and 100 Mb/s to mobile users. Frequencies used in the US are 1.71 – 1.755 GHz and 2.11 – 2.155 GHz
EV-DO Rel. 0 EV-DO Rev.A EV-DO Rev.B	3rd Generation Partnership Project 2	2.45 3.1 4.9xN	0.15 1.8 1.8xN	"N" is the number of 1.25 Mbps channels used. Frequencies used in the US are 1.71 – 1.755 GHz and 2.11 – 2.155 GHz
WiMAX release 1	IEEE 802.16	37 (10 MHz TDD)	17 (10 MHz TDD)	With 2x2 MIMO.
WiMAX release 1.5	IEEE 802.16	83 (20 MHz TDD) 141 (2x20 MHz FDD)	46 (20 MHz TDD) 138 (2x20 MHz FDD)	With 2x2 MIMO.Enhanced with 20Mhz channels in 802.16-2009

Common Name	Governing Specification	Downstream (Mbit/s)	Upstream (Mbit/s)	Notes and Frequency Band
WiMAX release 2	IEEE 802.16	2x2 MIMO 110 (20 MHz TDD) 183 (2x20 MHz FDD) 4x4 MIMO 219 (20 MHz TDD) 365 (2x20 MHz FDD)	2x2 MIMO 70 (20 MHz TDD) 188 (2x20 MHz FDD) 4x4 MIMO 140(20 MHz TDD) 376 (2x20 MHz FDD)	Also low mobility users[44] can aggregate multiple channels for data transmission rates up to 1Gbps.

Data Bus Summary and Recommendations

So what are the lessons in all of this?

1. Modern avionics has data bandwidth needs exceeding the capabilities of a single ARINC 429 data bus operating at "high speed". Some other high speed bus will be used by the avionics supplier to connect major compute resources. Today this is likely to be some proprietary variation of Ethernet operating at 10 Mbit/sec. In the future, there may be a migration to AFDX as airframers demand to interface to the central portions of the avionics through high speed data busses. This high speed connection is becoming important for diagnostic and testing purposes, as well as a convenient port to gather flight data[45] and perform nav data updates and device software updates.
2. ARINC 429 is still widely used to control and monitor radios and sensors (GPS/Air Data Computers/Smart Probes/AHRS, etc). ARINC 453 is used by the Honeywell RDR-2XXX weather radars. Some devices still use "low speed" ARINC 429.
3. RS-232/RS-422/RS-485 is now commonly seen in some peripheral boxes such as audio control centers. Most general purpose interface boxes from the avionics suppliers incorporate one or more of these serial interfaces. In general, using RS-422 is preferred to using RS-232 because of the increased noise rejection and higher tolerance to differences between signal ground voltage levels.

In summary, by piggy-backing onto commercially available circuitry manufactured in very large volume for the computer and information services industries, the cost of providing high bandwidth in the aviation environment has dropped from about 20 cents per Kbit/sec (ARINC 429) to about 0.03 cents per Kbit/sec (Ethernet) for just the node circuitry. The reduction in wire cost and weight has been almost as dramatic since the bandwidth for 10 Mbit/sec Ethernet is 100 times that of ARINC 429 using a similar cable. The dropping cost of high bandwidth data highways has matched our appetite for circulating more and more data to make operation and maintenance of the aircraft easier.

All of the wireless protocols discussed in this monograph provide a basic media link between devices that allows transmission of data packets in a common IP or UDP protocol. That allows changes in the underlying wireless

44. Low mobility users are ones that are relatively fixed in location so the cells don't have to do a handoff to another cell.
45. The desire to incorporate cockpit and exterior video capture is driving the bandwidth needs of flight data recorders ever higher. As memory prices continue their sharp decline, the cost of capturing and later recovering video images has become practical for the smallest business aircraft.

technology to not affect the structure of the data packets being transported. This is very important since the pace of development of the wireless technologies can proceed without requiring frequent major data structure changes.

Data Bus Error Detection and Data Integrity Assurance

Communications using data buses may be divided into two types. For the sake of this discussion, we'll call them non-robust and robust.

Non-robust communications assumes that errors seen by the receiving hardware are rare, and therefore inconsequential[46]. This approach is often used by older data bus designs, especially where the data is periodically repeated on the data bus. The concept is that a single wrong reading is quickly replaced by a good reading – often so quickly the wrong data is not seen by the pilot[47]. There is nothing wrong with this approach (widely used in ARINC 429, for example) but the consequences on stored data (such as a flight data recorder) must be understood for later accurate data retrieval[48]. There is no error detection requirement (and hence no retransmission mechanism needed) for non-robust communications, so the software is simpler.

By comparison, robust communications using data buses consist of two parts – the hardware circuitry and the software protocol. Both are important to providing an error-free means of transferring data from one location to another.

Rather than assuming the hardware works without errors[49], the robustness results from assuming the receiving hardware circuitry does not always receive what was sent by the sending hardware. Then the software makes up for any errors by detecting, then correcting any errors that have occurred.

There are two parts to what the software does – first detecting the error, then (in most cases), requesting a retransmission to replace the wrong data. The ability to request a transmission can only apply to directed communications, not broadcast communications[50].

The error detection means used includes some level of redundancy in the data. The simplest but least efficient method is to always send the data twice (for error detection by comparing the two copies) or three times (for error detection and error elimination by majority vote). Obviously this method causes a tremendous reduction in available bandwidth of the data bus. One early method of adding redundancy is the parity bit added to each

46. And not detected and/or corrected by any error detection mechanism.
47 Most ARINC 429 data formats have a "valid data" bit that is used by the receiver to suppress bad data. When used correctly, the pilot never sees bad data.
48. The program or user retrieving the stored data must understand that individual data points might be corrupted and should be discarded if they don't conform to the value or trend of surrounding data. This is acceptable for most applications – but should be understood.
49. Most of the time. The errors are not considered important enough to eliminate from the data stream.
50. Directed communications is where a message is sent to a single address. Broadcast communications is where a message is sent to multiple addresses. An example of directed communications is RS-422, where a single recipient is allowed. An example of broadcast communications is ARINC 429, where a single transmitter may be connected to as many as eight receivers.

byte[51]. Another early attempt to provide some redundancy is the checksum. Each data element in a message is added together and the sum is transmitted as an additional data element. While this method provides some ability to identify transmission errors and it eliminates most of the data duplication of the former method of simple data duplication, it is not very efficient at identifying the types of errors seen in typical data transmissions. A third method is the CRC or cyclic redundancy check.

CRC calculation may be implemented in hardware or software. The standard polynomials used in CRC-16 and CRC-32 have been implanted in readily available, tested software routines[52].

The effectiveness of a particular CRC method varies with message length. CRC-16 (which adds two bytes to every message) is generally used when the message lengths are less than 8192 bytes long. At message lengths of 8192 bytes, the probability that CRC-16 will discover a random error in the message is 50%. For messages longer than 8192 bytes, CRC-32 should be used (it adds four bytes to each message)[53].

Once an error in the data is detected, the error can be handled in two ways. If the data is repeated periodically, the particular data in error can be marked invalid (some data structures have a means for doing that). Or a request can be made to have the data re-transmitted. Most modern data protocols have a re-transmission capability, but it does create extra complexity when considering if a data protocol meets the requirements to be deterministic (predictable).

Data Retrieval

There are several ways data may be transferred from a data collector to a ground-based data center. Since a data collector may have one or more means of retrieving data while in the air and on the ground, the following sections just discuss some of the issues that might be encountered.

SatCom/Wireless Data Access

Wireless data access to the collected data may occur in-flight and/or on the ground after a flight is completed.

Wireless access during flight is limited to SatCom (Iridium, Orbcomm, Globalstar, Inmarsat) or one of the companies offering a hybrid satellite/ground station service aimed at Internet service on the airplane (currently

51. The parity bit was used to cause the number of bits set in each byte (including the parity bit) to be either odd ("odd parity") or even ("even parity"). Early computers from IBM included a parity bit for each memory location, but that practice fell out of favor as the memory chips became more reliable and users realized that most popular applications did not include any process for correcting the erroneous data byte. Parity in general has fallen out of favor because of the odd number of bits it causes. An exception is ARINC 429 which includes a parity bit in its 32 bit data structure.
52. However, it is vitally important to verify that the exact same polynomial is being used at both ends of a communications bus! This verification of error detection scheme and (if CRC is used) the exact polynomial used should be a documented task for each piece of equipment on any communications bus, especially if two of the pieces of equipment are supplied by different companies.
53. The message length where an even larger CRC is needed is about 537 Megabytes.

only Aircell through its Gogo brand). It is our opinion that eventually this latter category will develop faster as it is driven by the consumer desire of the passengers rather than by the commercial needs of the aircraft carriers.

Some of the challenges of using SatCom during flight are selection of what data to send, where to send it (your data center, or a special operations center just for in-flight data?), and what to do with it once it makes its way to the ground-based destination.

For ground-based data retrieval, an attractive solution (in addition to Wi-Fi) is wireless service through one of the cell phone carriers. While it is impossible to find one technology that provides world-wide service today; that will change in the future as consumer's cellular data demands require large investments in developing faster and more reliable technologies. Just as the US carriers are coalescing around LTE as their favorite technology, eventually the world will settle on software controlled radios that will accommodate all of the popular technologies.

Of course, Wi-Fi should be considered. But whose access point do you use? And who pays for the Internet access? Eclipse Aviation (in 2007) considered giving free access points to the FBOs, but even that would not guarantee coverage where it would be needed.

Some challenges are spotty coverage (a problem with both Wi-Fi and cellular), how long a data transfer might take, retry algorithms to resend data that didn't get through the first time, and how to manage security over Wi-Fi links (virtually all public Wi-Fi access points require some kind of manual sign-in the first time you visit).

Sneakernet

Sneakernet is any data transport mechanism that requires a human to carry a data storage device (such as a USB memory stick or a compact flash card) from one location to another[54]. While Sneakernet may have greater bandwidth than other transport means, it does rely on actions initiated by a human.

Often the Sneakernet solution requires removal of a compact flash card from the data collector and then inserting the compact flash card into a laptop which is connected to the Internet.

In addition to prodding a person to download the data manually, there is the issue of how do you get the data to your data center. If you are at your home base, this is not a problem. But if you happen to have landed at a remote site, then you might have to rely on someone else's Internet connection.

AeroVoodoo recommends Sneakernet as a backup to one of the wireless solutions.

54. Some might prefer the term "Guccinet".

Data Encryption

Should the aircraft data be encrypted? There are two possible places where data might be encrypted. One is data stored in the data collector, and the other is while the data is being transported from the aircraft to the data center. If used, encryption would need to be incorporated into the software/firmware of the data collector box on the aircraft and in any ground-based location that analyzes the data[55]. Unless a specific requirement exists (military aircraft, or information about ITARS-restricted technology, for example) AeroVoodoo recommends that encryption not be used.

Data Storage

Ok ok, you retrieved the data from the aircraft. Now what?

Based on decisions you made during the data system architecture design, you might want to accumulate the data in a single location; you might want to use the data to create wizzy graphs and charts for your aircraft owners; you might want to save the data for later analysis; you might want to distribute some of the data to others. In all cases, you probably want the data to reach a single location where additional data processing can occur. That single location is called the data center in this monograph.

Data Center

The data center for the collected data has these main functions:

- **Data Transport In.** The data center is the natural destination for all data retrieved from the data collector. Methods for receiving the data (including firewalls and malware protection) will be part of the data center.
- **Data Storage.** This is probably best implemented in a data base. MySQL for small amounts of data, Oracle for large amounts of data.
- **Web Hosting for Aircraft Owners (and others).** Since the data will reside in this data center, it makes sense to host the web site for aircraft owners (if one exists) from within this data center. Since this function represents an external data portal, appropriate firewall and malware protection will exist for this function). In addition to the aircraft owners, OEM engineering and third-party mainatenance organization have a valid need to access the aircraft data, and web hosting provides a convenient method to satisfy all of these needs.
- **Data Transport Out.** Data sent to others will be handled by the data center. Even for small amounts of collected data, automated scripts should exist to minimize required human intervention. Since

[55]. Recall the argument in the "Data System Architecture" section about "doing little to the data while still in the air".

this function also represents an external data portal, appropriate firewall and malware protection will exist for this function).

- **Maintenance Functions.** These functions include dat backup, data archive, media translation, disaster recovery, and software program updates.

The data center for the collected data may be a single computer tucked away in a closet at your company or may occupy racks and racks of space in an environmentally-controlled data center run by your company or by others.

Regardless of physical size of the data center, there are some operational and disaster recovery issues that should be addressed. These are presented as a list just to use as a reminder of issues to be addressed. The detailed solutions will depend upon how much data will be collected, how long the data will be stored, what user web site processing will be done, and how portions of the data will be distributed to others.

This list is oriented to larger data centers, but some of the items should be decided for the smallest of data centers (such as disaster recovery strategy).

- Data Center Disaster Recovery Strategy
- Strategy for Mitigating Denial of Building Access (Fire, Earthquake, Union Strikes, etc.)
- Strategy for Mitigating Power Outages
- Strategy for Mitigating Communication Outages
- Strategy for Mitigating Equipment Failures
- Notification Matrix at Company
- Recovery Times after Disasters Strike
- Backup and Archive Strategy for the Data Center
- Strategy and Frequency of Backups
- Archive Media
- Obsolete Media Strategy
- Location of Archive
- Archive Data Refresh
- Data Recovery Strategy
- Routine Maintenance Timing and Duration
- Uptime Availability Goal
- Physical Security
- Operator/Janitor Screening

The Cloud

A hot topic in consumer circles is “the Cloud”. One way to think of the Cloud and how it applies to aircraft data is the Cloud represents a third-party data center. The larger third-party data centers (famous examples are Google and Amazon) provide secure storage plus on-demand compute power. One advantage of using the third-party computer power is that you only pay for it while you use it.

Data Analyses

Once the collected aircraft data is safely on the ground, sophisticated software tools may be used to extract important information. These tools may be focused on improvements in maintenance diagnostics, improvements on training, or general analysis of flights as an added benefit to the aircraft user. Part of the benefits provided might include a web site where the aircraft user can display information about flights in his own aircraft. The site might also have a convenient means for uploading flight data (retrieved via Sneakernet) to the data center.

Another possibility is the forwarding of a portion of the data to satisfy an engine manufacturer's requirement for performance data to be submitted.

Some companies specialize in data analysis for FOQA (Flight Operational Quality Assurance) or for engine maintenance, etc. Of course use of third parties to analyze data may raise issues of data ownership and right to use the data for other purposes.

Table 3 below shows the many potential uses for data collected from an aircraft before, during, and after flight. This table assumes an always-active data collection system interfaced with a SatCom transceiver is present.

Table 3 – Potential Uses for Collected Aircraft Data

Purpose of Data Collection	User	Notes
Training Adjustments	Aircraft Manufacturer	Feedback to adjust overall training program or individual recurrent training
Engine Trending Data Collection	Engine Supplier	Likely Engine Product Support Agreement requirement
Flight Operational Quality Assurance (FOQA)	Aircraft Manufacturer and Fleet Operators	Includes performance improvements to reduce tire wear, for example. Accurate time-stamping (sub second granularity) is important.
Event Investigations/Reconstructions	Aircraft Manufacturer, FAA, and NTSB	Same analysis as FOQA. Note that the aircraft data may be a credible defense against unfounded claims of poor aircraft design.
Comprehensive Maintenance Billing	Aircraft Manufacturer	Could be basis for billing customers for Comprehensive Maintenance service
Fractional Operator Billing	Fractional Operators	Could be basis for automated billing of individual fractional share owners
Fleet Operator Dispatch and Flight Planning Adjustments	Fleet Operators	Adjust flight plans for delays. Know how much fuel is on board, for example.
Fleet Operator Fuel Management	Fleet Operators	Build statistics for fuel consumption on same routes

Purpose of Data Collection	User	Notes
LRU Fault Determination	Aircraft Manufacturer, Fleet Operators, and Engine Supplier	Service Centers
Real Time Critical Event Code (and CAS message) Monitoring	Aircraft Manufacturer	Initiate parts and technician positioning for up-coming AOG situations
Maintenance Efficiency Enhancements	Aircraft Manufacturer	Knowledge-based diagnostics
Customer Notification of Upcoming Scheduled Service Events	Aircraft Manufacturer	Might drive web-based information service
Stolen Aircraft Location	Aircraft Manufacturer, Law Enforcement	SatCom sends position information

Monetize the Data

As mentioned earlier in this monograph, the extreme reductions in cost for both airborne data collection and ground-based data analyses provide new opportunities for monetizing the data. Each organization will have differing needs, thus differing tactics for monetization of data – but a common theme is to be thinking of benefits to the user, the OEM, and third parties.

When thinking through the data collection architecture, the potential for monetization may drive some of the business and legal decisions discussed near the beginning of this paper.