Netw Spat Econ (2006) 6: 235–251
DOI 10.1007/s11067-006-9282-y

Improving Airline Network Robustness and Operational Reliability by Sequential Optimisation Algorithms

Cheng-Lung Wu

Abstract A sequential optimisation algorithm is developed to improve the operational reliability of airline schedules. Simulation results show that departure delays are reduced by 30% after optimisation by using extra 260 min buffer times in the schedule. This also increases the network-wide schedule reliability from 37 to 52% and an estimated delay cost saving of $20 million dollars per annum for a small airline network. The advantage of sequential optimisation is that it considers the delay/punctuality propagation in airline networks, so to prevent airlines from planning excessive buffer times to individual flights by considering aircraft rotation as a whole process.

Keywords Airline schedule robustness · Airline schedule reliability · Sequential optimisation · Aircraft routing · Delay propagation

1 Airline Schedule Planning and Operations

Airline schedule planning typically involves four steps from schedule design, fleet assignment, aircraft routing to crew pairing/rostering. At the stage of aircraft routing, schedule planning involves the optimisation of aircraft routing by formulating aircraft routing as integer programming problems such as the work by Arguello et al. (1998), Barnhart et al. (1998), Luo and Yu (1997), Rexing et al. (2000), Teodorovic and Stojkovic (1995) and Yan and Young (1996). Sophisticated optimisation algorithms have been developed to solve complex integer programming problems and satisfactory results can be achieved within reasonable computing times and resources nowadays. In the real practice, aircraft routing problems are hardly solved by a single step. Instead, the routing optimisation and problem solving process may involve intervention from experienced airline schedulers, e.g., manually adjusting routing

C.-L. Wu (✉)
Department of Aviation, University of New South Wales, Sydney NSW 2052, Australia
e-mail: C.L.Wu@unsw.edu.au

plans to reflect some operating constraints such as aircraft turnaround times at specific airports and flight punctuality targets of key feeder flights. This 'fine-tuning' procedure is widely used by airlines because the current approach in aircraft routing optimisation, e.g., integer programming is not an operation-oriented approach, but more optimisation focused. The lack of consideration in aircraft routing optimisation to reflect real operational issues may result in lower schedule robustness and reliability in daily operations (Mederer and Frank, 2002). The observable consequences of lower schedule reliability are flight delays and potential delay propagation in an airline's network.

Typical aircraft routing optimisation is aimed at searching for the optimal solution of an integer program with side constraints without explicit consideration of some issues in airline operations such as turnaround efficiency at airports, the allocation of buffer times in schedules and the reliability/robustness of schedule operations. In day-to-day airline operations, we can observe that some flights tend to arrive late and need to be turned around within a shorter time and consequently result in departure delays and delay propagation in the network (Mukherjee et al., 2004; Watterson and De Proost, 1999). Given the current aircraft routing optimisation approach, the robustness and reliability of airline schedules is not well considered against stochastic disruptions in operations. Therefore, experienced airline schedulers may manually adjust the preliminary results of aircraft routing to achieve higher schedule reliability and/or higher punctuality. In addition, this task is usually conducted under certain operational constraints, e.g., limited use of buffer times, airport slot availability, achieving punctuality targets and reducing delay propagation in the network. Regarding the algorithms to optimally adjust airline schedules at this stage of planning, most schedulers rely heavily on individual experience rather than on sound theoretical backgrounds and algorithms.

The delay propagation in airline schedules is well known to airlines and has been studied recently (Beatty et al., 1998; Mederer and Frank, 2002; Wu, 2005). To mitigate the impact of delay propagation and to improve on-time performance, airlines embed buffer times in schedules (in turnaround times or block times), so to absorb accumulated and unexpected delays in aircraft rotation. This measure may improve on-time performance, but the drawback is to consume longer aircraft turn times and possibly longer block times for flights (Howarth and O'Tool, 2005; Sunday Times, 2000). Delays may propagate in a network via aircraft rotation, passenger transfer and crew rostering, if embedded buffer times are not enough to absorb delays. By the same token, flight punctuality also propagates and this is the case which one can observe in a normal day of airline operations. In order to improve the robustness of schedules and reduce the likelihood of delay propagation, airline schedulers may manually adjust the optimised aircraft routing results. However, two problems are often encountered in this 'fine-tune' process including: (1) determining the optimal usage of expensive aircraft times as buffer times; and (2) establishing the punctuality benchmark with which new schedules are designed and compared.

To solve these two problems in the schedule fine-tune process, sequential optimisation algorithms are proposed in this paper. Algorithms are developed to apply sequential optimisation in fine-tuning aircraft routing plans based on the preliminary results of aircraft routing planning. To evaluate the effectiveness of embedded buffer times in draft airline schedules, a simulation model is used to generate the benchmark required to effectively allocate scarce aircraft times. In the following parts of the paper, it starts with the airline schedule simulation model, followed by the introduction of schedule reliability indices developed to evaluate the operational reliability of airline schedules. Sequential optimisation algorithms are developed to fine-tune aircraft routing schedules under certain constraints. Developed algorithms are then applied in a case study to demonstrate the effectiveness of the proposed sequential optimisation algorithm in more detail.

2 Schedule Simulation and Schedule Reliability Index

2.1 Schedule Simulation

Two issues are often encountered by schedulers during the process of aircraft routing optimisation including: a) how to determine the optimal use of schedule buffer times and, and b) how effectively schedule buffer times can improve operational reliability and save delays. Since an aircraft is assigned to carry out a few flights in a flight cycle (being one day for domestic operations and one week for international operations), these two issues may become extremely complicated and hard to solve in a big airline network, if the only evaluation tool for schedulers is their experience in scheduling. In order to evaluate the punctuality/delay status of a flight schedule at planning stages, a simulation model developed earlier is used to provide the estimated delays of draft schedules in various planning scenarios (Wu, 2005).

An airline schedule is composed of a number of flight sectors to be carried out by airline fleets. A flight sector is defined to start from the on-block of an aircraft at the origin airport gate until the on-block of the aircraft at the destination airport gate. Hence, a flight sector comprises turnaround operations at the origin airport and airborne flight operations in the air space between the origin and destination airport. Consequently in the simulation model, a flight sector is modelled by two components, namely the turnaround module and the enroute module. Individual flight sectors are then inter-connected in the schedule simulation model according to the given aircraft routing schedules. The simulation of schedule operation is carried out by Monte Carlo techniques to consider stochastic disruptions in airline operations and the stochastic nature of operational activities in airline operations. In a simulation run, each flight of the schedule is simulated by randomly selected samples from stochastic functions derived from historical airline data. The schedule is then simulated for 1,000 times (with 1,000 samples for each flight of the schedule) so to control the simulation noise level.

The turnaround module comprises a number of parallel workflows, which are conducted simultaneously and parallelaly on the airport ramps to turn around an aircraft for a following flight. Major workflows include passenger de-boarding/boarding (also includes crewing and cabin cleaning), cargo and baggage unloading/loading, catering unloading/loading and other independent work such as aircraft engineering check and fuelling. Certain procedures must be followed for activities in the same workflow and delays to an activity may delay following activities in the workflow and possibly the departure time. For instance, cabin cleaning does not start until all passengers have left the aircraft and passengers will not start boarding until cabin cleaning is finished. Given the sequential nature of activities in workflows and stochastic disruptions within workflows, Semi-Markov Chains are used to model these workflows. Since the operating times of activities in workflows vary according to a few factors such as available human resources and work loading, the use of Markov Chains also reflects the stochastic aspects of aircraft turnaround operations. Disrupting events in workflows, e.g., missing check-in passengers or late baggage loading, are modelled as 'disrupting states' in the Markov model with proper transition probability linking to normal operations (normal 'states').

2.1.1 Turnaround Module

Let t_i^{ATD} be the actual time of departure of flight i ($\forall i \in N$, the set of all flights in a study schedule), which forms a probability density function (PDF) of flight i and is denoted by $f_i^{ATD}(t)$. S_i^D denotes the given scheduled departure time of flight i, and departure delays are defined by Eq. (1) as follows:

$$D_i^D = t_i^{ATD} - S_i^D \tag{1}$$

where D_i^D denotes the departure delay of flight i from Airports I to J ($\forall I \neq J$). t_i^{ATD} is a dependent variable influenced by two other variables, called: the actual time of arrival of the previous flight m ($m \neq i$ and $\forall m \in N$), denoted by t_m^{ATA}, and the stochastic turnaround operation time of flight i, denoted by T_i^{OP}. T_i^{OP} is the (longest) time required to finish all turnaround activities including two major turnaround processes (passenger processing and cargo/baggage processing), delays from disruptions, and other aircraft service activities as described in Eq. (2).

$$t_i^{ATD} = t_m^{ATA} + T_i^{OP} = t_m^{ATA} + \max\left[T_i^{cargo}, T_i^{pax}, T_i^{events}\right] \tag{2}$$

For instance, T_i^{cargo} is the time required to finish cargo and baggage processing for flight i. A total of Ω activities need to be carried out in this process and each activity k ($k \in \Omega$) has a stochastic operating time and an expected operating time, ε_k as given in Eq. (3).

$$T_i^{cargo} = \sum_{k=1}^{\Omega} \varepsilon_k \tag{3}$$

Activity k is modelled as a Markovian state, which transits to any state at any time t with a state transient probability function, $A_k(t)$ as shown by Eq. (4). By the same token, the time required to finish passenger processing, i.e., T_i^{pax} is also modelled as a Markov Chain.

$$\sum_{k=1}^{\Omega} \varepsilon_k = \sum_{k=1}^{\Omega} (E_k[t]) = \sum_{k=1}^{\Omega} \left(\int_0^{\infty} tA_k(t)dt \right) \tag{4}$$

Discrete events, which operate independently from the above mentioned workflows might delay aircraft departure only if the finish time of an event exceeds the scheduled departure time. Discrete events are modelled as stochastic variables in Eq. (5) with a PDF denoted by $\Phi_q^e(t)$. ε_q^e in Eq. (5) denotes the expected disruption time of event q ($q \in Q$, the set of all potential events) and P_q^e denotes the occurrence probability of such an event.

$$T_i^{events} = \max\left[\varepsilon_q^e\right] = \max\left[P_q^e E_q[t]\right] = \max\left[P_q^e \int_0^\infty t\Phi_q^e(t)dt\right] \tag{5}$$

2.1.2 Enroute Module

The enroute module describes the airborne flight operations between airports. It is a complex procedure for flight operations from pushing back at the gate, taxing, taking off, airborne operations and landing at the destination airport. To simplify the simulation model and focus on airline operations on the ground, the enroute module aggregately describes the airborne flight operations by stochastic functions. Hence, considerations in airport layouts, runway congestion, queues on taxiways and enroute air traffic control delays are modelled aggregately by stochastic functions with model parameters calculated from historical block times between airports. Although this module is not as sophisticated as the turnaround module, it closely reflects the overall flight operation patterns in aircraft rotation, subject to current capacity constraints in the airport and air traffic control systems. Further more, airlines around the world hardly have any influence on airport and air traffic control procedures, except changing the hubbing patterns of flights at some hub airports at which the airline is the dominant hubbing carrier.

The enroute module is described by Eqs. (6) and (7). t_i^{ATA} is the actual time of arrival of flight i at the destination Airport J, which forms a PDF, denoted by $f_i^{ATA}(t)$. t_i^{ATA} is influenced by the other two stochastic variables: the actual time of departure of flight i at Airport I, i.e., t_i^{ATD} and the expected en-route flight time between Airports I and J, denoted by ε_i^{ER} in Eq. (6). ε_i^{ER} is derived from the PDF function of en-route flight time of flight i, denoted by $f_i^{ER}(t)$. Hence, the arrival delay of flight i is modelled by D_i^A in Eq. (7), where S_i^A denotes the given scheduled arrival time at the destination Airport J.

$$t_i^{ATA} = t_i^{ATD} + \varepsilon_i^{ER} = t_i^{ATD} + \int_0^\infty tf_i^{ER}(t)dt \tag{6}$$

$$D_i^A = t_i^{ATA} - S_i^A \tag{7}$$

Delays due to constrains of airport capacity and airspace congestion are modelled aggregately by the stochastic function of enroute operation times (t_i^{ER} from $f_i^{ER}(t)$) in Eq. (6), which counts from the time an aircraft is pushed back at a gate until the time an aircraft is on chock at the destination airport. Accordingly, the actual turnaround time of an aircraft is the time between on and off chock at a gate. The departure/arrival delay, i.e., D_i^D and D_i^A in Eqs. (1) and (7) could sometimes be negative values due to early departure/arrival of flights. Given the current capacity constraints at major airports and the pressure to reduce the scheduled turnaround times for most airlines, early departure/arrival operations are limited. Turnaround operations may start early or finish early, but this also puts pressure on ground resources allocation as well because early starts also disrupt resource allocation.

Historical punctuality data is obtained from an airline to calculate those parameters required to calibrate and run the schedule simulation model. The most significant advantage of such a simulation model is that we are able to probe the "expected status'' of a draft schedule plan before operation. It is well known that buffer times are usually embedded in airline schedules, so to provide airlines with some protection against the stochasticity in operations. Given designed buffer times, delays still exist in airline operations and the resulting delays are the complex interaction among three key factors: turnaround operations, schedules and stochastic disruptions in operations. Among three factors, the airline schedule is a fixed timetable, while the other two factors are stochastic in nature.

With embedded buffer times, airlines expect those buffer times to take effect and absorb some delays in operations. Since aircraft times have high opportunity costs (as revenue-making times), it is hardly seen that an airline would buffer its schedule to reduce delays to nearly the zero-delay level (called the *Perfect Case*) as illustrated in Fig. 1 (Wu, 2005). Instead, airlines tend to use limited buffer times scattering among flight sectors, hoping that these buffer times will control delay propagation in the network to a satisfactory extent. By using the simulation model together with appropriate parameters (usually by using the standard operating procedure data in aircraft operations), we are able to probe the "expected delays (on-time performance)'' before schedule operations, which closely reflect schedule planning philo-

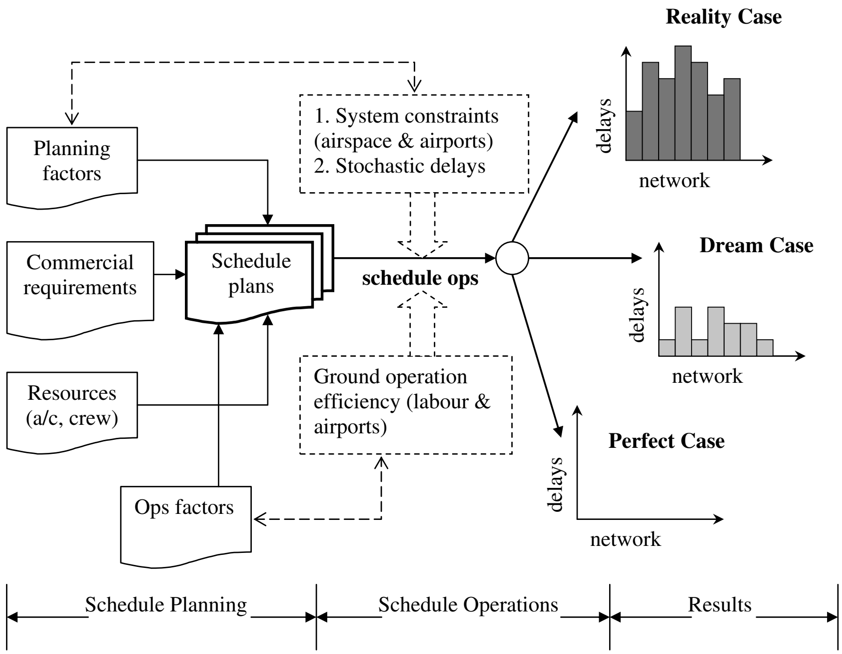

Fig. 1 Inherent delays of airline operations (source: Wu, 2005)

sophy, scheduling constraints and trade-offs. The expected delays (called *inherent delays*) of such an optimised schedule with limited buffer (called the *Dream Case*) are usually higher than the *Perfect Case*, because airline operations are always subject to disruptions from various sources, e.g., passengers and airport ground congestion. Very often it is found by airlines that the real delays after schedule operations (called the *Reality Case*) are higher than the delay level of the Dream Case, mostly due to inadequate schedule buffer times (or the reluctance to design more buffer times due to trade offs) and stochastic disruptions in airline operations. Figure 1 below illustrates the relationship among three possible situations in schedule operations.

A comprehensive model calibration process is conducted earlier to adjust parameters used in the simulation model, so to mimic the current flight operations in the real environment (Wu, 2005). A set of parameters is obtained from an airline to generate the inherent delays of the Dream Case. The same set of parameters is then further calibrated to generate simulation results, which are close to current results. When the flight schedule is changed, corresponding model parameters are modified and fed into the simulation model to generate simulated schedule operation results including mean delays and on-time performance (OTP) figures. Flight schedules are run by the simulation model for 1,000 times (representing 1,000 days operation), so to reduce the influence of simulation noises on results.

2.2 Schedule Reliability Index

At the stage of aircraft routing optimisation, airlines always encounter the trade off between limiting the use of expensive buffer times in schedules and reducing delays (or improving on-time performance) by using more buffer times. In order to minimise the use of expensive buffer times in a schedule, a benchmark is required to maintain the balance of the trade off between the use of buffer times and flight delays. A set of schedule reliability indices is proposed to serve as benchmarks in schedule planning. The *schedule reliability* of flight i is defined by comparing the inherent delays generated from the Dream Case with the real delays from the Reality Case as formulated by Eq. (8) below.

$$R_i^D = \frac{ED_i^D}{D_i^D} \quad R_i^A = \frac{ED_i^A}{D_i^A} \quad R_i = \frac{\left(ED_i^D + ED_i^A\right)}{\left(D_i^D + D_i^A\right)} \tag{8}$$

where R_i^D / R_i^A denotes the departure/arrival reliability of flight i, respectively; ED_i^D / ED_i^A represents the inherent departure/arrival delay of flight i, while D_i^D and D_i^A the actual departure and arrival delay of flight i. Hence, R_i is used to evaluate the overall operational reliability of flight i.

The concept behind this schedule reliability index is to benchmark real schedule delays against the inherent delays so airlines can evaluate how close the schedule operation is to the expected delay levels in schedule planning. By this index, the schedule reliability reflects both the schedule planning philosophy, e.g., the willingness to buffer schedule, and the situation in which real operations are conducted by the airline. For instance, if R_i^D has a value larger than 100%, it means that the real departure delays of flight $i\left(D_i^D\right)$ is less than the inherent delays $\left(ED_i^D\right)$. It implies that the real operation outperforms the expectation and the scheduled ground times may be shortened so to save aircraft times, if needed. However, quite often we see R_i^D less than 100%, meaning that the real delay level is higher than the expected one. This may result from inadequate buffer times for turnarounds, delay propagation in aircraft rotations, less robust turnaround operations, or excessive disruptions occurred in turnarounds.

3 Sequential Optimisation and Algorithms

The biggest challenge to improve the robustness of aircraft routing is not about the optimisation techniques one can employ when solving integer programming problems. Rather, it is more about how an airline can improve the robustness of schedule operations. Currently, the integer programming approach is widely used by airlines to solve aircraft routing problems. However, given the nature of integer programming, optimisation tends to produce tight schedules under cost minimisation objectives without thorough consideration of delay propagation effects in networks. Very often it is seen from airline operations at airports that some particular flights need more turnaround times than other flights due to the efficiency of turnaround at certain airports, connecting passengers inbound/outbound the flights, work loading of ground staff at certain times and so forth. The integer-programming approach to optimise aircraft routing may also cause some problems in real-world operations such as delay propagation, frequent delays to specific flights and tight turnarounds, if further manual schedule fine-tune is not conducted thoroughly. Given the current aircraft routing optimisation and schedule adjustment procedures, sequential optimisation algorithms are proposed in this paper to improve the efficiency and effectiveness of the manual schedule tuning process and to supplement existing integer programming algorithms. The goal of sequential optimisation is to improve the schedule robustness and reliability in daily airline operations and minimise delays.

After initial aircraft routing, a routing plan is generated which consists M rotation patterns to cover N flights in a network. A rotation pattern j (also called a route) consists of a number of flight sectors for an aircraft to conduct in a rotation cycle. Flight i in rotation j (denoted by flight (i,j)) is defined to start from the on-block time of the aircraft at the origin airport gate until the on-block time of the same aircraft at the destination airport gate, namely the actual time of arrival, t_{ij}^{ATA}. The actual time of departure of flight (i,j) is defined as the off-block time of the aircraft at the origin airport gate and is denoted by t_{ij}^{ATD}. The scheduled times of departure and arrival for flight (i,j) are denoted by S_{ij}^D and S_{ij}^A, respectively. Hence, the departure delay of flight (i,j) is $\left(t_{ij}^{ATD} - S_{ij}^D\right)$ and denoted by D_{ij}^D; the arrival delay of flight (i,j) is $\left(t_{ij}^{ATA} - S_{ij}^A\right)$ and denoted by D_{ij}^A.

Based on 'schedule reliability' defined earlier, the goal to improve the reliability of flight (i,j) is to control the real delay levels as close as possible to the benchmark standard, i.e., the inherent delays. Hence, to optimise the operational reliability of a schedule, the given aircraft routing patterns are further relaxed by allocating extra buffer times. Since buffer times are expensive costs to airlines, the use of buffer times in the relaxation process is limited to achieve a chosen performance target, which is measured by delays in this model. A reliability target, denoted by R^{TAR}, can be chosen arbitrarily by an airline as the schedule operation target. According to the definition of the reliability index given earlier by Eq. (8), the resulting target departure delays for flight (i,j) is:

$$D_{ij}^{TAR} = \frac{ED_{ij}^{D}}{R^{TAR}} \tag{9}$$

Accordingly, the schedule adjustment for flight (i,j), SA_{ij}^{D} is expressed by Eq. (10) below, where D_{ij}^{R} is the estimated delays of flight (i,j) from simulation. It is noted that if SA_{ij}^{D} is larger than zero, it implies that the current delay level of flight (i,j) is higher than the target level and vice versa. This also implies that an airline can save excessive aircraft times from those flights which have SA_{ij}^{D} less than zero and use the saved aircraft times on those flights which have SA_{ij}^{D} larger than zero.

$$SA_{ij}^{D} = D_{ij}^{R} - D_{ij}^{TAR} \tag{10}$$

Therefore, the total available schedule adjustment times for rotation j (denoted by SA_{j}^{D}) can be expressed by Eq. (11), where K_i is the number of flights in rotation j.

$$SA_{j}^{D} = \sum_{i=1}^{K_i} SA_{ij}^{D} \tag{11}$$

Since delays and punctuality propagate along aircraft rotations in a network, the relaxation of ground times of earlier flights in rotation j will reduce the delays of following flights in the same rotation. Accordingly, the final schedule adjustment times for flight (i,j) (denoted by ${}_{F}SA_{ij}^{D}$) will be less than or equal to the available schedule adjustment time estimates, i.e., SA_{ij}^{D} as shown by Eq. (12). Hence, the final total schedule adjustment times for rotation j (denoted by ${}_{F}SA_{j}^{D}$) will be less than or equal to the initial schedule adjustment time estimates, i.e., SA_{j}^{D} in Eq. (13).

$${}_{F}SA_{ij}^{D} \le SA_{ij}^{D} \quad \forall i \;\; 1 \le i \le K_i \text{ and } \forall j \;\; 1 \le j \le M \tag{12}$$

$${}_{F}SA_{j}^{D} \le SA_{j}^{D} \quad \forall j \;\; 1 \le j \le M \tag{13}$$

Given the unique delay/punctuality propagation effect, the optimisation of aircraft routing schedule is conducted sequentially within each rotation. The objective function of sequential optimisation can be expressed by Eq. (14) as follows:

To minimise:

$$\sum_{j=1}^{M} \sum_{i=1}^{K_i} D_{ij}^{R} \tag{14}$$

Subject to:

$$ {}_{F}SA_j^D \leq SA_j^D \quad \forall j\ 1 \leq j \leq M \tag{15} $$

$$ \left\{ D_{ij}^R \leq 15\,\text{mins} \right\} \quad \text{or} \quad \left\{ D_{ij}^R \leq D_{ij}^{TAR} \right\} \tag{16} $$

where Eq. (15) constrains the usage of total adjustment times of rotation j less than or equal to the first estimation, i.e., SA_j^D; M stands for the total number of rotations of the study schedule, which is equal to the fleet size. Equation (16) limits the estimated delay times of flight (i,j) after schedule adjustment under 15 min (which is the industry delay threshold) or under the chosen target delay level (D_{ij}^{TAR}), whichever requires less buffer times. To conduct the optimisation, an algorithm is developed as follows:

Step 1: For rotation j of the schedule, $1 \leq j \leq M$:

Step 2: Run the simulation model and generate the inherent delay estimates, i.e., ED_{ij}^D / ED_{ij}^A for each flight (i,j) in rotation j.
Calculate current schedule reliability, i.e., R_{ij}^D / R_{ij}^A for each flight (i,j) in rotation j and compare these with the chosen target schedule reliability index, R^{TAR}.

Step 3: Calculate the initial estimate of schedule adjustment, SA_{ij}^D by $(SA_{ij}^D = D_{ij}^R - D_{ij}^{TAR})$ and calculate the total available schedule adjustment for rotation j, SA_j^D by $(SA_j^D = \sum_{i=1}^{K_i} SA_{ij}^D)$.

Step 4: For flight (i,j) $(1 \leq i \leq K_i)$ in rotation j:

If $SA_{ij}^D > 0$ and $D_{ij}^R > 15$, add SA_{ij}^D minutes (in the multiple of 5 min) to the ground time of flight (i,j).
If $SA_{ij}^D < 0$, deduct SA_{ij}^D minutes (in the multiple of 5 min) from the ground time of flight (i,j).
Repeat Steps 2 and 3 for all flights in rotation j.

Step 5: Repeat Steps 1–4 for all rotations.

The optimisation algorithm above improves schedule reliability of each flight in a rotation sequentially according to the chosen reliability target, the historical operating delays, the total available schedule adjustment times and the impact of delay/punctuality propagation in the rotation. The sequential nature of the optimisation algorithm fully utilises the concept of 'punctuality/delay propagation,' so to avoid excessive use of aircraft times in optimisation.

4 Application in Schedule Planning

4.1 Research Data and Case Study Network

To demonstrate the effectiveness of the proposed sequential optimisation algorithm on schedule planning, a case study is conducted by using schedule information and punctuality data of an anonymous airline, denoted by 'Airline X.' A selected fleet with 17 narrow-body aircraft flying to 20 destinations is used in the following case study.

Schedules and operational data, i.e., the standard operating times of activities in aircraft turnaround processes adopted by the carrier, are used in the schedule simulation model to estimate the level of inherent delays of the study schedule (the *Dream Case*). This result is used together with historical operating delays to calculate the current operational reliability of each flights/rotations in the schedule. The simulation model is then calibrated based on existing model parameters of the *Dream Case* and real operating statistics in the previous years. Calibration results are summarised by Fig. 2 below. Parameters used in the turnaround module and the enroute module of the simulation model are calibrated to reflect the current operating delay levels. Disrupting events in aircraft turnarounds, e.g., missing passengers and late baggage loading are simulated by the real occurrence probability derived from historical IATA delay codes recorded by the airline. The average durations of disrupting events are calibrated accordingly to reduce the gap between simulation results and the *Reality Case*. It is seen from Fig. 2 that calibration results are close to real delay levels. Calibration errors between two cases are kept under 5% during the calibration exercise. The calibrated simulation model is then employed to conduct schedule simulation based on given draft schedules and generate simulated delays after corresponding schedule changes.

4.2 Sequential Optimisation Results

The chosen reliability target of schedule optimisation in the following case study is to meet either one of the following two criteria: (1) 70% individual reliability for

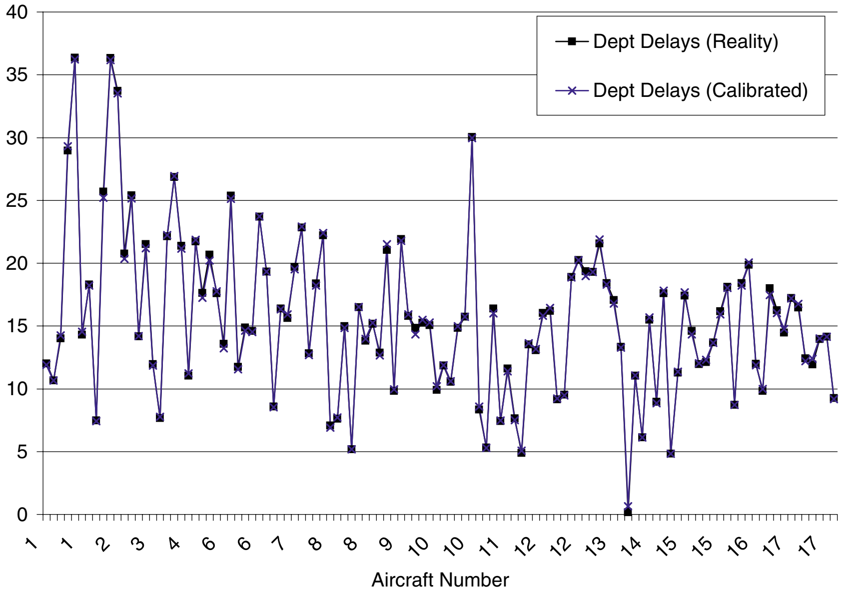

Fig. 2 Comparison between calibration results and real delays

each flight in the network; or (2) less than 15 min mean delays, whichever requires the least buffer times. Schedule adjustments in optimisation are only made to the scheduled ground times of flights. In other words, extra buffer times are embedded in the ground times (turnaround times) of flights, so to control departure delays. When schedule adjustments are required for a flight, a unit of schedule adjustment is 5 min. This prevents the final schedule from having unrealistic departure times, e.g., 9:13 departure. This practice also complies with the current policy of airport slot allocation and coordination at those slot-constrained airports (IATA, 2004).

The developed sequential optimisation algorithm is applied to Airline X's network and the summary results are shown in Fig. 3. The usage of buffer times for each rotation (operated individually by one aircraft) is compared with the estimated saving of delay times after sequential optimisation. Additional 260 min are used to relax aircraft routing plans according to optimisation criteria given above, and this also generates an estimated saving of 540 min delay network-wide. Delays of some flights in some rotations are significantly reduced after optimisation such as aircraft 2, 3 and 6, while the impact of minor schedule changes on some rotations such as aircraft 8 and 11 is less than others. The significant delay reduction effect (being the surrogate of punctuality improvement) after sequential optimisation also reveals the extent to which delay propagation may impact airline operations. A rough estimate of the impact scale of delay propagation from results above is 1-min delay might result in 2-min delay network-wide for Airline X.

If we assume that the monetary cost of one unit buffer time and delay time is \$200 per minute, then the estimated impact of the sequential optimisation on the new schedule is equivalent to an extra expenditure for Airline X by \$19 million dollars per annum (calculated from \$200/min@260 min/day@365 days/year) but a delay cost saving by \$39 million dollars per annum (calculated from \$200/min@540 min/day@365 days/year), resulting in \$20 million dollars net saving on operating costs per annum. This estimated cost saving could be significant enough for the study regional fleet and could have a positive impact on the profitability of an airline.

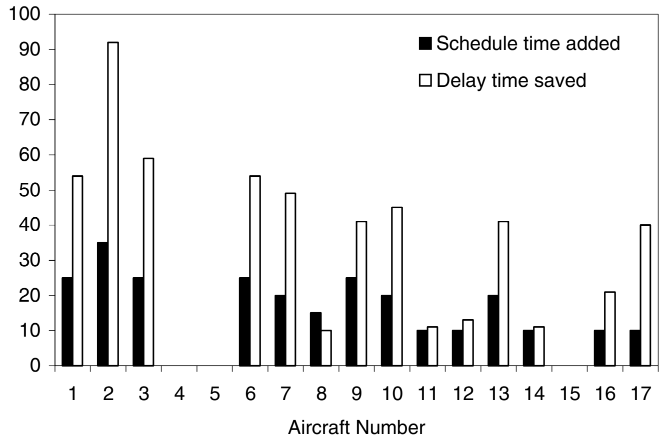

Fig. 3 Schedule adjustment and the impact on network-wide delay time saving

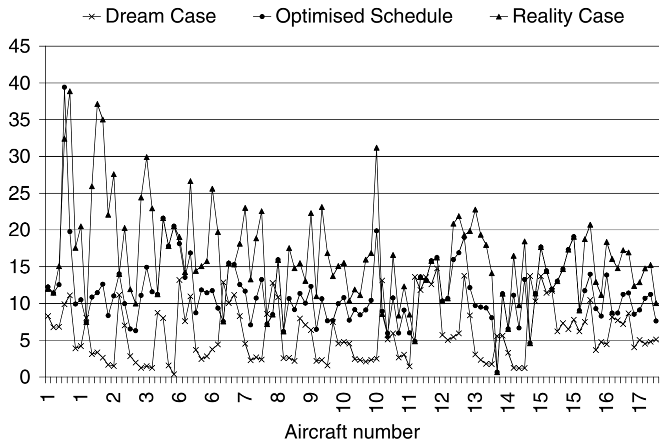

Fig. 4 Delay reduction by the optimised schedule

The optimised schedule is tested by the schedule simulator to evaluate how the optimised schedule may react to current operating environment faced by Airline X. Simulation results of the optimised schedule are compared with the inherent delays of the Dream Case and the current delays of the Reality Case in Fig. 4 below. Although the delay level of the optimised schedule is only as good as 70% of the expected level, i.e., inherent delays, the optimised schedule effectively controls the overall delays across the network to the required target level. Total departure delays of the original schedule are 1,816 min, which is reduced to 1,278 min after optimisation. This result also reflects on the increase of average network-wide schedule reliability from 37 to 52% after optimisation.

To demonstrate the impact of sequential optimisation on improving schedule reliability and reducing delays, the optimisation results of Aircraft 1 are given in Fig. 5 comparing the reliability index before and after optimisation. It's seen that the original reliability of the rotation decreases from 70% at the start of the cycle to 25% at the end of the cycle. It implies that the real delay level of this rotation is higher than the inherent delay level expected from Airline X's schedule. Either the

Fig. 5 Reliability of rotation by Aircraft 1 (before/after optimisation)

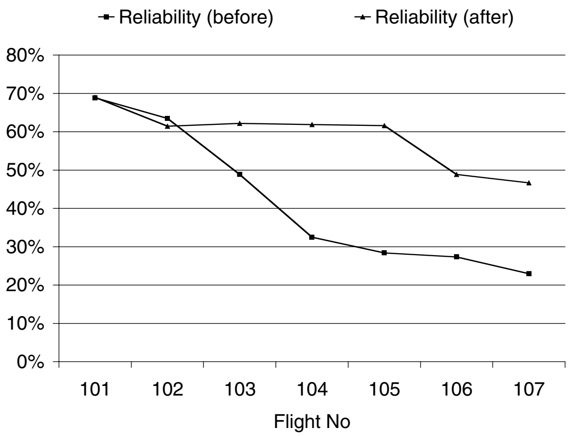

inherent delay level is too low because Airline X is too optimistic to the operation of these flights in the rotation, or the inherent delay level is close enough but the operations of these flights are not reliable enough.

When delays before optimisation are compared with those after optimisation for this rotation in Fig. 6, it can be seen that delays after optimisation are better controlled below 15–20 min. Schedule adjustments to this rotation cycle include: 5 min to Flight 103, 15 min to Flight 104 and 5 min to Flight 105, totalling extra buffer times of 25 min for Aircraft 1. The usage of buffer times in optimisation reflects closely the current operating status as shown in Fig. 6, so more buffer times are deployed for those flights that suffer from higher delays than others. Given the cost of deploying more buffer times, the estimated average delay time saving after optimisation is 54 min for this rotation.

The major advantage of the proposed sequential optimisation algorithm is to utilise the delay/punctuality propagation in a close network so to utilise limited schedule times and prevent from using excessive buffer times during the optimisation process. To demonstrate how punctuality propagation can help sequential optimisation and to illustrate the convergence of the proposed optimisation algorithm, interim results of the optimisation of Aircraft 1 are extracted and shown in Fig. 7. Current operating delays in this rotation hike significantly starting from Flight 104, so subsequent flights in the rotation suffer from severe delay propagation. The 'target delay level' is calculated based on the target reliability (70%) and the inherent delays of this rotation, which are strongly influenced by current ground operations and scheduling policy of Airline X.

The optimisation process starts from Flight 101. The initial estimate of the total schedule adjustment times for all flights in this rotation is calculated by (11) as 66 min. The optimisation procedure first makes 5-min change to Flight 103. From Fig. 7 we can see the simulated delay of the rotation after the first schedule change is reduced from the current level and consequently the re-calculated total schedule adjustment for following flights drops to 50 min, thanks for punctuality propagation.

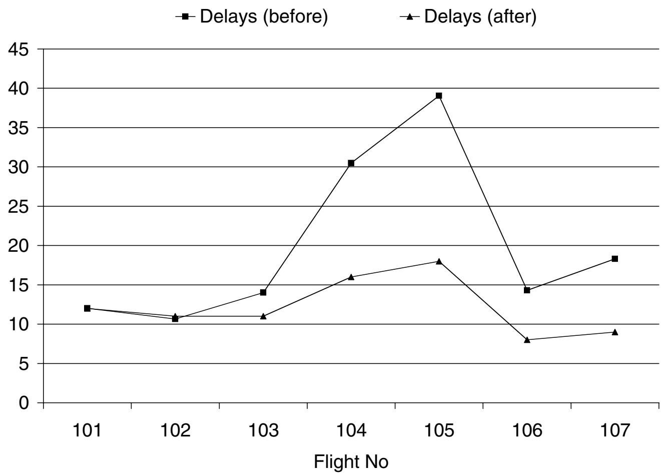

Fig. 6 Delays before/after optimisation for Aircraft 1

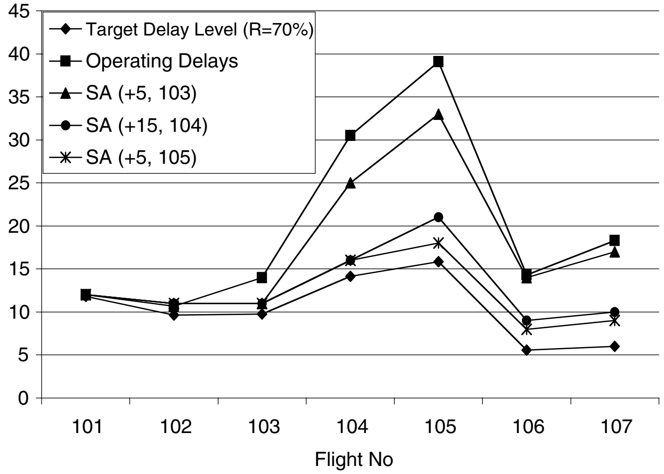

Fig. 7 Comparison of interim results of sequential optimisation for Aircraft 1

After adding 15 min to Flight 104, the delay level of the rotation drops further to a lower level, causing the required total schedule adjustment times for following flights by only 17 min. The impact of punctuality propagation in this rotation is best seen by Flight 104, and we can also observe how significantly delays of the following flights after 104 are reduced.

The optimisation process stops after adding 5 min to Flight 105. At the end, the used schedule times in the optimisation is 25 min, which is significantly lower than the first estimation in the beginning of optimisation, i.e., 66 min. Without the sequential optimisation algorithm, the application of 66 min to the rotation will improve the reliability of each flights much more significantly, but will result in spending extra 41 min when compared with results of sequential optimisation. By this optimisation algorithm, airlines can optimally utilise valuable aircraft times on critical flights in the network to effectively achieve the target schedule reliability.

4.3 Implications and Limitations

In practice, the immediate challenge before finalising aircraft routing plans is to match scheduled departure/arrival times of flights with airport slots available to an airline. For flights departing/arriving at those airports frequently served by an airline, the airport slot availability issue is a relatively minor concern as the airline can always swap slots among its own flights. Since the magnitude of flight time adjustment in the above optimisation is generally below 15 min, this change would not impose too much pressure on airport slot availability, even though an airline does not enjoy a dominant position at some airports.

The further implication of airport slots is the ‘saleability’ of flights to passengers. Since the sequential optimisation or any manual changes during aircraft routing planning always involves altering the schedule, the optimised schedule times are not necessarily the most saleable times for potential passengers. There is a possibility that marketing/sales decisions may overturn some results of aircraft routing with specified departure/arrival time requests for certain flights. This practice is seen quite often in the airline industry especially for airlines operating hubbing schedules due to the need to consider connecting times and connecting opportunities for passengers at hub airports.

The present sequential optimisation algorithm has not yet fully considered the constraints airlines may face when dealing with hubbing and schedule synchronisation across the network. Due to this limit, the applicability of the current algorithm seems more beneficial to those airlines, which do not operate highly synchronised hubbing schedules. However, the proposed algorithm can be generalised to consider hubbing schedules in future work as side constraints, if needed. Considering the operation of strong hubbing activities, there has been growing concerns over the economic benefits/costs for strong hubbing schedules and the congestion costs imposed on airports and air passengers (Goolsbee, 2005). The de-peaking of hubbing schedules recently by American Airlines at Chicago O'Haire and Dallas-Fort Worth (DFW), by Lufthansa at Frankfurt, and by Delta Air Lines at Atlanta, demonstrate the vulnerability of heavy hubbing operations in terms of operational reliability of airline schedules and the hidden costs to operate such schedules (Field, 2005; Goedeking and Sala, 2003). The trend in airline scheduling has gradually shifted from maximising connecting opportunities and network traffic in the past to maximising the operational reliability of airline schedules under the influence of disruptions (Kang, 2004). Hence, airlines are nowadays more willing to trade off hubbing and aircraft utilisation with reliable and robust "weak-hubbing" or more point-to-point schedule operations (Mederer and Frank, 2002; Wu, 2005). Given this trend, the proposed sequential algorithm will gradually see its benefit in airline schedule planning in the future.

An essential advantage of the proposed sequential optimisation is to combine schedule simulation models with schedule optimisation algorithms. The advantage is to provide airline schedulers with immediate feedback during schedule planning. However, the argument on simulation applications often roots in the validity of the simulation model in representing the real-world operations. The simulation model used here is developed from the perspective of airline operations and highly focuses on aircraft turnaround operations and the network effects of stochastic disruption in airline operations. Since airline operations are fully bounded and strongly influenced by airport operations and air traffic control, an alternative approach is to simulate the airport system by network queuing models such as the MEANS model (MIT, 2005).

5 Conclusions

A schedule optimisation algorithm is proposed and tested in this paper by using sequential optimisation techniques. The objective of the sequential optimisation is to optimally deploy limited and valuable aircraft times as buffer in a network and achieve a target schedule reliability of 70% or a mean departure delay of 15 min, whichever requires the least extra buffer times. Simulation results of the optimised schedule show that the total estimated departure delays of the optimised schedule are 1,278 min, which is reduced from 1,816 min before optimisation. This reflects on the increase of network-wide schedule reliability from 37 to 52% after optimisation. Results also show that the optimal use of schedule buffer times by 260 min can save up to 540 min delay. In monetary terms, it sums up to a net estimated saving of $20 million dollars per annum by a unit delay cost of $200 per minute for a small-size network.

The proposed optimisation algorithm can be used by airlines during aircraft routing planning, especially in the schedule fine-tune process. The use of schedule simulation models in schedule planning provides airlines with immediate feedback upon schedule alternations and the visualisation of possible results of schedule operations. The major advantage of sequential optimisation is that it considers operational characteristics of aircraft rotation, i.e., the mostly noted delay/punctuality propagation phenomenon in a network. This unique attribute distinguishes the proposed optimisation algorithm from other schedule optimisation techniques such as integer programming, which hardly considers stochastic operational factors in airline operations. Consequently, sequential optimisation prevents airlines from allocating excessive buffer times to individual flights, and meanwhile maintain the required schedule reliability targets. Therefore, expensive aircraft times are only used for critical flights, allowing punctuality to propagate before planning more buffer times for later flights in the same rotation.

References

Arguello M, Bard J, Yu G (1998) Models and methods for managing airline irregular operations. In: G Yu (eds) Operations research in the airline industry. Kluwer Academic: Boston

Barnhart C, Boland N, Clarke L, Johnson E, Nemhauser G, Shenoi R (1998) Flight string models for aircraft fleeting and routing. Transp Sci 32(3):208–220

Beatty R, Hsu R, Berry L, Rome J (1998) Preliminary evaluation of flight delay propagation through an airline schedule. 2nd USA/Europe Air Traffic Management R&D Seminar, Orlando

Field D (2005) Hubs attract low-cost trade. Airline Business, March issue, 15

Goedeking P, Sala S (2003) Breaking the bank. Airline Business, September issue, 93–97

Goolsbee A (2005) Tragedy of the airport. http://slate.msn.com/id/2123240

Howarth G, O'Tool K (2005) Expect delays. Airline Business, July issue, 62–64

International Air Transportation Association (IATA) (2004) Worldwide scheduling guideline

Kang L (2004) Degradable airline scheduling: an approach to improve operational robustness and differentiate service quality. PhD thesis, Massachusetts Institute of Technology

Luo S, Yu G (1997) On the airline schedule perturbation problem caused by the ground delay program. Transp Sci 31(4):298–311

Mederer M, Frank M (2002). Increasing robustness of flight schedules through stochastic modeling of planning parameters. AGIFORS Airline Operations Conference (Downloadable from website: http://www.agifors.org)

MIT (2005) MIT extensible air network simulation. http://means.mit.edu/means

Mukherjee A, Lovell D, Ball M, Odoni A, Zerbib G (2004) Modeling delays and cancellation probabilities to support strategic simulations. NEXTOR (Downloadable from website: http://www.nextor.org)

Rexing B, Barnhart C, Kniker T, Jarrah A, Krishnamurthy N (2000) Airline fleet assignment with time windows. Transp Sci 34(1):1–20

Sunday Times (2000, August 13) Airline extend flight times to conceal growing delays

Teodorovic D, Stojkovic G (1995) Model to reduce airline schedule disturbances. J Transp Eng 121:324–331

Watterson A, De Proost D (1999) On-time performance: operational challenges for airlines. Mercer on Transport (Downloadable from website: http://www.mercermc.com)

Wu C-L (2005) Inherent delays and operational reliability of airline schedules. J Air Transp Manag 11:273–282

Yan S, Young H-F (1996) A decision support framework for multi-fleet routing and multi-stop flight scheduling. Transp Res 30A:379–398