Although revenue management systems have evolved extensively over the past 10 years, there are many enhancements and techniques coming available that airlines can use to fine tune the revenue management process. These techniques and related systems are explored.

Dynamic optimisation & other revenue management techniques

The different elements of the passenger sales process are constantly being refined and improved. Revenue management (RM) is one that attracts of lot of attention. RM techniques and systems have become increasingly advanced over the decades, with the current generation of RM systems using more sophisticated algorithms and more sources of data, and reacting more dynamically to try to generate maximum revenue from the RM process. The ability to have dynamic RM and to re-optimise each flight after a booking is taken has only become possible in the past year. Many other methods are also emerging that can fine tune the RM process.

Current RM

There are two basic types of RM systems: leg- or segment-based systems that forecast and optimise on a leg basis; and origin & destination (O&D) systems, which typically forecast and optimise for the whole network. It is claimed that leg-based systems add 3-10% to revenues, while the more complex O&D RM systems are said to add 1-2% for those carriers with the complexity of network and sophistication to use them. Most Tier 1 and a larger number of smaller airlines have moved to O&D RM systems. There are many medium-sized and small carriers that find leg-based systems more than adequate.

Current RM systems work on the basis of a planned schedule of re-optimisations prior to each flight. These re-optimisations can start as much as a year before the flight departs. The simpler leg-based systems can be optimised nightly and frequently on demand. The more complex O&D systems typically operate on a schedule of re-optimisation, but are re-optimised daily, usually during the night, in the last one or two weeks prior to a flight.

A flight is re-optimised by using different criteria to calculate RM recommendations or a bid-price curve. RM solutions make demand forecasts using inputs from the reservation system, such as bookings taken, available seats, schedules and, in some cases, published fares.

The recommendations or bid-price curve generated during an optimisation will determine the number of target seats to be made available at each booking class until the next re-optimisation takes place. "Using this approach, the lowest booking class the airline will accept for the remaining seats goes up in steps," explains Larry Michaels, a consultant and former director of decision solutions at ICF SH&E, whose RM product is marketed independently and by SITA as part of the Horizon portfolio.

Alternative systems, such as Navitaire's SkyPrice, produce optimal demand and price points, which are translated into optimal seat allocations.

Earlier generation RM systems did not consider sell-up of passengers from one booking class to another. In the current environment where it is popular to offer fares with no rules or restrictions, RM systems have to consider buyers' sensitivity to higher fares. This sensitivity, or ability to 'sell-up', is considered through forecasts which estimate how many passengers will buy a higher booking class when the lower one has been closed.

The next optimisation can be one or several days later. There may be 100 seats still available, for example, and five booking classes left, for example classes A, B, C, D and E, with A being the class with the highest fare. In modern RM strategies airlines typically make the booking class with the lowest fare, E, available to the first buyers, unless it is a period where demand is known to be high.

In the case of low-cost carriers (LCCs), which have a pure LCC fare structure, the lowest fares are simply made available to the earliest buyers on a first-come, first-served basis. As the seats allocated to each booking class are sold, the RM system closes that booking class and the inventory control system makes the next highest booking class available to subsequent buyers. The booking class available during the day therefore increases as seats are sold.

Many traditional, hybrid and LCC airlines aim to generate higher yield mixes by selling higher booking classes to business travellers. This is frequently done through a combination of fare rules and inventory control. Fare rules are used to 'segment the market' by placing onerous conditions, such as a Saturday night stay, on the cheapest fares. The airline will close the lower booking classes, particularly on flights during known periods of high demand.

The buyer's required travel dates and times are compared to the rules of each booking class in the reservation system, in addition to which booking classes are still available, to determine what booking class should be offered to the buyer.

A low booking class, such as an E class, will have unappealing conditions to a business passenger looking for a day return fare, even if they are one of the first buyers after a flight has been re-optimised, and seats are still available in E class. Their dates and times of travel may mean that the lowest fare available to them is actually a B class.

One recent development in the RM process is the development of algorithms to predict a buyer's willingness to pay a higher fare.

The RM system recommendations are sent to the inventory control system to determine what class to make available in real-time to buyers making reservations through the various sales channels. Another element of controlling inventory is the process of 'nesting'. The bid-price curve, for example, will calculate how many of the 100 remaining seats on a flight should be available at each of the five booking classes from A to E. "The RM system may forecast that it can sell 10 of the E class, 30 of the D class, 20 of the C class, 30 of the B class, and 10 of the A class," explains Hari Subramanian, director of revenue management solutions at Sabre Airline Solutions. "Nesting means that the airline does not want to sell an E class seat to a business passenger that wants a day-return trip, just because they are one of the first buyers. Fare rules mean that the cheapest booking class this buyer should be offered is a B class. It does not make sense for the system to offer the buyer an E class fare, when they are likely to accept a B class. This is why nesting is required. Although there is an expected demand for 30 D class and 10 E class fares, the system allows 40 of D class to be sold if 40 buyers want travel times that match the rules of the D class fare, and there are simultaneously no buyers whose requirements match the rules of E class. In the example above then, the process of nesting means that: 10 seats will be available at E class; up to 40 at D class; up to 60 at C class; up to 90 at B class; and up to 100 at A class. This way as many as 100 of the remaining 100 seats could be sold at the highest class if there is higher than forecast demand from buyers whose travel requirements all match the rules of this booking class. Similarly, up to 90 could be sold at B class if actual demand exceeds forecast demand by 60."

Inherent weaknesses

Current systems have many limitations. "The first of these is that the bid-price curve calculation and the inventory control processes are done at different times, when ideally they should be done at the same time," says Subramanian.

"Another limitation of current systems is that while demand is for a flight is a function of competitor fares offered in the marketplace, many RM systems do not consider the influence of competitor fares in demand forecasts," continues Subramanian. "Sabre's RM system provides business rules that an airline analyst can use to influence inventory controls based on competitor fares. The influence of competitor's fares on passenger demand forecast is an active area of research."

A further weakness of current RM systems is that they make assumptions relating to a buyer's willingness to pay a higher fare. These assumptions are that a buyer is willing to pay the fare in accordance with their travel dates and times. Buyers are sensitive to changes in the availability of booking classes and fares, and an increase in the fare available will depress demand.

Another issue is with flights selling at a pace slower or faster than forecast. If the forecast is off, the RM recommendations provided by the RM system are likely to be off as well. "A lot of legacy systems used by traditional airlines cache the bid-price curve, so the number of seats available at each fare class does not respond quickly enough to the changes in demand," explains Michaels.

The problem with caching the bid-price curve is that less than optimal revenue can be generated. A slower than forecast rate of bookings will pose a danger of seats remaining unsold because the booking class available dissuaded buyers. A faster than forecast rate of bookings will mean that too many seats were made available in lower booking classes, when the number available in the higher booking classes should have been increased. The actual amount of deviation may or may not be important. Large deviations can be detected by the system, which then triggers a re-forecast and a re-optimisation.

RM enhancements

One refinement of the overall RM process relates to the algorithm used to calculate the number of seats allocated to each booking class. The largest elements of this algorithm are forecasting the demand for the remaining sales period for the remaining available seats and the price-sensitivity of buyers to different fares.

As described, the algorithms of legacy systems generally do not have any ability to calculate the sensitivity of buyers to available fares.

"The RM systems built in the 1980s and 1990s did not directly model the impact that price has on demand," says Barbara Vega, product direct revenue management at Navitaire. "That is, the legacy system relies on demand for yesterday's prices to forecast future demand. It also relies on inherent class or fare rules to force sell-up to a higher booking class, and does not worry about how the fare affects demand. Many airlines have relaxed or eliminated fare rules, and so instead have several price points for a booking class with the same rules and restrictions. In this case, price is the only

Many RM systems do not consider, or are not able to consider the influence of competitor's fares in their demand forecasts.

differentiator which determines demand so it is critical that the price-demand relationship be modelled within the RM solution." Navitaire's SkyPrice solution was one of the first RM solutions to model the price-demand relationship.

Typical legacy systems use the traditional method that a passenger is only interested in one booking class. "Although passengers actually want the lowest fare, some may be prepared to go up one or more booking classes if the booking class they want is closed," says Michaels. "A hybrid RM system calculates the elasticity or sensitivity of a passenger going up a certain number of booking classes. This is an enhancement in the forecasting and optimisation process, in systems such as SITA's current hybrid RM system."

Algorithms have evolved considerably to calculate a buyer's sensitivity to price changes. "Ultimately airlines would like to calculate people's willingness to pay more for higher booking classes," says Subramanian. "This is essential to fine-tune the process of calculating the bid-price curve whenever an optimisation is carried out. The algorithm in the RM system has to calculate how many seats to make available at each booking class. Allocating the correct number of seats to each class will generate the maximum revenue.

"The problem with calculating price-sensitivity is that the airline only has the booking pattern to date, and the previous booking and sales history of the particular flight," continues Subramanian. "Historical booking data play a crucial role in determining price sensitivity. The prime source of historical data will be for the same flight number that has been operated on the same day of the week as the flight being analysed. An airline may operate several times a day on a busy business route, such as Washington DC-New York or London-Frankfurt. If the airline is calculating price-sensitivity for a flight that leaves at 7:45am on a Monday, it will be primarily interested in historical booking and sales data for the same flight number on Mondays going back over several years. Of more interest will be the same flight number on Mondays that were operated in the same week and at a similar time of the year. Additional historical data of interest will be flights that operated on other days of the week and at other times of the year. The flight numbers that are for departures just before and after the flight at 7:45am are also analysed. Sabre employs passenger choice models to forecast demand, which takes into account the influence of other flights on the demand for a 7:45am departure."

Sabre's passenger choice models forecast demand on schedule and price attributes. These models mimic the passenger selection process in the demand forecasting. A passenger's willingness to pay is calibrated on the airline booking data. A higher willingness to pay implies a price inelasticity.

Lufthansa Systems has also developed algorithms within its RM module to help determine buyers' price sensitivity. "The algorithms are used to determine demand, spill and re-capture rates of passengers accepted in lower booking classes for each booking class," says Stephen Bayer, vice president of commercial solutions at Lufthansa Systems. "These systems simulate market and customer behaviour to understand different types of buying behaviour. The RM modules we offer are the ProfitLine Yield Rembrandt system for point-to-point services, and the O&D version for O&D operations."

Amadeus has developed systems for a better understanding of price-sensitivity. "We have implemented a technique called marginal revenue transformation (MRT) to develop a Q forecasting algorithm. Q is used because the Q booking class is the lowest revenue class in the economy cabin," says Jean-Michel Sauvage, senior manager of airline IT development at Amadeus. "All passengers want to buy the lowest fare possible, so the Q forecasting algorithm predicts or estimates the probability of a passenger buying a higher fare if a lower booking class is closed off by the RM system. This algorithm is based on a combination of the buyer's willingness to pay the higher fare and the elasticity of demand. For example if there are 100 seats left to fill and two fare classes of $100 and $150, the algorithm will calculate how many people will buy the $150 fare if the $100 fare is closed off."

Competitor intelligence

Another element of improved forecasting is information relating to competitors' activity. While RM systems require data feeds from various modules of an airline's own sales process, improved forecasting and optimisation are achieved through regular feeds of information relating to the fares and booking classes being offered by competitors.

Airlines use four main sales channels: call centres, travel agents, third-party websites and their own websites. It is not possible for airlines to get real-time information on the fares their competitors are offering through their call centres and via travel agents. Fares offered by competitors through their own and via third-party websites can be browsed periodically, however. This is made simpler and performed faster with robotic webcrawlers, which have been available for seven or eight years. "Fare management systems, such as SITA's Airfare Insight, include competitive monitoring and are another element of the forecasting and RM process that can help refine and optimise RM," says Michaels.

Fares offered by competitors change during the day between re-optimisations. Although RM systems re-optimise automatically at night, yield managers at

Dynamic optimisation for leg-based RM systems is a relatively simple process, and one that can be completed in a few seconds or just minutes. Dynamic optimisation for an O&D RM system, for an airline's complete network, requires a large amount of computing power, and most carriers lack the required capacity.

airlines can manually trigger re-optimisations during the day. This will lead to changes in fares being offered on their websites, and so generate a need for other airlines to adjust the fares they are making available. This then raises the issue of how frequently airlines should re-optimise their flights in the final few days before departure. "While more information can be gained to fine tune RM, such as near real-time data on competitors' fares, the airline still needs faster technology to make use of it," comments Michaels.

Re-optimisation triggers

A main vulnerability of legacy RM and Inventory Control systems is the caching of the bid-price curve between re-optimisations. Fixing the number of seats allocated to each booking class during the day, or for several days, makes the airline vulnerable to sales varying from forecasts.

While it is hard to re-optimise at high frequencies during the day in the last few days prior to a flight, it is possible to put in automatic triggers to make the system re-optimise the flight if the level of sales activity falls below, or rises above, certain sales level thresholds. These triggers can work in several steps, such as sales activity being 10%, 20% and 30% below or above the usual level triggering re-optimisations so that the recommendations are changed during the day to improve yield mix and overall revenue. "There can be a manual trigger, so that a single re-optimisation is requested during the day by an analyst," says Vega. "There can also be multiple scheduled re-optimisation thresholds so that several re-optimisations can be done over the course of the day based on defined parameters or thresholds. We will be introducing re-optimisation triggers, both manual and scheduled, in 2012."

Navitaire has always offered daily optimisations for every flight, once per night. This is an advantage over other solutions which have defined optimisation days that can often be weeks apart.

Lufthansa Systems has had the ability to re-optimise flights during the day for about 15 years, while Sabre has offered it for about 10 years. "The trigger to start a re-optimisation can be a manual one from an airline's revenue managers, but we have been offering programmed triggers for the past four years now," says Subramanian.

Periodic re-optimisation

The inevitable variation of actual demand levels and rate of sales from forecast or predicted levels raises the issue of merits and benefits of periodic re-optimisation. "Some airlines feel there would be a benefit to revenues if they re-optimised more often. In the extreme, that could mean a re-optimisation after every booking is taken," says Michaels. "It has to be remembered that first this is only worth it and possible if a very accurate forecasting module and algorithm is available. Re-optimising also requires a large calculation. The simpler optimisation in leg-based RM systems could be done periodically during the day, or even in response to each booking. The network optimisation of O&D RM is considerably more complex, taking in the order of minutes," continues Michaels. "O&D RM systems may be able to do partial re-optimisations during the day in response to triggers, such as unanticipated demand or schedule changes." In O&D RM, the entire network is optimised in one go, because segments cannot be separated anymore. So an optimisation can take many hours for larger airlines. Partitioning the network into groups of routes helps a bit. A re-optimisation is therefore not possible after each booking in O&D RM without reverting back to segment control, which risks neglecting network effects. Real-time network optimisation cannot exist without making large and questionable simplifications.

"Besides the inventory control and RM modules, airlines also need a fares management or pricing module, webcrawler robots and databases of historical booking and sales information as the main modules to complete a sales system," says Michaels. "The problem with this is that the communication of data between the different modules is not 100%. Many airlines have developed and evolved the different modules of their legacy systems at different times and with different technology standards and languages. The accuracy with which different modules speak to each other has to be considered. Vendors now offer integrated modules and complete solutions."

Amadeus has built an integrated inventory control and RM system. "The old traditional systems needed a lot of manual intervention between modules, in terms of feeding data from one to the other. This is because modules of legacy systems have data in different language formats. Exchanging data between them requires a coded third language. This coding is costly, time-consuming and inflexible," says Sauvage. "This made it very difficult to re-optimise a flight more than once a day. Our new integrated inventory control and RM system makes it possible to use dynamic revenue management. This integrated system means that inventory control and optimisation can operate in real-time. If systems are not integrated there needs to be a data file exchange, which takes time. This is makes it impossible to optimise in real-time every time a booking is made. An integrated system would dynamically adjust inventory levels after every booking to ensure maximum revenue.

"This was introduced in March 2011, and the new Altea RM system can re-optimise a flight in about one second," continues Sauvage. "The system was launched by Icelandair and Air Mauritius,

One aim of RM systems in the future might be to prioritise buyers that are known to be regular buyers of ancillary products over other buyers.

and will soon be implemented by another two airlines, with four more implementing it in 2012. All the relevant systems and modules of a sales system are hosted together at an airline's headquarters, so the modules interface with each other.

"A dynamic RM system means that revenue managers get alerts early enough to say if flights are selling faster or slower than usual," adds Sauvage. "An essential element of dynamic RM is to have specific algorithms, and we have implemented Dynamic Inventory Control (DICO) in our Altea system. This is an advanced mathematical model to ensure optimisation across the entire network. This means it does not calculate optimal prices for a single flight in isolation. It considers other onwards flights a passenger might be taking in the journey. A simple example is an airline providing connecting services between Nice and Paris, and Paris and New York. It may decide to block the last few remaining seats for customers taking the more profitable Nice-Paris-New York journey, in preference to passengers only wanting the short-haul Nice-Paris trip. Every time DICO is applied, the most sensible revenue decisions are made for the carrier across the network factoring all possible connections a passenger could wish to make."

Priority sales

Another method to fine tune RM is to analyse the sales achieved through each sales channel. The ultimate objective would be to give preferential treatment to the sales channels that cost less, or have highest percentage of sales in the highest booking classes. In other words, isolate the buyers that have the same booking characteristics. "Inventory systems with O&D capabilities, such as SITA's Horizon inventory system, have the capability to monitor sales generated through each sales channel, and use this information to optimise revenue generation," says Michaels. "When deciding whether or not to accept a particular passenger, the systems compare the average revenue by O&D, point of sale and sales channel, and booking class to the bid prices on the legs the passenger would traverse. The average revenues could also consider the cost of the different sales channels to optimise profit per passenger."

Amadeus says it can now offer real-time network optimisation by applying DICO across the whole airline network when each booking is made. It helps to ensure that flights are not viewed by the RM system in isolation, which could mean revenue is lost.

A further need for enhancements to the RM process is the emergence of ancillary products and the unbundled products offered by airlines in addition to full-service products. The objective is for airlines to sell additional ancillary products to buyers that have bought unbundled products, and which select the ancillary products from a menu of options. The danger to airlines is that not all buyers purchase enough ancillary products. This raises a need for the RM system to identify which passengers buy a large number of ancillary products, and then give them priority over other passengers.

Ultimately, priority would be given to buyers which generate the highest contributing revenue, rather than just the highest overall revenue. This would require an analysis of the cost of sales through each channel and a true identification of the net revenue, both from tickets and ancillary products. "The only way to identify those passengers that buy a large number of ancillary products is through historical sales records. These would be stored in the frequent flyer (FF) and customer relationship manager (CRM) modules of an airline's sales system," says Sauvage. "First, the sales channel would need to identify the buyer, and relate them to their previous sales record, so there would need to be a link between the CRM and RM modules."

It is still some time before this will actually be possible. "The current systems are not sufficiently integrated to predict which buyers of particular fare classes are more likely to buy ancillary products," says Michaels. "The data points and values exist within our system, however, so SITA's vision and roadmap foresees differentiating availability based on total customer value."

Navitaire's RM SkyPrice solution does allow for the manual input of ancillary revenue data in the optimisation solution to balance demand and price. The higher the ancillary revenue that has been generated, the more the system favours achieving high load factors over favouring higher booking classes and prices.

Lufthansa Systems offers a connection to the FF and CRM modules. "It is therefore possible to quote individual pricing for a buyer that has a history of buying ancillary products. The question is whether or not this makes sense," says Bayer.

A final issue in RM fine tuning is the feedback of actual achieved sales data to the RM system for optimising flights in the future. AC