The Effect of the Planning Horizon and the Freezing Time on Take-off Sequencing

Jason A. D. Atkin and Edmund K. Burke
School of Computer Science and I.T.
University Of Nottingham, Jubilee Campus
Nottingham, NG8 1BB, England
Email: jaa/ekb@cs.nott.ac.uk

John Greenwood
NATS CTC
4000 Parkway, Whiteley, Fareham
Hampshire, PO15 7FLT, England

Dale Reeson
National Air Traffic Services
Heathrow Airport, Hounslow
Middlesex, TW6 1JJ, England

***Abstract*— In this paper we consider the problem faced by runway controllers at London Heathrow Airport as they attempt to determine the best order for aircraft to take off. The order in which aircraft take off can have a large effect upon the throughput of the runway and the consequent delay for aircraft. Although two runways are available for use at Heathrow, in order to control noise for residents on the flight paths, only one is used for departures at any time. The other is used for arrivals. As this is one of the busiest international airports in the world, and the runway is the bottleneck for the departure system, it is important to attain high runway throughput. A runway controller re-arranges the aircraft for take-off within holding points at the ends of the runways, currently performing this task manually. The decision about the take-off order has to be made with very limited decision time, precipitating investigations into the development of an on-line decision support system to aid in this task. We have developed a model for such a decision support system and a simulation of the departure system, and discuss these in this paper. Our experiments predict that our system can provide suggestions fast enough to be of use in practice, while also being able to consider more aircraft than a real runway controller can. Thus it can obtain consequent benefits from highlighting potential problems before they occur. In order to maximise the potential benefits of a decision support system, it is important to understand how the decisions are affected by the planning horizon, how far ahead the system considers aircraft. From this information, we can better understand the inputs that the system would need to be able to fulfil its role. The position of an aircraft in the take-off order has to be frozen at some point prior to take-off. Our investigations revealed a trade-off between the time at which the take-off order was frozen and the planning horizon necessary to get the best results from the system. This paper considers this trade-off, evaluates where the planning horizon needs to be, and shows the relationship with the point at which the take-off schedule is frozen. We present results which show that there is a substantial delay benefit from including taxiing aircraft. Furthermore, we show that, with a schedule frozen for two minutes before take-off, the vast majority of the benefit can be gained from a knowledge of only those aircraft which will arrive at the holding point within the next eight minutes, and that aircraft can be ignored until they push-back from their stands. Finally, the results also show how the planning horizon must increase if the time for which the schedule is frozen increases.**

I. INTRODUCTION

A schematic diagram of Heathrow can be seen in figure 1. The airport currently has two runways, north and south, and both can be used in either direction. Each runway is named according to the current direction of use and whether it is on the left or right from that direction. There are currently four terminals, labelled T1 to T4 in figure 1, although a fifth is being built. Near the ends of each runway, the taxiways form holding points, labelled HP in the diagram, within which a runway controller manually performs the complex task of ordering the aircraft for take-off. Heathrow's proximity to housing means that only one runway can be used for departures at any time, with the other being utilised by arrivals. This makes it vital to achieve a high throughput on the one available departure runway.

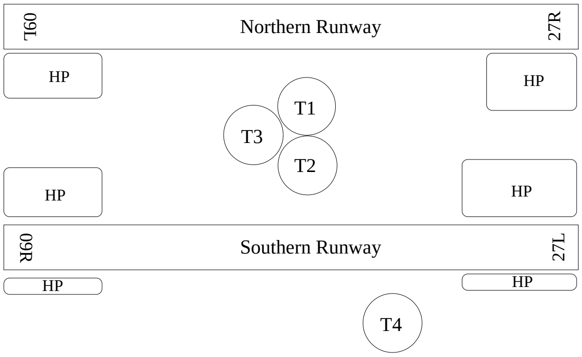

Fig. 1. The layout of London Heathrow Airport

We have designed a decision support system for the runway controller, to aid in the difficult task of ordering the aircraft for take-off. In order to make the best use of this, it is vital to understand the effect of other parts of the departure system upon it. In this paper we consider the advantages of moving the planning horizon, giving the system knowledge of some or all of the aircraft on the taxiways, allowing these to be considered in the schedule, before they actually come under the control of the runway controller. We also consider the effect of changing the time at which the schedule is frozen, effectively changing the other end of the planning window, and consider the relationship between this and the planning horizon.

In this paper, we first present the take-off scheduling problem and discuss the planning horizon and schedule freezing issues. We then present our decision support system and departure system simulation, which we use to investigate the effects of varying the planning horizon and the time at which the schedule is frozen. We end the paper by presenting the results of our experiments and by drawing conclusions about the implications.

II. THE TAKE-OFF SCHEDULING PROBLEM

The take-off scheduling problem can be summarised as finding a take-off order with low delay, low workload for pilots and controllers and maximal compliance to the various temporal and physical constraints upon the aircraft in the schedule. It is important, as the delay in the schedule has obvious implications for airlines, passengers and the airport.

When aircraft take off, a minimum separation time is required between them. Unnecessarily large separations delay the following aircraft from taking off, causing unnecessary fuel burn, with the associated financial and environmental costs. The runway controller's aim is to order the aircraft so that the number of larger separations that are required is reduced, consequently reducing the delay for the aircraft awaiting take-off.

Mandatory separation times are required between aircraft at take-off, to ensure that the wake vortices created by the previous aircraft to take off have had time to dissipate. This *'wake vortex separation'* depends upon the weight classes of the current and preceding aircraft. Every time the take-off schedule has a lighter class aircraft following a heavier class aircraft, a larger separation will be required, adversely affecting runway throughput and delaying later take-offs.

Aircraft take off along fixed routes, called *'Standard Instrument Departure'* (SID) routes. A further mandatory separation time is required to ensure that in-flight separation distances will be attained. This *'SID separation'* is based upon the relative SID routes on which the aircraft are departing. Aircraft departing on the same or similar SIDs may need a larger separation to be applied at take-off to ensure that the in-flight separation is attained. This separation is further modified according to the speed groups of the aircraft, to allow for the fact that a faster or slower following aircraft will decrease or increase the separation distance. A good schedule will, therefore, often ensure that aircraft with similar SIDs are not in adjacent positions in the take-off order.

Some aircraft have a *'Calculated Time Of Take-off'* (CTOT) assigned to them. This designates a target take-off time and is assigned to avoid congestion, en-route or at destination airports, by staggering the arrivals to the congested position. As aircraft are permitted to take off up to five minutes before or ten minutes after the CTOT time, it effectively designates a fifteen minute take-off window. CTOTs are assigned without regard to the source airport, so can be difficult, or impossible, for the runway controllers to achieve at busy periods. To allow for this, a limited number of five-minute extensions can be used, but as few as possible should be employed.

Take-off scheduling has been considered in a number of research papers in the past. Idris et al. considered the departure system at Boston Logon airport in [1] and [2], concluding that the runway represents the bottleneck. Anagnostakis et al. developed a two-stage departure planner, which they presented in [3] and [4]. A constraint satisfaction based model was applied to the take-off scheduling problem by van Leeuwen et al in [5]. Trivizas used a maximum position shift approach in [6], with dynamic programming.

There are similarities between arrival scheduling and departure scheduling, but also important differences. Abela et al, [7], Beasley et al, [8] and [9], and Ernst et al, [10], have all considered the arrival scheduling problem and applied different methods to solving the problem. However, although there are similarities in the wake vortex separation rules, the downstream constraints and the constraints inflicted by the holding point structure do not apply. Thus, the direct application of such methods to our problem is not possible.

The departure problem was considered by Bianco et al. as a special case of the cumulative asymmetric travelling salesman problem (ATSP) with release dates, in [11]. However, the equivalency does not hold for Heathrow, as the SID separation rules do not obey the triangle inequality, so it is not always sufficient to only ensure adequate separations between adjacent take-offs.

At Heathrow, aircraft are directed around the taxiways by Ground Movement Controllers (GMC) to holding points near the end of the runway, within which a runway controller aims to reorder them for take-off. Due to the complexity of the roles of the controllers involved, and the time constraints upon them, there is often little co-ordination between them for individual aircraft. The GMC usually determines to which of a number of holding point entrances each aircraft should be delivered, considering issues such as ease of reaching the entrance, congestion in the holding area and on the taxiways, and any requests the runway controller may have made.

Performing the overtaking within the holding points, in the last few minutes before take-off, ensures that the uncertainties in the take-off schedule are minimised. The controller does not need to consider the variability of taxi-times nor the contention with arriving aircraft for stands that would occur if the scheduling was performed at the stands. However, the structure of the holding points restricts how much overtaking can be performed. To be overtaken, there must be a position in which an aircraft can wait and a clear path for the overtaking aircraft to go past.

As, at Heathrow, the overtaking to achieve the take-off schedule is performed within the holding points, the constraints they inflict upon what overtaking is possible must be considered in the solution method chosen. The applicability of the aforementioned research to the take-off scheduling at Heathrow is, therefore, limited. In [12], Craig et al. considered the effect of a simplified holding point structure at Heathrow, applying a dynamic programming approach to solve it. The position of the aircraft have to be included as they have a great effect upon the feasibility of further reordering. The holding points are actually much more complicated than in [12], however, so the feasibility of the dynamic programming approach quickly breaks down as the number of possible positions for aircraft increases. We presented an alternative approach in [13], further developed the approach in [14] and apply it in this paper to investigate the planning horizon and take-off freezing time.

III. PLANNING HORIZON AND TAKE-OFF FREEZING

Departing aircraft can be considered to pass through the following states:

1) At the stand.
2) On the taxiway.
3) At the holding point.
4) At the holding point in a frozen take-off order.
5) Taken off.

A runway controller will usually only consider the aircraft that are already within the holding point when determining the take-off order. The first aspect that we consider in this paper is the effect of increasing the planning horizon beyond the holding point arrival time, so that at least some of the aircraft on the taxiways, and possibly also some that are still at their stands, are considered in the take-off schedule.

The earliest take-off time is a vital piece of information for any take-off scheduling system. In practice, however, the variability in the taxi times means that there is a degree of uncertainty in the earliest time at which an aircraft could reach the runway. This uncertainty decreases as the aircraft gets closer to the holding points and the runway.

It is useful to understand the effects of the planning horizon as it is easier to obtain shorter term predictions of taxi times than longer term ones. The task of a taxi time prediction system (providing these to a decision support system) would, consequently, be simplified. Understanding the effects can avoid the need for an unnecessarily large planning horizon. Additionally, larger planning horizons mean incorporating more aircraft in the search, so the search space is much larger, making the job of the decision support system much harder. Finally, it is useful to know whether aircraft need to be considered before they have pushed-back from their stands, as there will always be much more uncertainty involved in the holding point arrival time until that point.

The second aspect we wish to consider is the effect of take-off freezing. At some point before take-off the aircraft will have been given instructions for take-off or actually be lining up for take-off. At this point it is impractical to change the position of that aircraft in the schedule, except in extreme circumstances such as an inability to take-off for some reason.

The runway controller will usually give some conditional clearances to pilots, telling them to line up for take-off following another specified aircraft. Conditional clearances effectively freeze that part of the take-off schedule. They are used as they have workload advantages for a runway controller, as well as giving the pilots more visibility of their planned take-off time. The take-off order could still be changed after conditional clearances have been given, but doing so will involve more work than before they are given.

Additionally, a controller will normally have a planned take-off order in mind for the aircraft in the holding points, probably with gaps later in the schedule that future arrivals could fill, if possible. Reconsidering the order for these aircraft requires time and effort, so is not always practical due to the extremely busy workload of the controller. Therefore, often, more of the schedule, beyond that where conditional clearances have been given, can be considered as frozen, for all practical purposes, as it may be frozen from the point of view of the controller making the decisions. Of course, if necessary, changing this part of the schedule is less costly than changing the schedule where conditional clearances have been given, it will just not necessarily be considered.

We can consider a conditional clearance to move an aircraft from state 3 to state 4 in the list above. Conditional clearances given by controllers are not actually as constraining as freezing a part of the take-off schedule is. It is often possible, with only limited extra work, for a controller to fit an extra take-off into an existing large separation in the take-off schedule, especially if the take-off times for the other aircraft are unaffected, assuming, of course, that the new aircraft can perform the required overtaking within the holding points.

To investigate the effects of early freezing of the schedule, and of the planning horizon, we used a decision support system we had designed, together with an abstract simulation of the departure system of the airport. To ensure that the experiments were realistic, we used real, historic data provided by National Air Traffic Services. This data included temporal information such as the times at which the aircraft pushed back from their stands, arrived at the holding points and took off, as well as the details of the aircraft involved, such as the weight classes, SID routes and speed groups. In the following sections we explain the decision support system, simulation and experiments.

IV. DECISION SUPPORT SYSTEM

A decision support system for a runway controller must make decisions about the desirable take-off order using only the information available at the time. The responsiveness of an on-line decision support system to a changing situation is determined by the search time required to make each decision. Our system has a search time of around a second on a 2.4GHz pentium 4, making it fast enough for use as a real-time system.

A decision support system will have to solve a sequence of problems over time, each of which consists of only a snapshot of the daily schedule. A system would be running constantly, re-deciding what to do as the situation evolves, responding to changing circumstances such as new aircraft pushing back from stands or the effects of previous decisions. The aim, however, is to obtain the best overall schedule for the day, rather than necessarily the solution of highest throughput at any particular instant in time. The suggested schedule that the system returns must, therefore, not cause problems for the scheduling of later take-offs.

Furthermore, as the overtaking required to achieve a take-off scheduling takes place within the holding points at the end of the runway, our decision support system not only considers whether a take-off order is desirable, but also whether it is achievable within the holding points.

Our decision support system has three main parts. The first part uses a tabu search to investigate the possible take-off orders and to seek a high quality one. The second part is a system to verify whether the overtaking required to achieve a desired take-off schedule is possible. The third part determines how good a take-off schedule is by predicting take-off times for aircraft and evaluating a consequent cost for the schedule. Each of these parts is described below.

A. Tabu search

Tabu search was first introduced by Glover in [15] and has been applied to many different kinds of optimisation problems. Further information about Tabu Search can be found in [16]. The flow of our tabu search is illustrated in figure 2. Starting at an initial, achievable, take-off schedule (called a solution) the search makes progressive changes, seeking better and better schedules.

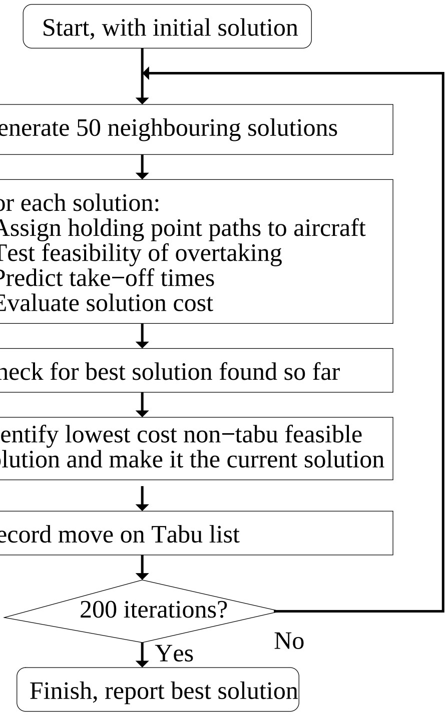

Fig. 2. The tabu search

Each neighbouring solution is created by applying a random move. Available moves include exchanging the positions of two aircraft in the schedule or moving up to five aircraft forward or backward in the schedule. In this way, fifty different schedules are created at a time. The search then selects the best of these and adopts it as the new current solution.

In order to inhibit the search from cycling between a few good schedules, a tabu list is applied. Whenever a move is made, the details of the old positions of the aircraft that were moved in the schedule are stored. The next ten moves after that one are explicitly prohibited from returning all of these aircraft to the positions from which they came. Any move which would do so is declared to be tabu and will not be accepted as the new current schedule. This avoids the search returning to schedules that have been recently evaluated, aiding it in escaping some of the local optima. Even if a solution is declared to be tabu, however, it is still recorded as the best if it has a lower cost than the best solution found so far.

A check is performed for each solution to verify that the required overtaking is achievable within the holding point. If it is not, the solution is declared as infeasible and will not be adopted. During the search the best solution found so far is always maintained. This is returned as the suggested schedule at the end of the search.

B. Overtaking within the holding points

It is important to verify that the required overtaking can be performed within the holding points. To ensure this we have a two-stage process. First sensible paths are assigned to aircraft, then the holding point model is used to verify that the overtaking can all take place.

We use a directed graph model of the current holding point structure for this verification. An example graph, for the 27R holding point, is given in figure 3. The graph has a node for each valid waiting position for an aircraft and arcs for transitions aircraft can make between nodes as they move through the holding point.

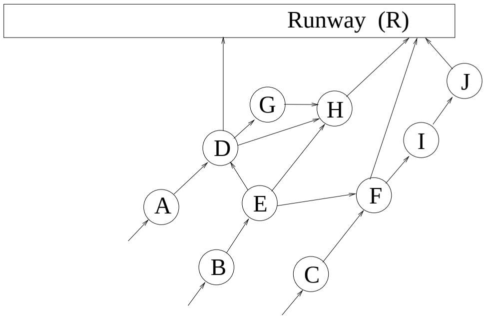

Fig. 3. The 27R holding point network

A path assignment heuristic is used to allocate paths through the holding point to aircraft. Paths are assigned based upon the overtaking that is required. Each path can be uniquely identified by the set of nodes it passes. Fast paths are assigned to aircraft which overtake and slower ones to those which are overtaken. For example, with the holding point given in figure 3, if two aircraft that arrive at entrance A must reverse their order the overtaking one will be assigned the path ADHR and the overtaken one ADGHR. The heuristic method ensures that the path allocation is sensible, from the point of view of the runway controller. This is important as, otherwise, a solution would be immediately rejected, and the decision support system would be worthless.

Once paths have been allocated, a feasibility check is performed to verify whether the required overtaking is possible. The holding point graph for the current holding point is used to determine this.

Initially, each aircraft that is already in the holding point is placed in the node related to its current position. If the aircraft is between nodes it is placed at the node it will next enter. Aircraft that are currently on the taxiways are placed in queues, in predicted arrival order, at the holding point entrance at which they are predicted to arrive. If any entrance node is empty then the first aircraft from that entrance queue is placed in the node and removed from the queue.

The overtaking test is then performed by moving aircraft, one node at a time, through the holding point graph towards the runway. Each aircraft can only enter the next node on its path if the node is empty, and entering it will not block another aircraft from reaching the runway on time. A fast method has been developed for verifying that the latter condition has been met. This works by building partial take-off orders for each node and tracking the number of free nodes that an aircraft could use to move out of the way of an aircraft that should take off before it.

If all aircraft can enter the runway node in the desired take-off order then the order is achievable within the holding point. If not, then the schedule is discarded as being unachievable, or infeasible.

C. Schedule cost

The tabu search requires an objective measure of the cost of a schedule, its aim being to find a low cost solution. To measure the cost of a schedule, take-off times are predicted for all aircraft, then the total delay is measured and the number of missed CTOTs are counted.

Predicting take-off times requires that the earliest time of take-off is calculated for each aircraft, allowing sufficient time for the aircraft to reach the runway and ensuring that all required separations are maintained.

The time at which an aircraft will be able to reach the runway is calculated by adding an expected traversal time of the holding point to the predicted holding point arrival time. The traversal times used are pessimistic to ensure they are achievable.

The earliest time an aircraft can take off while obeying separation rules can be calculated by considering each previous take-off in turn. Adding the higher of the required wake vortex or SID separations to the take-off time of the earlier take-off gives an earliest take-off time for this aircraft.

If the aircraft has a CTOT then it is not permitted to take off more than 5 minutes before the CTOT time. This forms an additional constraint for the earliest take-off time.

As there is never an advantage in delaying the take-off, aircraft are assumed to take off as early as they can, so the predicted take-off time is assumed to be the earliest time which allows the aircraft to physically reach the runway, obey any CTOT and fulfil all required separation rules.

Once the take-off times have been predicted, a cost for the schedule can be calculated. The cost takes account of the number of CTOT extensions required, the total time for which the aircraft are in the holding points and the deviation of the schedule from the first-come-first-served schedule. The cost for the schedule is a combination of these three factors, weighted so that reducing the CTOT extensions is the primary criterion, reducing delay is secondary and reducing deviation is tertiary. If any schedule involves an aircraft which is scheduled too late even for a CTOT extension it is given a prohibitively high cost, preventing the search from adopting these schedules.

We refer to the time in the holding points as the delay for aircraft. We use a *delay-based* rather than *throughput-based* measure for the cost of the solutions to the problems. Both delay-based and throughput-based measures share the characteristic of penalising larger separations. However, a delay-based measure penalises these separations more if they are earlier in the schedule, as they delay more aircraft. It, therefore, favours schedules where the large separations are later on in the schedule and are more likely to be utilised by later aircraft. This often aids later iterations of the search to perform better.

V. SIMULATION

Our departure system simulation allows us to investigate the consequences of adopting the take-off sequence recommended by the decision support system. The simulation maintains information about the expected arrival time at the holding point for aircraft on the taxiways and the position within the holding point of any aircraft currently there.

The simulation uses a one minute timestep, presenting the decision support system with a sequence of problems consisting of only the information an on-line decision support system would have. At each step, the decision support system is given the opportunity to order the aircraft for take-off, providing the simulation with information about the desired take-off order and how it will be achieved. The simulation is responsible for using this information to update the aircraft positions before presenting the next problem to the system to solve.

As mentioned in section III, some or all of the aircraft on the taxiways may be included in the problem the simulation presents to the system. For these aircraft, the system is given a predicted arrival time at the holding point. In our tests we assumed no uncertainty in the arrival time, as we wished to evaluate the effects of varying the planning horizon without having to account for the effects of the uncertainty. The prediction is, therefore, always accurate in these tests.

The decision support system also needs to know which holding point entrance (for example, A, B or C in figure 3) the aircraft is due to arrive at, in order to test the feasibility of performing the required overtaking. For these experiments we assumed the Ground Movement Controller (GMC) delivered each aircraft to the nearest entrance to its stand. Experiments have shown that better results can be obtained by delivering the aircraft more evenly across entrances, allowing more flexibility in rescheduling. However, assigning aircraft to the nearest entrance is the only allocation method that is guaranteed not to involve prohibitive work from the GMC under any circumstances.

At the end of each experiment, all aircraft will have been scheduled for take-off and the final schedule can be evaluated to determine the CTOT compliance and delay for aircraft.

TABLE I

DATASET DETAILS AND MANUAL RESULTS

Dataset	Number of aircraft				Manual results	
	Heavy	Medium	Light	CTOTs	CTOTs	Delay (s)
1	90	239	1	100	5	366
2	100	244	1	98	5	312
3	64	193	2	42	4	372

TABLE II

CTOT COMPLIANCE AND DELAY FOR VARIABLE TAXI KNOWLEDGE

Taxi knowledge(s)	Dataset 1		Dataset 2		Dataset 3	
	CTOT	Delay	CTOT	Delay	CTOT	Delay
0	5	307	2	281	4	312
30	4	294	1	255	4	307
60	4	274	2	241	4	274
90	3	266	2	237	4	262
120	3	262	2	233	4	264
180	3	259	1	223	4	257
240	3	248	1	216	4	251
300	3	249	1	217	4	252
360	3	246	1	215	4	248
420	3	246	1	213	4	244
480	3	242	1	213	4	244
540	3	242	1	213	4	244
600	3	242	1	213	4	244
660	3	242	1	213	4	244
720	3	242	1	212	4	244
780	3	242	1	212	4	244

TABLE III

CTOT COMPLIANCE AND DELAY FOR VARIABLE TAXI KNOWLEDGE, IGNORING STAND LEAVING TIME

Taxi knowledge(s)	Dataset 1		Dataset 2		Dataset 3	
	CTOT	Delay	CTOT	Delay	CTOT	Delay
0	5	307	2	281	4	312
30	4	294	1	255	4	307
60	4	274	2	241	4	274
90	3	266	2	237	4	262
120	3	262	2	232	4	264
180	3	259	1	223	4	257
240	3	248	1	216	4	251
300	3	249	1	217	4	252
360	3	246	1	214	4	248
420	3	246	1	213	4	244
480	3	241	1	213	4	244
540	3	241	1	213	4	244
600	3	241	1	213	4	244
660	3	241	1	213	4	244
720	3	240	1	211	4	243
780	3	240	1	208	4	243

VI. RESULTS

We performed a series of experiments to evaluate the effects of varying the planning horizon and varying the schedule freezing time. Each experiment was performed ten times, for each of the same three datasets, because of the stochastic nature of the tabu search. The mean results are shown in the tables. In fact, in every case, the results for each of the ten executions were very close to each other, the CTOT compliance being the same in all cases and the delay varying only slightly. The number of aircraft of each weight class and the number with CTOTs are shown in table I, along with the performance of the real controllers in terms of the number of CTOTs missed and the average delay per aircraft.

In our first experiment we evaluated the effect of varying the planning horizon by altering the amount of knowledge the system has about taxiing aircraft. Aircraft are included in the simulation if they will both arrive at the holding point within a given time window and have already left their stand. For example, if this taxi knowledge time is 0, then aircraft are only considered once they have reached the holding point. Table II presents details of the performance of the system as the taxi knowledge time is changed. The number of CTOTs missed and the total delay for the aircraft in the system is given for each of the three datasets. In this experiment, the decision support system froze the take-off order for aircraft within the holding point two minutes before take-off and penalised changes made within three minutes of take-off.

We can clearly see from table II that the delay is significantly improved by giving the decision support system some knowledge of the aircraft currently on the taxiways, and that, in general, the system performed better with more knowledge. In some cases, the CTOT compliance was also improved. In all cases, however, the delay seems to plateau at around eight minutes taxi knowledge.

We note, however, that aircraft were only added once they had left their stands, so there is a limit to how early the simulation will include each aircraft, regardless of how far ahead the system is permitted to know about aircraft. To determine whether this was the reason for the plateau we performed a second experiment, where the aircraft were added to the simulation a given number of minutes before arrival at the holding point, even if they had not left their stands at that time.

Table III shows the results of this second experiment. The results are very similar to those in table II. We conclude that the plateau in performance is not entirely due to the push-back times. We will return to the reason for this plateau later in this section.

We also wished to consider the effect of varying the point at which the take-off order is frozen. When freezing the take-off order we usually assume that the order can only be frozen for aircraft actually at the holding point. This reflects the fact that controllers can only give conditional clearances for take-off based on aircraft actually in the holding point. It also avoids a possible feasibility problem, as the holding point arrival order is not fixed for aircraft still on the taxiway, as aircraft could push back from stands in front of them. As the arrival order is not fixed, the required overtaking to achieve a schedule may cease being possible, which is a problem if that part of the schedule has been frozen.

We performed experiments that varied the length of time for which the take-off order was frozen. In these experiments we assumed that the system was aware of all aircraft on the taxiways at that time. However, the results were remarkably similar to the results when the take-off order was not frozen, so are not documented here. Investigation revealed the reason for this. The schedules produced at any stage often had aircraft that were still on the taxiways placed in early positions in the schedule, for instance to fit a taxiing aircraft between two at the holding point that need a large separation. This significantly limited how much of the schedule was actually frozen. Combined with this, as the next experiment will show, the system was often scheduling the aircraft before they reached the holding point, so fixing the schedule for aircraft at the holding point made not difference.

In the absence of taxi time uncertainty, there is an obvious relationship between the effects of varying the planning horizon and varying the amount of time the schedule is frozen for, as both vary the number of aircraft that are available for scheduling. To examine the relationship further, and to better understand the results for the previous experiments, we performed a series of experiments where we varied both the planning horizon and the time for which the take-off order was frozen. To change the planning horizon we changed the time before arrival at the holding point at which each aircraft was considered by the search. The position of an aircraft in the take-off schedule was frozen a given number of minutes before take-off, regardless of whether it was at the holding point at that time. In order to be able to freeze the schedule without introducing the aforementioned feasibility problem, aircraft were added to the simulation the given number of minutes before holding point arrival, regardless of whether they had pushed-back or not.

The results are presented for dataset 1 in tables IV, V and VI. Table IV shows the number of CTOTs missed for taxi knowledge of up to 5 minutes. Beyond 5 minutes the number of CTOTs missed is 3 in all but one case (for frozen time 540 seconds and knowledge 360 seconds it is 4) so the results are not shown here. Tables V and VI show how the total delay in the schedule changes as the taxi knowledge and frozen time vary. The results for the other datasets show similar trends.

The diagonal pattern in the tables clearly shows the expected relationship between freezing the take-off order earlier and knowing earlier about taxiing aircraft that will soon arrive at the holding point. These affect opposite ends of the planning window and, as expected, have similar effects. Exact symmetry would not be expected, even if using an exact solution method to find the best schedule, as, even though the planning window may be a similar size the problems presented to the decision support system at each iteration will involve different aircraft, and sometimes even different numbers of aircraft.

TABLE IV

FROZEN TIME VS TAXI KNOWLEDGE - CTOTs

Frozen	Taxi knowledge for :					
time(s)	0s	60s	120s	180s	240s	300s
0	3	3	3	3	3	3
60	3	3	3	3	3	3
120	4	3	3	3	3	3
180	5	4	3	3	3	3
240	6	4	3	3	3	3
300	6	6	4	3	3	3
360	6	6	6	4	3	3
420	8	6	6	6	4	3
480	11	7	6	6	6	4
540	14	9	8	6	6	6

TABLE V

FROZEN TIME VS TAXI KNOWLEDGE - DELAY

Frozen	Taxi knowledge for :					
time(s)	0s	60s	120s	180s	240s	300s
0	284	257	255	251	247	245
60	285	262	255	251	247	246
120	287	261	260	252	246	247
180	321	278	262	259	248	249
240	345	306	273	259	252	249
300	391	337	305	276	252	252
360	409	374	332	302	272	253
420	489	421	373	332	302	272
480	520	463	412	373	331	302
540	574	513	473	410	373	331

For a larger taxi time knowledge, freezing the schedule earlier does not change the delay for the aircraft. This implies that the take-off order for these aircraft had already been determined by that point. Examination of the last time at which the take-off order was changed for aircraft supports this interpretation. This explains the earlier results, where varying the frozen time for which the schedule made little effect when the system had full knowledge of the taxiing aircraft.

Additionally, the fact that aircraft are scheduled earlier also tells us that sufficient knowledge was available to be able to do this much earlier than was actually necessary. The results obtained for very large taxi knowledge match the value at which tables II and III were plateauing. This, finally, explains the plateau in tables II and III. With around eight minutes warning about aircraft that will later arrive at the holding point, the system had enough knowledge to perform the scheduling, so additional information helped very little.

So, in summary, there is a benefit from accurate knowledge of taxiing aircraft, up to around eight minutes before holding point arrival, beyond which there is little benefit to delay. But, as tables II and III are similar, there is little gain from knowing about aircraft before they are ready for push-back. Additionally, tables IV, V and VI show that fixing the schedule for longer before take-off requires more knowledge of taxiing aircraft if similar performance is to be obtained from the system.

TABLE VI

FROZEN TIME VS TAXI KNOWLEDGE - DELAY

Frozen time(s)	Taxi knowledge for :					
	360s	420s	480s	540s	600s	660s
0	244	244	241	241	241	241
60	244	244	241	241	241	241
120	246	244	241	241	241	241
180	246	246	241	241	241	241
240	247	247	243	241	241	241
300	249	247	244	243	241	241
360	252	249	250	244	243	241
420	253	257	247	250	244	243
480	270	258	254	247	250	244
540	302	270	255	254	247	250

VII. CONCLUSION

Take-off scheduling within the holding points at Heathrow is a very complicated problem. Due to the amount of time runway controllers have to spend communicating with pilots and monitoring the airport, there is limited time available for them to consider the take-off scheduling task. It is common to have little knowledge of the aircraft on the taxiways, usually only considering them if large separations are seen within the take-off schedule for aircraft already at the holding points. The result of this is that problematic sequences of aircraft on the taxiways are not always foreseen in time to do something about them.

In this paper we have evaluated the effect of changing the planning horizon, including knowledge of taxiing aircraft that will soon arrive at the holding points. To do this, we used a simulation of the departure system, with a real-time decision support system that we designed taking the place of a runway controller. We have shown that there are definite delay and CTOT compliance benefits from having knowledge of taxiing aircraft, and we feel that this is where a decision support system will be of most benefit to the controllers.

We have shown that, for a two-minute frozen take-off order, as long as the system can accurately predict the holding point arrival times of aircraft that will arrive within the next eight minutes the vast majority of the benefit can be gained. There is little need for making longer term predictions, or for predicting the arrival of aircraft that have not yet left their stands.

However, the frozen parts of the discovered schedules often contained aircraft that were still on the taxiways. These schedules would not be possible if the take-off schedule was frozen considering only aircraft in the holding points at that time, as has to be done for conditional clearances. There is, therefore, a consequent danger of missing better schedules if too many conditional clearances are given.

Finally, we have seen that there is a correlation between the time at which the take-off order is frozen and the required planning horizon. If the schedule is frozen for longer, the planning horizon will need to be moved, requiring a longer term taxi time prediction.

ACKNOWLEDGMENT

This work was supported by EPSRC (The Engineering and Physical Sciences Research Council) and NATS (National Air Traffic Services) Ltd. from a grant awarded via the Smith Institute for Industrial Mathematics and Systems Engineering.

REFERENCES

[1] H. R. Idris, B. Delcaire, I. Anagnostakis, W. D. Hall, N. Pujet, E. Feron, R. J. Hansman, J.-P. Clarke, and A. Odoni, "Identification of flow constraint and control points in departure operations at airport systems," *AIAA Guidance, Navigation and Control conference, Boston, MA, August*, 1998.

[2] H. R. Idris, I. Anagnostakis, B. Delcaire, R. J. Hansman, J.-P. Clarke, E. Feron, and A. R. Odoni, "Observations of departure processes at logan airport to support the development of departure planning tools," *Air Traffic Control Quarterly*, vol. 7, pp. 229–257, 1999.

[3] I. Anagnostakis and J.-P. Clarke, "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 [24 February 2006].

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

[5] P. van Leeuwen, H. Hesselink, and J. Rohling, "Scheduling aircraft using constraint satisfaction," *Electronic Notes in Theoretical Computer Science*, no. 76, 2002.

[6] D. A. Trivizas, "Optimal scheduling with maximum position shift (mps) constraints: A runway scheduling application," *Journal of Navigation*, vol. 51, pp. 250–266, 1998.

[7] J. Abela, D. Abramson, M. Krishnamoothy, A. de Silva, and G. Mills, "Computing optimal schedules for landing aircraft," *The 12th National Conference of the Australian Society for Operations Research, Adelaide, July 7-9*, 1993, pp. 71–90, available at: http://www.csse.monash.edu.au/∼davida/papers/asorpaper.pdf [24 February 2006].

[8] J. E. Beasley, M. Krishnamoorthy, Y. M. Sharaiha, and D. Abramson, "Scheduling aircraft landings - the static case," *Transportation Science*, vol. 34, no. 2, pp. 180–197, May 2000.

[9] J. E. Beasley, J. Sonander, and P. Havelock, "Scheduling aircraft landings at london heathrow using a population heuristic," *Journal of the Operational Research Society*, vol. 52, pp. 483–493, 2001.

[10] A. T. Ernst, M. Krishnamoorthy, and R. H. Storer, "Heuristic and exact algorithms for scheduling aircraft landings," *Networks*, vol. 34, no. 3, pp. 229–241, 1999.

[11] L. Bianco, P. Dell'Olma, and S. Giordani, "Minimizing total completion time subject to release dates and sequence-dependent processing times," *Annals of Operations Research*, vol. 86, pp. 393–416, 1999.

[12] A. Craig, R. Ketzscer, R. A. Leese, S. D. Noble, K. Parrott, J. Preater, R. E. Wilson, and D. A. Wood, "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 [24 February 2006].

[13] J. A. D. Atkin, E. K. Burke, J. S. Greenwood, and D. Reeson, "A meta-heuristic approach to departure scheduling at London Heathrow airport," *The 9th International Conference on Computer-Aided Scheduling of Public Transport (CASPT), San Diego, USA, August 9-11*, 2004. Available at: http://fugazi.engr.arizona.edu/caspt/atkin.pdf

[14] J. A. D. Atkin, E. K. Burke, J. S. Greenwood, and D. Reeson, "Hybrid meta-heuristics to aid runway scheduling at London Heathrow airport," accepted for publication in *Transportation Science*, to appear 2006.

[15] F. Glover, "Future paths for integer programming and links to artificial intelligence," *Computers and Operations Research*, vol. 13, pp. 533–549, 1986.

[16] F. Glover and M. Laguna, *Tabu Search*. Kluwer Academic Publishers, 1998.