Big data analytics and predictive maintenance is a relatively new system of managing aircraft technical faults. Designed to detect the degradation of components and systems in advance of failure, a few examples in practical airline operations are investigated.

Examples of airline use & experience with big data & predictive maintenance

Prognostics and predictive maintenance, which is carried out by processing large volumes of flight and aircraft systems data, using big data analytics and machine learning, aims to predict when rotable and system components on an aircraft will fail or reach a certain point of degraded performance. This will change the approach to managing these components from reactive to predictive.

It has been several years since several industry original equipment manufacturers (OEMs) and system providers introduced the concept of big data analytics and predictive maintenance to the market. There are now a few examples of these systems being fully implemented in airline operations and maintenance departments. The results of these new processes can now be seen.

Current system

The current system for dealing with technical faults and failures on aircraft systems involves a mix of scheduled and unscheduled maintenance.

Scheduled maintenance consists of several types of tasks and inspections at pre-determined intervals. Many of these are structural and zonal inspections, and so can easily be managed within a fixed inspection and test interval. Others are aircraft system inspections and functional tests. While these may also be managed easily in a fixed interval system, system functional tests performed during aircraft maintenance checks can reveal a system or component malfunction. This often involves replacing a rotable or repairable component with one from an inventory of spare components.

Unscheduled maintenance deals with the malfunction or even failure of aircraft components or systems that occur on a random basis during aircraft operation. These have been detected up to now by the relatively simple built-in test equipment (BITE), which sends fault messages to the aircraft's central maintenance computer (CMC). The CMC then issues a fault code, which can either be transmitted automatically to the ground in-flight, or reported to line mechanics and engineering departments via a technical log. IT systems have evolved to provide a faster, and more sophisticated and accurate method of fault message analysis. These include Airbus's Aircraft Maintenance Analysis (AIRMAN) system and now the Airbus real-time health monitoring (AiRTHM) system, which was introduced on the A380. Both provide an electronic link between the CMC fault codes transmitted to the ground and an electronic version of the troubleshooting manual (TSM) for automatic and speedy diagnosis of the fault. AiRTHM can also send requests to the aircraft for further information to assist with the diagnosis.

While these systems reduce the time needed to diagnose system malfunctions and failures, they still rely on aircraft CMC fault messages, and BITE to detect a component fault and malfunction. This reactive way of identifying a malfunction in-flight means that it can only be detected when the component or system has degraded to a level that is easily detectable. In some cases fault messages are not issued by the CMC until the component has completely failed.

In addition to this fault detection system, airline engineering departments are also legally required to operate and maintain reliability programmes. They also have to file reliability statistics.

These reliability data statistics can first provide an indication of which systems and components should be given priority when it comes to using predictive maintenance to eliminate the biggest problems. Moreover, big data analytics can increase the wealth of reliability data and statistics for the engineering department.

Another main issue with the traditional system of detecting and managing faults on a reactive basis is the practical issues of recording faults. This is in paper or electronic technical logs, and informing various departments of the fault, and entering the data at different tims into different systems, and then coordinating responses. If component degradation can be recognised early, and then tracked and monitored, then the process of planning their inspection or removal ahead allows a more coordinated and controlled response.

Airlines therefore need to keep inventories of rotable and repairable components in several locations. The size of inventory held at each station is relatively large, because of the random nature of component failure. Their failure rates and intervals conform to statistical distribution curves, making it hard to predict where and when they will fail. This increases the risk of lengthy and expensive schedule disruptions and delays. The cost of owning or having access to large component inventories therefore has to be borne by airlines to protect their schedule reliability.

The traditional system of rotable component inventory management also

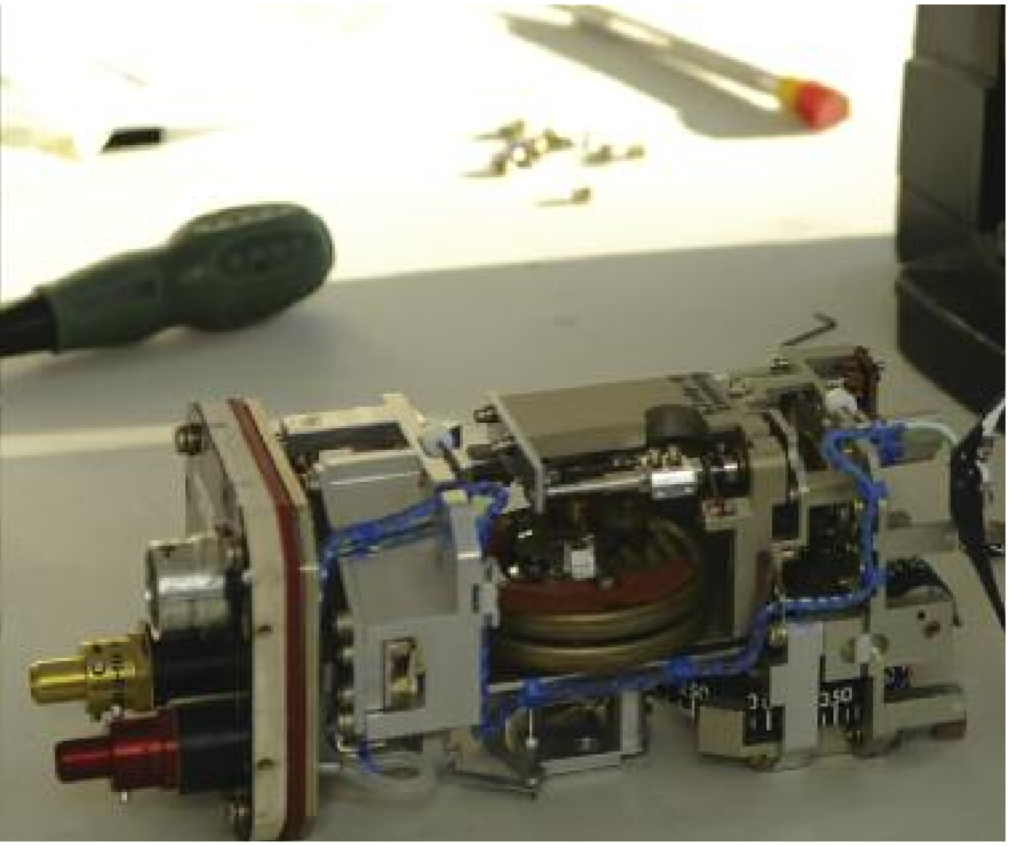

The basis of predictive maintenance revolves around changing the management of components from a reactive operation to a predictive system. The implications of this are savings from several categories of engineering, maintenance, and operations.

presents additional issues such as no fault found (NFF). This relates to the removal and testing of components suspected of malfunction or failure, and subsequently being found to have no fault. This is mainly caused by the fault and malfunction detection systems on the aircraft being relatively simple.

Dealing with unscheduled component and system failures can lead to events such as aircraft-on-ground (AOG) events incurring high costs. These occur after an unexpected system or component failure that causes a lengthy, and so expensive, delay that can involve the temporary removal of the aircraft from service and passengers being provided with hotel accommodation.

Predictive maintenance

Predictive maintenance aims to improve the monitoring of aircraft component and system function by monitoring a larger number of aircraft system and component parameters. The data is transferred to the ground, collated and then processed using big data analytics to provide an earlier indication of sub-component malfunction than the traditional system of detection.

This has several objectives. The first is that it more accurately detects the start of a component malfunction, at an earlier stage than has previously been possible. It also provides data on a larger number of parameters and sub-components.

The big data analytics process therefore analyses a large volume of data to provide more detailed and accurate trending data and analysis than the traditional trend monitoring system. It allows the performance and function of a component to be followed more closely and in more detail to give engineers an early indication of component degradation. The component can therefore be removed at a scheduled downtime for maintenance, such as an overnight or weekly check, which does not involve additional downtime for unscheduled maintenance. This will be a long time before the component has deteriorated further and is therefore likely to fail, causing unscheduled maintenance.

There are several advantages of this. The first is that it clearly avoids unscheduled maintenance downtime or AOG events. It also reduces the cost of component repairs, since the optimum time for removing a component with a particular malfunction can be chosen, rather than incurring higher maintenance costs following its failure. Another significant advantage is that early detection of component malfunction, and the ability to conveniently time removals for repair, can ultimately reduce the two major cost elements of component stock inventories and repairing parts.

OEM Services

In parallel to new generation aircraft generating large volumes of system behavioural and health data on a larger number of parameters and components than previous generation aircraft, OEMs are offering services to collect data from aircraft, collate it and apply big data analytics to offer predictive maintenance. Airbus announced its Skywise Predictive service in June 2017. This transfers all aircraft flight operations and maintenance data from an airline, and uses data analytics to process it. One result is to provide airlines a dashboard for their operations, maintenance control and engineering departments.

Skywise Predictive takes data from the aircraft through the aircraft condition monitoring system to generate a customised report. Airbus applies an algorithm to the data so that it provides a predictive maintenance capability. This allows the function and health of components to be monitored, and recommendations for removal and replacement to be made. The system can be configured to create alerts for component and system malfunction.

Airbus is also providing a connectivity and data transfer channel for the aircraft. This is the flight operations and maintenance exchanger (FOMAX) system. The router is placed in the aircraft, and it collects large volumes of aircraft. This differs from aircraft health monitoring (AHM) and engine health monitoring (EHM) data that has been transferred to the ground. FOMAX transfers about 24,000 times the amount of data, and it provides far better trend monitoring than the older generation AHM and EHM data.

The Airbus Skywise service and product is available for Airbus aircraft, and aircraft built by other manufacturers.

Skywise has already provided cost savings with respect to fuel burn and mitigated unscheduled maintenance on A320 family aircraft CFM56 engines.

In one case, a flag carrier in the Middle East and Africa achieved fuel savings on its 777 fleet. Traditionally airlines have relied on fixed centre of gravity (CoG) and zero fuel weight targets. These overlook critical parameters, which ultimately leads to inaccurate targets and overspend. Airlines use the Skywise system to create dynamic targets that consider all fuel-impacting parameters, such as data from the QAR, aircraft load parameters, and flight path. If all processed correctly, the aircraft's CoG can be reassessed and moved to generate a more fuel-efficient operation.

The airline in question used historical flight data to accurately re-calculate the

Using Airbus' Skywise predictive maintenance service, easyJet was able to avoid unscheduled engine maintenance for the fan module on its CFM56-5B fleet. As an initial benefit of the system, the airline avoided 18 of these incidents in 14 months for a fleet of 100 A319s.

CoG on its 777 fleet, and achieved fuel savings of $850,000 over a five-month period. Similarly, a major US airline identified fuel saving opportunities of $12 million after integrating total baggage, mail, freight and aircraft weights to improve zero fuel weight, and weight and balance predictions.

Skywise has been used by UK airline easyJet to accurately monitor the vibration of the intake fans of its CFM56-5B fleet. By detecting intake fan vibration, easyJet was able to fix the issue at an early stage and during planned maintenance, thereby avoiding unscheduled and expensive flight disruptions, maintenance, engine removals and repairs. easyJet avoided 18 such incidents over a 14-month period with a fleet of about 100 aircraft.

Airbus says that one of the biggest potential savings that can come from using Skywise is a more accurate prediction of the quantity of spare rotable and repairable components required. This will lead to large savings, particularly in inventories being held at outstations. Predictive maintenance can also help airlines avoid borrowing components from other airlines at outstations, or paying to be part of pooling agreements.

JAL experience

Japan Airlines (JAL) is one of the first major airlines to implement a big data analytics and predictive maintenance system, and realise positive results. JAL has moved from a react and repair to a prevent system for part of its fleet, and has done this with the cooperation of Boeing's Data Analytics.

The airline began operations with the 787 in 2012, and now operates a fleet of 40 aircraft. Like all airlines, JAL was subject to system and component failures, which caused delays and cancellations, and incurred high costs.

In cooperation with Boeing, JAL launched a zero zero 100 programme, aimed at using predictive maintenance to have an operation with: zero irregular operations; zero engine in-flight shutdowns and defects; and a 100% on-time departure record. Achieving this requires all recurring and non-recurring reliability problems to be dealt with.

JAL's second objective with the programme was to empower its maintenance and engineering (M&E) team with advanced tools to efficiently solve problems, and to let the engineering and maintenance technicians know that they are having a positive impact on the airline's operations. Up to now, JAL has operated in a reactive manner to technical issues, and aircraft component and system malfunctions. The problem, as with all airlines, is that maintenance does not prevent flight schedule interruptions.

JAL decided to change from a reactive to a preventative maintenance system for aircraft components and systems, without compromising safety. The basis for the project was the use of advanced data analytics and data science to provide predictive maintenance, the last line of defence to predict failures. It can also be used in addition to established maintenance programmes, such as design modifications of components to detect deterioration, and of assemblies to avoid disruptions.

The entire predictive maintenance system requires a team of experienced M&E technicians, in addition to the data science, data analytics and machine learning to provide all the necessary elements. JAL provides the M&E teams, while Boeing provides the data science and machine learning elements.

These teams decide what action to take, and which member of each team is responsible for a particular element of the maintenance. Boeing started working with JAL to devise the first element of predictive maintenance. Boeing used advanced data analytics to get additional insight and additional cost savings beyond traditional methods that had been applied in the past.

One of the first issues to be tackled was a recurring issue of reliability problems with power control units (PCUs) for the wing spoiler actuators. The root cause of this problem was found to be corrosion of motor wire coils inside the PCU unit. Spoiler activation was therefore not completely reliable. Diagnosing the problem was expensive, because the airline had to wait until the units had failed. The aircraft had to be grounded for unscheduled maintenance, about 30 man-hours (MH) of labour were used to inspect it, and spare units had to be fitted while removed units underwent an expensive repair.

With the 787 generating more data than previous generations of aircraft, Boeing and JAL had access to a huge amount of data, including data that just related to the activation of the spoilers. Using the 787's continuous parameter logging (CPL) data, JAL's big data analytics team collaborated with Boeing on using machine learning and advanced analytics to find the effective predictive algorithm for the spoiler failures.

The data used came from the continuous logging of thousands of aircraft system and component parameters from a fleet of 30 aircraft for three years. This provided more than 220 billion samples of information, which was filtered by the Boeing analytics team to try to find a key predictive signature that would allow JAL to predict spoiler degradation. This took more than two months of work.

The analytics team took 220 billion data points, and processed them with machine learning, to find a way to diagnose the aircraft's corrosion problem. The Boeing team developed an algorithm to detect when one of the two copper wires has corroded to the point when an open circuit failure develops. This compares to the established fault detection system; where there is only a CMC fault code displayed on the flightdeck MCDU when both wires have stopped working.

When one wire stops working this immediately causes a unit failure and the need for an expensive unscheduled maintenance event.

This recurring problem gave Boeing and JAL the impetus to build an algorithm to detect when the degradation of one or both wires starts. By monitoring the appropriate trends and function, it would be possible to devise an appropriate interval to manually inspect the PCU and coil wiring, and perform maintenance before the degradation becomes critical.

Without this new system of scheduled inspections, line mechanics would have to perform three inspections of the PCUs in situ on the wing structure, resulting in high maintenance costs.

By performing thousands of mathematical experiments, Boeing is able to identify operating patterns and component behavioural patterns that were previously invisible. This makes it possible to perform maintenance action before problems actually arise.

The first stage in this whole process to derive information from raw data is machine learning, and anomaly detection techniques to solve the spoiler PCU problem. Boeing found a pattern of abnormal unit operation out of the 220 billion data sets. These data sets include a reasonable number of patterns of data, which, when plotted and trended, reveal both normal and abnormal component behaviour. It is from these patterns that the point can be determined at which degradation and then malfunction starts to occur. The patterns that were established revealed a common feature that was shared by all the degraded spoilers, identified from a piece of flight data taken from an event that lasted for a few seconds on each flight.

The team was able to detect a difference in spoiler actuator behaviour of less than one-tenth of a millimetre (mm) in the actuator position. That is, there is a difference of 0.1mm between healthy and degraded spoilers. This difference is extremely difficult to detect with the naked eye, which is why machine learning is needed to identify behavioural patterns. Engineers can say which patterns are relevant to the technical problems that need to be solved. The large volume of data sets and parameters means that there will be a large number of data patterns that do not apply to particular problems.

In the entire process, there is an intersection of three important areas of expertise among the different teams of Boeing and JAL staff: design and engineering; operations knowledge; and data science.

The system found the needle in the haystack: data that gave Boeing and JAL the solution to detect abnormal spoiler behaviour.

To reach JAL's goal of zero zero 100, the airline needed to identify the source of the problem, but also needed to predict when it would arise and so prevent it from happening again.

Engineers know that most parts are neither in an operable nor failed state, but instead are on a continuous path of degradation, from the moment they are installed to the moment they fail, as degradation starts at a very low level almost from installation. The ultimate goal was to identify the optimum point

Japan Airlines cured the recurring problem of corroded wiring looms for the PCUs powering wing spoilers on its 787 fleet with the use of big data and predictive maintenance. This initial project can be extended to a range of reliability problems across a fleet.

on the degradation curve to remove the component for the most economic removal and repair.

By identifying this 'sweet spot', Boeing and JAL can identify which spoiler function requires attention and assess the risk of failure that would lead to an unscheduled maintenance event.

Overall the system allows proactive inspection of the degradation of a specific component or sub-component. This avoids random failures and non-scheduled maintenance. This was the appropriate time to test a new type of management for components and unscheduled maintenance events. JAL asked Boeing to use the devised model to identify which spoilers in the fleet were degrading or degraded, through regular inspections.

The algorithm that was developed in the process searched for units which had degraded functionality, and were therefore at risk from failures. The algorithm also predicted five new cases of degraded PCU coils on five different aircraft. The JAL maintenance team then inspected the relevant PCUs to see if the algorithm was accurate. A few weeks later the five predicted spoilers came back as confirmed as having been degraded, but not yet having developed a fault. This proved the system works.

If the degradation had not been detected, they would have caused PCU failures, unscheduled maintenance, and AOGs. Each one would result in 30MH of overnight inspections per aircraft. It is now also possible to target specific aircraft, and spoiler positions. This clearly predicts degradation on PCUs across the entire fleet, prevents all potential failures and all the subsequent consequences.

In this case, failures have been reduced by 68-85%. This item brings JAL closer to its goal of zero zero 100 operations. The added benefit is that the process has made the overall team more confident.

The project of predicting spoiler degradation with more than 90% accuracy was worked on for five weeks. Boeing has continued to work with JAL on predictive maintenance because of the need to find prognostic solutions.

The system adds another layer of complexity to the work, and highlights a big challenge: the ability to apply algorithms and results to a large number of other problems, including ones not seen before. This is referred to as the process of generalisation.

Despite this initial success, there is no silver bullet: airlines continue to collect flight data differently, in different operating environments, and to conduct maintenance differently. There is now a second version of the algorithm, which is used to predict inboard and outboard spoiler degradation. The algorithm has also been applied successfully to other fleets.

JAL is now working on many other areas. First, however, it has needed to build a foundation of trust so that its M&E engineering team feels confident in the predictive model. A result of this is that the engineers have the confidence to get the maintenance team to inspect the component.

In addition, faults detected are now confirmed as being correct, thereby eliminating the problem of NFF and generating savings.

Scheduling maintenance of a rotable or repairable component several weeks in advance has many implications, all related to financial savings.

Boeing now expects to get more real-time access to data such as quick access recorder (QAR) data, so it hopes that it will be able to investigate a lot more opportunities to do predictive maintenance. Ultimately the system will allow airlines to transfer from a repair to prevent component maintenance management system.

The future

Ultimately, once big data analytics has been developed to eliminate component and system failure and reliability problems, then virtually all unscheduled delays and AOGs could be a thing of the past. All components could be managed in a preventative way.

The full benefits will be smaller maintenance control and line maintenance departments, smaller rotable and repairable component inventories, reduced expenses to cover for AOG situations, smaller operational control and line maintenance teams and departments, and even higher aircraft utilisation and smaller fleets. The resulting cost savings are not yet fully understood or realised, but a few carriers are just beginning to get an indication.

In addition to the maintenance and engineering savings, there are also knock-on or consequential savings that relate to flight operations. The example of re-adjusting aircraft CoG to realise fuel burn savings gives an indication of what will be possible when carried across an entire fleet. Data analytics in flight operations can also provide more accurate cost information to operations and maintenance control departments.

More airlines are implementing big data analytics, and predictive maintenance technologies. The adaption of airline maintenance IT and flight operations systems to accommodate this change will be examined in a future issue. AC