Modern aircraft are producing increasing amounts of technical data that is analysed to predict and solve repetitive faults. But data comes at a cost data comes at a cost. Are airlines being saturated with large amounts of data that are not adding any tangible value to the operation?

The aircraft big data explosion: how much is enough?

One aspect of big data analytics is the practice of taking large quantities of data from a source and analysing and using it for business-related tasks, such as determining the root causes of aircraft component and system failures.

Passenger aircraft today have the ability to capture data from a multitude of onboard systems and components in flight. Once the data is offloaded from the aircraft, data engineers can analyse it to detect trends and identify fault patterns in components and systems.

It is considered that the larger the quantity of data available, the easier it is to detect a fault pattern, and predict when a component failure is likely to occur. Preventive maintenance can therefore be scheduled, causing minimal disruption to the aircraft's operation.

Scheduled preventive maintenance tasks can prevent aircraft-on-ground (AOG) situations and improve an operator's on-time performance, scheduling, and revenue. The many operational variables make it difficult to quantify how much an AOG event will cost an airline, but some estimates do cite an average cost of a delay at $150,000 per hour per flight.

Airbus's newest flight operations and maintenance exchanger (FOMAX) data router captures more than 20,000 real-time aircraft parameters. FOMAX sends this data to Airbus's Skywise data analytics platform, providing predictive maintenance capability to its operators and subscribers.

As useful as big data and predictive maintenance are to an airline, these do not come for free; the larger the volumes of data offloaded from the aircraft, the higher are the related costs.

Engineers will report that no amount of data is too much, because it offers the possibility to reduce the number of unscheduled maintenance events. However, for an airline's financial officer, monitoring an increasing number of aircraft component and system parameters represents an increasing cost in data management and analytics, and the possibility of only a marginal return.

Airline managers therefore need to determine the quantity of data required to noticeably streamline the day of operation, and the quantity where the cost negates any additional value or gain.

Taking stock

Air France-KLM claims to be the first airline group to develop predictive maintenance techniques following the A380's introduction to service.

"The A380 was one of the first aircraft to be fitted with a large number of sensors. Air France Industries KLM Engineering & Maintenance (AFI KLM E&M) department wanted to see how they could use the large volume of data coming from the aircraft to improve maintenance operations," says Philippe Bordel, digital services commercial manager at AFI KLM E&M.

When the A380 was introduced into service there was a significant number of aircraft disruptions, delays and AOG events. Once a recurring AOG related to a failing hydraulic pump was identified, the engineering department began analysing the data generated by the pump and the workshop reports. They realised that they could correlate the data patterns to predict an incoming failure.

There are now predictive models that cover the critical ATA chapters on both Airbus and Boeing aircraft, and these have been helping to significantly reduce the number of aircraft delays and AOGs. With the development of the A350 by Airbus, and the 787 by Boeing, modern aircraft are now increasingly rich in data and sensors.

"We cannot use all of the available sensors on the aircraft so it is pointless to monitor them all," says Bordel. "Some of the sensors are used by the pilot to fly the aircraft. What we are interested in is developing a purpose-built approach for making better use of the appropriate sensors and the parameters to develop our maintenance practices."

The A350 has about 6,000 sensors, although it is difficult to quantify the exact number, because the definition of a sensor can be interpreted differently. Nonetheless most aircraft sensors measure temperature, flow speed, torque positioning and global positioning system (GPS) data.

"We know that on each flight, the A350 generates about 2.5 terabytes (TB) of data. To monitor the aircraft's maintenance condition we typically harvest data from 4,000 sensors and download about half a gigabyte (GB) of data per flight," says Bordel.

The A380 has about 10,000 sensors because it has four engines and a comparatively large wing. The systems and sensor data generated onboard an aircraft is called asset machine data.

"It is the asset machine data that is associated with a lot of the hype relating to the high volumes of data an aircraft can produce," says William Cecil, managing director at Aircraft IoT Consulting &

Air France developed predictive maintenance techniques following the A380's introduction to service. The A380 was one of the first aircraft to generate a large volume of data. Its large wing and its four-engine configuration mean each A380 has about 10,000 sensors. The asset generated data yielded from the aircraft reduced many AOG events.

Solutions. "There have been claims that the 787 generates 500 GB every flight, or 844 TB during a 12-hour flight. Whatever the exact number, only a tiny fraction of the data that is transmitted off the aircraft and subsequently used."

Often less data than this is collected, and the reality is that a few GB per month per aircraft of time-series data is sufficient for predictive maintenance applications. Large volumes of engine data are available, but for the most part not wirelessly. It is mostly manually downloaded and on a sample basis. Not all data is collected, and much is over-written between downloads.

According to Cecil, no more than 5 GB of time-series data is collected post-flight from aircraft today, and this is primarily used for safety analysis. Real-time data transmission via the aircraft communications and addressing system (ACARS) is typically less than 50 MB per month, or no more than 1-2 MB per day.

"The amount of continuous time-series data automatically collected and offloaded from the aircraft via WiFi Gatelink or 3G cellular network after landing, is typically about 5MB for each flight hour. Therefore a 12-hour flight would result in a 60 MB download after landing," explains Cecil.

Prominent data is routinely collected by the aircraft quick access recorder (QAR) or a digital ACMS recorder (DAR). The data is usually the same as, or is an expansion of, that data acquired by the flight data recorder (FDR). Operators will typically collect about 1,000 parameters, but some carriers have made software upgrades to expand this number to around 3,000 parameters.

Hardware and/or software upgrades make it possible to acquire more data. The data acquisition system requires a software update and installing an Airbus FOMAX system makes it theoretically possible to download all 24,000 parameters available on the A320.

The volume of data captured from the aircraft is determined by factors other than just the total number of sensors. First is the rate at which each sensor records data which is measured in samples per second (or Hz i.e. frequency).

"The more data that is harvested, the more analysis that you can do. If operators choose to monitor a particular valve within the air conditioning system that has an ongoing fault, the more data samples that are taken per second, then the better the chance of understanding the reasons for the part failing," says Murray Skelton, senior director of sales at Teledyne. "Monitoring a parameter at a high Words per Second (WPS) data rate will improve the accuracy of predicting when a part is going to fail, than with a lower rate."

Secondly sensor technology has evolved over the years, and younger aircraft have a higher number of sensors fitted that generate larger volumes of data. Many newer aircraft such as the A320CEO/NEO have ACMS software (Aircraft Condition Monitoring System) that can analyse the RAW data in real time in flight, prognose a potential fault, and generate a report that is sent immediately via ACARS to airline engineering, without the need to offload data and analyse it on the ground beforehand.

Data

According to Steve Bogie, vice president of flight operations and technology at Drone Delivery Canada: "There is lots of talk about capturing and downloading high volumes of data from the aircraft. The real issue is how you use the data to add value to your operation. Which parameters do you want to monitor, and what do you want to achieve?"

Improving aircraft reliability is the airline's primary goal when using aircraft data. Each individual carrier may, however, have different requirements with respect to the parameters and components they want to monitor. Typically, this is dependent on their maintenance goals and standard operating procedures (SOP), and the operational focus of its operation.

Fundamentally, the ratio between flight hours (FH) and flight cycles (FC) must be considered, as well as the aircraft's typical operating environment. Noticeably different effects on some of the aircraft's systems and components could be caused, depending on whether the aircraft is operating in a desert or tropical environment.

Since the maintenance condition of many of the aircraft's systems and components is measured in terms of FH or FC, high FC operators will typically focus their attention on monitoring data from items such as landing gear and brakes.

Furthermore, high FC operations will increase the stress and wear on wing components such as ailerons and spoilers, including their associated hydraulic systems, meaning they will have a higher propensity to fail, compared to aircraft flying high FH-driven operations.

High FH operators' priorities are likely to be FH-driven components with parameters that relate to the performance and reliability of the engine.

High value components

According to Bogie, Air Canada has identified 200-250 major components on the 787 to monitor.

"This number mainly included critical components, so if one of these parts fails, the aircraft becomes unserviceable and must be grounded," explains Bogie. "Our approach was to focus on the high-value, high-yield items that have a direct impact on reliability, instead of monitoring all 5,500 parameters available. Therefore, we filtered out parameters that were not immediately and critically important to us, to avoid burying ourselves in data that has

Finnair are able to predict integrated drive generator (IDG) faults in up to a month in advance. Between 2018 and 2019 the number of maintenance tasks that were scheduled from results of predictive maintenance increased by 50%, and many faults repetitive faults can be resolved by monitoring a single parameter.

little value and does not add anything to the operation."

Once the value has been extracted from the high-value parameters, it is possible to look at adding more parameters, and label them in categories of value. Airbus and Boeing report the potential for their operators to monitor more than 20,000 parameters and download the data to the respective analytics platforms for analysis.

"Whether there are 5,500 or 26,000 parameters, the point remains that many of these items are not high-value and should not be an airline's initial focus," says Bogie. "Monitoring 200 is typical for a 787, depending on the operator's options and priorities. The focus is on the parts that are most important to the operation."

What is equally important is the aircraft utilisation data, which can be used to reduce the number of high-value components kept in storerooms throughout the airline's route network, and which are used to replace components that fail at random. It is better to manage and allocate replacement components in accordance with aircraft location by performing preventive maintenance before the fault occurs. This is an alternative to strategically placing expensive components at outstations, waiting for a fault that may never occur.

"The high-value parts are typically extremely important ones where you want to gain the maximum yield. You want to be able to provision the airline properly, without having a lot of surplus inventory sitting around," explains Bogie. "For example, a generator control unit (GCU) costs about $250,000, and is a 'no-go' item that will ground the aircraft in the event that it fails. So when you need one, you need it urgently. Yet it makes no sense to have an excessive number of spare items sitting in the storeroom. Predictive maintenance makes 'just-in-time' inventory possible."

Data storage & science

Real world use cases relate to reliability and monitoring high-value components to maximise yield and minimise maintenance costs and the capital inventory component. Capturing all the data that is produced by the aircraft, however, only increases the cost of transmitting, managing and storing it.

It is easy for operators to fall into negative equity with data storage, since large volumes of unused data accumulate and increase storage costs. Storage costs from data accrued by 24,000 recordable parameters per flight will be much higher than for data gained by a couple of hundred. Furthermore, recurring cloud-based storage costs will continue as long as the data remains stored within the cloud, so it must be analysed to be put to good use or it has no worth.

The 787 is more data rich than legacy aircraft types, such as the 767 and older generation A320s. When the 787 was launched, the concept of big data was well received, although the first analytics tools were limited, and did not have enough capability to manage and use the data. Data analytics tools have since improved, although much of airlines' current expertise does not actually rest with data science.

"If you do not have people physically analysing the data, then what value is it?" asks Bogie. "If the operator does not have the data scientists to do this work, then it can outsource the task to original equipment manufacturers (OEMs) or analytics companies, but the airlines will still need the analytics tools and the cloud storage. This will be an added cost to the airline."

A concern with OEM analytics tools is that they can only analyse data from their aircraft. Skywise is Airbus-centric, which means that operators with multi-OEM fleets must subscribe to more than one analytics platform, with each one only solving part of the problem.

Finnair

Finnair has been using the Skywise cloud-based platform since late 2018, and uses it to host predictive models that are partly built upon aircraft sensor data. Finnair reports that the system has proven to be an efficient tool in handling the large data sets needed. It also allows Finnair to combine other data sources, such as content management system (CMS) fault and maintenance messages with the sensor data that had previously been separated. By combining the data sets and monitoring them over a period of time, Finnair better understands the correlations between them and becomes more effective in identifying trends and patterns.

According to maintenance analytics engineer at Finnair, Oskari Nihtila: "We are using Skywise for all our Airbus fleet. This includes the A350, A330s and A320s. The amount of data that each type produces varies. The A350 provides a large amount of data compared to older aircraft types, so it allows more opportunities in that sense. We are not allergic to the term 'big data', but it is more about finding the meaningful streams in the data, than the amount of data itself."

The corrective data streams that typically prove to be insightful and most helpful in terms of predictive maintenance can actually be small. In some cases it is only a small amount of data that has enabled Finnair to resolve difficult fault cases.

For example, it is possible to monitor every cabin temperature control valve and know its position for every second of its operation. "This is not helpful," says Tero Polamo, operative engineer at Finnair.

Instead of monitoring all the many parameters available, Air Canada focused on the most operational critical and expensive components when it introduced the 787 into service. Doing so meant the airline was not buying and analysing data that did not add value to the operation. Monitoring high value components enabled Air Canada to streamline its part inventory.

"Actively monitoring these parameters is not a good use of time. Knowing a lot about the valve position does not mean anything. However, when the crew report that a zone within the cabin is too hot or too cold, then the data becomes useful."

It is possible to filter out unwanted parameters, which was useful when Finnair introduced the data-rich A350 to its fleet. At the beginning the airline's focus was to monitor data from the most troublesome parameters, and then add to them over time.

Following the launch of A350 operations, Finnair experienced recurring issues with the nose landing gear, whose actuator developed an internal leak and failed after repeated use.

"In some cases, it was impossible to retract the landing gear fully after take-off, so the aircraft had to return or divert," explains Polamo. "On the A350 there are plenty of datasets available, however, to track this fault if conventional aircraft condition monitoring system (ACMS) data is used. This includes data 'snapshots' from individual components, or small reports detailing the operation of a system."

By extracting landing gear retraction and extension reports, it was possible to time how long it took for the landing gear to retract and extend. Any significant change in these times indicates the retraction actuator is beginning to deteriorate, and the airline must schedule the installation of a replacement or have the system serviced before it fails.

Major aircraft components and systems, such as the undercarriage, must be included within the operator's minimum equipment list (MEL) and be approved by the national airworthiness authorities. An operator must not operate an aircraft that does not comply with the approved MEL, except with the explicit permission of the appropriate regulatory authority. Authorisation for such operations is typically granted to airlines for flight testing and aircraft positioning purposes only, and not for passenger operations or to generate revenue.

Components and systems that have the greatest risk potential to cause operational interruption are typically the items airlines want to monitor to prevent AOG situations.

"On the A330, we monitor the integrated drive generator (IDG), which is driven by the engine accessory gearbox. On a twin-engine aircraft like the A330 there are two IDGs. If one fails, we lose a significant portion of redundancy in the system," says Polamo. "We can dispatch the aircraft with one functional IDG, but the auxiliary power unit (APU) must be running for the entire flight as an alternative AC source."

Operating the APU for the duration of the flight increases fuel consumption, while operating with a single IDG means the aircraft will be indefinitely grounded if the remaining item fails. Finnair has almost cleared the aircraft of any unscheduled IDG replacements, helping to improve its operational efficiency.

The carrier can plan its maintenance schedule around predictive maintenance and therefore decides the optimal time slot to complete a repair. The aircraft does not generate revenue during maintenance, so it is important to optimise all aircraft ground time. It is always preferable to schedule ground time maintenance together with unscheduled ground time for maintenance.

"We have a few indicators that show how the predictive maintenance programme has been helping us," says Nihtila. "The number of preventive maintenance work orders, based on the results of predictive maintenance, increased by more than 50% in 2018 and 2019."

Predictive maintenance has also influenced Finnair's operation, since data on its A350 fleet shows that technical dispatch reliability improved in 2018 and 2019.

Predictive maintenance has also had an impact on repetitive fault defects on the A330. The engine bleed system diverts air from the engines to be used in sub-systems for air-conditioning, cabin pressurisation and ice protection.

"We had a lot of problems with the A330's two bleed systems. If one of them fails, then the remaining system must double its load," says Polamo. "It was difficult to work out why the system failed. Once we started using predictive maintenance to look into the data really carefully, we managed to almost entirely get rid of the bleed system faults and the operational interruptions that they caused."

There are several issues relating to the A350's complex flap and slat system. This changes the aerodynamic shape of the wing during flight, by altering its curvature and ultimately moving the centre of lift on the wing. Traditionally, this was done by moving the aircraft's centre of gravity (CoG) by transferring fuel to the horizontal stabiliser tank on the A330. Due to the importance and complexity of the system, Finnair monitors many wing components and subsystems.

Many of the items that are classified as 'no-go' items are time-limited, and must be repaired within a number of days, FC or

FH. Some MEL items increase aircraft form drag, meaning they incur an extra weight penalty that increases fuel burn.

Back to the future

For big data and predictive maintenance to be useful it must be able to detect a fault far enough into the future to allow the required preparation for carrying out remedial action. Finnair is now predicting IDG failures on the A330 almost a month in advance. The amount of notice given before a component fails, however, depends on the system that is being monitored.

"I analysed the data after an IDG had actually failed, and was able to see first signs of the upcoming failure one month before. But then we were just building that system and finding examples," says Polamo. "There are still some faults that might come out of nowhere, but usually they start developing pretty slowly so we can catch them. It is not easy, but we are able to catch most of those faults."

According to Finnair they cannot quantify the exact value-add of predictive maintenance to the operation, because it is difficult to quantify a 'what could have been' cost if the fault had remained undetected, and caused an AOG event.

"By doing a predictive repair and service we avoid the operational interruption so we cannot know exactly what the savings are. It is really hard to say what could have happened, and what the worst-case scenario could be," says Polamo. "I think if we have managed to avoid a number of cancellations and long delays per year, it gives an idea of the potential cost savings. Nevertheless, the monetary cost saving depends on many factors."

In an attempt to reduce overhead costs during the Coronavirus pandemic, some airlines are scaling down projects that do not have a definite value-add to their operation. Yet Finnair believes it is finding enough value from big data and predictive maintenance to continue applying it to influence the operation of its aircraft.

The legacy ACMS and aircraft health monitoring (AHM) systems are highly effective at capturing data and can monitor a broad spectrum of high-value parameters.

"We do not need data for every single point of time during the flight. It is possible to monitor just a single parameter, and this will give us enough data. One method used to solve the recurring A330 bleed air system fault was analysing its pressure and calculating a long-term trend," explains Polamo. "We use one data point to measure the pressure value per flight once the aircraft is in a stable cruise."

Once several months' worth of data has been harvested, then it is possible to identify the long-term trend.

PROGNOS✱

AFI KLM E&M provides operational support to the Air France and KLM fleets and other customers. As an airline-MRO, it has used its knowledge of aircraft systems gained from providing MRO services to a large number of different aircraft types, to develop its PROGNOS✱ predictive maintenance analysis tool on components, engines and APUs.

"Airbus uses one of the Skywise Predictive Maintenance (SPM) modules, which is comparable to PROGNOS✱," says Bordel. "We are developing an independent approach from the aircraft manufacturer like AVIATAR does. Many airlines will still only use the core basic modules of Skywise or AVIATAR without necessarily using SPM or AVIATAR Prediktor modules. Since we have access to large amounts of historical data from Air France and KLM, we know the value of each component and system for many

Finnair is using Skywise for all its Airbus fleet. This includes the A350, A330s and A320s, even though the amount of data that each type produces varies. The A350 provides a large amount of data compared to older aircraft types, yet many operational critical issues can be solved by ACMS data alone.

aircraft types."

Additionally, AFI KLM E&M archives repair data for components and systems. The PROGNOS✱ solution will monitor sensors within an aircraft's components, engines and APU. The most advanced data sets are from aircraft components, because data harvesting technology from the sensor is more advanced. High temperature and vibration make it difficult to place sensors in the engine or APU, and many need special coatings. According to Bordel, PROGNOS✱ focuses primarily on sensors located in the landing gear, brakes, flight controls, fuel, and hydraulic and electrical systems, from about 4,000 parameters.

"Thanks to our experience we can anticipate faults 30-50 flights ahead. This means it is possible to have three to 10 days' notice to prepare and rectify, depending on the aircraft utilisation rate and the number of FC it completes per day," says Bordel. "We have an operator in Southeast Asia, whose aircraft typically complete 10FC per day. At this rate of utilisation the operator has about two or three days' notice, but generally we predict the fault a week in advance."

AFI KLM E&M developed its first predictive models based on operational issues. To begin addressing these, they turn to the ATA chapter that is causing the most problems in terms of reliability of the aircraft and monitor those parameters.

It has been discovered that almost all operators report the same type and number of faults against the same ATA chapter, such as Chapter 32 – landing gear, Chapter 21- air conditioning, Chapter 27 – flight controls, and Chapter 36 – pneumatics.

"More often than not, we find the same 'troublemakers' in the ATA chapters, regardless of aircraft manufacture. If you take Chapter 32, which is landing gear and brakes, these components are regularly put under a lot of stress during repeated take-offs and landings. It is not unusual to find the same type of a problem recurring for all aircraft because fundamentally the basic operation of their systems is the same," says Bordel. "We have developed about 30 models on different ATA chapters and we can predict 80% of faults within them."

Next steps

Much of the useful data that is harvested from aircraft today correlates directly to the maintenance condition of major aircraft systems such as flight controls, undercarriage, engine and the APU. Moving forward, it is likely that operators and vendors will shift the focus of their predictive efforts to the cabin.

"Aircraft interiors are far more complex and automated than they used to be. It is possible to harvest data from seats, inflight entertainment (IFE) systems, and to use that data to improve passenger wellbeing," says Skelton. Teledyne Controls is researching sensor technology pertaining to cabin air toxicity and cabin air condition monitoring systems, by examining pressure, bleed flow and particulate gases.

Improvements to sensors and IoT technology mean that an increasing amount of aircraft data now relates to monitoring parameters in the aircraft passenger domain. Broken lavatory systems, lighting and IFE systems may not prevent operation, yet all are damaging in terms of passenger experience and airline reputation.

Because of limited access, many recurring passenger domain faults can take a long time to resolve, using up a lot of man-hours. It is believed that monitoring more cabin components will reduce mechanics' time in diagnosing, accessing, and repairing these types of faults. Then mechanics can spend more time focusing on repairs that directly impact the operation.

Summary

Airlines are now accumulating more and more data. The more data acquired, the more the associated cost to analyse it. If an airline cannot afford an army of data scientists to analyse the data, it is likely to be incurring data storage costs, for no added value

From a carrier's perspective, aircraft reliability and serviceability are part of the reason to choose a particular aircraft. The argument is that if an aircraft is unreliable and extra investment is needed to improve its serviceability, the airline should have chosen a better-quality product from a different OEM.

It is believed that the OEM should take ownership in improving the reliability of its products, so it should be the OEM that analyses the data and parameters, and then issues service bulletins (SB).

There are diminishing returns with higher data volumes. It will end in the middle with a flexible approach, with more data being collected than today, but less than the maximum possible or available. As more data collection becomes possible and transmission costs get even lower, and use cases for the data grow, then the data volumes collected will naturally increase.

It is assumed that few operators of A320 neo and 737 MAX family aircraft are offloading OEM-expanded data recordings.

The expanded data has value, but this value will not be realised until it is clear and compelling to the operators. There is a lot of talk about the high volumes of data. But what is the benefit to the airline? Do the costs outweigh the benefits? Where is this going to end?

Currently it is impossible to quantify the direct benefit by monitoring 20,000 or more parameters. Yet operational gains have been made by monitoring a select few. AC