How big data is influencing rotable part inventory and predictive analytics to minimise aircraft on ground situations. By analysing part failure rates and stock inventory operators can insure inventory is in the correct location and part failures are minimal.

Using big data analytics in rotable inventory

Big data systematically analyses data to extract trends from an appointed number of very large data sets, and then processes the information with statistics. This has the potential to monitor behavioural patterns of components and systems and can be used to predict when and where an aircraft component will fail.

This will transform the management of rotable components from a reactive process to a proactive one. It will allow airlines and operators to mitigate costly aircraft-on-ground (AOG) situations, and reduce aircraft downtime because a component change can be managed and planned before it fails. In addition, a higher level of part availability can be achieved without a disproportionately large quantity of stock having to be held.

Inventory management

Airlines and MROs plan their parts and component inventories at a particular level based on the probability of failures. High stock levels often have to be held at several locations to prevent protracted AOG situations. The aim is to balance most component failures without holding excessive stock.

The amount of rotable stock held is carefully calculated to minimise the unavailability of needed parts. It is estimated, however, that more than 40% of rotable inventory is surplus. Many airlines and organisations hold inactive and obsolete inventory items that have depreciated and hold no useful value.

There are several ways that big data analytics can reduce and optimise rotable inventory.

AJW

AJW provides component and rotable inventory support solutions.

It harvests data from its customer base and logistic suppliers, has access to data from repair vendors, and can source reference material from its own supply facility, AJW Technique in Canada.

AJW's access to a large number of data sets and data collection points from many vendors, means that it can analyse an extensive set of information. It also has access to historical data at a component and serial number (S/N) level that helps it analyse the past operational environment of a component.

"We work with the macro data sets provided by the airframes, the granular data sets for an individual component, and the criticality of that component and the disruption it causes to a network. If you have a seven-day turnaround to replace a component, because of large volumes of data to help forecasting, it takes the pressure off your logistics and spares environment," says Boris Wolstenholme, chief strategy officer.

"We have spent the past 25 years looking at this problem. Airlines have used different approaches, including outsourcing their supply chain solutions, flooding inventories with stock or developing in-house engineering solutions. This can result in an insular world because the data is often only from their own environment and therefore restricted. We benefit from seeing the macro data from many environments and operators. This allows us to spot trends."

The methodology is to use the many data points to accurately forecast component failures and then to use that information to develop an intelligent inventory-planning solution. To evolve the forecasting methodology further, what is happening in real time is compared to what is expected to happen.

If a problem with a component occurs four times when it is expected to only happen twice, more analysis is needed to determine what is driving its failure rate.

There are many instances where AJW supports the same aircraft through various operator ownerships or leases, so AJW holds more comprehensive data. This means AJW benefits from a better overview of that individual aircraft.

"The level of detail we have enables us to look at each aircraft's reliability level, and the service bulletin (SB) and modification status of each component on that aircraft," says Wolstenholme.

"We take the macro statistical analysis and compare it to what is expected to happen. The granular analysis can be compared to what is installed on the aircraft. This includes the modification standard, reliability, and when each part was last overhauled or repaired. Bringing all this together, and comparing it to what is happening in real-time is when predictive maintenance is used," adds Wolstenholme.

Wolstenholme explains that, for example, sisterships that have been in the same fleet since the day they came off the production line can behave completely differently.

"To overcome this we focus our data-sets on the criticality of the components, so that we are looking at the economics of what is impacting the operator. This includes the operator experience and the despatch reliability of the aircraft, what is impacting AJW's and the airline's maintenance cost base, and the ownership of inventory," says Wolstenholme. "Then we need to find out what is impacting the total inventory, and how much inventory needs to be owned."

Talking about the bigger picture, Greg Hoggett, technical director of AJW adds "A lot is affected by airlines' maintenance practices as well. You might have completely different maintenance practices in terms of planning for each of the similar products you are operating. Detailed analysis of the data sets and all the variable inputs going into the jigsaw need to be processed, before informed calculated decisions can be made.

"The system can then make decisions down to a component and S/N going into a workshop. It can go to the level of what degree of workscoping we require on each part, and what is the failure mode based on our historical data. With forward-looking data we can decide if we need that inventory," says Hoggett.

Minimising failed fits and premature failures is another utilisation of big data to ensure the right work scopes are

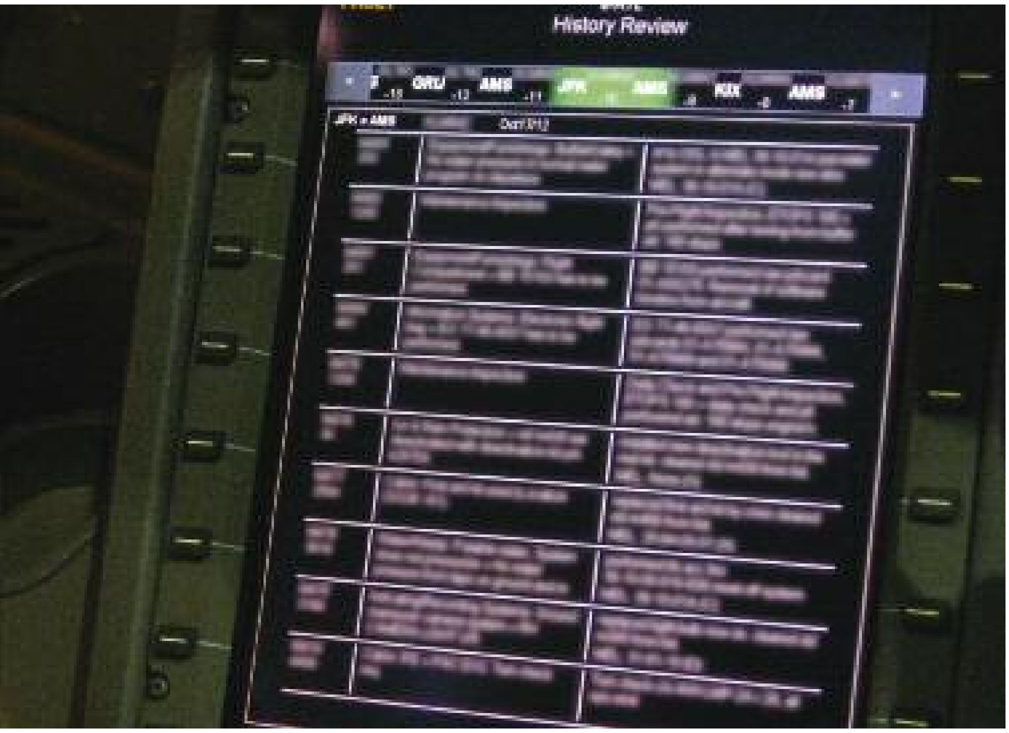

Large volumes of relevant health data from the aircraft assist in taking a pro-active approach to rotable component management.

completed to drive reliability into those components.

EU261, the penalties paid for disruption to operations, are another factor for commercial operators, so on-time performance is a critical key driver for them in terms of the partners they work with.

In that respect AJW has an operational disruption tool to enable it to support its customers' on-time performance by considering the fleet's operation, how many aircraft are due to retire, and the technologies on board replacement aircraft.

AOG events

Looking to the future of big data and today's constraints, Wolstenholme says that, "without the technology being embedded at production level or being proven to be effective in predictive maintenance capabilities, you are looking at a long time before we have zero AOG events."

Some AOG events are caused by an in-service effect like impact damage, which leads to an unexpected requirement for a component such as a fan-blade, a fan-cowl, or a thrust reverser. According to Wolstenholme, however, extreme AOG events can be regarded as an exception. The wider AOG challenge relates to components that fail outside of extreme environments.

From a rotable component perspective within the ATA chapters covered by ordinary flight hour (FH) contracts, Wolstenholme believes non-extreme failures can be minimised. "We monitor all the components within the ATA chapters that we support by triangulating three data points: customer data, airframe data and our data.

"Some parts we do not see in our own system, so we do not get the full picture. But we monitor all of them within the ATA chapters," continues Wolstenholme. "We have contracts where we have full supply chain support solutions. We have the benefit in these instances of seeing the full data set and what causes all the AOGs and critical events and despatch delays. It is also down to effective maintenance planning by the airlines'."

Service level

Many airlines are running complex route networks, with short turnaround times and complex line base maintenance activity. When they have a defective component, they have a limited time to remove and replace it, or repair the defect. How can they predict where the aircraft is going to be at any one point?

"What can happen is that the aircraft must be re-routed to another destination, so often you end up chasing defects around the network," says Wolstenholme. "Lots of our customers demand parts to be available within 24 or 48 hours, even though they have a routine seven days to rectify the defect. The 24 or 48-hour availability gives them the flexibility to change the component."

The higher the service level an operator stipulates, the higher the costs that are driven into logistics to meet expedited targets. This drives costs into stock ownership, so an assessment of an operator's network must be done.

"Airlines are trying to lower their costs while maximising component availability, and it is a fine balancing act," says Wolstenholme. "We see some airlines owning surplus inventory. Our systems, expertise, methodology, logistics, data and our technical capability revolve around the components rather than the airframe," Wolstenholme continues.

Looking forward

AJW will continue to collate historical data, and there will be more pressure on airlines to improve data quality, because airlines only focus on capturing the relevant data.

Then the data release becomes the challenge. Data capture and release can be poor, so there will be pressure to improve data capture and release quality. Evolution of the systems that analyse historical data will be important, and the artificial intelligence that supports that, and compares it to real-time events, will improve forecasting and planning.

Predictive maintenance gives real-time insight into what is about to fail, and information to avoid that despatch risk to the aircraft.

Breaking this down, all these elements complement each other. We therefore cannot say that the future is all predictive, because predictive relates to only a small sub-set of components, and a small sub-section of aircraft on today's fleet.

Predictive maintenance includes a degree of statistical modelling. Once applied, it will give a degree of predictability, which in the sense is a much smaller data set.

Nevertheless, it has an important role, and its ability to predict what will fail, in the terms of how the aircraft system is behaving, represents a real change in how to manage it. That is why it is so exciting, because this has been the missing link. The quicker that can progress the better.

"If a component is expected to stay on wing for, say 10,000FH, that will be three years in the operation of a narrowbody aircraft," says Wolstenholme. "Before you can start analytics, the part is predicted to come off. It comes off, and then the challenge is 'what do we do with it?' Because it has not failed at this point, what repair work scope should be used? So for this example do we do an overhaul?

"Then the part is installed on the aircraft, and we monitor how long it goes back on-wing again. We can then determine if that work scope has been effective. It could, however, be three years until we next see that S/N fail again," adds Wolstenholme.

Ideally, predictive maintenance indicates where and when a part is going to fail. It comes off, the work scope is small because you can just replace one part. Replacing the part before it is damaged, lowers your cost and maintains your reliability. You do not know if you can actually maintain your reliability until you track the part again for the next few years.

What is really required are some big fleet data sets to analyse. The evolution is to move away from tracking components at P/N level, to S/N level. For example, an S/N at 5,000FH may be removed from one aircraft, while a different S/N is reliable for 10,000FH.

"It is important to have the data sets, and collate and aggregate them, and get to the point where we are customising the work scope for those components at shop level," says Hoggett. "This is because you are adding the most value at the right time to that component to keep your maintenance events down to a minimum. This is introducing the reliability with the least amount of operational disruption."

Historically, efforts have been restricted to rogue unit analysis, when a 'rogue unit' just fails every time it is installed for no particular reason. It is important to understand each airline's rogue unit policy. This is how long or how many times the same S/N has been seen or how long it has been on-wing.

Determining what are premature failures, and what you do with such a component needs to be decided. Do you scrap it or does it go back to the OEM?

It will be interesting to see how that language changes as predictive technologies improve. This is because you are removing something from the aircraft before it fails. Therefore no fault found (NFF) metrics will have to be revised and reconsidered in terms of rogue unit policy. Technologies are now developing at such a rate that it could be possible to eliminate these events and problems.

The point is then when will this be supported across the airframe fleet? It is necessary to look at the retirements of the existing fleet that are due.

An aircraft is in operation for about 30 years. Aircraft built before this technology was enabled will not benefit from predictive maintenance.

"We are still working towards developing that philosophy and technology across our current data sets," says Hoggett. "There is a gap between a modern fleet's capability to give those predictive outputs, and what we do with the legacy fleet without the ability.

"Do we try to bridge that gap between this kind of nirvana of S/N level of predictive maintenance; and the current installed fleet, which still relies on us collating and aggregating data in the way we have always done it?" finishes Hoggett.

Wolstenholme says that: "Another benefit of predictability is it gives you that real-time assessment of the system's impact on a component failure.

"You have to remember that the component failures within an aircraft's system are independent of each other," continues Wolstenholme. "The benefit of predictive maintenance is that because it can monitor, for example, the hydraulic pressures or mounting temperatures or vibrations in real time, you can see what is affecting the inputs to that particular component.

"With traditional analysis, you can see that a valve has failed but you do not know why. Predictive maintenance gives that insight into why it might have failed," explains Wolstenholme. "Data is key to enabling you to assemble the puzzle from all the individual elements,

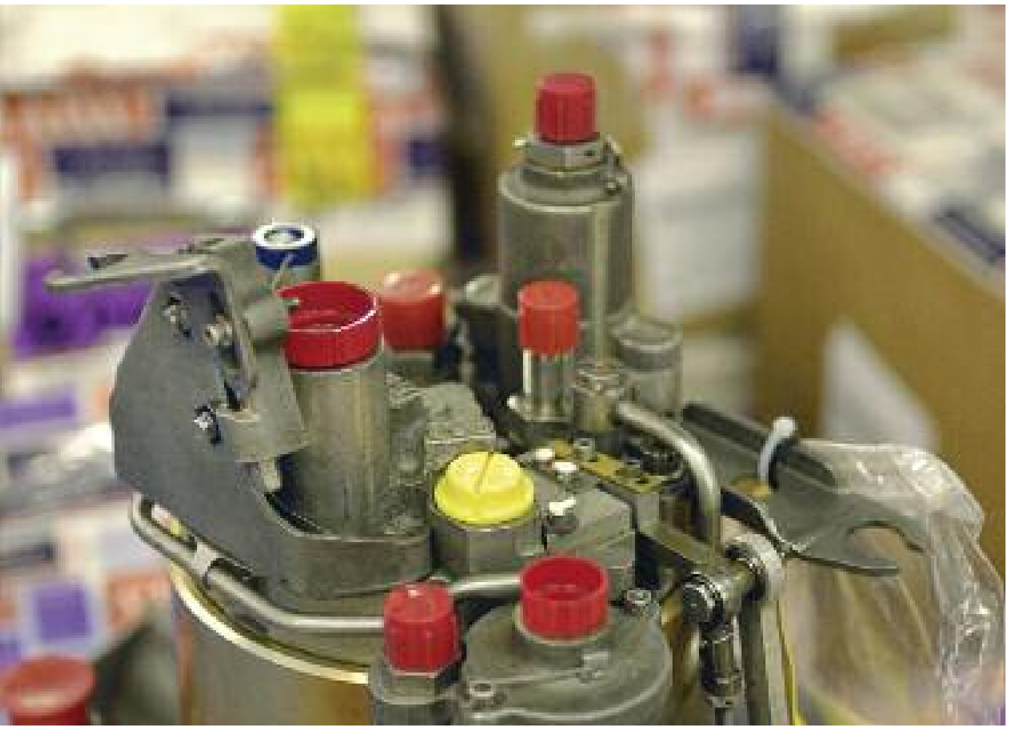

Big data analytics assists with planning rotable component removals, rather than them occurring on a random basis that results in an AOG situation.

such as the operational environment, the airline's engineering maintenance practices, the component's history, the repair work scope, and the systems that it is sitting in. Then it all comes together."

ARMAC

ARMAC has a rotable optimisation software product called RIOsys, which is also used for managing consumable, expendable and repairable inventories. It uses the data within a user's Maintenance and Engineering (M&E) IT system to make a decision support layer.

Since the M&E systems are configured to support transactional activities, RIOsys uses data to make decisions about planning inventory, which considers multiple stock locations.

"RIOsys does not purely set stock levels. The system can monitor various component activities. As parts are removed from aircraft and events occur, we can identify shortage, surplus and parts that become scrap," says Michael Armstrong, chief executive officer at ARMAC. "An airline may be below its ideal stock level, but it is also constrained from a cash perspective. We help customers make smart decisions.

"We have another product, called Inventory Planning Capability Framework (IPCF)," continues Armstrong. "This system puts inventory management planning capability into four pillars: strategy and policy; people and organisation and processes; systems and data; and finance control."

It is uneconomic to randomly move thousands of parts, because they represent a risk and a cost. There may be another part in another location that is best suited for a requirement. The problem is knowing which parts to expedite and which to repair if there are thousands of them in different locations.

It is important to avoid moving parts for marginal benefit. For example, if you could wait a week for the next part to come into repair. These are the things that IPCF can take into account.

"We do not do predictive analytics," says Armstrong. "We optimise inventory. Predictive analytics is all about taking telemetry from the aircraft and using it to predict and pre-empt a part's failure using historical information. What we do is we manage the four types of demand in aviation: planned, preventative, probabilistic and predictive demand.

"Predictive maintenance converts probabilistic demand to planned demand," continues Armstrong. "In other words, you are trying to make uncertainty more certain."

While the parts will still have a mean-time between removal (MTBR), big data analytics can be used to convert the probabilistic demand into a planned demand.

Looking at it from this perspective, and when talking about inventory planning, there are three distinct layers.

The first is top-level strategic planning: what is your supply network, what service levels do you want to deliver, what is your investment, how do you source materials, what level of stock do you pool, what percentage do you own, what do you lease, what are the levels of exchange on the market place, and what is your approach to all of this?

The second is tactical planning, which addresses imbalances in the network, such as a spike in demand that leads to lots of parts in repair. It is about managing your operation and preparing for these unexpected variances in demand.

The third is fulfilment: this is the planning horizon. You know you will need these parts in the next two weeks, so where are you going to source the items for the particular demand areas and requisition them on your system? It is about fulfilling real-life demand.

According to Armstrong, predictive maintenance falls into and interacts with the third level of fulfilment planning. "Will I have a predictive demand in the next two weeks? So where ARMAC sees the value of predictive maintenance is as enabling a more proactive approach to the supply chain in terms of fulfilment," says Armstrong. "That is where our solution 'dovetails' into the predictive maintenance solutions."

ARMAC uses big data to assess the rotable parts inventory requirement. There are other uses for these analytics and different types of big data, however. You have the information and data generated by the aircraft. Within the industry are the likes of Skywise and AVIATAR, which are revolutionising datalakes so that you can have accessible data in your enterprise resource planning (ERP) system, or your MRO IT system.

"We carry out analytics to optimise inventory at a strategic planning and tactical level. At a time of fulfilment, our solution provides a decision support layer for strategic tactical and fulfilment planning," says Armstrong. "We use big data because it is a highly complex supply chain that differs from a standard purchase-to-consume model.

"When you are looking at a repair to re-use a component, you cannot just use re-order level and MRP, which is stochastically used in manufacturing and production types of systems, to manage your stock," says Armstrong. "Big data analytics focuses on the supply chain. What parts do I have and where are they right now? Planning takes the different drivers of demand, for example the removal history per check type, per utilisation and per line maintenance demand. Even in a scheduled task was it routine or non-routine?"

How do you decide at a strategic, tactical and fulfilment level? With predictive maintenance, the consensus is that analytics can remove uncertainty, and reduce costs by creating a more foreseeable environment. This technology

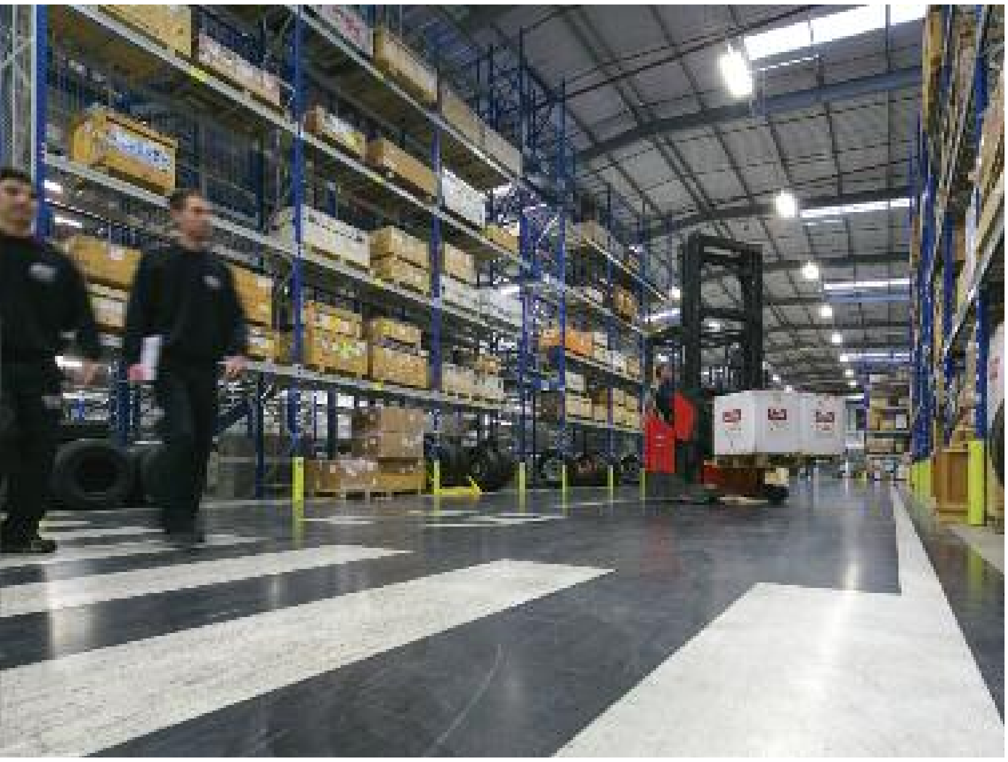

The use of big data analytics and predictives in rotable inventory management is still in its infancy, and used by a relatively small number of specialist providers.

is in its infancy, and few airlines use it.

Early adopters, like AVIATAR and Skywise, for example, have their support customers. The scope to carry out predictive analytics is still quite limited, given the range of parts that are in use.

"I do not believe we will get to a 100% success rate with predictive maintenance. We will always have those four Ps, and even today more than 80% of removals are probabilistic. Whether we will get to the point where the removal of most of those components is predictive, is uncertain. We are at the leading edge of this technology," says Armstrong.

"There is a need for more early adopters of the technology," he adds. "Component failures happen randomly, and result in an AOG or delay. If you can prepare for that in advance you have cost savings, EU261 savings, supply chain savings; and you avoid an AOG shipment, exchange or purchase.

"If need to buy, move or exchange a part urgently you incur more costs," continues Armstrong. "So anything you can do to get some advance warning will absolutely reduce your supply chain costs. If you have a part that will also deliver a high degree of predictability, you do not have to hold it as close to where the demand is going to occur, and you can also hold less stock."

Predictive maintenance does not change the fact that component removals occur. It gives you early warning that what was a probabilistic change becomes certain within a particular time.

"Our system can reduce rotable part inventory, even without the predictive maintenance," says Armstrong. "Just by using probabilistic demand, our system can dramatically reduce investment by streamlining the supply chain."

Armstrong adds that predictive maintenance allows further optimisation of the supply chain. For example, if you have a five-day warning of a failure, you can plan and set the parameters of an optimisation solution to say, 'how much do I need to hold on the shelf when my delivery time is X number of days'. This will adjust the level of stock that needs to be held.

Predictive maintenance will not stop the demand occurring, it just allows more time to react so investment levels can be reduced. It is believed that if operators and MROs use the data they have today by using probabilistic maintenance and planning better with it, investment can be reduced by 20-40%.

Another challenge is who owns the data, and how prepared are they to share it? Who owns the consequential analytical results of the processed data?

Armstrong says, "Johannes Bussmann, chairman of Lufthansa Technik, says we should not compete over the data. He says we should all share it, but democratise it so once it is shared companies compete over the analytics they carry out on that data."

Teledyne

Teledyne Controls provides systems that collect data from the aircraft. It is probably the biggest company that does it on more aircraft than anyone else.

The Quick Access Recorder (QAR) on the aircraft is a data concentrator that is connected to all the different systems and sensors. It collects that data and packages a minimum subset that is required by European Aviation Safety Agency (EASA) and the Federal Aviation Administration (FAA).

"Over the past 20-30 years, the QAR, and similar items on the aircraft have been acquiring data. We can record twice, four times or even eight times the volume of data on the QAR than on the flight recorder," says William Cecil, business development director, aircraft data services at Teledyne Controls.

Teledyne also pushes live snapshots in flight. It will download as much as 10 megabytes (MB) of data after landing. In comparison, small quantities of data are sent via the aircraft communications and reporting system (ACARS) connectivity system in flight.

"ACARS generally sends 2-10MB per aircraft," says Cecil. "Sending data via ACARS in flight is expensive.

"You do not need to receive data for predictives in real time," continues Cecil. "For example, two-thirds of aircraft have a wireless, ground connectivity system that uses WiFi or cellular connectivity. So instead of 2-10MB after each flight, aircraft are now downloading 500MB to 1 gigabyte (GB). It is typically 1,000 times more data, and the critical thing is it is time-series data, captured four, eight, and 16 times per second so it is not a snapshot-style text message," says Cecil.

It can take just a few minutes to download this data using a cellular network. This is an autonomous system.

Teledyne conceived the concept of the wireless QAR about 20 years ago, and added cellular technology to it.

The data that is offloaded from the aircraft is binary, and needs to be de-coded. Some airlines have got together with providers like Teledyne that do the decoding part for them. "It is like pre-processing your data. It needs to be done to make it understandable," says Cecil.

Some components, often on older aircraft, do not have any sensors, and so are unable send data. Teledyne's solution can expand recordable data sets.

Delta Airlines retrofitted sensors to its existing fleet to expand its predictive data analytic capabilities. Delta can now harvest, monitor and predict much more accurately, so it has more control running its own logistics and supply chain, giving it better forecasting capabilities.

Airlines are starting to open up to allow more data to be used outside of things like flight safety. For example, in

Germany government regulations are that all aircraft must have a wireless system, and that when data is transmitted to a flight safety department, copies must be despatched to the relevant government department too.

"The Chinese have a government flight safety analysis programme for the entire fleet of aircraft on the Chinese register, but it is unlikely that the FAA will do this any time soon," says Cecil.

"We are setting up mechanisms so it is easy for the airlines. 200 airlines have 1,400 aircraft equipped with our systems. Channelling this data to the relevant flight safety department allows airlines to re-direct or duplicate the flow of data to multiple third parties of their choice."

Expanded data set acquisitioning systems on the 737 MAX and the A320neo have been introduced by the OEMs via an on-board server. The servers acquire data at a much higher rate than the QAR. Once measured in 500MB to 1GB per month, more than that can now be achieved each day or each flight.

"The new systems increase the amount of data by a factor of 50. This is the OEM system," says Cecil. "There is hype about the amount of data generated by the aircraft. The reality is that most of that data lives and dies on the aircraft.

"There is a law of diminishing returns. If you jump from 5Kb per flight to 2MB, and then to 10MB, it is obviously better," continues Cecil. "You can jump from 10MB to 100MB and then to 10GB, or even one Terabyte. The incremental benefits of having higher levels of data are that there is a higher level of resolution. With predictive analytics you are predicting that something will fail a week, or even a month in advance. But how far in advance do you actually need?"

If you have already got a good line maintenance process system for troubleshooting, for example the pneumatics or bleed air system on the aircraft, then this simple process can continue to be used. "If there are four potential components that might have failed, and it only takes the line mechanics five minutes to find out which one it is, and you predict this problem a week in advance, do you then really need predictive maintenance to identify which one it really is?" asks Cecil.

If something is predicted to fail in the next couple of days or weeks, an airline will schedule line maintenance to look at the problem on an overnight stop. They are not panicking, because they do not have to turn the aircraft around or have an AOG.

One of the big benefits of predictive maintenance is that when a problem happens and a component is needed, airlines can position the aircraft to the airport where the part is located or schedule the maintenance to be done at an overnight stop where the parts are.

Another example is that airlines may have some maintenance locations with a specific set of expertise and specialist mechanics, while at others maintenance is outsourced so they do not know what they will get. Being able to predict a failure therefore allows you to move your aircraft to where specific talent is making for a more economic resolution.

Teledyne also develops software for ground applications, such as one which translates binary data into useful data. "We are already able to predict certain failures. An example is ATA Chapter 27, where there is a problem with flap symmetry, which is common on some aircraft," says Cecil. He adds that the most useful way forward is to focus on the highest value problems.

"The most problematic systems on a 737NG are pneumatics: ATA Chapter 36. So we are focused on this and four other chapters right now, and expanding more. We have done quite a bit of work on other chapters too," continues Cecil. "Not all aircraft are the same. There are different generations of aircraft, and different levels of data acquisition." AC