Departure Runway Scheduling at London Heathrow Airport Extended Abstract

Jason A. D. Atkin[1] Edmund K. Burke[1]
John S. Greenwood[2] Dale Reeson[3]

[1] {jaa,ekb}@cs.nott.ac.uk, School of Computer Science and Information Technology, University Of Nottingham, Jubilee Campus, Wollaton Road, Nottingham, NG8 2BB
[2] Analysis & Research, National Air Traffic Services Ltd, Spectrum House, Gatwick, West Sussex, RH6 OLG
[3] National Air Traffic Services, Heathrow Airport, Hounslow, Middlesex, TW6 1JJ

October, 2004

1 Introduction

London Heathrow airport is one of the busiest airports in the world as it is very popular with both airlines and passengers. The capacity of the departure system at Heathrow is limited by the capacity of the runway. It is therefore imperative to maximise the throughput of the single available departure runway. Aircraft taxi from the stands to the ends of the runway where they queue in holding points, awaiting instructions to take off. We will present a model for scheduling aircraft awaiting take-off at the runways at London Heathrow and evaluate the effectiveness of our approach to solving it.

In [11] and [12] Idris et al. analysed the departure flow at Logan airport and identified the runway as the key constraint. In [14], Newell provided a model for estimating the capacity of an airport in terms of number of arrivals and departures from the runways. The capacity of an airport was shown to be increased when all of the runways are used for alternating arrivals and departures.

In both the departure and arrival systems, separations need to be enforced between aircraft for reasons of safety. The departure problem involves finding a take-off sequence for which either the throughput of the runway is maximised or the delay upon the aircraft awaiting take-off is minimised. The arrivals problem involves finding a landing order for arriving aircraft so that either the total deviation from ideal landing times for aircraft or the landing time for the last aircraft is minimised. There are obvious similarities between the arrivals and departures problem. Both of these problems can be seen to be similar to the machine job scheduling problem with release times and sequence dependent processing times. The objective is often to minimise the total completion time or to minimise earliness and tardiness.

Bianco et al. showed in [8] that the single machine problem with sequence-dependent set-up times and release dates is equivalent to the cumulative asymmetric travelling salesman problem with release dates. A dynamic programming formulation was used to attain lower bounds for the problem and then heuristic algorithms were presented to find approximate solutions. This model and algorithms were then applied to the arrivals problem to sequence arriving aircraft.

Abela et al. [1] presented a genetic algorithm and branch and bound approaches for solving the arrivals problem. Beasley et al. used mixed integer 0-1 formulations to solve the arrivals problem in [6] and Beasley et al. used Genetic Algorithms in [7]. In [10] Ernst et al. presented heuristic and exact algorithms for solving the aircraft landing problem.

All of these solutions assumed that separations only needed to be maintained between adjacent arrivals, a condition which is not sufficient for the departure problem at Heathrow. Using the separation rules for Heathrow it is easy to generate take-off schedules for which all pairs of adjacent aircraft have the required separations but other pairs of aircraft do not.

In [2], Anagnostakis et al. gave a summary of the criteria that affect departure scheduling and specified a model for the runway usage, accounting for departures, arrivals and runway crossings. A search tree was described to solve the problem and branch and bound techniques or an A* algorithm were recommended. The departure process was also analysed and a departure planner was proposed by Anagnostakis et al. in [3], [4] and [5].

Trivizas suggested a dynamic program to solve the departure order problem in [15] by limiting the number of aircraft that are considered for any place in the schedule to simplify the search space. The results of our experiments and the real data we have available show that, at Heathrow, aircraft often move forward or backwards up to eight places in a schedule.

Van Leeuwen et al. presented a constraint satisfaction based model for the departure problem in [13], including the allocation of aircraft to runways. Heathrow, however, has a far greater departure rate than was considered in [13].

There are important constraints at Heathrow due to the physical holding point structures that are not normally considered in the academic departure problem. Craig et al. looked at the effects of a simplified holding point structure in [9] and gave a dynamic programming solution for scheduling take-offs. However, holding point structures are more flexible than this in practice and in some configurations aircraft can enter the runway from either side.

The various constraints upon the Heathrow departure problem are identified in the problem description below.

2 Problem description

London Heathrow has two runways. Local agreements with nearby residences mean that only one runway can be used for departures at any given time, preventing the more efficient utilisation of the runways for both arrivals and departures. The demand at the airport is not evenly spread throughout the day and at times there is severe congestion in the departure system, arrival system or both.

Aircraft taxi along the taxiways from the stands to holding points near the end of the current departure runway. At the holding points, a runway controller is responsible for sequencing the aircraft into the best order for take-off. The order of take-off can have a large effect upon the throughput of the runway. The lack of space for taxiways, due to the restricted size of the plot of land upon which Heathrow is situated, means that it is usually impractical to reorder aircraft before they reach the holding points.

There are a number of constraints upon when an aircraft can take off due to required separations between aircraft, take-off time slots and physical manoeuvring room.

A wake vortex separation is imposed between aircraft to ensure that the wake vortex from the first aircraft has dissipated before the second takes off. This separation depends upon the relative weight classes of the aircraft.

Aircraft leave an airport along predefined departure routes to limit the workload of pilots and controllers. To ensure that a safe distance is maintained between aircraft in flight, a time separation is imposed at take-off. This separation depends upon both the relative departure paths of the aircraft and their speeds. At times the departure route separation may be further increased for flights into congested airspace to control the congestion. This increased separation is called a Minimum Departure Interval, or MDI.

Some aircraft have a fifteen minute departure timeslot within which they must take off. This is called a Calculated Time of Take-off (CTOT) and is applied to control crowded airspace and congestion at busy destination airports. This will limit both how early and how late the aircraft can depart and will affect take-off schedules accordingly.

The physical holding point structure is an important constraint upon the departure system at Heathrow as it determines both how much reordering is possible in the holding point and the cost of achieving the reordering in terms of pilot and controller time and effort. With two physical runways usable in either direction there are, of course, four possible holding point configurations to consider. However, only three of them are usually used for departures because of agreements with local residents. The constraints imposed by the physical structure of the holding points are different for each holding point configuration.

The departure system is a dynamic process with new aircraft becoming available over time, aircraft that have already taken off leaving the system and the status of aircraft within the system changing over time. The runway controller will usually have visibility only of the aircraft already within the holding point and will usually have insufficient time to take much consideration of the aircraft that may arrive at the holding point later. The controller will therefore currently aim to leave enough flexibility in the reordering to be able to account for later problems that may arise.

The objective of our research is to increase the throughput of the departure runway subject to the real-world requirements while always maintaining the safe separations between aircraft, avoiding undue workload for controller and pilots, and ensuring aircraft take-off within their CTOT slots, where possible.

3 Solution method

In order to aid the throughput of a departure runway we are proposing a decision support system where take-off orders can be suggested to the runway controller in order to ease the task. A decision support system can take into account more aircraft than can a human controller. One of the goals is to avoid later problems that may not be visible to the controllers at the time. The full paper will present results showing the advantages of increased knowledge in both meeting CTOTs and decreasing the holding point delay.

To be of use a decision support system must react extremely quickly to changes in circumstances. We therefore assume a maximum search time of one second as the search time will directly reflect the responsiveness of the system. As the number of aircraft under consideration increases, an exhaustive search quickly becomes infeasible within the allowed timeframe and new search methods are necessary. Here we present hybrid metaheuristic searches that perform well even in the limited time available.

The holding point structure is a major constraint upon what reordering can be done in the holding point and is the main reason why previous research is not applicable to the problem at Heathrow. This physical structure must be incorporated into the solution method.

We will discuss the problems with incorporating the structure directly into the solution space in the full paper, including aspects such as the exponential increase in the size of search space and number of local optima. We will also discuss the varying value of the paths the aircraft use to traverse the holding point and the issues that this introduces if solutions generate the paths during exploration of the search space. For example, a dynamic programming approach has to consider the positions and movement of aircraft in the holding points as a part of the state, as these have an effect upon the reordering possible with later aircraft. This ensures that the number of states becomes prohibitively large for real holding point structures, limiting the value of the approach in the real world.

Our alternative model for the problem involves a two stage, hybrid, approach. The problem is decomposed by treating the holding point constraints separately to the main search. Metaheuristic searches are applied to find good take-off orders for the aircraft, evaluating the cost of a schedule without considering whether it is feasible or not. Each take-off order that is found by the search is tested for feasibility by heuristically assigning paths through the holding point to the aircraft and determining whether the desired reordering is possible. This heuristic assignment ensures that only good paths are allocated to aircraft and that the path taken reflects the amount of time the aircraft has available, eliminating the need to cope with these aspects in the objective function of the search itself.

In order for our model to work in a real-time system the feasibility check for whether the reordering is possible or not must be extremely fast. To achieve this we use preprocessing of the holding point structure and the possible traversal paths so that less calculation is required during the feasibility check itself. We will discuss the holding point feasibility check in the full paper, with details of how the model copes with aircraft already in the holding point.

Even though the searches are given only a very small execution time, we employ a large number of possible moves. These moves are designed around the characteristics of the problem and help to reduce the number of local optima discovered. In the full paper we will discuss the moves we allow and justify them in the context of the solution space for the problem.

4 Formal model of the problem

Let n be the number of aircraft currently under consideration. Let $i \in \{1, \ldots, n\}$ be an integer to represent an aircraft currently in the system. Let a_i represent the position of aircraft i in the arrival order at the holding point, so if i is the first aircraft to arrive then $a_i = 1$. Let c_i be the position of aircraft i in the departure order, so if i is the first aircraft to depart then $c_i = 1$. The positional delay of aircraft i in the take off schedule is then given by $c_i - a_i$.

Let d_i be the time of take-off for aircraft i and h_i be the time aircraft i arrives at the holding point. The holding point delay for aircraft i is then given by $d_i - h_i$

For reasons explained in the full paper the end of the CTOT slot cannot be implemented as a hard constraint. We define a function $C(d_i, b_i, l_i, h_i)$ to assign a penalty to aircraft i, depending upon its take-off time in relation to its holding point arrival time and its assigned CTOT slot, where for aircraft i, b_i is the time of the beginning of the CTOT slot and l_i is the time of the end of the CTOT slot. This is not a simple function as it has to account for different costs of delays and introduces factors to allow for schedule deviation from predicted times of take off. This function is fully and mathematically defined in the full paper.

We aim to minimise the total holding point delay, positional delay and CTOT non-compliance for all aircraft in the system. The objective function for our minimisation problem can be seen in expression 1, where the weights W_1 and W_2 give the relative importance of each factor. It should be noted that the feasibility check at the holding point is specifically excluded from this evaluation, as discussed earlier.

Minimise

$$\sum_{i=1}^{n} (W_1(d_i - h_i) + W_2(\max(0, c_i - a_i))^2 + C(d_i, b_i, l_i, h_i)) \tag{1}$$

Given a take-off schedule, it is necessary to predict the take-off times for aircraft in order to evaluate the cost of the schedule. To predict d_i for any aircraft i we assume that all aircraft take off as early as possible while fulfilling all required separation and physical movement time constraints.

The earliest time an aircraft can take off while maintaining all required separations can be calculated given the take-off times for all earlier flights.

Let v_i be the weight class of aircraft i, r_i be the route aircraft i will depart upon and s_i be the speed class of aircraft i. For aircraft i and j, where $c_j < c_i$, we define $V(v_j, v_i)$ to be a function to return the required wake vortex separation and $R(r_j, s_j, r_i, s_i)$ to be a function to return the required separation for the routes and speeds. Both of these are defined to return the standard separations in accordance with current regulations.

The earliest time, e'_i, at which aircraft i can take off, while maintaining required separations, is then given by equation 2.

Given a known path, t_i, by which aircraft i will traverse the holding point structure and a function $T(t_i)$ to predict the traversal time of the holding point for an aircraft using path t_i, the earliest time the aircraft could physically reach the runway, e''_i, can be predicted using equation 3.

Using the physical taxi-time and separation constraints and enforcing the start of the CTOT slot, b_i, as a hard constraint, the earliest departure time for aircraft i can be seen to be given by equation 4.

$$e'_i = \begin{cases} 0 & \text{if } c_i = 1 \\ \max\limits_{j \in \{1,\ldots,n\} | c_j < c_i} (d_j + \max(V(v_j, v_i), R(r_j, s_j, r_i, s_i))) & \text{if } c_i \geq 2 \end{cases} \tag{2}$$

$$e''_i = h_i + T(t_i) \tag{3}$$

$$d_i = \max(e'_i, e''_i, b_i) \tag{4}$$

5 Results

In the full paper we will present results from applying our model to six datasets containing real, historic data provided by National Air Traffic Services Ltd. In order to test the searches we determine what information would have been available to a decision support system at the time and provide only that information to the searches. The tests start at the beginning of the data period, with knowledge of all aircraft in the holding points and all aircraft taxiing around the runways towards the holding point at that time. The tests progress through the data period in steps of 15 seconds. Aircraft are added to the system as they leave their stands and start to taxi towards the holding points, or when they reach the holding points, depending on the test. Aircraft are removed from the system after take off, once they can have no more effect upon the take-off times of later aircraft.

We will show that with a level of knowledge and constraints similar to that which the controllers currently have the searches perform similarly to the controllers. By increasing the level of knowledge of the system to include the aircraft taxiing towards the holding point as well as those already there we see that the searches can perform significantly better than the controllers do at present. Indeed the model predicts a decrease in the delay of between 10% and 25%, depending upon the test dataset, while simultaneously increasing the CTOT compliance. We conclude from this that the employment of a decision support system could improve both CTOT compliance and delay.

We will show that even a first descent algorithm can perform well with the moves that we allow, once the holding point contraints have been moved to a separate evaluation. As it is not possible to remove all local optima, however, we also show that applying Tabu Search and Simulated Annealing algorithms to the problem gives much better consistency of results as well as better mean delays.

We therefore conclude that the hybrid metaheuristic approach with heuristic allocation of holding point traversal routes to aircraft is an appropriate approach to this problem, consistently giving good results for the data provided.

6 Concluding comments

The departure problem is a complicated one due to the many constraints upon the schedule and the sequence-dependent separations between aircraft. At present the runway controllers have to reorder aircraft manually under very tight time constraints. The goal of this work is to provide the methodologies which could underpin the development of an advisory system to help them to perform this exacting task.

Existing research into both the arrival and departure problems commonly checks only the separations between adjacent aircraft and assumes that required separations will therefore be achieved between all pairs of aircraft. This is, however, not the case for the normal required separations at Heathrow and certainly not the case in the presence of an MDI upon certain departure routes.

None of the techniques previously applied to the departure problem consider the physical constraints of the holding point structure that are present at Heathrow Airport. We can see, however, that for Heathrow the physical holding points are an important constraint which affect which take-off schedules are possible and the workload required.

To be of use to a controller, an advisory system must react very quickly to changes in the situation. Our hybrid metaheuristic approach is ideal in this kind of environment as it can very quickly produce high quality results and so react to changing circumstances very quickly.

We present a model which takes account of the real world constraints at Heathrow such as holding point structure, workload, separation rules and CTOTs. Moreover, we show, via simulations of the real situation, that by increasing the knowledge available to a decision support system, the holding point delay can be improved while simultaneously better meeting the other constraints. The results predict a large reduction in the delay suffered by aircraft as well as better compliance to the CTOT time slots.

References

[1] Abela J, Abramson D, Krishnamoothy M, de Silva A, Mills G (1993) Computing optimal schedules for landing aircraft. Proceedings of the 12th National Conference of the Australian Society for Operations Research, Adelaide, July 7-9, 1993, p71-90 Available at: http://www.csse.monash.edu.au/∼davida/papers/asorpaper.pdf [26 August 2004].

[2] Anagnostakis I, Clarke J-P, Böhme D, Völckers Uwe (2001) Runway operations planning and control, sequencing and scheduling. Proceedings of the 34th Hawaii International Conference on System Sciences (HICSS-34), Hawaii, January 3-6, 2001.

[3] Anagnostakis I, Idris HR, Clarke J-P, Feron E, Hansman RJ, Odoni AR, Hall WD (2000) A conceptual design of a departure planner decision aid. 3rd FAA/Eurocontrol International Air Traffic Management R & D seminar, ATM-2000, Naples, Italy, June 13-16, 2000. Available at: http://atm-seminar-2000.eurocontrol.fr/acceptedpapers/pdf/paper68.pdf [26 August 2004]

[4] Anagnostakis I, Clarke J-P (2003) Runway operations planning, a two-stage methodology. Proceedings of the 36th Hawaii International Conference on System Sciences (HICSS-36), Hawaii, January 6-9, 2003.

[5] 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. Available at: http://icat-server.mit.edu/Library/Download/167_paper0024.pdf [26 August 2004]

[6] Beasley JE, Krishnamoorthy M, Sharaiha YM, Abramson D (2000) Scheduling aircraft landings - the static case. Transportation Science 34:180-197

[7] Beasley JE, Sonander J, Havelok P (2001) Scheduling aircraft landings at London Heathrow using a population heuristic. Journal of the Operational Research Society 52:483-493

[8] 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

[9] Craig A, Ketzscer R, Leese R A, Noble S D, Parrott K, Preater J, Wilson R E, Wood D A (2001) The sequencing of aircraft departures. 40th European Study Group with Industry, Keele 2001. Available at: http://www.smithinst.ac.uk/Projects/ESGI40/ESGI40-NATS/Report/AircraftSequencing.pdf [26 August 2004]

[10] Ernst A T, Krishnamoorthy M, Storer R H (1999) Heuristic and Exact Algorithms for Scheduling Aircraft Landings Networks, Vol 34, Number 3, p229-241

[11] Idris HR, Delcaire B, Anagnostakis I, Hall WD, Clarke JP, Hansman RJ, Feron E, Odoni AR (1998a) Observations of departure processes at Logan airport to support the development of departure planning tools. Presented at the 2nd USA/Europe Air Traffic Management R&D Seminar ATM-98, Orlando, Florida, Dec 1st-4th 1998. Available at: http://atm-seminar-98.eurocontrol.fr/finalpapers/track2/idris1.pdf [26 August 2004]

[12] Idris HR, Delcaire B, Anagnostakis I, Hall WD, Pujet N, Feron E, Hansman RJ, Clarke JP, Odoni A (1998b) Identification of Flow Constraint and Control Points in Departure Operations at Airport Systems. Proceedings of the AIAA Guidance, Navigation and Control conference, Boston, MA, August 1998.

[13] van Leeuwen P, Hesselink H, Rohling J (2002) Scheduling Aircraft Using Constraint Satisfaction. Electronic Notes in Theoretical Computer Science 76.

[14] Newell GF (1979) Airport Capacity and Delays. Transportation Science 13:201-241

[15] Trivizas DA (1998) Optimal Scheduling with Maximum Position Shift (MPS) Constraints: A Runway Scheduling Application. Journal of Navigation 51:250-266