An examination of take-off scheduling constraints at London Heathrow Airport

Jason Atkin[1], Edmund Burke[1], John Greenwood[2], and Dale Reeson[3]

[1] School of Computer Science and Information Technology, University Of Nottingham, Jubilee Campus, Wollaton Road, Nottingham, NG8 1BB, England
`jaa/ekb@cs.nott.ac.uk`
[2] NATS CTC, 4000 Parkway, Whiteley, Fareham, Hampshire, PO15 7FL
[3] National Air Traffic Services, Heathrow Airport, Hounslow, Middlesex, TW6 1JJ

Abstract:

In this paper, we focus upon the departure system for London Heathrow airport, one of the busiest airports in the world. Decreasing the delay for aircraft awaiting take-off with their engines running would decrease fuel usage and have consequent cost and pollution benefits. We explain how the departure system at Heathrow currently works and we describe the various constraints that apply to take-off schedules. A model for the take-off order problem is presented from the point of view of the runway controller, the person who is responsible for the take-off scheduling. We investigate the effects of each constraint and combination of constraints, using a simulation of the Heathrow departure system. The role of the runway controller in the simulation is performed by a search which was designed to form the basis of an online decision support system. Both the simulation and the decision support system are fully described. We use the results to evaluate the effect upon delay that we would expect from various changes that could be made to the departure system. We end the paper by drawing conclusions about the predicted effectiveness of different changes that could be made to the departure system and focus upon a further opportunity for decision support research.

1 Introduction and problem description

At London Heathrow, there are two runways available. However, in order to control noise, only one may be used for arrivals and the other for departures at any time.

Departing aircraft taxi from the stand, by the terminal, around the taxiways to holding area structures near the end of the current departure runway. There they are reordered for take-off by a runway controller. These holding areas are different for each runway end and comprise one or more entrances

from the taxiways, some manoeuvring space and one or more exits onto the runway. The amount of reordering that is possible is limited and depends upon the structure of the holding areas in use at the time. Despite the constraints upon the re-sequencing that are enforced by the holding area structures, re-sequencing is performed there in order to avoid the intrinsic uncertainties associated with re-sequencing aircraft at the stands, such as contention between aircraft at pushback time and variations in taxi times. These issues are discussed in more detail in section 9.

Separations must be imposed between aircraft taking off, to ensure safety both at take-off and in-flight. These separations are sequence dependent, thus the take-off order can significantly affect both the throughput of the runway and the consequent delay for departing aircraft. Wake vortices are left behind aircraft at take-off; the size depending upon the size of the aircraft. Following aircraft may be affected by the wake vortices if they take off too soon; smaller aircraft being more affected than larger aircraft. A minimum separation time is therefore mandated between any pair of aircraft, being larger whenever a smaller class of aircraft follows a larger class of aircraft. For this reason it is often beneficial to consecutively schedule several aircraft of the same weight class, reducing the number of larger separations needed.

Aircraft that take off have to maintain a given safe separation distance in flight. To ensure that the in-flight separation distance will be attained, a second type of minimum separation time is also required between aircraft at take-off. Aircraft depart along specific, pre-defined departure routes called Standard Instrument Departure routes, or SIDs. This second separation depends upon the departure routes of the aircraft; a larger separation being needed if routes are similar than if they diverge very quickly. It can be worked out from the SIDs of the two aircraft and is then modified by the speed groups of the aircraft to allow for divergence or convergence in flight due to the following aircraft being slower or faster than the preceding one. The modification rules also depend upon the SIDs that are being used.

Wake vortex separation rules depend upon the weight class of the aircraft and are asymmetric. SID separations depend upon the departure routes and speeds of the aircraft and are not only asymmetric but also do not obey the triangle inequality. It is therefore not adequate to only verify the compliance of separations between adjacent take-offs.

Some aircraft are given a *'Calculated Time Of Take-off'* (CTOT), assigned to avoid congestion en-route or at destination airports. These aircraft must take off within a fifteen minute window, from five minutes before the CTOT to ten minutes after the CTOT. This can make the scheduling hard to achieve at times since aircraft that arrive late but have a tight CTOT could have to overtake aircraft that have been waiting longer.

The final major constraint is that of the holding area structure. Since the overtaking is performed within the holding area, near the end of the runway, there is a limit to the level of overtaking that can be performed. In many

cases, whether an aircraft can overtake another aircraft depends upon what other aircraft are doing and how congested the holding area is at the time.

The take-off sequencing problem that we address in this paper can be summarised as determining take-off orders which reduce the delay for aircraft while obeying separation rules, ensuring as many aircraft take off within their CTOT (if they have one) as possible and that the necessary reordering is actually achievable, without a prohibitive workload for pilots or controllers.

In this paper, we will be using a model we designed for a decision support system for the runway controller to evaluate the relative effects of the four major constraints upon the scheduling: departure route separations, wake vortex separations, CTOTs and the physical constraints imposed by the holding area structure. This research is motivated by a desire to better understand the benefits that could be gained from different types of changes to the departure system. This motivation is more fully explained in the conclusions section.

We will discuss previous related research in section 2, introduce our model for the departure problem in section 3, explain our solution system in section 4 and detail our simulation to validate the results in section 5. In section 6, we explain how we used our simulation to obtain the results we present in section 7. In section 8, we draw a number of conclusions from the results, explaining how they could lead to different improvements that could be made to the departure system. Finally, in section 9 we focus upon one change that could be made and discuss how it changes the problem for any decision support system.

2 Previous research

Idris examined the departure system at Boston Logan airport in [14] and [15], concluding that the runway is the primary bottleneck for the departure system. This is also the case at Heathrow. Newell presented a model for airport capacity in [18], showing that mixed mode, where each runway is used for both arrivals and departures, is the more efficient mode of operation for two-runway airports like Heathrow. The fact that this is severely limited at Heathrow due to noise reduction agreements with local residents further adds to the bottleneck problem.

The scheduling is not performed in holding areas at all airports so the physical constraints imposed by the holding area structures do not then apply. Anagnostakis and Clarke in [3] and [4] suggested a two-stage departure planner, where the first stage ignores the downstream constraints such as the SID-based separation rules and CTOTs. We will see in the results section of this paper that the downstream constraints are much more important at Heathrow given its position in very busy airspace.

Anagnostakis et al described a solution method using a search tree and an A* algorithm or branch and bound approach in [2], van Leeuwen et al presented a constraint satisfaction based model in [17] and Trivizas suggested a

dynamic programming approach utilising a maximum position shift to further reduce the size of the search space in [19]. Unfortunately, when the overtaking is performed in the holding areas, the physical locations of the aircraft constrain the schedules which can be produced and change the value of partial solutions, thus they have to be a part of the state considered in the dynamic programming solution.

Craig et al considered a model of a simplified holding area structure in [10] and applied a dynamic programming solution. However, the real holding area structures are much more complex than the example used, making the state space prohibitively large for solving an online problem in this way. We presented an initial model for this problem, including the holding area structures, in [6], and have since improved it to include CTOT constraints; this is the model described later in this paper. Further details about the improved system can be found in [5].

Bianco et al considered the departure problem as a special case of the cumulative asymmetric travelling salesman problem with release dates in [9]. The equivalence breaks down for Heathrow, however, as the separation rules do not obey the triangle inequality, as explained in section 3.3.

There are similarities to the arrivals problem described and studied by Abela et al in [1], Beasley et al in [7] and [8] and Ernst et al in [11], since both have sequence-dependent separations. However, the differences in the detail of the problems, particularly with the physical holding area constraints for departures and the potential gain in landing later and gaining a better fuel economy for arrivals, mean that solutions to one problem are far from being immediately applicable to the other.

Similarities can be seen between the receding horizon approach to scheduling landings in [13] and the online simulation that we use in this paper, the main difference being that our horizon is defined implicitly by the times aircraft left the stands. A similar approach to handling a dynamic problem, which was applied to the vehicle routing problem, can also be seen in [12].

3 Simulation model

Our simulation aims to reproduce a realistic model of the departure system at Heathrow from the point of view of the runway controller. Real, historic data is used for the simulations, including details of aircraft, when they actually left the stands, arrived at the holding area and took off. The simulation includes different levels of detail as aircraft move through the departure system from the stands, along the taxiways through the holding areas to the runways.

Our solution system is intended to form the basis for a decision support system for the runway controller, so we assume that anything outside of the control of the runway controller is outside of the control of the automated support system and will occur in the same way as it did historically. In particular, we assume that, until aircraft reach the holding area, they perform

in exactly the same way as they did in the historic datasets. Once within the holding area the simulation has to determine and enact the effects of decisions made by the runway controller.

The controller has responsibility for all aircraft at the holding area so the simulation has to know the identity and position of all of these aircraft. The taxiways form the inputs for the holding area so information about all aircraft on the taxiways is also maintained. The separation rules mean that aircraft that have already taken off can influence how early another aircraft can take off, so they have to remain in the system until this is no longer the case.

Our simulation model has a number of parts, or sub-models. We will now discuss the different parts in more detail and then in section 5 we will give more detail of how the simulation uses the model to test the solution system.

3.1 Abstract taxiway model

The abstract taxiway model is the simplest part of the simulation, as it models the ground movement of the aircraft around the taxiways. This can remain at a high level of abstraction as the runway controller has no control over it. The taxiways form the inputs to the holding area and the taxiway model must provide details of when aircraft will arrive at the holding area and at which holding area entrance they will arrive. The arrival order can, of course, be determined from the arrival times.

As we use real historic data, the actual holding area arrival times are available to the model and are used. Until the time that the aircraft actually reached the holding area, there would have been a degree of uncertainty in the arrival time at the holding area due to the impossibility of accurately predicting taxi speeds for individual aircraft and of delays for contention between aircraft on the taxiways. The level of this uncertainty would decrease as the aircraft approached the holding area. The taxiway model can simulate this uncertainty by adding appropriate prediction errors to the timings in the data, but we do not do this in the tests performed for this paper as we wish to assess the effects of the constraints and do not want these effects to be obfuscated by any introduced prediction errors.

The taxiway model is responsible for determining the holding area entrance at which aircraft arrive. This will, in practice, be a decision made by the ground movement controller and will depend upon taxiway contention and holding area status. The simplest case assumes that all aircraft enter the holding area by the nearest entrance to their stand, taking the simplest route through the taxiways and hence causing the least congestion. Previous experiments have shown that the overtaking in the holding area can be simplified by adding more intelligence to the entrance allocation method, making improved take-off schedules possible. In this paper, we assume that aircraft are sent to the nearest entrance, even though this may mean that the results we achieve will be slightly pessimistic when compared to having a real ground movement controller making the decisions.

3.2 Holding area model

Once aircraft are in the holding area, they are under the control of the runway controller and could be influenced by decisions suggested by a decision support system. At this point, a less abstract model is required to properly model the real life restrictions. The holding area model consists of a directed graph representation of each holding area structure, where the nodes represent the valid positions at which aircraft could be located and the arcs represent the transitions that aircraft can make between nodes. An example holding area network can be seen in figure 1. The NATS Airport Analysis and Playback Tool (APT) can be used to play back the real movement of aircraft around Heathrow, showing where aircraft actually were at any instant in time. Figure 2 presents an extract of a screen shot of APT, showing the layout of the same holding area as the graph in figure 1, including the real instantaneous positions of aircraft in the holding area. Different runways have different holding area structures and so inflict different constraints upon the overtaking which may be achieved.
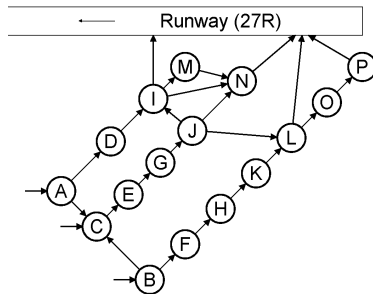

Fig. 1. Directed graph representation of the 27R Holding Area

Fig. 2. Picture of aircraft in the 27R Holding Area

The physical constraints upon the reordering of aircraft can easily be checked by initially positioning each aircraft at the node it is currently at, or will enter next if it is between nodes, and verifying that all aircraft can exit the graph in the desired take-off order via the runway.

We consider the taxiing aircraft in the predicted take-off schedule, so we also need to consider the feasibility of reordering these. To check this we build a queue for each holding area entrance then add aircraft to the queue for the entrance at which they are predicted to arrive, in the predicted arrival order.

Whenever the entrance node is empty, the next aircraft from the queue is placed at that entrance and allowed to move through the graph in the same way as any other aircraft. Again, the aircraft must be able to enter the runway in the allocated take-off order to declare the reordering as feasible.

3.3 Take-off time prediction

A formal model is used to predict take-off times for all aircraft, based upon the assumption that aircraft take off as early as they can. This assumption is only invalid at very quiet times, when delays to take-off will neither affect throughput nor delay other aircraft, but at those times decision support is not necessary anyway. Unlike the analogous situation in arrival scheduling, it is never of benefit to an aircraft to take off later than the earliest take-off time.

Given any two aircraft, i followed by j, let V_{ij} be the minimum required wake vortex separation and R_{ij} be the minimum required SID separation. The wake vortex and SID separations are mandated by the airport regulations and can be easily calculated via a lookup based on weight class or SID route and speed group, respectively. Note that $V_{ij} = V_{ji}$ only if the weight classes of i and j are the same and R_{ij} is not always the same as R_{ji}, even if the speed group is the same for both aircraft and rarely holds if not. Furthermore, the triangle inequality does not hold, so $R_{ij} + R_{jk} \geq R_{ik}$ may not be true for all aircraft which take off in the order i, j, k.

Throughout this model we define times as the number of seconds from the start time of the dataset under consideration. Let c_j be the position of aircraft, j, in the take-off order, so that $c_j = 1$ for the first aircraft to take off. Let d_i be the real or predicted take-off time for aircraft i. The earliest take-off time, e_j, for aircraft j that meets the separation constraints can then be calculated from equation 1.

$$e_j = \begin{cases} 0 & \text{if } c_j = 1 \\ \max_{i \in \{1,\dots,n\} | c_i < c_j} (d_i + \max(V_{ij}, R_{ij})) & \text{if } c_j \geq 2 \end{cases} \tag{1}$$

$$d_j = \max(e_j, h_j + T(t_j), b_j) \tag{2}$$

If aircraft j has a CTOT, then let b_j denote the start time of the slot, before which the aircraft should not take off, otherwise let $b_j = 0$. Let h_j be the time aircraft j arrived at the holding area and t_j be the traversal path it will follow through the holding area. Let $T(t_j)$ be a function to determine a traversal time for an aircraft following path t_j. Assuming aircraft j takes off as early as possible, its take-off time can be predicted using equation 2.

The estimated traversal time function $T(t_j)$ depends upon the route taken through the holding area, which is determined as a part of the feasibility check, so it is important to perform the feasibility check prior to predicting take-off times.

4 Solution method

Given the model of the problem specified above, our solution method has to find a take-off schedule that meets all of the constraints, has a low delay for the aircraft and does not require unnecessary or excessive workload from either the pilots or controller.

In this section we first specify many of our objectives mathematically in the form of an objective function to minimise, then we explain how we solve the model to find a good take-off order. Finally, we explain how the remaining objectives are covered.

4.1 Objective function

Let h_j be the arrival time of aircraft j at the holding area, let c_j be the position of the aircraft in the take-off schedule and a_j be the position in the arrival order at the holding area. Let v_j be the weight class of aircraft j and s_j be the speed group for aircraft j. If aircraft j has a CTOT then let b_j and l_j denote the start and end times, respectively, of the fifteen minute slot, otherwise let b_j be 0 and let l_j be a large number, larger than the total number of seconds the dataset covers.

Expression 3 specifies our objectives as a formula to be minimised where n is the number of aircraft in the schedule and $C(d_i, b_i, l_i, h_i)$ and $P(c_i, a_i, v_i, s_i)$ are functions which will be explained later. α, β, γ and δ are constants to ensure the correct prioritisation of the objectives.

Minimise:

$$\sum_{i=1}^{n}(\alpha C(d_i, b_i, l_i, h_i) + \beta(d_i - h_i) + \gamma(\max(0, c_i - a_i))^2 + \delta P(c_i, a_i, v_i, s_i)) \quad (3)$$

Where aircraft have a CTOT, this could be thought of as a hard constraint, specifying a 15 minute window within which the aircraft must take off. However, in practice the workload of the departure system at Heathrow can prohibit some aircraft from being able to achieve their CTOT time and the controllers are actually permitted to miss a small number, gaining five minute extensions for these aircraft. It is very important though to minimise the number missed, so we make this our primary objective (the first term in formula 3).

Our secondary objective is to minimise the total delay for aircraft as this contributes directly to both pollution and dissatisfaction. This is the second term in formula 3. Minimising delay will also maximise runway throughput whenever there are aircraft waiting, since wasted runway throughput would inflict the wasted time as additional delay upon later aircraft. Delay-based measures have an additional benefit over throughput based measures as they shift bad separations later in the schedule, where they delay less aircraft and are, importantly, more likely to be usable by aircraft that will soon leave their stands.

Our tertiary objective is to reduce the 'unfairness' in the schedule. The third and fourth terms in our objective function refer to this. The third term is the simplest and aims to reduce the sum of the squares of the positional delay of the aircraft. The fourth term is a function which applies penalties to schedules which delay certain types of aircraft and aims to avoid penalising these aircraft. For example, this function penalises schedules which put slow or light aircraft at the end of the schedule. When one of these aircraft takes off, a larger separation is required, either before or after it, so a delay based measure will tend to push the aircraft later on into the schedule. Often, there is a benefit in pushing large separations to positions which are later in the schedule but, in this case there are so few of these types of aircraft that there is unlikely to be a benefit in delaying them. The penalty from $P(c_i, a_i, v_i, s_i)$ avoids the unnecessary penalising of this type of aircraft.

$$
C(d_i, b_i, l_i, h_i) =
\begin{cases}
W_1((d_i - l_i)^{1.1}) + W_2 & \text{if } d_i \geq (l_i + 300) & (i) \\
W_3((d_i - l_i)^{1.1}) + W_4 & \text{if } (l_i + 300) > d_i > max((h_i + F_H), l_i) & (ii) \\
W_5(d_i - l_i) + W_4 & \text{if } (h_i + F_H) \geq d_i > l_i & (iii) \\
W_6(F_L + d_i - l_i) & \text{if } l_i \geq d_i > (l_i - F_L) & (iv) \\
0 & \text{if } (l_i - F_L) \geq d_i > (b_i + F_B) & (v) \\
W_7(b_i + F_B - d_i) & \text{if } (b_i + F_B) \geq d_i & (vi)
\end{cases}
\tag{4}
$$

The function $C(d_i, b_i, l_i, h_i)$ is designed to promote CTOT compliance and has a number of cases depending upon when an aircraft is scheduled to take off in comparison with its CTOT. It is fully defined in [5] and is summarised here. W_1 to W_7 are constant weights to ensure a correct prioritisation of objectives.

Case (v) applies to aircraft without a CTOT, or scheduled to take-off well within their CTOT. Cases (iv) and (vi) apply to aircraft that are scheduled within their CTOT, but close to one end of it. It is useful to have some flexibility in case the schedule slips forwards or backwards so these cases apply a small penalty to this sort of less flexible schedule. W_6 and W_7 have low values and F_B and F_L are set to 60 and 120 seconds respectively.

If an aircraft is more than five minutes late for its CTOT slot then even an extension will not allow the aircraft to take off within it. In this case, an aircraft would have to wait for a new CTOT to be allocated; a situation to be avoided and covered by case (i). W_1 and W_2 are the largest weights in the model and ensure that any schedule which has aircraft taking off too late even for an extension is extremely heavily penalised.

Case (ii) applies when an aircraft is scheduled to take off too late for its CTOT but within an allowable extension. W_3 and W_4 are large to ensure that these schedules will be worse than schedules where all aircraft are within CTOT but are much smaller than W_1 and W_2 so that multiple extensions are preferred to a single aircraft being too late for an extension. The large size of W_3 ensures that schedules where the delay is not as pronounced will be

preferred and the non-linear power of the delay ensures that multiple short extensions are preferred to fewer, longer extensions.

If an aircraft arrives at the holding area too late to be able to take off within its CTOT then it is not always wise to schedule it to take off immediately if that would cause problems for the other aircraft. Case (iii) covers this circumstance, where F_H is a constant denoting the minimum number of seconds it is reasonable to expect between the holding area arrival time of an aircraft and the take-off time of the aircraft. W_5 is relatively small, to allow more flexibility to schedule the aircraft within this flexible time period while still slightly favouring schedules where it takes off as early as it can. Without case (iii), aircraft in this situation would be scheduled as early as they could be (due to the large size of W_3 and W_4), even if doing so meant a loss of throughput on the runway.

4.2 Tabu search

To be of use to an online decision support system in a real world environment, the search must execute extremely quickly. We work on a one second limit to the search time, which reflects a one second response time to situational changes in a real departure system. The search we use must, therefore, be very quick. A tabu search algorithm is used to find good take-off orders and another heuristic method is used to validate that each is achievable within the current holding area. The formal model is then used to predict take-off times and determine a cost for the schedule. Given the take-off times, the cost of the solution can be determined from the objective function and the tabu search can then use it to determine how to move through the search space.

The tabu search is started from an initial feasible solution. We use the first-come-first-served order as an initial solution at the start, then seed future searches with the best schedule found by the previous search, modified accordingly for new aircraft which entered the simulation, old aircraft which left the simulation and changes in predicted arrival order at the holding area. The initial solution is then evaluated to obtain an initial cost.

The search steps below are then iterated 200 times and at the end the best solution found during the search is returned as the solution to suggest to the controller.

1. Fifty neighbouring solutions are generated by applying random moves to the current solution, cs. Let NS denote the set of neighbouring solutions which are generated from cs, so $|NS| = 50$.
2. Each solution is evaluated to verify feasibility, take-off times are predicted and the cost of each is determined.
3. The non-tabu feasible solution (if there is one), which has the lowest cost is then accepted as the new current solution and the algorithm returns to step 1. More formally, let $FS \subseteq NS$ denote the set of feasible solutions, and $COST_s$ denote the cost of a solution $s \in FS$. If $FS \neq \emptyset$, then let

$cs = a$ for some solution a such that $COST_a = \min_{f \in FS} COST_f$. If multiple solutions, a, meet the minimum value criterion then one of the minimum value solutions is selected at random to be used as the new current solution.

A tabu solution is here defined to be one which reverses the effects of a previous move. Whenever a move is adopted, the previous positions of all aircraft which were moved are recorded on a tabu list. This tabu list is checked for each solution that may be adopted, and if all of these aircraft are found to be back in the positions they were moved from the solution is declared to be tabu and will not be adopted. This list records these aircraft for the last ten moves, preventing the reversal of any move made within ten moves of making it.

The neighbourhood is designed to provide a wide variety of moves based around the characteristics of a good schedule. Three types of moves are used, the actual move being chosen at random each time. There is a 50% chance for any move that a shift move will be selected. This will move a randomly selected group (of from one to five aircraft) to a new position in the take-off schedule. There is a 30% chance of using the swap move, choosing two aircraft at random and swapping their positions in the take-off schedule. Otherwise, a randomisation move will be used, selecting a random group of three to five aircraft from sequential positions in the take-off sequence and randomly reordering them.

4.3 Holding area constraints

When considering the number of options for each aircraft to move through the holding area, the search space is prohibitively large to solve in a reasonable time. Moreover, there will be many functionally equivalent ways to reorder aircraft and many ways that have the same results but involve more or less work or movement from the aircraft or controller. However, an important objective for our solution system is to control the workload of controllers and pilots.

We allocate paths through the holding area to aircraft prior to testing whether the re-ordering is possible. This has the dual benefit of vastly reducing the search space while also ensuring that the more sensible paths are always used, reducing the workload for pilots. Additionally, the paths assigned to aircraft in the holding area are remembered by the simulation so that they are not changed, reflecting the fact that, even when it is actually possible, changing instructions to pilots is costly from a workload point of view.

Paths are assigned to aircraft according to whether they are overtaken by other aircraft from the same holding area entrance. It is easy to overtake aircraft at other entrances when interleaving the arrival queues from each entrance. However, overtaking aircraft from the same entrance requires the aircraft to follow different paths. For each entrance, there is at least one path

which has somewhere for an aircraft to wait to be overtaken. The path allocation heuristic assigns these paths to aircraft that are overtaken as they are slightly longer. It allocates the direct path through the holding area to other aircraft. In this way we can be sure that all aircraft follow routes that are as good as they can be, and so control the load.

Pre-processing is performed upon the paths through the holding area prior to any search being performed and the points of convergence and number of nodes beyond points of divergence are recorded. When a take-off order is to be checked, this pre-processed information is used to form partial orders for aircraft at each node which, together with counts of aircraft that have passed each node, can be used to determine whether moving any aircraft will block aircraft on another path. This facilitates a very fast feasibility check, avoiding the need for backtracking. More detail of the fast feasibility check and the path allocation algorithm used can be found in [5].

5 Simulation tests

In this section we discuss the simulation that is executed. In section 6, we consider how it is varied in order to evaluate the effects of the different constraints. Each simulation we perform uses one of the real datasets provided by National Air Traffic Services. These datasets cover around half a day each and include all of the information that the simulation requires.

The simulation is responsible for building the static problems for the solution system to solve, taking the resulting solution and enacting it, then re-presenting the new, consequent, situation to the solution system to be re-solved.

We assume for these tests that the runway controller always does what the solution system suggests so that the simulation does not require the presence of a real runway controller. This is possible for the simulation only because there is no communication needed with pilots and no unexpected input from outside of the simulation. This does mean, however, that many of the benefits that controllers can gain, for example from discretionary changes to separation rules or adoption of more work intensive manoeuvres, will not be gained. We therefore expect the results of using the solution system like this to be pessimistic compared to performance in the presence of a real controller.

The simulation process has the following steps:

1. Create initial problem, for the start of the dataset. The simulator adds aircraft to the problem as soon as they leave their stand. Until this time the level of uncertainty in ready time makes it impractical to include the aircraft.
2. Pass the problem to the solution system to determine what to do. The problem passed consists of:
    - The characteristics of all aircraft in the problem, including weight class, departure route and speed group.

- The predicted arrival times at the holding area for the aircraft on the taxiways.
- The current position (node), currently allocated traversal path and real holding area arrival time for all aircraft at the holding area.
- The take off time for all aircraft in the problem that have taken off.

3. Retrieve the proposed solution, comprising a suggested take-off order, the holding area paths assigned and the predicted take-off times.
4. Advance time by an increment of one minute and update data accordingly.
   - Aircraft in the holding areas move along the paths that have been allocated to them.
   - Aircraft on the taxiways move closer to the holding area and any prediction errors are decreased accordingly.
   - An aircraft may reach its predicted take-off time, in which case it is assumed to take off. (The minimum separation time between aircraft is one minute so it is not possible for more than one aircraft to take off in this time.)
   - Any aircraft that left their stand in the time increment are added to the problem.
   - Once sufficient time has passed after the take-off of an aircraft so that it can no longer have any effect upon the other aircraft in the problem it can be removed from the problem.
5. Return to step 2 if there are still aircraft in the dataset that have not taken off, otherwise end the simulation and evaluate the final take-off schedule.

Some information in the problem passed to the solution system is sometimes considered to be static.

- The take-off time and position for an aircraft is assumed to be static for at least two minutes before take-off. This ensures that pilots are given sufficient warning time to prepare and line-up.
- The holding area traversal path is assumed to be fixed for all aircraft in the holding area. This is over-restrictive but aims to avoid excessive rescheduling which would involve giving pilots new taxi instructions. In fact, the solution system could easily allow the path allocation heuristic to reallocate paths until the point where the aircraft has passed the node where the paths diverge, but this may involve more work for controller and pilot so is not assumed to be possible.

At the end of the simulation, the final take-off order is examined, considering the take-off times of all aircraft. The total holding area delay and number of CTOTs that were missed is calculated for the schedule.

6 Comparing the effects of constraints

The aim of this paper is to evaluate the effects of different constraints upon the quality of the take-off schedules that can be produced. It is vital to understand

these in order to better predict the effects of possible improvements to the departure system at Heathrow. We discuss possibilities for such manipulation in section 8 and focus upon a single possibility where decision support could help in section 9.

To do this we used a previously developed simulation of the departure system. A solution system designed as a decision support system for the runway controller took the place of the runway controller in the tests. We have explained the model, solution approach and simulation in some detail above.

We knew from previous results that the solution system performed well with all constraints in place, but we wished to examine how much better it could perform in the absence of each constraint or the absence of multiple constraints. This allowed us to better understand the effects of the constraints and how manipulation of them could be used to further decrease the delay at Heathrow.

To test the effects of the constraints we performed a simulation with and without each of the constraints and compared the resulting schedules. We used four different datasets to perform these tests. Each contained historic data provided by National Air Traffic Services. As the tabu search has a stochastic element in the neighbour selection, each simulation was performed ten times with each combination of constraints for each dataset, in order to minimise the effect of this. The mean results are presented. The ten executions used the same ten different random seeds for each constraint combination and dataset.

7 Results

Table 1 gives details about the test datasets and shows the results obtained by the real controllers. The 'number of aircraft' column identifies how many aircraft were in the dataset. The next two columns show the number of CTOTs the real controllers missed and the total number of seconds of holding area delay the real schedules had. These are calculated using the real holding area arrival times, take-off times and CTOTs.

Table 1. Dataset size and manual, real world results

Dataset	Number of aircraft	Real CTOT misses	Real delay
1	341	6	117894
2	330	5	120893
3	345	5	107786
4	259	4	96235

In table 2, the '#C.' columns specify the number of CTOT misses and the 'delay' columns give the total delay (in seconds) that the simulation predicts

for the schedules it produces. Each row shows the effects of removing different combinations of constraints. Any row labelled 'HA' has had the physical holding area constraints removed, so any sequence is achievable. 'WV' refers to removing the weight class based wake vortex constraints, so one minute separations are possible regardless of the weight classes of the aircraft. 'SIDS' refers to removing the departure route based constraints, so one minute separations are possible regardless of departure route and speed. Finally, rows labelled 'CTOT' have had the CTOT constraints removed. Obviously, in these cases the number of CTOTs missed is 0. The first row shows the results for normal operation, with all constraints present, and the final row shows the results with all constraints removed. The mean values for CTOT compliance and total delay obtained from ten runs are shown in the table. In each experiment, the CTOT compliance was identical across the different runs, since a large cost is associated with CTOT misses. The delay in the generated schedules often (but not always) varied slightly between runs, with up to four different sequences being adopted across the ten runs for some of the experiments. The total delay across the alternative sequences almost always differed by less than five minutes (300 seconds), usually by less than two minutes and was sometimes identical. The mean delay was, thus, representative of the effects of the constraints upon the results which could be obtained by the system.

A number of observations are possible from these results. The constraints imposed by the holding area structure are the least restrictive of all of the constraints, but usually have some effect on restricting good schedules. These tests were conducted with the 27R holding area structure which is much more flexible than the 09R holding area structure so we would expect more effect from using that structure instead.

Wake vortex constraints make a significant difference but far less than either CTOT or SID constraints. It is surprising that the CTOT constraints have such a great effect upon the delay. Examination of what is happening reveals that our objective function ensures that aircraft which arrive late for their CTOT overtake aircraft that are already in the holding area. From the point of view of CTOT compliance this is a good thing, however, not only is it extremely unfair but it also forces the rest of the aircraft to move into positions where they can be overtaken and occasionally forces them to take off in a less efficient order as the overtaking needed to obtain the more efficient order is no longer possible. Removing this constraint ensures that the better take-off orders can be kept and so reduces the delay for aircraft.

The SID constraints are the most restrictive in all cases. However, in all cases a considerable further reduction in delay can be obtained by also removing the wake vortex constraints. We see, therefore, that both wake vortex separations and SID separations matter and appear to be fairly independent.

Each dataset tends to have a surplus of either northbounds or southbounds, depending upon the current flight paths over the atlantic, but even in datasets with a surplus of one direction there are still times when there is a deficit at the holding area of aircraft going that direction and consequent wasted through-

Table 2. Effects of removing constraints

Removed constraints	Dataset 1		Dataset 2		Dataset 3		Dataset 4	
	#C.	Delay	#C.	Delay	#C.	Delay	#C.	Delay
None	3	83581	3	79857	1	73057	4	63153
HA	3	81998	3	79880	1	73057	4	61998
WV	3	73533	3	75854	1	65031	4	53703
SIDS	3	64742	3	63251	1	58294	4	43639
CTOT	0	68925	0	65469	0	60732	0	55835
HA, WV	3	72989	3	75562	1	65043	4	53691
HA, SIDS	3	64595	3	62126	1	58294	4	43639
WV, SIDS	3	55962	3	54650	1	51181	4	38635
HA, CTOT	0	68817	0	65421	0	60748	0	54301
WV, CTOT	0	61248	0	59658	0	54314	0	47697
SIDS, CTOT	0	54625	0	47692	0	46104	0	39293
HA, WV, SIDS	3	55331	3	54650	1	51181	4	38635
HA, WV, CTOT	0	61212	0	59658	0	54314	0	47666
HA, SIDS, CTOT	0	54565	0	47640	0	46117	0	39393
WV, SIDS, CTOT	0	45225	0	39932	0	40157	0	33370
All	0	45225	0	39932	0	40157	0	33370

put. The generated schedules show that there is significant spare capacity for aircraft to take-off from Heathrow, provided that they take off at the right times.

Finally, it should be noted that, with the future introduction of larger aircraft, for example the airbus A380, and the consequent expected increase in wake vortex separations, the wake vortex separations are likely to start to have more effect than they currently do. However, we are unable to evaluate the full effects of these until the separation rules have been finalised.

8 Conclusions

In this section, we draw conclusions about what these results mean to a variety of potential modifications that could be made to the departure system. We begin by considering the effects of physical changes that could be made and then we consider the effects of changes to the way the departure system currently operates, including the usage of the runways and the point at which the take-off scheduling takes place. Some changes would also have the effect of increasing the capacity of the departure system, however, these effects cannot be concluded from the results of the tests we documented in this paper and are, in some ways, more obvious.

8.1 Physical changes

By laying down more concrete, the holding area structure could be modified. This would make more overtaking possible and schedules that are not currently achievable may become so. Similarly, work on the taxiways could increase the possibilities for reordering the aircraft before they reach the holding area. Both of these approaches are expensive, however, in terms of both cost and disruption while the work is being performed and are prohibited at present by the lack of physical space within which to expand. Our tests show that the physical constraints imposed by the 27R holding area structure are not actually as restrictive as many of the other constraints and suggests that changes elsewhere may have more effect. The 09R holding area, which is considerably more restrictive, is already being modified as a consequence of the work on the new terminal 5.

Another option is to build another runway, which has obvious potential throughput benefits. If the number of aircraft remained the same, the effect on delay would be similar to that of implementing mixed mode, explained below, since more holding areas would be available for reordering and there could, conceivably, be a runway available to use for departures to utilise the wasted time due to larger wake vortex separations on the other runway.

8.2 Procedural changes

Adopting mixed mode for the runways has been discussed in the 2003 White Paper on aviation [20], so may become possible at some times of the day. Using the runways in mixed mode changes the problem considerably as the aim is then to interleave arrivals and departures in such a way that the total throughput of the airport is increased, decreasing the delay for both arrivals and departures. This brings in coordination issues not only between the arrivals and departures on each runway, but also between the departures on the two runways, due to the downstream constraints. In terms of delay for departing aircraft, we could expect the holding area structure to be much less constraining since there would be two holding area structures available, thus much less overtaking should be needed in each. We could also expect the wake vortex separations to be less constraining, with many being eliminated, since an increased wake vortex separation on one runway could be utilised by an arrival landing on that runway while a departure took place from the other runway. As long as the coordination issues could be adequately addressed, we could expect a decrease in the delay of up to the amount predicted by the results for eliminating both the holding area constraints and the wake vortex separations (the downstream constraints of SID separations and CTOTs would still apply). This would be a significant gain from the point of view of passengers, airlines and pollution. There is currently no way to deal with wake vortices other than to avoid them (for instance by using the other runway) or to wait for them to dissipate. The effects of wake vortex separations cannot,

therefore, be eliminated while the operating procedures require the use of a single runway.

The delay decreases for mixed mode are possible even without changing the CTOT or departure route constraints. Furthermore, the delay for arrivals should also be decreased by the implementation of either mixed mode or a third runway, facilitating further pollution and cost benefits, as arrivals are also affected by wake vortex separations.

Another option is to perform the overtaking somewhere other than in the holding area. The possibilities for overtaking on the taxiways are very limited due to issues of controller workload and taxiway congestion. The ground movement at Heathrow is already a complicated task and the taxiways are sufficiently busy to ensure that it is far more practical to have aircraft overtaking in the holding areas which were built for that purpose. The exception to this, which was mentioned earlier, is that it is often possible for the ground movement controller to pick a holding area entrance to deliver an aircraft to, which can have benefits in reducing the overtaking necessary in the holding area. It is possible that a decision support system may be of use to the ground movement controller in performing some of the overtaking on the taxiways but it is unlikely to gain great benefits due to the limited structure of the taxiways.

Conversely, stand holding has great potential for benefits in the departure system. This involves aircraft being held at the stands and released at the correct time for take-off, moving some of the delay the aircraft suffer from waiting at the holding area with engines running to waiting at the stand without the engines running. There should be consequent air and noise pollution benefits as well as cost savings for the airlines from doing this. In theory, releasing aircraft from their stands in time to reach the runway in the correct order would, ideally, ensure that no overtaking was required at the holding areas, so the physical constraints imposed by the holding area structure would vanish, with additional delay and CTOT compliance benefits as we saw earlier.

The downstream SID separation constraints have the most effect upon the delay in the take-off schedule. None of the changes discussed above would affect these, only a change to the structure of the downstream airspace, or limitations to the other airports feeding into this airspace, would do so. A stand holding system would, however, mean that the delayed aircraft were burning less fuel and so would have the pollution and fuel cost benefits. We discuss the issues of stand holding, how it changes the take-off problem and whether the holding area is still relevant in the next section.

9 Future work : the stand holding problem

The stand holding problem has many similarities to the problem of scheduling take-offs at the holding areas. The main differences are that the scheduling decisions have to be made much earlier, there must be a far greater emphasis on the taxi times and the holding area structure should make less difference

to the produced schedules. A stand holding system could work essentially the same as the solution system we defined earlier, with changes in these three areas.

It is important to know the earliest time at which each aircraft could push back from its stand. Investigation is already taking place into Collaborative Decision Making [16], which should make the earliest push-back times available sooner. An expected taxi time to the runway can be added to the earliest push-back time to obtain a prediction for the earliest take-off time. These times could then be used to determine a take-off order and take-off times, exactly as we described above. In an ideal situation, the estimated taxi times can then be deducted from the predicted take-off times to determine the times at which aircraft should leave the stands.

In practice, issues of stand contention and taxi time variability mean that it is not so easy. Stand contention occurs when an arriving aircraft requires the stand to unload passengers, or when two aircraft wish to leave nearby stands. To avoid delay to the arriving aircraft it may be necessary for an aircraft to push-back earlier than desired. Overtaking at the holding area would then be required to recover the desired take-off order.

Calculation of the stand release time requires accurate taxi time estimates. Taxi time will vary depending upon taxiway contention so is influenced by the leave stand time. Taxi time also includes intrinsic variability in the processes of pushing back from the stand, starting engines, and the pilot's pre-flight checks. The simplest method to ensure aircraft reach the runway on time is for them to push-back early, allowing for longer taxi times than expected. This moves some of the delay back to the holding area and so reduces some of the potential benefit of stand holding in terms of fuel burn and constrains possible sequences. Since the ground movement controller would know the desired take-off order, appropriate prioritisation of aircraft on the taxiways could alleviate some of the problems with taxi time inaccuracy.

The holding area structure matters whenever there is overtaking required in the holding area. Both early push-back due to stand contention and variable taxi times can cause the aircraft to arrive at the holding area in a different order to that in which they should take off. However, the overtaking problem should be much simpler than it is at the moment. A stand holding system would be expected to build robust schedules, allowing for stand contention and the effects of variable taxi times and ensuring that good schedules can be recovered by the runway controller at the holding area.

We intend to use our solution system to further investigate the stand holding problem, the effects of the various elements and the robustness issues. To do this we will modify our solution system to work as a decision support system for stand holding and use the current solution system to perform the schedule recovery at the holding area.

10 Acknowledgements

References

1. Abela J, Abramson D, Krishnamoothy M, de Silva A, Mills G (1993) Computing optimal schedules for landing aircraft. The 12th National Conference of the Australian Society for Operations Research, Adelaide, July 7-9, 1993. Available at: `http://www.csse.monash.edu.au/~davida/papers/asorpaper.pdf` [14 November 2006]
2. Anagnostakis I, Clarke J-P, Böhme D, Völckers U (2001) Runway operations planning and control, sequencing and scheduling. The 34th Hawaii International Conference on System Sciences (HICSS-34), Hawaii, January 3-6, 2001.
3. Anagnostakis I, Clarke J-P (2003) Runway operations planning, a two-stage methodology. The 36th Hawaii International Conference on System Sciences (HICSS-36), Hawaii, January 6-9, 2003.
4. Anagnostakis I, Clarke, J-P (2002) Runway operations planning, a two-stage heuristic algorithm. AIAA Aircraft, Technology, Integration and Operations Forum, Los Angeles, CA, October 1st-3rd, 2002.
5. Atkin JAD, Burke EK, Greenwood JS and Reeson D (2007) Hybrid meta-heuristics to aid runway scheduling at London Heathrow airport. Transportation Science 41:90–106
6. Atkin JAD, Burke EK, Greenwood JS and Reeson D (2008) A meta-heuristic approach to departure scheduling at London Heathrow airport. Computer Aided Systems of Public Transport, M. Hickman, P. Mirchandani, and S. Voss, eds., Lecture Notes in Economics and Mathematical Systems, Springer, Berlin.
7. Beasley JE, Krishnamoorthy M, Sharaiha YM, Abramson D (2000) Scheduling aircraft landings - the static case. Transportation Science 34:180–197
8. Beasley JE, Sonander J, Havelock P (2001) Scheduling aircraft landings at London Heathrow using a population heuristic. Journal of the Operational Research Society 52:483-493
9. Bianco L, Dell'Olma P, Giordani S (1999) Minimizing total completion time subject to release dates and sequence-dependent processing times. Annals of Operations Research 86:393-416
10. Craig A, Ketzscer R, Leese RA, Noble SD, Parrott K, Preater J, Wilson RE, Wood DA (2001) The sequencing of aircraft departures. 40th European Study Group with Industry, Keele 2001.
11. Ernst AT, Krishnamoorthy M, Storer RH (1999) Heuristic and Exact Algorithms for Scheduling Aircraft Landings. Networks 34:229–241
12. Gendreau M, Guertin F, Potvin J-Y, Taillard E (1999) Tabu Search for Real-Time Vehicle Routing and Dispatching. Transportation Science 33:381–390.
13. Hu, X-B, Chen, W-H (2005) Genetic algorithm based on receding horizon control for arrival sequencing and scheduling. Engineering Applications of Artifical Intelligence 18:633–642

14. Idris HR, Anagnostakis I, Delcaire B, Hansman RJ, Clarke J-P, Feron E, Odoni AR (1999) Observations of Departure Processes at Logan Airport to Support the Development of Departure Planning Tools. Air Traffic Control Quarterly 7:229-257
15. Idris HR, Delcaire B, Anagnostakis I, Hall WD, Pujet N, Feron E, Hansman RJ, Clarke J-P, Odoni A (1998) Identification of Flow Constraint and Control Points in Departure Operations at Airport Systems. AIAA Guidance, Navigation and Control conference, Boston, MA, August 1998.
16. Lane R, Hogg D (2005) London Heathrow CDM WP1. Eurocontrol Experimental Centre Notes. Available at: `http://www.eurocontrol.int/eec/public/standard_page/2005_note_03.html` [14 November 2006]
17. van Leeuwen P, Hesselink H, Rohling J (2002) Scheduling Aircraft Using Constraint Satisfaction. Electronic Notes in Theoretical Computer Science 76.
18. Newell GF (1979) Airport Capacity and Delays. Transportation Science 13:201-241
19. Trivizas DA (1998) Optimal Scheduling with Maximum Position Shift (MPS) Constraints: A Runway Scheduling Application. Journal of Navigation 51:250-266
20. Department for Transport (2003) White Paper. The Future of Air Transport. Appraisal framework for airports in South East and East of England. Available at: `http://www.dft.gov.uk/stellent/groups/dft_aviation/documents/pdf/dft_aviation_pdf_031504.pdf` [14 November 2006]