Aircraft produce an increasing amount of data which can be analysed for predictive maintenance purposes to improve the operating efficiency of an aircraft. The solutions to offload this data, its usefulness and the economics to broadcast it are examined.

Economic solutions for offloading aircraft data

Airlines are now counting the number of messages they send over the Aircraft Communications Addressing and Reporting System (ACARS) via satcom to minimise their communication costs.

ACARS is a flightdeck message communication system, designed by ARINC and first deployed in 1978. ACARS can theoretically use any connectivity network available on the aircraft, so it can be physically enabled by communications networks such as high frequency (HF), and very high frequency (VHF) radios and Internet Protocol (I.P.) over satellite communications (satcom).

ACARS' ability to communicate reliably across a number of global communications networks means that an airline can transmit data to improve the safety, efficiency and reliability of its operation.

In addition, there is now an improved potential for aircraft to harvest data, yet it is debatable which is the most effective and economical way to offload it and exploit its usefulness.

SATCOM

Satcom is becoming more common, although there are some aircraft without it that broadcast over HF and VHF. Typically all long-range, trans-oceanic aircraft rely on ACARs to communicate.

Types of communications include basic air traffic control (ATC) instructions, operational and administrative data, and safety-related and aircraft performance data.

As ACARS is a character-based system, its messages are small, ranging from 10 to 200 bytes. For example, a 10-character message will equate to 10 bytes.

ACARS must have one redundant back-up channel as part of its reliability requirements. An aircraft will therefore need to have either a secondary HF channel or a satcom system in case the primary usage communication channel fails.

Compared to a high frequency data link (HFDL), satcom makes it possible to reliably send and receive ACARS messages over longer distances. Its higher frequency means higher data rates can be achieved.

Satellite frequencies range from one to 40 gigahertz (GHz). The L-band spectrum operates at 1-2 GHz, while the newer generation of Ku and Ka networks operate at 12-40 GHz.

All mobile phones operate across the L-band spectrum, partly because of the manageable size of the transmitters, receivers and antennas. Their wavelength of just a few centimetres makes it possible to build small L-band receivers that are efficient and suitable for aviation applications.

Inmarsat developed a service called Swift Broadband (SBB) over its L-band network, where data rates of 400 Kilobits per second (Kbps) are achievable.

L-band's robustness means the network is essential for many global safety services, but the reliability of the spectrum and its ageing, high-maintenance infrastructure mean that data transmission costs can be high.

Inmarsat Swift Broadband Safety (SBS) uses the L-band network. This offers a reasonable amount of bandwidth, although the total bandwidth used by operators is small because of its high usage cost.

Much of the safety data that is sent across SBS will consume only a percentage of the offered bandwidth, so much of it will be unused. In addition, there is an increasing trend to offload more aircraft data across the aircraft's existing in-flight entertainment (IFE) Ka-Ku channel.

The Ka-Ku channel is typically used for passenger in-flight connectivity. It is expensive to install, so not all aircraft have it.

A problem with the L-band network is the limited bandwidth available across the spectrum. This means there is an ongoing struggle for spectrum capacity between mobile phone operators and other users, such as marine operators and the military.

According to Hannes Griebel, Griebel Aerospace Consulting: "Some parts of the radio spectrum such as L-band are as desirable as a property in downtown LA or the City of London. High cost means that many transmissions that are not safety-related are gradually moving away from the L-Band network and transmitted across the less expensive Ka-Ku networks."

Inmarsat made a huge investment in launching its Global Xpress (GX) spacecraft, which operates on the Ka spectrum with high throughput capabilities. Low price per kilobit (Kbit) rates are achieved because of the wide bandwidth that is available.

"Because the data rate is so much higher, it is possible to sell more bandwidth and more data to more consumers at a much lower price per-bit. Furthermore, not all Satcom communications need the robustness afforded to the L-band network, meaning the Ka-Ku spectrum is a more economic alternative," says Griebel.

Rain fade

Certain parts of the radio spectrum are more prone to suffering from an effect called moisture attenuation. The L-band's radio spectrum is less susceptible to it, so optimum data rates are achievable, even with a dense amount of cloud cover.

Transmissions across the Ka-Ku network are more susceptible to rain-fade and are therefore less reliable in inclement weather.

Even beneath a haze of Cirrus clouds, it is believed for the most part that the Ka-

ViaSat-2 is the second Ka-band satellite launched by ViaSat after ViaSat-1. The $600 million satellite was launched by Arianespace from the Guiana Space Centre in French Guiana. ViaSat-2 uses Ka-Band frequencies with a maximum throughput of 300Gbits/s.

Ku network is reliable enough for a host of aviation applications.

According to Murray Skelton, Teledyne's senior director of aircraft solutions: "Moisture attenuation must be extremely high before it starts to affect transmission. There is always a possibility that rain fade can disrupt the system, but not to an extent to where it will halt the broadcast. Technological advancements have made the system inherently more reliable."

Safety-related data must be sent across an approved network such as Inmarsat's or Iridium's L-band services. Once a message has been sent by ACARS over an I.P. node to the standard ACARS host, the host is currently unable to identify the type of network used. Furthermore, there are discrepancies between operators on what data is classed as being safety related, such as weather-related data.

Economics

Only very few operators like ViaSat, Iridium and Inmarsat offer mobile satellite services across spacecraft they own and operate themselves.

Meaning many resellers rent capacity from other networks and from third parties.

There are arguments about what is more economic for the client. Is it best to operate your own spacecraft, or lease transponder space at a mark-up? Or to reduce overhead costs by not having the expense of procuring, launching and operating expensive spacecraft?

Air-to-ground networks are based on 4G cellular Long Term Evolution (LTE) technology that uses ground stations to the aircraft flying overhead.

The air-to-ground network is the least expensive in terms of data rates because the ground-based infrastructure system can use existing cell towers to tie into an existing data network.

The aircraft's antenna is small and draws lower power than typical satcom systems.

However, coverage is limited to the range of cell towers and therefore cannot stretch across larger bodies of water or mountain ranges.

The widespread use of data-hungry smartphones means that there is a massive data transmission infrastructure in place that aircraft air-to-ground (ATG) networks can access.

ACARS over I.P.

The average cost of ACARS data varies between airlines.

With ACARS data transmitted over traditional satcom and VHF services, the cost can easily amount to $2 million per year for operators with a fleet of 90 aircraft. For larger airlines it is possible to reduce the usage cost per aircraft.

A recent large customer with Teledyne, who successfully implemented the Teledyne ACARS over I.P (AoIP) GroundLink DataLink solution, used AoIP to justify the investment in terms of kit and capability on the basis that its aircraft condition monitoring system (ACMS) and aircraft health monitoring (AHM) data traffic would effectively become free.

The Teledyne AoIP solution leverages both existing broadband connectivity onboard the aircraft to transmit ACARS data during flight, and cellular networks to transfer non-critical messages when the aircraft is on the ground, at a fraction of the cost of traditional satcom networks.

An aircraft can generate lots of real-time AHM-ACMS reports, and send all this data across the L-band network, which is expensive. Therefore, smaller airlines will typically download AHM-ACMS data for analysis more economically in batch mode once the aircraft is on the ground, or only send the most important notifications and reports across L-band in real time during flight.

Instead of using L-band, AoIP enables AHM-ACMS data to be broadcast more economically over the Ka-Ku network, or via any I.P. air-to-ground channel. Ka-Ku vendors typically offer fixed-rate data contracts that can include an unlimited data usage allowance. However, many of these contracts may have a fair usage clause written in them, and extremely high data usage could incur a financial penalty.

"ACARS over I.P. will autonomously choose the most cost-effective channel, such as the cellular network or the passenger IFE satcom. Then the data transmission costs will be absorbed within the airline's LTE and Ka-Ku band data," explains Skelton. "If an airline quadrupled the amount of ACARS traffic sent across its passenger network, the impact in terms of data usage on these systems would be insignificant. This is because the extra ACARS data sent will amount to Kbits and not Mbits."

Baseline data

Baseline AHM-ACMS reports are standard on all new aircraft and record elemental data, such as out-off on-in (OOOI) messages that are automatically transmitted via ACARS. The airline can activate, deactivate and customise AHM-ACMS reports.

Once an aircraft has taken off, and throttle settings are set from take-off go-around (TOGA) to cruise, the aircraft will automatically generate and store a baseline report on the engine's performance.

The AHM-ACMS report can be sent immediately if it is required by the airline's line maintenance department. Alternatively the data can be offloaded at the end of each flight across the airport's 4G and LTE networks.

For an aircraft to receive a Ka-Ku band transmission, it must be equipped with an Ka-Ku antenna that has a receiver designed to track a selected satellite. For aircraft applications, it was necessary to shrink these large antennas that are also associated terrestrial and maritime operation.

Flight operations quality assurance (FOQA) data mandates that several datasets are to be recorded and monitored. The mandate does not require the data to be offloaded in real-time, however, so it is possible to collate these reports at the end of each day or flight.

"FOQA data is mandated in most parts of the world except the US. What is not mandated is how this data is used, and how often it needs to be checked. It needs to be made available for regulators such as the Federal Aviation Administration (FAA) or European Aviation Safety Agency (EASA), but how quickly the data must be received by the airline is not specified," says Skelton.

The FOQA report will contain data and text, referring to why the report was generated. These will include reasons such as an exceedence of one of the aircraft's operating parameters. Included within the report are any snapshots of the data that the aircraft has looked at.

Quick Access Recorder

The quick access recorder (QAR) records ARINC 717 data and ARINC 429 data. ARINC 717 is a protocol that defines a standard of data communication between the flight data recorder (FDR) and the data acquisition unit (DAU). ARINC 429 data standard relates to the transfer of data between aircraft avionics and systems.

Avionica LLC has expanded the capabilities of its Mini QAR product. The remote data concentrator (avRDC), is similar to the Mini QAR, because it can record all the mandated 717 data. However, the avRDC can harvest a larger amount of data from more of the aircraft's 429 channels.

"Our Mini QAR is a small and lightweight piece of hardware, certified on more than 300 aircraft types. It records the standard FOQA data and some 429 data. The avRDC has all the functionality of the Mini QAR, but the RDC can record data from more of the 429 buses," says Sean Reilly, vice president business development, Avionica.

Access to 429 channels makes it possible to harvest information from almost any system on the aircraft that produces data, including from the engine and hydraulics. By collating data from the aircraft's full authority digital engine control (FADEC) system it will be possible to capture enhanced engine parameters.

Once the data is recorded, a cellular device on the QAR or avRDC, which gives the unit global cellular connectivity, enables the transmission of data. When the aircraft lands the RDC will package it into datasets and send it over the cellular network.

Once the aircraft lands the hardware connects to Avionica's avSYNC cloud-based programme to route the data anywhere it is needed. After the aircraft data has been collated, and verified as being received, it can then be forwarded to an airline's in-house facility for processing or to a contracted third party.

"Two years ago, we entered into a joint partnership agreement with GE Aviation. GE has a division that can perform analytics on the data if the airline chooses. Nevertheless, we are agnostic to where the customer needs to send the data for post-flight processing," explains Reilly. "The key to this technology is that it is possible to capture almost any parameter on the aircraft, provided we are connected to that data source, and record it, depending on the customer's individual requirements, and then transmit it to where it needs to go for post-flight analysis."

Typical service level agreements stipulate that all the data from the aircraft will be extracted within 24 hours. Nevertheless, 80% or more of the data is usually extracted within 20 minutes of connecting to the cellular network. When the aircraft lands, the data is transferred by 4G or LTE as soon as possible. However, the exact time taken to fully extract the data will depend on the flight duration and the number of recordable parameters, and how much data the aircraft generates.

For example, some DAUs record at a data frame rate of 512 words per second (WPS), and others at 1024 WPS. Newer aircraft can record data at a higher frame rate, so the data files are bigger, and therefore take longer to transmit.

According to Reilly, an aircraft will typically generate about 2Mb of data per flight. "It is not a gigantic amount of data, but it contains a significant amount of information. The flight length and aircraft type need to be taken into consideration. Technological advancements mean it is possible to record more data sets from contemporary 787 aircraft than a legacy 777-200."

The DAU will typically take a snapshot of data at given time intervals. If there are eight fuel nozzles on an engine, the DAU will record the average flow rates of all eight fuel nozzles, and record this information in a snapshot. If connected to a FADEC more detailed information can be recorded.

The FADEC can produce 4Mb per hour of data alone. Previously a long flight could yield a total of 1Mb of data, yet a three-hour flight, including FADEC data, could produce a 15Mb data file.

Collating a larger amount of raw data makes it possible to process the results with higher accuracy, and then turn the results into something tangible and beneficial to the airline.

When an airline records a higher number of its 429 buses, it is possible to get data from every fuel nozzle, so if one fuel nozzle is injecting more fuel than all the rest, it can be replaced before a hot-spot on an engine develops that will incur the high cost of a removal and shop visit.

Furthermore, instead of getting a data snapshot, it is possible to record the enhanced data and observe trends and perform predictive maintenance, meaning faults will be identified before they occur.

Edge processing

Distributed computing is a field of computer science where the components are located on a different network. Computers communicate and coordinate their actions by passing messages to one another to achieve a common goal.

Edge processing is a distributed computing paradigm that brings computation and data storage closer to the location where it is needed, to improve response time and save on bandwidth.

It is possible to install an application onto the aircraft server that analyses the 429 databuses in real time. The server-based application will compare the data against a predetermined set of parameters. If one of the data parameters on the 429 channel exceeds its pre-defined minima, the system will send an alert to the ground.

The system can identify fault trends, so that AHM notifications can be sent directly to the airline's line maintenance facility before a fault light is illuminated on the flightdeck.

Once the line maintenance operator on the ground receives the notification, they can objectively look at the problem and proactively formulate a solution. Remedies could include a mechanic or a part waiting at the airport when the aircraft lands.

"The system will allow proactive maintenance actions to be taken, instead of waiting for the aircraft to land and then downloading the data to identify the fault before prescribing a remedy. The proactive maintenance approach keeps data costs down and the aircraft flying," says Reilly.

Next Gen edge processing

Next generation edge processing means that instead of sending a snapshot of data, the system sends a 15-minute summary-block of data on fluctuations developing on the aircraft for on-ground analysis.

If a ground maintenance action needs more data to pin-point the exact nature of the fault, an additional data set can be downloaded from the aircraft. Depending on the severity of the anomaly, it is possible to monitor the situation in almost real time and even direct the crew to change some of the aircraft settings.

Yet if the aircraft's onboard systems are functioning within their set parameters, the edge processing app will not broadcast data that is not required, saving the airline a lot in bandwidth and data costs.

Once the snapshot has been received, a decision can be made to wait and to off-load a more comprehensive dataset more cost-effectively on the ground via the LTE network.

"We are seeing that the cost of moving data during flight is becoming less expensive. At Avionica we are looking at the SBB network by Inmarsat. SBS is still under the SBB network, but I am not worried about the safety part of the SBS system because this is for regulated air traffic control (ATC) and safety services data. I just need it to move data as economically as possible," adds Reilly.

"Using the broadband pipe means it is also possible to use passenger connectivity. There are some security issues and FAA regulations to be overcome. However, many aircraft do not have a Ka system installed," says Reilly.

Installing a Ka-Ku system can cost up to $750,000 per aircraft. There are concerns, however, about rain-fade across

Predictive maintenance will reduce the chance of an aircraft-on-ground situation. Monitoring 150-160 aircraft data parameters will allow the operator to perform prescriptive and predictive maintenance across 10 ATA chapters, such as wheels and brakes, mechanical systems and the APU amongst others.

this spectrum. The SBB service on the other hand offers a more reliable connection and is not subjected to moisture attenuation.

Sending data only when an anomaly is detected by edge processing techniques means that data transmissions are small, and use a small amount of bandwidth.

"Compared to Ka-band, it is a little more expensive to send data across the L-band network, although the added value is that the ground is guaranteed to pick up the message regardless of the weather conditions. This could justify the cost difference," says Reilly.

Next generation edge processing is a controlled way to offload data from the aircraft. Furthermore, it is possible to evaluate problems in-flight.

"The value is how to keep the aircraft operational and operating to schedule. Does it make sense to pay $100 extra in data charges to safeguard the aircraft against an on-ground event that could cost $100,000 in lost revenue?" asks Reilly.

There are many different views on how to perform the analytics and the amount of data needed to ensure aircraft reliability. There will be new technologies that require expanded data sets to improve the airline's efficiency.

For the most part, airlines are now starting to see the potential of next generation edge processing. Initially airlines were unsure about the benefit of the technology, but they are noticing how edge processing and big data analytics can make a difference to their operation.

"The Avionica platform has been designed as an open architecture system that uses third-party software, such as Ultramain Logbooks. The system's open architecture adds flexibility to the solution, making it possible for the subscriber to simply adjust any of the aircraft's recording parameters it is recording," says Reilly.

This is because the app within the edge processor server uses docker technology that makes it possible to facilitate the many recordable parameters to be monitored, individually, in real-time.

Low cost routing

To be most efficient, the software also needs to be smart enough to understand all the connectivity options to offload data from the aircraft to the ground. Application processing interface (API) software on the server will pick the best and most efficient communications link to the ground as set by the operator.

For example, high priority messages can be sent across SBB, while lower priority messages can use the ground cellular network.

The key is that acknowledgement (ACK) reports mean the system will also know if a data packet has been received in its entirety and free of errors. If the system sends a message across the Ka network and it does not reach its destination because of weather conditions, it could transmit the data again, or re-send the data packet over the reliable SBB network.

The system would automatically use air-to-ground methods, depending on the priority rules set by the airline and its wants and needs.

Using a priority message type system that picks the most efficient datalink connection will minimise the impact on the airline. Low-cost routing helps airlines maintain control of their data delivery and its impact on costs.

Data frames

Honeywell has developed solutions that allow airlines to capture and transmit aircraft data, which can be analysed for valuable insights. It is important to discuss not only how data is transferred but also how the data sets are selected.

"Flight data has two different dimensions," explains John Peterson, vice president and general manager, connectivity services, Honeywell.

The first dimension is the data frame, which can include parameters such as the date and time of day, as well as altitude, speed and heading, and engine- and sensor-related data.

Modern aircraft are designed with more sensors, so they record a lot more data in their data frame.

The second dimension is the frame rate. "This is the rate at which the data set is captured," says Peterson. "People talk about the total amount of data an aircraft can produce, such as 1TB per hour. Yet to harvest data at 1Tb per hour could mean collecting the aircraft's entire data frame, at a data rate that is not adding real value and driving unnecessary costs." Honeywell understands how to balance the optimum frame rate to the data that the airline wants to analyse, optimising the data's value and minimising the cost expended.

Labels

Honeywell refers to the data contained within a data frame as a label. A label is the recordable parameters and the data sets of information.

"So how many labels and parameters does an airline need?" asks Peterson. "Honeywell has broken these parameters down, meaning it is possible to explain to the airline what the tangible results are for any number of labels. Alternatively, an airline can ask what results can be achieved by analysing 50 or 175 or 350 labels."

Monitoring 10-20 labels allows an operator to perform predictive and prescriptive maintenance on the aircraft's auxiliary power unit (APU). Monitoring 150-160 labels allows the operator to carry out prescriptive and predictive maintenance across 10 ATA chapters, such as wheels and brakes, mechanical systems, the APU and many pumps, generators and alternators. Monitoring data across 350 labels will allow the operator to perform flight optimisation, fuel efficiency and predictive and prescriptive maintenance.

The amount of data harvested across 350 labels after an average 1 hour 45 min flight time will be about 35Mb. The cost of sending this across an LTE network connection once the aircraft is on the ground is very low.

"Once you have captured the data across 350 labels at a specified frame rate, it is important to ask if there is a global network to route the data to a cloud-based server or directly to the airline's server," says Peterson.

The accuracy of the data is defined by their recordable frame rate. The logic is that the higher the frame rate the higher the degree of accuracy, in terms of being notified of an issue. The higher the frame rate and the amount of data captured, the higher the transmission and analytical costs.

"We advise operators which databuses you need to get data from, such as the 717 bus, the 429 bus and EPIC ASCB bus," explains Peterson. "We have product gateways and edge nodes that can read data from all of these buses. Then we can calculate what can be achieved with the data. Many in the industry claim to be able to transfer flight data, so then the question becomes what do you want to achieve with the data? We have a long track record of working with major airlines on flight efficiency and predictive maintenance initiatives that leverage flight data."

Not all airlines have invested in the hardware and infrastructure to process the data as quickly as possible.

The possibilities depend on how the operator wants to retrieve data from the aircraft. If this has to be done manually with a laptop that is plugged into the aircraft, it can only be done when the aircraft is on the ground, and a maintenance person is available to work the software. This makes it likely that the airline will receive the data typically once a week to once a month.

Another possibility is to use a datalink subscription. The configuration file can be updated on the communication management unit (CMU), and a quantity of data can be extracted from the aircraft while it is in a datalink environment.

When the aircraft is on the ground, it can transmit the data via a 4G and LTE network. If that is not available, it is also possible to send data from the aircraft via an L-band system, such as SBB or Iridium services, no matter where the aircraft is in the world.

Finally, it is possible to use Ka and Ku services or low-cost air-to-ground connectivity solutions. Almost all aircraft that transition oceanic airspace will have an L-band system installed. Many aircraft that have been commissioned into service in the past five years have been line-fitted with a Ka and Ku antenna.

Yet the value that the data provides has to be greater than the cost of acquiring it. In a typical set-up, it is possible to buy an edge node and a network hub for less than $10,000. Installation costs are approximately $5,000, so a total of $15,000. So for an airline with a fleet of 300 aircraft, this will cost millions of dollars.

Then there are the complexities between the capital cost and the recurring costs for transferring the data. For example, the typical airline uses thousands of cellphones for its crew and staff.

"We are working hard to combine technologies, to reduce an airline's data acquisition, and we expect the data costs per aircraft to be no higher than of a cellphone," says Johnson. "This is more palatable for an airline, because it can combine this cost within its cellular call plans."

As the cost of transmitting data decreases, it is conceivable that downloading data in real time will become increasingly more affordable for airlines. It is expected that the costs in unforeseen maintenance will be much greater than the extra cost in data incurred by streaming real-time data.

"For example, we can predict when the toilet vacuum pumps are going to fail and

By analysing aircraft data, it is possible to predict when a toilet vacuum pump is about to fail and needs to be replaced. Changing the toilet vacuum pump before it fails will not only keep all the aircraft's toilets in service, but it will also save time and cleaning costs.

need to be replaced. When a toilet vacuum pump fails it creates a mess in the toilet, putting it out of action and delaying the flight," says Peterson.

A key part of a complete connected aircraft strategy is savings in fuel used and the better management of disruption on the day of operations that improves aircraft on-time performance.

Real-time data acquisition enables airlines to enhance the data models, and then make it possible to calculate the most efficient amount of data to be streamed from the aircraft. In addition. real-time data will become more attractive as the cost of sending it across the Ka and Ka band comes down.

"We will get to a point when downloading specific data in real-time will generate enough value, meaning the value will outweigh the acquisition cost. We ultimately believe the analytic services will pay for the aircraft connectivity itself, and passenger connectivity purchases will be part of ancillary revenue alongside eCommerce" explains Peterson.

Combining real time data infrastructure with an advance analytics platform, it becomes possible for an airline to select the right set of data that provides immediate opportunities for optimisation, thereby reducing data transfer costs while maximising positive outcomes.

Telemetry

Teledyne has recently completed telemetry trials with its GroundLink Comm+ system, which has been used for the last two decades by airlines worldwide to offload their aircraft data upon landing.

This trial consisted of offloading aircraft data in-flight and in real-time. "There is much talk about aircraft real-time telemetry, and we have successfully demonstrated that it is possible to transmit a reliable datastream in real-time," says Skelton. "The challenge with real-time telemetry, however, is that there are few available ground-based applications that can effectively process the data streams in real-time."

The telemetry system currently collates data from all the ARINC 717 acquisition unit labels, and can send a stream of up to 250 user selected labels. The maximum frame rate is once every second, which is the minimum delay rate possible.

The advantage of streaming data in real time is that the application on the ground is continually monitoring the aircraft. This means that maintenance gets an immediate acknowledgement that something is unusual with the aircraft that initiates a monitor and tracking scenario.

Therefore, it is possible to watch an engine core temperature, and monitor it for any exceedences. Unlike a static report, it is possible to look at the ongoing effect and make better decisions because of the granularity of the real-time data that is streamed to the ground.

If a core temperature remains high, it is possible to turn the aircraft around. If the core temperature drops back to normal, then it is possible to make other assumptions, such as the possibility that foreign object debris (FOD) has been ingested by the engine.

Furthermore, the data can be forwarded directly to the engine original equipment manufacturers (OEMs,) such as Rolls-Royce or GE.

With edge processing techniques, if an anomaly is not detected, then all the aircraft's telemetry data is simply offloaded after the flight. This results in a time delay, however, before any trending and predictive maintenance analysis can be performed on that data.

Snapshots of data are used for proactive maintenance and monitoring. You need to access all the main aircraft data sets, however, before it is possible to perform predictive maintenance trending. If the data is being continually streamed, the airline can constantly process trending analytics on all the datasets while the aircraft is in flight.

Therefore, telemetry data is more suited for predictive maintenance analysis than edge processing techniques.

Predictive considerations

Procedures that rely on predictive analytics must be integrated piecemeal to avoid overwhelming the airline or industry. Certification issues within the industry mean that the aviation sector is slow to embrace the new technology, so the cost of implementing it is important.

Before it is possible to put a newly developed piece of hardware on an aircraft, it needs to meet a minimum specified certifiable requirement, and requires a supplemental type certificate (STC) to install it. The cost of developing an STC can range from $100,000 to $300,000, depending on the impact of the system on the aircraft type's original design. Certification costs mean the barrier for entry for new technology is high.

Summary

Next generation satellite waveforms and edge processing systems are exponentially more efficient than their forerunners. It is therefore likely that in-flight maintenance will become a reality.

Capacity and technological improvements across the Ka and Ku network are making the spectrum more secure and reliable. It is possible in the future that more safety-related data will be permitted to be sent across the network.

Furthermore, the possibility to record more data sets and offload the information more efficiently will lead to more comprehensive predictive and preventative maintenance. This will translate to a safer, more reliable and better optimisation of aircraft within an airline's inventory. AC