MISTA 2009

A comparison of two methods for reducing take-off delay at London Heathrow airport

Jason A. D. Atkin • Edmund K. Burke • John S Greenwood

Abstract This paper describes recent research into the departure process at London Heathrow, with the primary motivation of reducing the amount of fuel used, improving both the economic and environmental cost. Two proposals are considered here. The first proposal considers the practicality and potential benefits of aiding the controller to improve the take-off sequencing. The second proposal aims to absorb some of the inevitable delay for aircraft at the stands, before the engines are started, but also involves a take-off sequencing aspect. Models for the two take-off sequencing problems are presented in this paper, the second of which includes an additional pushback time (or TSAT) allocation sub-problem which has to be solved subsequently. These models have distinctive differences from the models for the take-off and arrival sequencing problems which are usually considered in the literature, since they take into account necessary constraints imposed due to the control problem (whether a sequence can actually be achieved, and how) in each case. As discussed in this paper, the control problem cannot be ignored by the controllers at Heathrow, so cannot be ignored by any realistic system to aid them in their tasks. Comparative take-off sequencing results are presented for the two systems and the potential benefits from providing decision support to the runway controllers or improved TSAT allocation at the stands are considered. The paper ends with some overall conclusions from the research, showing the large potential benefits of these systems. The TSAT allocation system which is discussed in this paper has been developed for implementation at London Heathrow as one element of the Collaborative Decision Making project.

1 Introduction

This paper considers potential improvements to the departure process at London Heathrow airport in order to reduce the delay for aircraft awaiting take-off, thus reducing the amount of fuel burnt, leading to cost and environmental benefits. Models are presented in this paper for two problems at different stages of the departure system and solution systems are proposed and

Jason A. D. Atkin and Edmund K. Burke
School of Computer Science, The University of Nottingham, Jubilee Campus,
Nottingham, NG8 1BB, UK
E-mail: jaa@cs.nott.ac.uk, ekb@cs.nott.ac.uk

John S. Greenwood
NATS, NATS CTC, 4000 Parkway, Whiteley, Fareham, Hampshire, PO15 7FL, UK

evaluated. The first system considers the possibility of improving the take-off sequencing by providing a decision support system to the controller who performs that task. The second system considers the potential to absorb some of the necessary delay for aircraft at the stand so that the fuel burn is further reduced. It has similarities to the first system in that it must predict take-off sequences and it may indirectly aid the controller to improve the take-off sequence by reducing congestion in the holding area and providing a more appropriate selection of aircraft from which to select. The potential benefits from both of these systems are evaluated and contrasted in this paper.

From the point of view of an automated system to help the controllers in the control tower, the departure process has a number of stages. When granted permission to do so, aircraft push back from their stands, start their engines and enter the taxiways. Once on the taxiways, one or more ground controllers will direct the aircraft to a holding area by the current departure runway, where they queue awaiting their turn to enter the runway and take off. At some point prior to take-off, a take-off sequence must be determined. Under present operating procedures, a runway controller will re-sequence the aircraft within the holding area, attempting to improve the take-off sequence and reduce delay for aircraft.

The first system which is described in this paper has been designed as a decision support system for the runway controller. It attempts to make decisions about the take-off sequence to adopt by considering known controller preferences and operating methods as well as the physical and regulatory constraints upon the sequencing problem. One of the main requirements of such a system is that it must react virtually instantaneously to changes in the situation at the runway, thus results have to be determined within a second. Obviously, there are human-computer-interface (HCI) issues to resolve in order to incorporate such a system for the runway controller, since it is vital (for safety-related reasons) that it does not distract the controller or increase the workload. So far, this research has been useful for providing insight into the role of the runway controller, the constraints upon the departure system and the likely effects of any changes to the way in which the departure system operates.

The second system which is described in this paper has been designed to assign Target Start-At Times (TSATs) to aircraft while they are still at the stands. TSATs define the times at which an aircraft should be given permission to push back from the stands, start their engines and commence their journeys towards the take-off runway. A take-off sequencing element is again involved in this problem, but decision times can be longer. This second system is currently being integrated into Heathrow.

2 Take-off sequencing

The runway forms the bottleneck for the departure system at London Heathrow. Although it is such a popular and busy airport, Heathrow currently has only two runways and, despite the fact that mixed mode operations (using both runways for interleaved take-offs and landings) would be more efficient, usually only one of the runways is available for take-offs at any time of the day - the other being used simultaneously for landings. This restriction is designed to control the noise for residents on the flight paths. Since the runway is normally used purely for take-offs, landings do not need to be considered in the models presented here.

A minimum separation is required between consecutive take-offs to allow wake vortices to dissipate. This separation depends upon the relative weight classes of the two aircraft and is asymmetric. An additional minimum separation must also be met between aircraft which depart along the same or similar departure routes, even if they are not consecutive take-offs (the separations do not obey the triangle inequality), in order to control the rate of flow along each route. This route-based separation can depend upon the time of day (temporary increases may be applied at times) and will change depending upon the runway that is in use and the direction of take-off. It also may need to be increased, or can sometimes be shortened, according to the relative speed groups of the aircraft involved (allowing for the convergence or divergence of aircraft at take-off or shortly afterwards). It is possible to determine the minimum separation that is required between any two aircraft given the weight classes, speed groups, departure routes and allocated runway(s), and the take-off time of the first of the aircraft (so that time-dependent separations can be considered). Any practical solution approach for the take-off problem must not assume that these separations are constants, even for a specific pair of aircraft.

There are distinct similarities between arrival sequencing problems [1], [2], [3] and take-off sequencing problems. However, it was shown in [4], [5], [6] that the control problem at Heathrow cannot be ignored. This situation is not unique to Heathrow [7]. The required separations, and hence the runway throughput, are very sensitive to the selection of aircraft which are available for take-off at that time, and especially to the routes upon which aircraft are due to depart. The fact that the route separations and take-off time-slots have major effects upon the delay for aircraft, [5], [6], means that approaches which categorise aircraft by weight category, such as [8], [9], [10] are not likely to be successful at Heathrow, even if the control problem could be ignored.

Controllers have to consider at least four objectives when considering which take-off sequences to enact. Firstly, runway throughput, or mean take-off delay, has to be considered. Secondly, the sequence cannot be prohibitively difficult to enact, in terms of the overtaking that is required and the time and workload involved in giving the necessary instructions to pilots. This is critical for the sake of ensuring safety. Thirdly, take-off time-slots apply to some aircraft and must be met. Finally, controllers must ensure that any single aircraft is not greatly penalised in the sequencing. This is important since characteristics of some aircraft (for example unusually small or slow aircraft) will necessarily require larger separations either before or after them. The re-sequencing must ensure that these aircraft are not just moved to the end of the take-off queue, where the larger separation would affect the take-off times of fewer aircraft.

Delay is used here as a measure of performance rather than throughput, since it is not only related to cost (for airlines and passengers, in terms of time as well as money) but also measures improvements that throughput measures can hide. The experiments performed for this paper consider half-day periods. The overall take-off throughput of the runway over a long period of time cannot exceed the rate at which departures are released from the stand. Even when busy, there can be times when the runway is starved of the correct types of aircraft which would be required to attain a good runway throughput. For example, if all waiting aircraft are northbound then no amount of sequencing ability will be able to reduce the separations between aircraft below two minutes. Throughput measures tend to consider only the take-off time of the last aircraft to take-off, rather than considering the take-off times of all aircraft, whereas delay-based measures take consideration to the sequencing of all aircraft. The advantages of measuring delay rather than throughput were discussed further in [6].

Due to congestion (or adverse weather conditions) in busy airspace or at busy destination airports, take-off time-slots exist for some aircraft. These are formulated as a Calculated Time Of Take-off (CTOT). An aircraft is not permitted to take off more than five minutes before this time and should not take off more than ten minutes after it. A limited number of further five minute extensions are available to controllers in order to reduce the number of renegotiations of CTOTs that are required. Thus, these CTOTs normally generate fifteen minute time-slots within which aircraft must take off, but on occasions can be stretched to twenty minutes, where necessary, by utilising an extension. Since extensions are limited and it is usually unknown how many extensions will be required later in the day, it is important for any sequencing to use as few as possible.

In order to predict when an aircraft would take off for a given take-off sequence, it is necessary to determine the earliest time at which it can do so. Predicting this prior to the aircraft reaching the holding area requires knowing the time at which an aircraft will be ready to push back from the stand and the taxi time from the stand to the runway. Due to the difficulty of predicting the time at which an aircraft will be ready for push back and the time required to reach the holding area, the current mode of operations is to release aircraft from the stands as soon as the workload for the ground controller will permit, get them to the runway holding area, and allow the runway controller to perform the take-off sequencing once they get there. However, only a small minority of the runway controller's time can actually be spent considering the take-off sequence; the majority being spent communicating with pilots. Fast (mental) heuristic methods have to be used, therefore. These heuristics usually work extremely well, but sometimes throughput improvements can be obtained from considering more aircraft. In particular, it is sometimes possible to avoid future larger separations for aircraft which are still taxiing by knowing about them in advance and appropriately sequencing the aircraft which are already in the holding area.

3 A model for the sequencing problem

Let d_j denote the (actual or predicted) take-off time for aircraft j. Let et_j denote the earliest time at which aircraft j can take off given the earliest time it is reasonable to expect it to be able to reach the runway. The derivation of this value differs between the two problems which are considered in this paper and is discussed below. Let ec_j denote the earliest time at which aircraft j can take off within its CTOT time-slot, or a low enough value to not restrict the take-off time if j has no CTOT. For any ordered pair of aircraft i and j, where i takes off before j, let RS_{ij} denote the minimum required take-off time separation.

Let ts_j denote the position of aircraft j in the planned take-off sequence and let as_j denote the position of aircraft j in the First-Come-First-Served (FCFS) take-off sequence. The FCFS sequence can be considered to be that which would be generated if aircraft took off in the order in which they arrived at the holding area, there was no contention at the cul-de-sac or on the taxiways, and no stand holds were applied. This can be generated by considering the earliest pushback times for each aircraft then adding on the pushback durations and the (predicted or real) taxi durations (from pushback completion to reaching the holding area).

Measuring delay requires a base-line from which to measure. The base-line time, bt_j for aircraft j, is here defined as the time at which aircraft j historically reached the holding area, and the delay is defined as the duration from holding area arrival to take-off, $d_j - bt_j$.

The sequencing problem involves finding a good take-off sequence such that the constraints are met and the value of the objective function is minimised. There are three main constraints upon the take-off time for aircraft j. Firstly, j cannot take off before it can reach the runway, as shown by (1). Secondly, it cannot take off before the start of any CTOT time-slot, as shown by (2). Thirdly, all runway separation rules must be obeyed, as shown by (3).

$$d_j \geq et_j \qquad (1)$$

$$d_j \geq ec_j \qquad (2)$$

$$d_j \geq d_i + RS_{ij} \quad \forall\, i \;\text{ s.t }\; ts_i < ts_j \qquad (3)$$

The objectives for the sequencing element of the problem are threefold: reduce the delay for aircraft (or, equivalently, keep the runway throughput high), control inequity in the sequencing (i.e. do not allow any aircraft to be unduly penalised) and ensure that CTOT take-off slots are met where possible (and that extensions are met otherwise). Of course, maintaining safety is imperative through all of this. The separation rules ensure safety at take-off and the solution method for the control problem ensures that the achievement of any predicted sequence will not involve too much work since a controller's time is a limited resource so a controller would reject any sequence which took too long to achieve.

Let $C(j,d_j)$ denote a function to determine the cost associated with any failure to comply with CTOT time-slots if aircraft j takes off at time d_j. Function $C(j,d_j)$ will return 0 if d_j is within the CTOT for aircraft j, a penalty cost if d_j is not within the CTOT for d_j but is within an extension, and a much greater cost if d_j is too late to be within even a CTOT extension (so that any schedule for which this is the case will not be adopted).

Let $E(ts_j,as_j)$ denote a function to penalise positional inequity in the take-off sequence, for example penalising the positional delay (defined as $\max(ts_j - as_j, 0)$) or both positional advancement and delay (for example, $\max(ts_j - as_j, as_j - ts_j)$). The effects of both of these factors were studied in [6], where factors related to the square of the positional delay and of the positional deviation were included in the objective function.

The objectives of the take-off sequencing can then be formalised by (4), where W_1, W_2 and W_3 are weights for the three components, set so that CTOT compliance is of primary importance and delay is of secondary importance. The weight for equity is chosen so that small deviations are of little importance but large deviations are greatly penalised. For the purposes of this paper, the overall delay and CTOT compliance results are more important than the objective function which was used by the sequencing algorithm to obtain the results. Details of functions $C(j,d_j)$ and $E(ts_j,as_j)$ can be found in [6], [11] and [12] along with values for the weights which were used in each case.

$$W_1\, C(j,d_j) + W_2\,(d_j - bt_j) + W_3\, E(ts_j,as_j) \qquad (4)$$

A linear penalty for delay is shown in the objective function (4). Although this form was used for the runway sequencing decision support system, a non-linear penalty is currently being used for the TSAT allocation system, in addition to the factor $E(ts_j,as_j)$, to promote equity in the sequencing. When re-sequencing within the holding areas, experiments showed that the constraints imposed by the holding area structures favoured equitable schedules, since holding an aircraft for too long can block a part of the holding area, thus limiting the improvements that can be made to the sequencing of other aircraft. When take-off sequencing is performed by the TSAT allocation the system (in the absence of the equity-aiding constraints from the holding area structure), the delay is raised to the power 1.5 in the objective function, to further encourage the spreading of the delay more equitably between aircraft. Results are presented in this paper for the TSAT system with both linear and non-linear objective functions.

4 Take-off time prediction, for both systems

The same take-off time prediction method is used by both of the systems described in this paper. Aircraft in a partial or full take-off sequence are considered one at a time, in take-off order and the earliest time which will satisfy constraints (1), (2) and (3) is assigned to each aircraft. The order of evaluation is important since RS_{ij} can depend upon the time of day so, before d_j can be determined, it is necessary to know d_i for each aircraft i which takes off earlier than j. This method for take-off time prediction is known to be pessimistic (as discussed in [6] and [11]) in that it may predict take-off times which are later than would actually happen since the controllers are permitted to reduce some non-safety-related separations at their discretion.

5 Sequencing within the holding area

The physical layout of the holding area structure can be a major problem for re-sequencing aircraft once they reach the holding area. In some airports the physical layout can be restrictive enough to simplify the take-off sequencing problem [7], however the Heathrow holding areas are flexible enough to allow much re-sequencing, but not so flexible that all desirable re-sequencing is achievable (in which case they be ignored by the sequencing algorithm). It was identified in [5] that the removal of the constraints imposed by the holding area could make better take-off sequences achievable, and thus improve the overall delay at take-off.

The runway controller has to consider the locations at which the aircraft are currently situated, the instructions which they have been given so far, and the desirability of the various take-off sequences, then use this information to select a take-off sequence to enact. The holding area movement problem was thoroughly described in [6] and the system that was used in [4], [6], [11], [12] has been utilised for the experiments in this paper which use the runway controller decision support system.

To understand the effects of the holding area, it is useful to consider the way in which a directed graph model of the holding area can be used to determine the feasibility of re-sequencing. One such graph, for the 27R holding area which was used in these experiments, is presented in Fig. 1. The 27R holding area can be considered to consist of three parallel queues, with interchange points between them. Determining the feasibility involves positioning aircraft which are within the holding area at the node to which they are closest. Aircraft which are still on the taxiways are added to queues for the appropriate entrance nodes. The re-sequencing is deemed feasible if aircraft can enter the runway node(s) in the take-off order having moved one arc at a time without any node being occupied by more than one aircraft at any time.

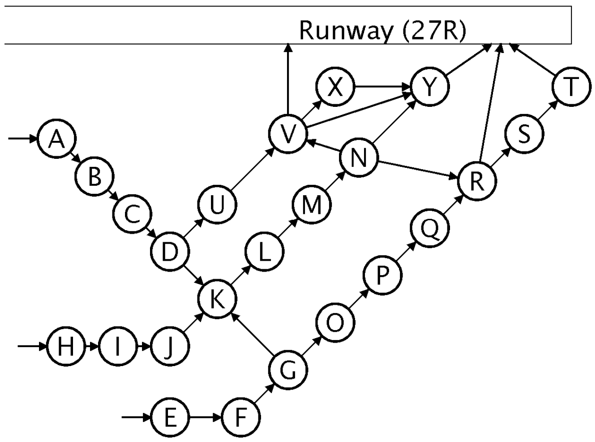

Figure 1. Directed graph of the 27R Holding area, taken from [6]

The sequencing problem faced by the runway controller is complicated by a number of factors. Firstly, the value of et_j can change depending upon the way in which the control problem is handled. For example, allocating a longer path through the holding area to aircraft j would be expected to increase the value of et_j, since aircraft will take longer to reach the runway. The value of et_j can be considered to be a composition of the holding area arrival time and the length of time necessary to traverse the holding area and line up for the runway and can, therefore, only be determined for aircraft once the control problem has been considered. However, the paths which should be allocated to aircraft will depend upon the target take-off sequence, since short paths need to be allocated to those aircraft which will have less time available to traverse the holding area. Furthermore, it is important for sequencing to be achievable in a manner which will be sensible to the runway controller and keep the workload low.

6 Solution method for sequencing at the holding area

The solution method for this problem was described in [4], [6], [11] and [12]. A tabu search meta-heuristic is used to investigate the potential take-off sequences, utilising moves to swap the positions of aircraft, shift multiple aircraft forwards or backwards in the sequence, or randomly re-sequence a continuous sub-sequence of two to five aircraft. Each sequence is evaluated individually. Paths through the holding area are heuristically assigned to aircraft according to the overtaking which they need to perform. This is performed in such a manner that et_j will usually be as low as possible for any aircraft j for which et_j is the binding constraint upon the take-off time. Partial sequences of movement are generated for the nodes close to the entrances and exits of the holding area, and counts of aircraft and spare nodes are maintained. A look-ahead algorithm is then used to move aircraft within the holding area graph, taking consideration of the aircraft counts in order to avoid blocking. The sequence is rejected if either it is infeasible or the controller is unlikely to accept the holding area movement which would be required to achieve it (for example, if the required paths would not appear sensible, or the sequencing of movement is too complex). This ensures that any adopted take-off sequence will be relatively easy to achieve.

The outputs of the control problem stage consist of an allocated path for each aircraft, an earliest take-off time and full details of how the re-sequencing would be achieved. The derivation of the earliest take-off time (et_j for aircraft j) takes into consideration the position of each aircraft, the allocated paths, the arrival times of aircraft at the holding area and any delay which is necessary due to waiting for other aircraft to move out of the way. The movement details are used by the simulation which evaluates the performance of the system in order to predict the positions of aircraft over time. They are also very useful for watching playback of the movement, to ensure that the system is working as intended and that all movement within the holding area is both as expected and acceptable. Once a take-off sequence has been determined to be feasible, and the et_j values for each aircraft have been determined, the sequence is evaluated to determine the take-off times (where d_j is assumed to be the earliest value which will satisfy (1), (2) and (3)) and the consequent cost of the sequence, as measured by (4).

7 Considering the sequencing earlier in the departure system

Given the difficulties of re-sequencing at the holding area, it is worth considering whether the re-sequencing can be performed any earlier in the departure system. To determine a take-off sequence and take-off times, it is necessary to know the earliest times at which each aircraft can reach the runway (et_j for each aircraft j). This requires a prediction for each aircraft for both the taxi time from the stand to the runway, and an earliest pushback time from the stand. Is has not previously been feasible to determine earliest pushback times, but recent innovations at Heathrow will make it more so. A Collaborative Decision Making (CDM) project is currently being implemented at Heathrow. For the first time, this will automate the sharing of information between airlines, ground handlers and controllers, potentially enabling all parties involved to make better decisions. For the purposes of this research, this means that earliest pushback times (called TOBTs, Target Off-Block Times, the times at which the airline believes the aircraft will be ready for pushback) will be provided by airlines to controllers. In combination with taxi durations (discussed below), this will make earliest take-off times more predictable. Controllers can then use the predicted take-off times to generate TSATs (Target Start-At Times) to pass back to airlines specifying the time at which the aircraft will actually be permitted to push back. The CDM project means that airlines gain greater visibility of expected delays, and controllers gain the knowledge that is necessary to predict take-off delays and assign appropriate stand holds.

7.1 Taxi-out duration

Taxi-out duration (or time) often refers to the time from leaving the stand to line-up on the runway, and may be used as a measure of the delay that aircraft should expect from stand/gate to take-off [13]. Since the queueing time can form the major element of the delay for aircraft, this measure of taxi time can be highly dependent upon the state of the airport at the time. This definition is, thus, not necessarily the most useful measure of the taxi time if the aim is to determine the earliest time at which a take-off could be scheduled.

At Heathrow, the taxi duration can be considered to have three main components: a delay from wishing to pushback to actually having completed pushback (consisting of a pushback duration plus any waiting time due to contention with other aircraft, as discussed later), the travel time from the pushed back position to the holding area queue, and some queueing and lining up time at the runway. The queueing time at the runway will depend upon the take-off sequence and is thus relatively predictable for a known take-off sequence, as long as the arrival time at the holding area is known and sufficient time is allowed for aircraft to traverse the holding area and line up for take-off. The taxi time from cul-de-sac to runway is believed to be relatively predictable, once the delays around the stands and at the runway are removed from the problem. Thus, there is a need to predict the final component of the taxi duration, the delay around the stand.

7.2 Cul-de-sac contention

Most stands at Heathrow are grouped around the terminal piers, as illustrated in Fig. 2. Cul-de-sacs are formed between the piers, with one end of the cul-de-sac opening out onto the taxiways and the other very close to the main part of the terminal. Due to these cul-de-sac structures, aircraft may block other aircraft from pushing back or from leaving the cul-de-sac until they have started their engines and moved out of the way. This causes delays for the blocked aircraft, which can traditionally appear as unpredictable increases in the taxi times. By explicitly modelling the contention within the cul-de-sac, the unpredictability of these cul-de-sac delays can be greatly reduced.

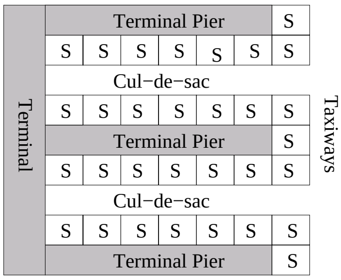

Figure 2. Stylised diagram of how cul-de-sacs are formed between the stands (labelled S) at the terminal piers

8 TSAT allocation

The TSAT allocation problem involves assigning pushback times (Target Start-At Times, TSATs) to aircraft such that excess delay can be absorbed at the stand without reducing the runway throughput. The solution method that has been implemented for this problem involves predicting a take-off sequence for aircraft while they are still at the stand, then using this to determine how much stand hold to assign to each aircraft. The main elements of such a system are a method to determine the amount of cul-de-sac delay that aircraft will expect (by modelling the contention between aircraft which are in the same cul-de-sac or at stands which are close together), a consideration of the take-off sequencing problem (to determine a good take-off sequence), and a method to determine the appropriate stand hold to allocate to an aircraft given a predicted take-off time (so that excess delay is absorbed at the stand but enough slack is available to allow for variation between real and actual taxi times).

The aim is to absorb any delay beyond a specified threshold value (labelled IRH_j for aircraft j in the model presented later) at the stand instead of the runway, rather than to attempt to absorb all of the delay at the stand. The characteristics of good take-off sequences mean that they tend to group aircraft by weight class, and that aircraft alternate departure routes. This means that aircraft in alternate positions in the take-off sequence often have similar (but not usually exactly the same) characteristics. Alternative good (but not necessarily quite as good) take-off sequences can often be generated from a good take-off sequence by switching the positions of aircraft which take off at similar times, especially those which take off two positions apart in the take-off sequence.

The TSAT system can only have an adverse effect upon the potential take-off sequencing if aircraft are held at the stands to the point where the runway controller is starved of the appropriate aircraft to achieve a good take-off sequence. At quieter times, when delay at the runway is low, stand holds will very rarely get allocated, in which case TSAT allocation will have no adverse effect upon the potential sequences. As the departure system get busier and delay increases, the benefits of the system become more apparent as stand holds increase. However, the threshold value for the runway hold ensures that alternative aircraft are usually available when unexpectedly long taxi times mean that a specific aircraft does not arrive in time. Thus, good alternative take-off sequences should be available when the predicted take-off sequence cannot be achieved. The potential for the achieved take-off sequence to deviate from the predicted take-off sequence means that the TSAT system needs to find good take-off sequences, but there is not necessarily a benefit from finding the optimal take-off sequence while aircraft are still at the stands.

The primary benefit of the TSAT allocation system is to reduce the waiting time with engines running while ensuring that the better take-off sequences are no less likely to be achieved than without the stand holds. However, a potential secondary benefit is to simplify the problem for the runway controller. The reduction in aircraft at the holding area at busy times should reduce the complexity of the sequencing problem, reducing the amount of blocking within the holding area and enabling the controllers to achieve the better take-off sequences which may not normally be available due to the restrictions of the physical layouts of the holding areas.

9 Take-off sequencing with cul-de-sac contention

The take-off sequencing aspect of the problem is very similar to that considered by the runway controller. The same objectives (as expressed by (4)) and constraints (expressed by (1), (2) and (3)) apply to the sequencing element of the problem. The main differences in this case are, firstly, that the control problem at the holding area is ignored (on the basis that so few aircraft will actually be at the holding area that any re-sequencing there will be achievable) and, secondly, that the earliest take-off time will depend upon the taxi duration and the time at which the aircraft leaves the cul-de-sac, which will itself depend upon the contention within the cul-de-sacs. With the earlier definition for bt_j as the historic holding area arrival time, $d_j - bt_j$ will be the sum of the stand delay and the time spent in the holding area. In a real system, in the absence of historic holding area arrival times, the equation $bt_j = ept_j + pd_j + td_j$ would be used to predict a holding area arrival time with no stand hold, with the following definitions for ept_j, pd_j and td_j.

Let pt_j denote the known or predicted pushback time of aircraft j. Let pd_j denote the known or predicted pushback duration for aircraft j. This is the time taken to push back, start engines, and be ready to taxi towards the runway. Let td_j denote the taxi duration for aircraft j, from the cul-de-sac to the holding area. Let ept_j denote the earliest pushback time for aircraft j (the TOBT). Let ct_i denote the cul-de-sac time for aircraft i, defined as the time at which i will be able to commence its taxi towards the runway without being blocked from doing so by other aircraft. Similarly, let ct_j denote the cul-de-sac time for aircraft j. Let cs_i and cs_j denote the positions of aircraft i and j (respectively) in the sequence in which aircraft commence their taxi towards the holding area, so that aircraft with earlier cul-de-sac times have lower values for cs_i.

$$et_j = ct_j + td_j + MRH_j \quad (5)$$

$$ct_j \geq ept_j + pd_j \quad (6)$$

$$ct_j \geq ct_i + MS_{ij} \quad \forall\, i \quad \text{s.t} \quad cs_i < cs_j \quad (7)$$

$$pt_j = ct_j - pd_j \quad (8)$$

Inequalities (1), (2) and (3) still apply to the take-off times. The value of et_j is now related to the cul-de-sac time of aircraft j, as shown by (5), where MRH_j is a minimum runway hold that can be allocated to aircraft j. The value of MRH_j should be large enough to allow for the time to traverse the holding area and line up (this normally requires less than a minute) and may also include some slack to allow for taxi time variability. Given that this value is used to determine how early an aircraft can be scheduled to take off, it is important not to set it to too great a value. Doing so would involve either an unnecessary wait for take-off if the sequence was adopted or, more likely, the controller would adopt a (better) sequence which would avoid the unnecessary delay.

The contention within the cul-de-sac can be modelled as minimum separation requirements between the cul-de-sac times for aircraft, so that each aircraft can be moved out of the way before another aircraft has to use that section of a cul-de-sac, as described in [14]. Let MS_{ij} denote a minimum cul-de-sac time separation that is required between aircraft i and j, when i leaves the cul-de-sac prior to j. The constraints upon cul-de-sac times can then be described by (7). The cul-de-sac times are related to the earliest pushback times (or TOBTs) by (6). Finally, given a cul-de-sac time, it is sensible for an aircraft to push back just in time to meet its calculated cul-de-sac time, so (8) can be used to determine a pushback time from a cul-de-sac time. The output of the take-off sequencing stage is, thus, a take-off time and feasible pushback time for each aircraft.

10 The push-back time allocation sub-problem

The pushback time allocation problem can be considered as one of assigning the optimal cul-de-sac times given the predicted take-off times. Rather than assign stand holds to absorb all of the delay, in which case achieving the target take-off schedule would depend upon all aircraft achieving the predicted pushback times and taxi times, the aim is to absorb all delay beyond a specific threshold value. An ideal runway hold, IRH_j for aircraft j, is determined, which specifies the target time prior to take-off at which j should be planned to arrive at the holding area - providing that sufficient time is available. In other words, for any aircraft j, IRH_j is the threshold duration beyond which hold should be absorbed as stand hold, and below which it should be absorbed as runway hold. Obviously, the values that are used for IRH_j will affect the balance between the expected runway and stand hold.

$$ict_j = \max(\, ept_j + pd_j,\ d_j - IRH_j - td_j\,) \tag{9}$$

$$\sum_j (\, 100 \max(\, 0,\ ct_j - ict_j\,)^{1.1} + \max(\, 0,\ ict_j - ct_j\,)^{1.1}\,) \tag{10}$$

Given an ideal runway hold, an ideal cul-de-sac time, ict_j, can be determined for each aircraft j using (9). A cost can then be associated with any deviation of the cul-de-sac time for aircraft j from the ideal time, and the objective for the pushback time allocation stage is to reduce the sum of these costs for all aircraft under consideration (10). The nonlinear cost of deviation is again used to slightly favour more equitable distributions of any necessary deviations. Weights of 100 and 1 are used for the two elements in (10) so that the algorithm favours aircraft pushing back early (increasing the runway hold time) rather than late.

10.1 Ideal and minimum runway hold values

The values that are used for MRH_j and IRH_j should be determined according to the reliability of the taxi time data. IRH_j and MRH_j have different but complementary effects. MRH_j can be used to ensure slack for the taxi times, while IRH_j can be used to favour schedules with more slack where possible, without prohibiting good schedules which do not allow so much slack. For the purpose of this paper, a minimum hold of two minutes was chosen so that the earliest take-off times are compatible with those for the runway sequencing decision support system values. An ideal runway hold of five minutes was chosen for all aircraft. This is short enough that some stand hold will be allocated to aircraft but large enough to provide a pool of a few aircraft at the holding area at most times.

11 Solution method for the TSAT problem

The implemented solution method for the TSAT allocation problem first generates a take-off sequence and good feasible take-off times, then uses this to assign appropriate stand holds to aircraft. The two sub-problems are solved separately, although (1) and (5) mean that cul-de-sac times have to be considered when considering the take-off sequencing, so that the second stage has at least one feasible solution. The cul-de-sac times depend, of course, upon the cul-de-sac sequence (due to (7)). The fact that cul-de-sac separations are both asymmetric and do not obey the triangle inequality means that the best cul-de-sac sequence does not necessarily match the planned take-off sequence, however, only the feasible rather than optimal cul-de-sac times are required at the take-off sequencing stage.

The developed solution method first heuristically generates an initial take-off sequence (based upon allocated CTOTs and TOBTs), then incrementally improves this using a branch-and-bound algorithm. The branch-and-bound algorithm is applied within a rolling window through the sequence (from earliest take-off to latest take-off), finding the optimal sequence of the aircraft within the current window. It assumes that the aircraft prior to the window have a fixed take-off sequence (and take-off times) and ignores aircraft beyond the end of the window. This solution approach was described in more detail in [14].

The branch and bound algorithm that is applied each window considers partial take-off sequences, starting from the first take-off within the window and consisting of consecutive take-offs. Aircraft are added to the sequence one at a time. As aircraft are added, a take-off time is determined for the aircraft, along with a feasible cul-de-sac time. If the cul-de-sac time is the limiting factor upon the take-off time, the cul-de-sac sequencing problem is solved to determine the cul-de-sac times which will result in the earliest take-off time for the current aircraft. In this way, the take-off times for aircraft in partial take-off sequences can be determined, along with a consequent cost for the partial sequence. A lower bound is then determined for the cost of the remaining aircraft in the window which have not yet been added. If the resulting lower bound for the cost of any sequence with this partial take-off sequence as a prefix exceeds the cost of a previously evaluated full take-off sequence, the partial sequence is pruned from further consideration.

The fact that exchanging two aircraft in a take-off sequence can often generate a very similarly valued take-off sequence is a useful property in that it allows recovery from unpredictable taxi times, but it is a problem for the branch and bound algorithm. There will usually be a large variety of similarly costing schedules with different initial partial take-off sequences, thus the potential for early pruning of partial schedules is limited. However, take-off times can only be predicted accurately when a partial sequence consists of a consecutive set of aircraft with some known starting time, so it is necessary to build-up sequences in this way. Although the branch and bound significantly reduces the number of sequences which must be considered, the solution time still rises rapidly as the window size increases.

In the presence of CTOTs it may be necessary for specific aircraft to move a long way in the take-off sequence (as shown in [6]) and, although the heuristic pre-sequencing helps with this, it is not always sufficient. In order to allow aircraft to be advanced a large number of places when necessary, the window size is kept relatively low and the rolling window is applied multiple times to provide further incremental improvements.

Once take-off times and feasible cul-de-sac times have been determined, a second stage algorithm is applied to optimise the cul-de-sac times, as defined by (10). Each aircraft will have an ideal cul-de-sac time (calculated from its take-off time, using (9)), an earliest cul-de-sac time (calculated from its TOBT, using (6)) and a latest cul-de-sac time (provided as input information or calculated from the take-off time as the latest time which would allow (5) to hold). Despite the non-linear costs of deviation from ideal times, this is usually a relatively simple problem to solve, since it can be decomposed into small sub-problems of mutually dependent aircraft. Since TSATs must lie on minute boundaries, the number of possible TSATs is usually very limited. A simple branch and bound algorithm is again used.

12 Comparing solution methods

The solution methods for the two take-off sequencing problems differ. Sequencing at the runway is performed using a meta-heuristic search. Sequencing at the stands is performed using a branch and bound method within a rolling window. This is due to fundamental differences in the problems. Firstly, sequencing at the runway has to be extremely fast, so a heuristic method is required, whereas the longer time available for TSAT allocation (and the ability to ignore the structure of the holding area when sequencing at the stands) makes the branch-and-bound possible within a window. In a dynamic situation, the TSAT allocation problem will have to cope with latest TSAT constraints as well as earliest TSAT constraints. These windows can make it very difficult to find feasible solutions that meet the constraints, limiting the effectiveness of a meta-heuristic search. Although the branch-and-bound solution method works within a window, in the presence of latest TSAT constraints it has to also take into consideration the aircraft outside of the current window in order to ensure feasibility of later TSAT allocations. These considerations are irrelevant for the experiments described in this paper, but important motivations for the system design when considering the real problem.

The second reason for the difference is the value of a partial take-off sequence. Issues related to the control problem at the holding area mean that partial take-off sequences are of limited value for the decision support system for the runway controller. Adding new aircraft to the sequence may require path changes for previous take-offs, which may alter take-off times. Furthermore, the movement of previous take-offs may greatly restrict the possibilities for re-sequencing later take-offs which had to be overtaken. In contrast, when sequencing at the stands, the value of a partial take-off sequence is a good indication of the value of a full take-off sequence, thus evaluating partial sequences within the branch-and-bound algorithm can avoid a lot of re-calculation of cul-de-sac times.

13 Comparative results

Experiments were performed using two datasets which had sufficient information to be used by both the TSAT allocation system and the decision support system for the runway controller. All experimental data was for the 27R holding area and avoided runway change-over times. The first dataset contained details of 345 consecutive take-offs and the second 341 consecutive take-offs.

The decision support system for the runway controller was executed within the simulation of the departure system that was described in [4], [6], [11]. Each evaluation involved a number of iterations of the simulation. In each iteration, the system was presented with a problem to solve before the time was advanced by one minute, and the new, consequent, problem was presented in the next iteration. Once all aircraft in the dataset had taken off, the resulting overall take-off sequence was evaluated. The decision support system has a stochastic element, so each evaluation was executed ten times and the mean results are presented. It should be noted that the results for each evaluation were very similar.

For comparison, the TSAT allocation system was applied to the static problem, considering the entire dataset of over 340 aircraft at once. Rolling windows of various window sizes were applied and the resulting take-off sequences and TSATs were evaluated. Of course, there is no intention to communicate to the controllers the predicted take-off times or sequences, but consideration of the internally predicted sequences is useful for comparing the sequencing at the stands against that at the runway.

For the runway controller decision support system, historic holding area arrival times were used and the system was permitted to schedule aircraft to take-off from two minutes after the historic holding area arrival time. For the TSAT allocation system, historic pushback times were used as earliest pushback times. Historical taxi times were also used, so that an aircraft without any stand hold, and that pushed back as soon as it could, would have the historic holding area arrival time. A two-minute minimum runway hold was applied, so that an aircraft could not take off within two minutes of its predicted holding area arrival time. Thus, with no stand hold, the two systems had the same earliest take-off times for aircraft, although stand contention had to be considered for the TSAT allocation system.

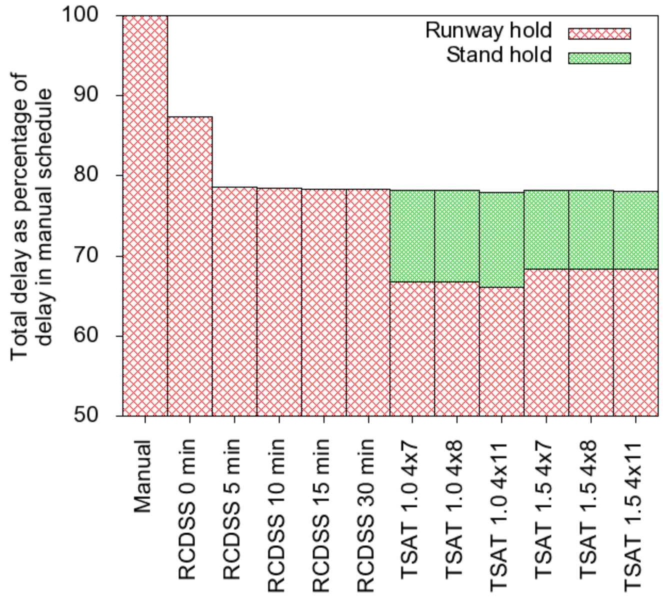

Figure 3. Comparative delays, dataset 1

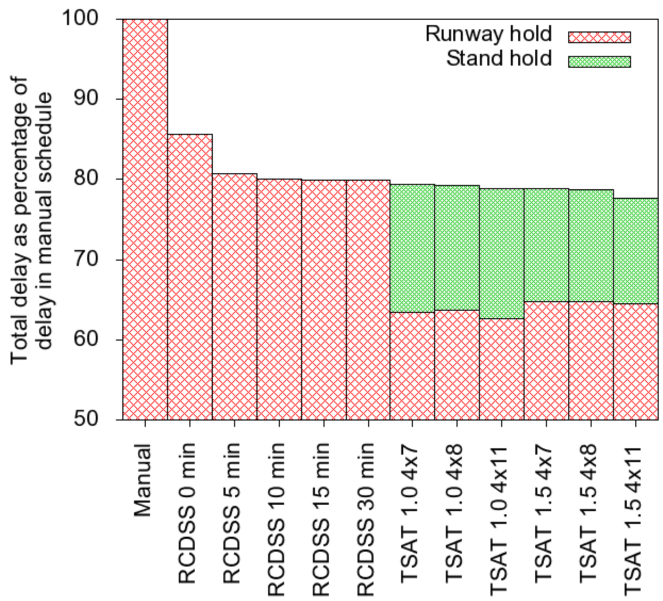

Figure 4. Comparative delays, dataset 2

There were three aims for these experiments. Firstly to obtain an idea of how much potential benefit could be gained from a decision support system which was able to act upon more information than a controller could be expected to cope with. Secondly, to identify how the take-off sequencing elements of the two systems performed when compared against each other. Thirdly, to obtain insight into how much of any prospective delay could reasonably be absorbed at the stands prior to starting the engines.

Figs. 3 and 4 illustrate the performance of the decision support system for the runway controller for each of two datasets. All values are shown as percentages of the delay obtained by the controllers for the manually produced take-off sequences with real take-off times. For reference, the manual results had mean delays per aircraft of 311s and 340s for datasets 1 and 2 respectively. In each case, the pessimistic take-off time prediction system was used to generate the take-off times for the automatically generated take-off sequences. i.e. the earliest take-off time which will meet constraints (1), (2) and (3) is assumed for each aircraft.

The results labelled 'Manual' show the real delay in the manually produced sequence. The 'RCDSS 0 min', 'RCDSS 5 min', 'RCDSS 10 min', 'RCDSS 15 min' and 'RCDSS 30 min' results show the performance of the decision support system for the runway controller with knowledge of aircraft 0, 5, 10, 15 and 30 minutes respectively prior to the aircraft arriving at the holding area. The 'TSAT 1.0 4x7', 'TSAT 1.0 4x8' and 'TSAT 1.0 4x11' results show the performance of the TSAT allocation system with linear costs for the delay, 4 passes of the rolling window and window sizes of 7, 8 and 11 aircraft respectively. The 'TSAT 1.5 4x7', 'TSAT 1.5 4x8' and 'TSAT 1.5 4x11' are identical to the 'TSAT 1.0 4x7', 'TSAT 1.0 4x8' and 'TSAT 1.0 4x11' experiments respectively except that a non-linear cost was applied for delay (the number of seconds delay in objective function (4) was raised to the power of 1.5).

The performance of the decision support system for the runway controller depends upon how early the aircraft enter the system and how many aircraft are considered. This was discussed further in [6] and [15]. In particular, the tested configuration of the system assigned paths to aircraft as soon as they arrived at the holding area. This meant that the system performance when it was only made aware of aircraft at the point at which they arrived at the holding area (i.e. 'RCDSS 0 min') was significantly worse than when the system had knowledge of them prior to holding area arrival. System performance was slightly better when more aircraft could be considered (i.e. when aircraft entered the system earlier) but the differences were relatively small. The simulations predicted that the decision support system was able to find schedules which reduced the delay by over 20%, a significant saving in time and fuel.

The performance of the take-off sequencing element of the TSAT allocation system was slightly better than that of the decision support system for the runway controller, even with low window sizes. This was expected for two reasons. Firstly, consideration of the entire static problem provides the TSAT allocation system with greater knowledge of aircraft than is available to the runway controller's decision support system. Secondly, under this formulation, it is possible to hold aircraft at the stand for arbitrarily long if necessary without reducing the sequencing flexibility, whereas the decision support system for the runway controller has to work within the holding area constraints and it was previously shown in [5] that the holding area constraints affect the performance of the sequencing.

Since stand delay does not contribute to fuel usage, the results for the TSAT system are significantly more environmentally friendly than those for the runway sequencing. Not only are similar total delay reductions achieved, but a significant amount of the delay can be achieved as stand hold rather than delay at the holding area. For the experiments performed for this paper, this resulted in an additional 9.9% to 16.1% of the full delay of the manual schedule being absorbed at the stand.

Interestingly, considering equity within the objective function by using a non-linear cost for delay actually slightly reduces the overall delay in the schedule for dataset 2, rather than increasing it. More importantly, the amount of delay which can be absorbed at the stand is higher for the linear delay cost. The mean delay values for datasets 1 and 2 were 4.1 and 4.5 minutes respectively. Only aircraft with greater than a 5 minute runway hold will receive any stand hold. A less equitable take-off sequence will have more deviation of delays away from the mean, so more aircraft with a delay greater than five minutes, thus more aircraft will have stand holds. However, these schedules will also have more aircraft with lower runway holds so would be expected to be less robust due to the lower slack and smaller pool of aircraft at the holding area. This robustness effect is another reason for promoting equity of delay and would comprise as interesting area for further research. Tuning of these parameters for the specific situation at Heathrow is a planned future stage of the TSAT implementation project.

It should be noted that these delay improvements in the automated schedules were not at the expense of CTOT compliance. The manually produced schedules missed 5 and 6 CTOTS respectively for datasets 1 and 2. In all cases, the automated schedules missed 1 CTOT for dataset 1, and between 3 and 5 for dataset 2. 'RCDSS 0 min' missed 5 CTOTs for dataset 2. 'RCDSS 5 min' and 'RCDSS 10 min' missed 4 CTOTs for dataset 2. 'RCDSS 15 min', 'RCDSS 30 min' and all of the TSAT results missed 3 CTOTs for dataset 2. Thus, in all cases, the automated schedules required fewer CTOT extensions than the manual schedules.

14 Potentially linking the two systems

Although the two models and systems could be used independently, better results could reasonably be expected from using both simultaneously. The primary benefit of the decision support system for the runway controller is that it provides visibility of the taxiing aircraft, and enables sequencing decisions to take into account more aircraft - with consequent reductions to the total delay. The presence of a TSAT allocation system should ensure that the holding area is not congested with aircraft which have long waits. Given that a good potential take-off sequence was originally planned by the TSAT system, there may be some benefit from giving visibility of this intended sequence to the runway controller. For example, a link between the two systems could be used to provide the decision support system for the runway controller with an initial sequence based upon that generated by the TSAT system, helping to guide the meta-heuristic search. The decision support system for the runway controller would then take into account the deviations between the predictions that the TSAT system made and what happened in reality and produce a recommended take-off sequence for the runway controller. The two systems would, therefore, potentially complement each other.

15 Conclusions

This paper presented a summary of some of the recent research involving the departure system that has taken place for London Heathrow airport. Two examples of take-off sequencing were considered. Each was applied at a different part of the departure process and, consequently, had different constraints upon it. In both cases the problem was complicated by not physically being able to achieve all desired take-off sequences due to contention between aircraft at the runway holding areas or in the cul-de-sac.

The first problem considered take-off sequencing from the point of view of the runway controller, and the second considered performing predictive sequencing earlier in the departure system, while aircraft are still at the stands, and (rather than doing anything with the sequence itself) using this to allocate an appropriate stand hold to aircraft so that an appropriate amount of any expected delay is absorbed prior to starting the engines.

The sequencing elements of the TSAT system and the decision support system for the runway controller were finding take-off sequences with very similar total delays, although the TSAT system performed slightly better. Of course, these results assumed predictable taxi times and pushback times. A pessimistic take-off time prediction system was used, and unnecessarily large minimum runway hold values were specified, so a real controller could potentially perform even better if given the appropriate information in a timely manner.

These results show that either of the two described systems could potentially provide great benefits at Heathrow, which is precisely why the TSAT system was developed. Moreover, the benefits of the TSAT system are more than just a potential for a reduction in the total delay, since a considerable amount of the remaining delay could potentially be absorbed at the stand. The improved selection of aircraft to provide to the runway controller which should result from the implementation of a TSAT system should aid the runway controllers by reducing the congestion which can prevent better sequences from being enacted.

Acknowledgements The authors wish to thank NATS (formerly National Air Traffic Services Ltd) for partially funding this work and for providing time with their controllers and analysts to ensure that the work is relevant for the real situation at Heathrow. We would also like to thank the Engineering and Physical Sciences Research Council for providing the majority of the funding.

References

1. A. T. Ernst, M. Krishnamoorthy, and R. H. Storer, "Heuristic and exact algorithms for scheduling aircraft landings," Networks, vol. 34, pp. 229– 241 (1999)
2. J. E. Beasley, M. Krishnamoorthy, Y. M. Sharaiha, and D. Abramson, "Scheduling aircraft landings -the static case," Transportation Science, vol. 34, pp. 180–197 (2000)
3. H. Pinol and J. Beasley, "Scatter search and bionomic algorithms for the aircraft landing problem," European Journal of Operational Research, vol. 171, pp. 439–462 (2006)
4. J. A. D. Atkin, E. K. Burke, J. S. Greenwood, and D. Reeson, "A meta-heuristic approach to aircraft departure scheduling at London Heathrow airport," in Computer Aided Systems of Public Transport, ser. Lecture Notes in Economics and Mathematical Systems, M. Hickman, P. Mirchandani, and S. Voss, Eds. Springer, Berlin (2008)
5. J. A. D. Atkin, E. K. Burke, J. S. Greenwood, and D. Reeson, "An examination of take-off scheduling constraints at London Heathrow airport," in proceedings of the 10th International Conference on Scheduling of Public Transport (CASPT), Leeds, UK (2006)
6. J. A. D. Atkin, "On-line decision support for take-off runway scheduling at London Heathrow airport," Ph.D. dissertation, University of Nottingham (2008)
7. M. A. Bolender and G. L. Slater, "Analysis and optimization of departure sequences," in proceedings of the AIAA Guidance, Navigation and Control Conference and Exhibit, Denver, CO, August 2000, pp. 1672– 1683
8. I. Anagnostakis and J.-P. Clarke, "Runway operations planning, a two-stage heuristic algorithm," in proceedings of the AIAA Aircraft, Technology, Integration and Operations Forum, Los Angeles, CA, October 2002, report AIAA-2002-5886
9. I. Anagnostakis and J.-P. Clarke, "Runway operations planning: a two-stage solution methodology," in proceedings of the 36th Hawaii International Conference on System Sciences (HICSS-36), Hawaii, vol. 3, January 2003
10. H. N. Psaraftis, "A dynamic programming approach for sequencing groups of identical jobs," Operations Research, vol. 28, no. 6, pp. 1347– 1359 (1980)
11. J. A. D. Atkin, E. K. Burke, J. S. Greenwood, and D. Reeson, "Hybrid meta-heuristics to aid runway scheduling at London Heathrow airport," Transportation Science, vol. 41, no. 1, pp. 90–106 (2007)
12. J. A. D. Atkin, E. K. Burke, J. S. Greenwood, and D. Reeson, "On-line decision support for take-off runway scheduling with uncertain taxi times at London Heathrow airport," The Journal of Scheduling, vol. 11, no. 5, pp. 323–346 (2008)
13. H. R. Idris, J.-P. Clarke, R. Bhuva, and L. Kang, "Queuing model for taxi-out time estimation," Air Traffic Control Quarterly (ATCA Publications), vol. 10, no. 1, pp. 1–22 (2002)
14. J. A. D. Atkin, E. K. Burke, and J. S. Greenwood, "The TSAT Allocation System at London Heathrow: The Relationship Between Slot Compliance, Throughput and Equity,",in proceedings of the 11th Conference on Advanced Systems for Public Transport (CASPT), Hong Kong, July 2009.
15. J. A. D. Atkin, E. K. Burke, J. S. Greenwood, and D. Reeson, "The effect of the planning horizon and the freezing time on take-off sequencing," in proceedings of the 2nd International Conference on Research in Air Transportation (ICRAT), Belgrade, Serbia and Montenegro, June 2006.