The complex task of manually planning schedules, assigning individual aircraft to flights and managing flights on the day operations is now made simpler with specialised software. These systems allow airline schedulers to plan operations months in advance and make automated changes to aircraft assignments.

Systems for aircraft scheduling & flight operations

Planning aircraft operating schedules and managing day-to-day aircraft operations is a complex task. This was traditionally done manually, but software products are available to handle this process. Airlines clearly seek to improve aircraft and maintenance-check interval utilisation, despatch reliability, on-time performance, flight connections, crew productivity and passenger service, as well as to reduce delays, schedule disruption and operating costs. Aircraft schedule management is a three-phase process: long- and medium-term planning of fleet assignment prior to operations; tail assignment of individual aircraft to specific routes a few days prior to operations; and real-time management of aircraft operations as they occur.

There are several suppliers of software for these two main tasks, including SITA, Lufthansa Systems and Sabre. Planned and actual aircraft operations and schedules clearly have an impact on several other airline departments and functions, such as fleet planning, maintenance and engineering, crew scheduling and rostering, and aircraft dispatch. Software to manage aircraft operating schedules clearly has to interface with other airline departments.

Planning schedules

Planning aircraft and fleet operating schedules is complex because of the many factors that influence and limit which of the individual aircraft in its fleet that an airline can use for each route.
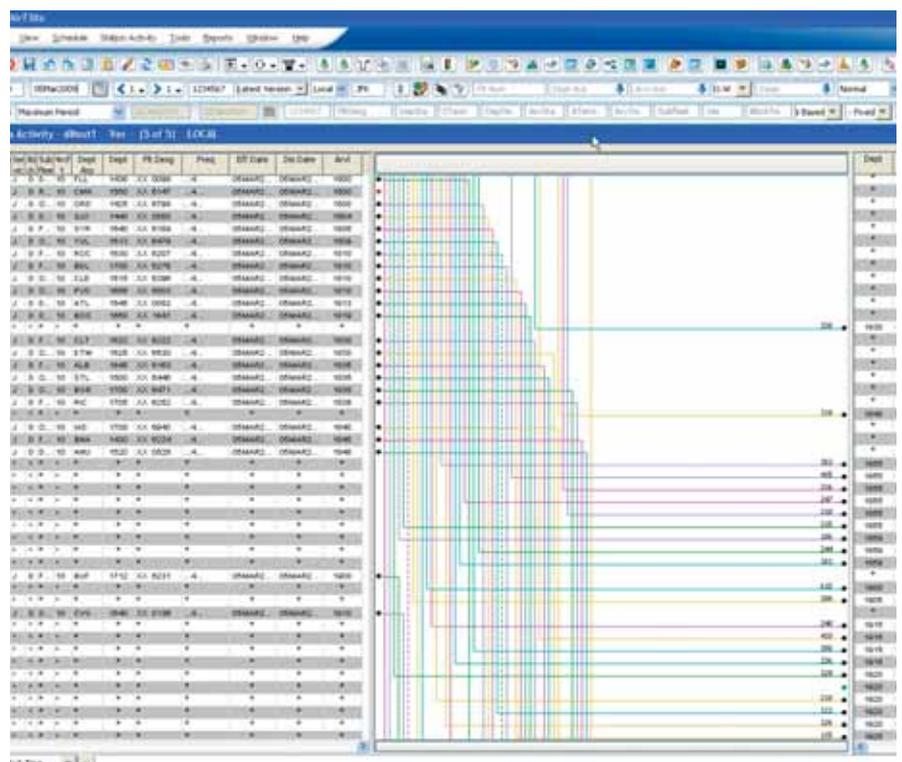
Airlines clearly have long-term fleet plans regarding the number of different types they require in their fleet, and the seating configuration of each one. This determines the availability of appropriate-sized aircraft in relation to expected demand. Airline schedulers assign individual aircraft, or tail numbers, to specific flights closer to the day of operations, according to the expected passenger loads on each flight.

There are often two or several sub-variants of the same model with different seating configurations in each fleet: aircraft with smaller and larger premium-class cabins to cater for different demand levels across the network, and therefore different economy-class cabins. Sub-types and sub-fleets mean that different groups of aircraft will be limited to certain routes.

Airline schedulers then have to consider the operational characteristics of the airports on the routes of their network, including: the airport elevation; the length and slope of runways; the certified landing category of each runway; the historical record of ambient temperatures throughout the year; and any special noise and emissions requirements (for example, Scandinavian airports charging higher landing fees for NOx emissions above certain levels).

Airline schedulers then also take into account the length of the routes on their network, as well as typical flight, taxi and en-route delay or holding times, which will vary for each aircraft type.

All these factors then have to be considered against the operating characteristics and performance of each individual aircraft in the fleet, so that

Sabre's AirVision scheduling software has the facility to display flights arriving and departing at a hub airport. This allows the scheduler to plan a connecting hub operation that allows sufficient connection time between all flights and sufficient turnaround time for all aircraft.

airline schedulers know the routes on which each aircraft can and cannot be operated. Even though a sub-fleet of an aircraft type may have the same seat configuration, individual aircraft may have different levels of operating performance, because they may have a lower-rated engine model, reduced exhaust gas temperature (EGT) margin, lower weight specification or fuel capacity, or some other characteristic that limits operating performance.

Once schedulers have grouped aircraft into sub-fleets, and know the routes on which each aircraft can and cannot operate, they then have to build the operating schedule, according to the requirements specified by the airline's commercial and network forecasting departments. The basic schedule will be planned for a season, but the process can start several years prior to operations. The plan will include the flight number and airport codes for each flight, the probable aircraft type and sub-type used, its seating configuration, its departure and arrival times, the days of the week it is operated, and the airport terminals used.

A schedule plan has to consider turn times at airports, attractive departure and arrival times for passengers, actual available slot times at airports, and the airport terminals that the airline wishes to use. The plan also has to allow for the inevitable delays and disruptions that will occur during actual operations, aircraft-on-ground (AOG) situations and delays caused by technical problems.

As the date of the flights gets closer, airline schedulers will update the plan according to forecast passenger loads, so the aircraft type assigned to the flights may then be changed.

Schedulers will then assign tails to flight numbers only one or two days prior to actual operations, because only then will they know for certain at which airport each aircraft will be located at the start of the day's operations. Individual aircraft flight hours (FHs) and flight cycles (FCs) will differ from those forecast, which will influence timing of maintenance events and checks, and therefore the availability of each aircraft.

Scheduling software

Sabre's schedule definition and long-term planning software is AirVision Network. "Airlines first have to define their schedule using basic data for each flight number relating to airport codes, aircraft fleet and sub-fleet, departure and arrival times, aircraft configuration, days of the week the aircraft will operate, and months or season that are applicable," says Jeremy Million, product manager, AirVision Network suite, at Sabre Airline Solutions. "This is usually done for a few months or for a whole season in advance. These data are entered line by line for each flight number.

"The input data is then transferred to a gantt chart for planning an operating schedule," continues Million. "The gantt chart has a horizontal timescale and individual aircraft on the vertical scale. Each flight is shown as coloured rectangles or 'pucks', which have the flight number in the middle and airport codes at either end. There is a line of flights for each aircraft in each sub-fleet. Although individual tails are not assigned to flight numbers, there is a line of flights for each aircraft so that schedulers can see the number of flights that each aircraft is physically capable of operating.

Schedulers need to consider several issues when determining departure and arrival times, including minimum turnaround times that are practical at each airport and for each aircraft type, and convenient connection times for passengers. "AirVision has the functionality of displaying flights departing and arriving at an individual airport. This is particularly useful for a hub operation with lots of connections," says Million. "The schedule manager shows a day's flights on the screen. Flights arriving are on the left, flights departing later out of the hub are on the right.

There is a blank space between the two in the middle of the screen, which shows the connection time between the two banks of flights. Each aircraft has to be allocated to an arriving flight and a subsequent departing flight. There is a coloured line for each aircraft which indicates which two flights it will operate. Since all aircraft cannot land or depart at the same time, flights will arrive over a period of 30-60 minutes, followed by a quiet ground period while passengers disembark and make their connections, and then by another 30-60 minute busy period when all flights leave. Schedulers have to consider practical aircraft turn and passenger connecting times, and the probability and length of delays. Airlines can choose schedule aircraft in a 'first in, first out' (FIFO) system, or 'last in, last out' (LIFO) system. The timescale can be changed to several days, weeks or months."

SITA's product for long-term schedule planning is FleetPlan, which manages schedules, allocates aircraft to them and co-ordinates airport slot availability. "An airline's proposed flight schedule must be approved by each airport for its own operational purposes, especially by busy airports," explains Paul Rainford, product manager for FleetWatch Services, at SITA.

FleetPlan plans schedules by listing individual aircraft in the fleet down the left side on a gantt chart, in a similar fashion to AirVision. The fleet can have dividing lines for each aircraft type and sub-fleet. Flights are shown as coloured pucks, and airport three-letter codes are shown at either end. A series of flights is shown for each aircraft.

"Each flight has its flight number, and departure and arrival times. These are shown in GMT, but the time can be switched to local if necessary," says Rainford. "All the relevant information about the airports for each route can also be examined, including its time zone, and details about the terminals, runway lengths and certified landing category."

FleetPlan has a colour code system for indicating the status of each flight: a blue box means that all requirements are met; a green one means that slots are required, with a symbol showing whether the slot is needed at the departure or arrival airport. Slots can be acquired at International Air Transport Association (IATA) slot meetings, or the airline can change the timing of the flight to where slots are available at both ends.

FleetPlan also allows the user to make a one-off change to an aircraft's schedule, and operate on a different flight number. In this case the flight is indicated in a different colour.

"FleetPlan can show the flight number, type of flight, aircraft version and seat configuration, and schedule of times," explains Rainford. "The software also has the ability to read ad-hoc and scheduled messages about changes to the planned schedule.

"The operating schedule clearly affects several other airline departments, including ground services, maintenance operations control, crew scheduling, maintenance scheduling, and flight dispatch, all of which need to view the system," continues Rainford**. "**FleetPlan can also create a maintenance timetable, especially for line maintenance, which clearly needs to be co-ordinated with the operating schedule for each aircraft. Time in the operating schedule also has to be blocked off by schedulers for base maintenance."

FleetPlan clearly cannot function without a lot of background information. 'This includes specification and performance data for each aircraft, and the turn times each aircraft needs at each airport depending on the type of cabin service it is offering and what is physically possible at each airport. Rules can also be set up to limit aircraft to particular routes based on their performance. If an aircraft that cannot perform a certain route is scheduled to operate, the system alerts the user.
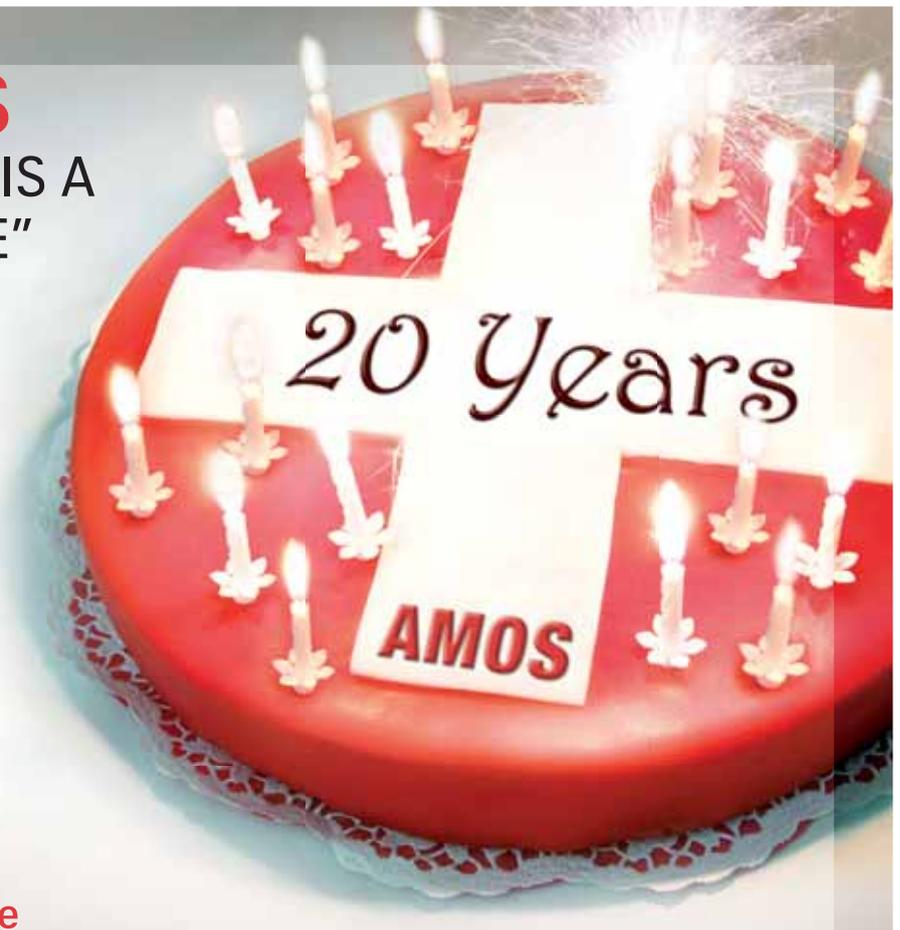
"Other technical details, known as

hard rules, can also be programmed into the system," continues Rainford. "This includes when an auxiliary power unit (APU) cannot be used and a ground power unit (GPU) is required, or if a toilet is unserviceable which limits the aircraft to a short flight. Soft rules can also be defined. This is when certain aircraft are preferred, but are not absolutely necessary, at an airport."

When a flight is created, the version of the aircraft is dictated. The timetable can already be stored in the system, together with all the aircraft that can be used. "FleetPlan automatically puts in arrival times, and chooses airport terminals," says Rainford. "It also checks that the same flight number does not already exist for that day. The timetable is also communicated to the crew-rostering system for long-term planning."

Schedule fine tuning

Long-term schedule planning clearly needs to be adjusted regularly. Demand levels are forecast when schedules are initially defined. Actual passenger numbers become clearer as forward bookings are taken. Forecasts can also be provided by revenue management. Airlines may therefore want to allocate different aircraft types to routes. "This will occur at regular intervals in the

months and weeks leading up to operations," says Million. "Our FleetManager product can edit a schedule, summarise the changes, including any cost implications, and match the aircraft type and seat numbers with projected demand to provide the best fleet mix for the schedule. The schedule and aircraft type assignment can be reiterated periodically every six weeks or so leading up to actual operations. Our ProfitManager product analyses the overall profitability of an operation."

This need to optimise aircraft size and the planned schedule in accordance with projected demand explains the popularity of the A320 and 737NG families. Both are essentially one type, with the same pilot-type ratings but available in four different sizes of 105 to 190 seats. "FleetManager can determine the number of each sub-fleet for each day's operation, and family type can be switched to another without any disruptions to crew rosters or timetables," says Million.

SITA's FleetPlan, Sabre's FleetManager and Lufthansa Systems' Netline Optimisation solutions can all optimise schedules within several constraints. Examples are a minimum of utilisation for each fleet type. Systems can even determine the number of each type of a family an airline needs, although this depends on long-term passenger numbers.

Another level of schedule optimisation offered by SITA's FleetPlan is that it can analyse whether additional flights can be fitted in using the given fleet, or if a given schedule can be operated with a smaller fleet. This can be done for just one day, or over an extended period.

Tail assignment

As the day of operations approaches and the schedule has been fine tuned and the aircraft type for each flight has been decided, airline schedulers will be ready to assign individual tails to flights. "The long- and medium-term planning assigns only virtual aircraft in the fleet to a flight, not actual aircraft," explains Rainford. "So although it is known that there are 50 A319s in a fleet, and the flights for each of these 50 aircraft have been selected, the actual tail numbers have not. This only occurs one or two days before operations. The schedule then passes from FleetPlan to FleetWatch, which assigns tail numbers to flights. Here the schedulers can use FIFO or LIFO algorithms for assigning aircraft, but there are advantages and disadvantages to both. LIFO, for example, means that the last aircraft to land has the shortest turnaround time.

One issue to consider when

optimising tail assignment is the standard reimbursement made to passengers who are denied boarding. Lufthansa Systems' Tail xOPT used for tail assignment takes into consideration the number of passengers booked and selects the right size of aircraft accordingly.

SITA's FleetWatch has a colour coding system to show the status of flights. Blue indicates that the flight is scheduled, but has not yet departed.

Sabre's AirCentre Movement Manager assigns tails to flights. This takes the planned schedule from FleetManager, and continues with the gantt display. Tail assignment is a process that is done nightly in many airlines' cases. The system needs to interface with several other systems, including the airline's maintenance system. Many maintenance events and checks have intervals in FH, although line checks are carried out prior to every flight, and either daily, every second day and weekly. Daily, twice-daily and weekly checks are usually scheduled during the night for short-haul operations, but they often have to be scheduled during the day between flights for long-haul operations. Sufficient downtime therefore has to be reserved by the maintenance department, and the scheduling system must be notified of this. Checks with FH interval have to be scheduled, and require several days' or weeks' downtime.

Lufthansa Systems' Netline Ops Tail xOPT optimises aircraft utilisation and plan maintenance checks just in time. A buffer of just 1-4% remaining maintenance interval is used to time checks. The timing of the check and its downtime can be displayed as coloured pucks. Netline Ops Tail xOPT optimises the maintenance interval for each check for each aircraft by automatically planning the checks as they come due.

Managing daily operations

Once tails have been assigned to flights, FleetWatch and other systems can be used to manage daily operations and flights as they occur.

"FleetWatch manages each puck for each flight. The airport codes and flight numbers are maintained, and a new colour-coding system is used to indicate a flight's status," says Rainford.

The screen again is a gantt chart, and usually displays a two- or three-day timescale. The middle of the screen has a bright green line to indicate the current time. Pucks that are black with green hashes, from the day before and earlier in the current day, mean that a flight has been completed. Green means that the flight is in progress, but not completed, while blue means that the flight is scheduled, but has not yet departed. Pink, yellow and brown are used to indicate

scheduled short and extended maintenance events, or aircraft in an AOG situation. A light blue colour indicates there is a possible delay, and flights departing late have a red block at the front of the blue, the length of which shows the extent of the delay. The actual departure time is indicated, and a vertical white marker shows the scheduled arrival time. Movement Manager and xOPT have a similar colour-coding system.

In FleetWatch, Movement Manager and xOPT a long delay to a flight that

clashes with the scheduled departure time of the subsequent flight is indicated by the puck for the next flight moving down, so that it is no longer in-line with other flights scheduled for the aircraft. This brings the problem visually to the attention of the controllers.

As delays or possible delays occur, controllers and operations staff have to manage flights by re-assigning them to other aircraft, which can have many consequences. Crews will be rostered onto particular flights, and in many cases

cannot change to other aircraft types. Passengers may miss their connections, and a delay can have a knock-on effect on all other flights scheduled for the day. This is why airlines need to have spare aircraft available, which can also be colour coded on the system.

There first has to be a system for communicating information about the progress and status of each flight, through type B messages, the internet, or the aircraft's aircraft crew and communicating system (ACARS). Basic data for flights are: out (O) from the gate; off (O) for take-off; on (O) for landing; and in (I) for arriving at the gate. These OOOI messages sent by the aircraft send the basic time information for each flight.

Operations staff need to know the nature and cause of the delay, such as air traffic control, connecting passengers, weather or slots being missed, in order to estimate the probable length of the delay. It may be harder to assess the length of technical delays, so knowing the nature of a delay is important.

The system uses icons to display the type and length of the delay, once details of it have been entered.

When flights are delayed, system users then need to re-assign flights to other aircraft. "This is simply done by dragging and dropping the flight into another aircraft's schedule. FleetWatch will show the cost of making this change to another aircraft," says Rainford. "The system will also show in colours the consequence of making changes. If these are not possible, the system will alert the user.

"As well as manually, the system can also re-assign flights automatically, and offers several options for replanning the schedule and the effects of each one," continues Rainford. "It can also give a simulation of each replanned schedule and fleet assignment. Re-assigned flights are shown by white hashes in the blue pucks."

Systems can show the details of each flight, including: the crew and number of passengers on board; the cause of any delay; the anticipated arrival time; and the number of connecting passengers and the flights they are connecting to.

Operations control

Besides dealing with individual delays, airline operations-control departments have to manage entire daily operations. There are numerous knock-on effects to even a single flight being delayed. For example, if an incoming flight is late, crews scheduled to change to a new aircraft when they arrive at an airport will be late boarding their new aircraft, thereby delaying its flight. Passenger mis-connections are a large consideration.

Lufthansa Systems has developed its Integrated Operations Control Center (IOCC) platform to monitor all flights. The systems tracks up to 30 different key performance indicators (KPIs) and key disruption indicators (KDIs). Several of the most important can be displayed on a single screen. Examples are the percentage of flights arriving on time, the number of AOG situations at specific time intervals, the number of standby crews available, and the number of flightdeck and cabin crews that are due to change aircraft. Each of these can be displayed graphically as coloured bars. Two colours can be used to indicate how critical each of these parameters are. The display of this information, and for other parameters, helps the maintenance control department anticipate problems by acting on situations before they get critical. The system allows the user to drill down on each KPI and KDI to get more information about each problem.

Lufthansa Systems explains that the IOCC platform helps airlines by: increasing revenue when re-optimising fleet assignment, improving punctuality and crew productivity, and reduced fuel consumption.   **AC**