The TSAT Allocation System at London Heathrow: The Relationship Between Slot Compliance, Throughput and Equity

Jason A. D. Atkin[1], Edmund K. Burke[1], John S. Greenwood[2]

[1] School of Computer Science, University Of Nottingham, Jubilee Campus, Wollaton Road, Nottingham, NG8 1BB, England. E-mail: {jaa/ekb}@cs.nott.ac.uk

[2] NATS, NATS CTC, 4000 Parkway, Whiteley, Fareham, Hampshire, PO15 7FL

ABSTRACT

London Heathrow is an extremely popular airport, where efficient take-off sequencing is important for ensuring a high runway throughput. Aircraft are usually released from the stands as soon as possible, providing the maximum pool of aircraft at the runway from which to choose from when sequencing. This paper considers the task of predicting the delay that aircraft will experience, so that some of it can be absorbed at the stand before the engines are started, reducing the fuel burn, with consequent environmental and economic benefits. Doing this requires determining the value of a take-off sequence, in order to identify good sequences. This paper considers the trade-off between three mutually conflicting objectives – minimising delay, complying with network departure constraints, and treating aircraft equitably - and provides important insights for use in tuning the system to controller preferences. It also indicates potential benefits for controllers from the implementation of such a system.

1 INTRODUCTION

London Heathrow is one of the busiest airports in the world, but it normally has only one runway available for departures, with the other being used for arrivals. Demand exceeds capacity, so it is important that runway capacity is not wasted. The take-off sequence can have a large effect upon the runway throughput since the separations which are required between aircraft at take-off depend upon the sequence in which aircraft take off. Re-sequencing can significantly increase the throughput of the runway, thus appropriate sequencing is extremely important, however obtaining the best overall take-off sequence is far from simple. Minimum separation requirements apply not only to allow wake vortices to dissipate to safe levels, but also to control the frequency of departures along each departure route or in each direction. Noise considerations and the complexity of airspace over London give rise to complex rules for the required time separations for departure routes, so that separations depend upon both the departure routes and speed groups of the two aircraft involved in addition to their wake vortex categories (based upon their weight classes).

The departure system involves each aircraft pushing back from a stand (either direct onto a taxiway or into a cul-de-sac between the piers of a terminal), then starting its engines, taxiing to a holding area near the current departure runway and awaiting its turn to enter the runway and line up for take-off. The complex take-off sequencing task is currently performed by a runway controller once aircraft have reached a holding area adjacent to the current departure runway. The runway controller has to sequence the departures so as to attain as high a throughput as possible (and, consequently, as low a delay as possible) while ensuring that all constraints upon the take-off sequence are attained, including those imposed by capacity limitations in the airspace and at destination airports. The task of the runway controller at Heathrow was considered by Atkin *et al.* (2007, 2008a, 2008b) and Atkin (2008) where a design for a decision support system was proposed which would be able to help the runway controller, while working within the current operating environment.

In order to ensure maximum flexibility for the runway controller (by maximising the number of aircraft at the holding area from which to choose), current operating procedures mean that departing aircraft are usually released from stands as soon as possible. Any take-off delay for an aircraft therefore occurs while it is queued at the holding area adjacent to the runway, with its engines running. Where it is possible to do so without having a detrimental effect upon the runway sequencing, it would be preferable to absorb as much of this delay as is practical at the stand, prior to starting the engines. This would lead to a reduction in the amount of fuel burnt, with both economic and environmental benefits. Since the load on different departure routes is not usually even, so some aircraft will necessarily have longer delays than others, adding a global hold to aircraft risks starving the runway of the types of aircraft which could actually take off at that time, and thus reducing runway throughput. In order to determine how much delay it is appropriate to absorb at the stand, it is important to be able to predict the delay that aircraft would otherwise have at the holding area.

This paper discusses a system to allocate TSATs (Target Start-At Times) to aircraft at Heathrow, specifying a target time at which an aircraft should commence its push back and engine start-up operation. This system includes an algorithm to predict take-off times for aircraft while they are still at the stand, and to then use these predictions to allocate appropriate stand holds to aircraft.

The objectives of the take-off sequencing that controllers perform are actually more than merely to minimise delay, since it is also important to consider the equity of delay between aircraft and avoid penalising aircraft with specific characteristics. It is, furthermore, important to meet the take-off time-slots which are allocated to around 30% to 40% of the Heathrow departures by the the Central Flow Management Unit of Eurocontrol in order to control downstream congestion. These three objectives (delay, equity and time-slot compliance) are mutually interacting. Sometimes constructively to guide a search towards good overall schedules, but this paper will show that they usually in conflict for the best schedules. In order to predict the take-off schedule that will be adopted, a three-part objective function is used to evaluate each schedule, where each part evaluates the compliance with one of the three objectives. In order to produce a more useful and accurate take-off time prediction system, and generate an objective function which will more accurately model the preferences of controllers, it is important to understand the way in which these objectives actually interact. Investigation of the interaction between these three objectives is the focus of this paper. In particular, this paper considers the effect of the weights used in the objective function and the ways in which the three objectives interact when used in the TSAT (i.e. stand-hold) allocation system described in this paper.

2 TSAT ALLOCATION AND TAKE-OFF SEQUENCING

The research described in this paper considers the TSAT (Target Start-At Time) allocation project, which is one element of the CDM (Collaborative Decision Making) project (Eurocontrol Experimental Centre, 2005) now being implemented at Heathrow. TSAT allocation involves determining a time for each aircraft at which its engine start-up should take place, so that it can push back from the stand, taxi to the holding area near the runway and arrive not long before it should take off. The risk in changing the operations in this way is that a bad allocation of stand hold time could prevent the controller from having available the aircraft that are necessary to achieve the better take-off sequences, thus wasting runway capacity. The system described in this paper, that is currently being implemented by NATS for use at Heathrow, avoids this problem by ensuring, as far as possible, that the best take-off sequence it can find will be achievable, while also making it likely that good alternative, schedules will also be available if necessary.

Not only are there obvious benefits from waiting at the stand, in terms of the fuel used, but there are also benefits in terms of increasing the visibility for airlines of when an aircraft is expected to be released from the stand, potentially facilitating better allocation strategies for resources needed during turnround to clean and re-fuel the aircraft, to load passengers, and the tugs which push aircraft back from stand. Finally, the strategic allocation of stand holds should reduce the congestion at the holding areas, since the delay for aircraft which need exceptionally long holds is usually predictable. Congestion at the holding area can limit the amount of re-sequencing that can be performed, as described by Atkin *et al.* (2006a), and mean that the best sequences become unavailable.

There are two stages to TSAT allocation. The first is to predict a take-off time for aircraft, and the second is to use this take-off time to determine an ideal pushback time. The ideal pushback times are then used in conjunction with a consideration of the contention between the aircraft at the stands to determine TSATs. Take-off sequence prediction utilises the fact that the controllers at Heathrow are highly skilled and will find good take-off sequences providing that the correct aircraft are available to them. Thus, sequence prediction involves assuming that the best sequence which is found by the algorithm is the one which the controllers will adopt, however this obviously assumes that a measure of the value of a take-off sequence (in terms of the desirability to a runway controller) is available. Providing a better understanding of the relative effects of the three main objectives, to obtain the necessary insight to tune an objective function, is the focus of this paper.

The take-off sequencing stage is of primary interest in this paper, but an understanding of stand contention at Heathrow is also required. The contention rules at the stands are modelled as minimum separation requirements between pushback times for aircraft. Runway separations and pushback time separations are both sequence-dependent, asymmetric and do not obey the triangle inequality, so it is not sufficient just to ensure compliance with separations for aircraft which are in consecutive positions in a push-back or take-off sequences. The pushback separations are obviously important for the TSAT allocation stage, to ensure that allocated TSATs are achievable, but also have to be considered in the take-off sequencing stage, since contention around the stands can delay the earliest take-off times for aircraft. Solution of the take-off sequencing problem therefore requires a simultaneous consideration of two sequence-dependent-separation sequencing problems (at the stands and the runway), which are linked together by the taxi times of the aircraft.

One of the methods that is utilised to handle downstream congestion is the Calculated Time Of Take-off (CTOT), assigned by the Eurocontrol Central Flow Management Unit. Aircraft with CTOTs have a fifteen minute window within which take-off must occur. CTOTs ensure that there is a limit upon the number of aircraft entering any busy airspace sector or busy destination airport at the same time. Complying with these CTOTs is important; however, in order to reduce the number which must be re-negotiated due to slight delays, a limited number of five-minute extensions are available to the controllers. Although these extensions are permitted, they should be avoided when possible, and limiting the number of extensions which are used is one of the objectives of take-off sequencing.

The take-off sequencing problem is multi-objective. The first objective is, obviously, to keep delay down for aircraft, the second is to comply with CTOT windows, and the third is that a schedule must not excessively penalise individual aircraft. Despite the multi-objective nature of the problem, however, a single objective solution approach has to be used since only a single TSAT should be presented for each aircraft rather than a set of solutions from which a controller should select. The take-off sequence prediction element is merely an internal algorithm, and the intention is only to reveal the TSAT, not the take-off time. Forcing controllers to select from a set of internal take-off sequences associated with the TSATs is not appropriate.

By necessity, the objective function has to consider all three objectives. The aim of the research described in this paper is to evaluate the effect upon the TSAT allocation system of varying the weights in the objective function that are used to determine the balance between take-off time-slot compliance, equity of delay and overall delay. For example, by reducing the importance of take-off time-slots, more extensions will be required, but the delay may potentially be improved by enabling sequences with a better throughput to be adopted. It may, thus, be better to take advantage of the permitted CTOT extensions rather than always trying to minimise the number used. It is also important to understand the way in which the overall delay is penalised by potential measures to ensure a degree of equity of delay between aircraft. In order to make an informed decision about the appropriate objective function for the real world problem, such an evaluation is vital. It is imperative that, at the very least, any increase in delay to account for CTOT compliance or limit any inequity between aircraft has been justified.

Previous departure sequencing research at the airport by Atkin *et al.* (2007,2008b) and Atkin (2008) utilised an objective function where avoiding CTOT extensions was given an extremely high weight. Comparison with controller schedules showed that the selected objectives of the system were more concerned with avoiding the use of CTOT extensions than a human controller was. For example, there were cases where the overall delay for aircraft was (perhaps unnecessarily) increased in order to meet a tight CTOT. Nevertheless, significant delay benefits were still predicted from the use of a decision support system even with this extreme focus upon meeting CTOTs. The results in this paper aim to give an idea of the additional benefits which could be obtained by increasing the emphasis on delay reduction and allowing more extensions to be used. The aim is to provide the insights necessary to better determine the objective function to use in practice to predict trade-off sequences that the controller would select.

Research into arrival and departure sequencing elsewhere (which are both instances of sequence-dependent separation problems) does not usually explicitly consider the trade-off between equity and delay. However, it is important for the inequity in aircraft delays to be managed, and there will always be a trade-off to be made between the total delay for aircraft and the equity of the division of that delay between the aircraft. It is common for considerations of equity to be handled by applying a limitation upon the movement of any specific aircraft (for instance using a maximum position shift approach (Trivizas, 1998). Furthermore, the concept of a time-slot which should be attained where possible, but which has extensions which may be utilised if necessary, is not often observed in the academic models. Time-slots are usually hard constraints where they are implemented at all, for example by the use of hard time-windows around the landing/take-off time for each aircraft (Beasley *et al.*, 2000, 2004, Ernst *et al.*, 1999). Similarly, the trade-off between equity of delay and total delay is not usually explicitly examined, although Beasley *et al.* (2001) did recognise the need for some degree of equity by using a non-linear objective function. In this research, we wish to perform the hard task of modelling the real-world operations, and thus we need to understand the trade-off that is actually taking place in terms of the increased delay for various levels of commitment to avoiding the use of CTOT extensions and for ensuring a more equitable allocation of delay.

3 PROBLEM MODEL AND DESCRIPTION

3.1 Definitions

The following definitions for constants and variables are used throughout the model described in this paper:

- pt_j — Allocated pushback time for aircraft *j*. This will be assigned as the TSAT.
- ept_j — Earliest pushback time for aircraft *j*, an input to the algorithm.
- pd_j — Predicted pushback duration for aircraft *j*. The number of seconds that the system should assume that *j* will take to push back.
- ct_i, ct_j — Allocated cul-de-sac time for aircraft *i* or *j* (respectively). The time at which *i* or *j* will be ready to commence the taxi to the runway, without being delayed by other pushbacks.
- td_j — Predicted taxi duration for aircraft *j*. The number of seconds that the aircraft is expected to (or did) take to taxi from the position at which it was when it started its engines to the holding area by the current departure runway. This specifically excludes

any runway hold delay, and any time spent pushing back and starting engines since these are covered elsewhere in the model.

bt_j A base-line time for aircraft *j* from which to measure delay. This provides a method to measure and penalise delay. In this paper this is set to the earliest take-off time for *j* if it was in isolation in the departure system (i.e. no queuing for take-off and no cul-de-sac delay).

ec_j Earliest take-off time for aircraft *j* which will meet any requirements for the allocated CTOT. This is set to an extremely early time if the aircraft has no CTOT so that no constraint is implied.

lc_j Latest take-off time for aircraft *j* which will meet any requirements for the allocated CTOT. This is set to an extremely late time if the aircraft has no CTOT so that no constraint is implied.

d_i, d_j Predicted or actual take-off time for aircraft *i* or *j* respectively.

cs_i, cs_j Position of aircraft *i* or *j* (respectively) in the cul-de-sac sequence – the sequence in which aircraft commence their taxi to the runway, having pushed back and completed their engine start-up. 0 for the first aircraft to set off. If $cs_i < cs_j$ then *i* commences its taxi operation prior to *j* doing so.

as_i,as_j The position of aircraft *i* or *j* (respectively) in the first-come-first-served take-off sequence. This is taken to be the sequence implied by sorting into ascending order of earliest take-off times of all aircraft, assuming no delays from contention with any other aircraft in the system. i.e. sorting the aircraft in ascending order of bt_j.

ts_i, ts_j The position of aircraft *i* or *j* (respectively) in the take-off sequence. 0 for the first take-off. If $ts_i < ts_j$ then *i* takes off before *j* does so.

RS_{ij} Minimum runway separation. The minimum number of seconds which must elapse between the take-off times of aircraft *i* and *j*, when *i* takes off before *j*. This value actually depends upon the time of day and the allocated runway as well as the characteristics of the pair of aircraft involved (i.e. should be a function of d_i as well as *i* and *j*), but can be treated as a constant for each pair of aircraft in the experiments described here.

MS_{ij} Minimum runway separation. The minimum number of seconds which must elapse between the take-off times of aircraft *i* and *j*, when *i* takes off before *j*. This value depends upon the time of day and the allocated runway but can be treated as a constant for each pair of aircraft in the experiments described here.

IRH_j Ideal runway hold for aircraft *j*. Can depend upon the predicted take-off time, but is treated as a constant (with value 300 seconds) for all aircraft in the experiments described here.

MRH_j Minimum runway hold for aircraft *j*. Can depend upon the predicted take-off time, but is treated as a constant (with value 60 seconds) for all aircraft in the experiments described here.

3.2 Objectives and Constraints

The model for the two stages of the TSAT allocation problem is described in sections 3.3 and 3.4 and can be expressed by the following equations, inequalities and formulae:

$$\mathrm{W}_1\, C(lc_j, d_j) + \mathrm{W}_2\, D(d_j, bt_j) + \mathrm{W}_3\, E(ts_j, as_j) \quad (1)$$

$$ct_j \geq ept_j + pd_j \quad (2)$$

$$ct_j \geq ct_i + MS_{ij} \quad \forall\, i \quad \text{s.t} \quad cs_i < cs_j \quad (3)$$

$$d_j \geq ct_j + td_j + MRH_j \quad (4)$$

$$d_j \geq ec_j \quad (5)$$

$$d_j \geq d_i + RS_{ij} \quad \forall\, i \quad \text{s.t} \quad ts_i < ts_j \quad (6)$$

$$pt_j = ct_j - pd_j \quad (7)$$

$$ict_j = \max(\, ept_j + pd_j,\, d_j - IRH_j - td_j\,) \quad (8)$$

$$\sum_j (\, 100 \max(\, 0,\, ct_j - ict_j\,)^{1.1} + \max(\, 0,\, ict_j - ct_j\,)^{1.1}\,) \quad (9)$$

TSAT allocation has two stages, which are described below. In the first stage involves predicting the take-off sequence, in order to predict take-off times and total delay for each aircraft. The second stage involves allocating achievable push-back times to aircraft, considering the cul-de-sac contention, so that an appropriate amount of the total delay is absorbed as stand hold.

3.3 Take-off sequence prediction

Take-off sequence prediction involves an assumption that the controllers will adopt the best take-off sequence, then attempting to find this sequence. This involves developing an evaluation function for a take-off sequence. The contribution of a single aircraft *j* to the overall cost can be expressed by Formula 1. Formula 1 has three components. The first component (the function $C(lc_j, d_j)$) penalizes schedules which fail to meet take-off time-slots. Function $C(lc_j, d_j)$ returns a cost of zero if *j* is within the allocated CTOT time-slot (i.e. $d_j \leq lc_j$), a penalty cost of $(500 + d_j - lc_j)$ if *j* is within a valid (five minute) CTOT extension (i.e. $lc_j \leq d_j \leq lc_j+300$) or a large penalty cost of $(50{,}000 + 10*(d_j - lc_j))$ if *j* is scheduled to take off too late for even an extension.

The second component of Formula 1 (the function $D(d_j, bt_j)$) penalises schedules according to the delay experienced by aircraft. Function $D(d_j, bt_j)$ returns a cost equal to $(d_j - bt_j)^{\alpha}$ where the value of 1.0 was usually used for α, but the effects of using α=1.5 and α=2.0 are considered in Section 5. The third component (the function $E(ts_j, as_j)$) explicitly penalises positional movement of an aircraft within the sequence, favouring the first-come-first-served sequence. The function $E(ts_j, as_j)$ was defined to return $(ts_j - as_j)^2$ in the experiments performed for this paper.

The aim of this paper is primarily to investigate the interaction of the weights W_1, W_2 and W_3 which alter the relative effects of the different components of the objective function. In particular, the aim is to provide the insights which are necessary in order to be able to select appropriate values for the weights and to understand the effects of the interaction between the objectives.

Given the contribution to the cost for each individual aircraft, the overall cost of a take-off sequence can be measured as the sum of the costs for the individual aircraft in the schedule. Take-off sequence prediction then becomes a case of finding a take-off sequence for which the objective function, expressed by Formula 1, is minimised when summed over all aircraft *j* in the take-off sequence. This involves knowing for each aircraft the position in the take-off sequence and a predicted take-off time (which may involve knowing a feasible cul-de-sac sequence).

Once a take-off sequence is known, it is possible to determine take-off times by assuming that each aircraft will take off as early as it can, given the various constraints upon the take-off times. To determine the earliest take-off time for an aircraft while it is still at the stands it is necessary to determine the earliest time at which the aircraft can complete its pushback. This earliest pushback time is limited by the contention with other aircraft which push back from stands which are close by, or by stands further away but on the same cul-de-sac. The cul-de-sac time for an aircraft is here defined to be the time at which it has completed its pushback and can freely commence its taxi towards the holding area, without being delayed by other aircraft which are also pushing back. Airlines will declare a TOBT (Target Off-Block Time) for each aircraft, specifying the time at which they believe the aircraft will be ready to push back. The TSAT allocation algorithm should use this as an earliest pushback time (ept_j for aircraft j), and seek to allocate a TSAT which is no earlier than the TOBT. This constraint upon the TSAT can be treated as a constraint upon the cul-de-sac time by adding on the pushback duration, as expressed by Inequality 2.

It is possible to model the contention rules within cul-de-sacs using the cul-de-sac times of the aircraft. Cul-de-sac contention is of two forms. Firstly, aircraft that are on a cul-de-sac and near to the taxiways can block aircraft which push back further from the taxiways from leaving the cul-de-sac until they have themselves moved. In this case, the cul-de-sac time of the second aircraft must be after the cul-de-sac time of the first aircraft and there may be a minimum separation time that must be applied to maintain a gap between the aircraft. Secondly, in some cases an aircraft cannot even commence its pushback until another aircraft has left the cul-de-sac. This often occurs when aircraft are on stands which are close together. In this case a minimum separation equal to the pushback time of the second aircraft plus any time delay between the first aircraft leaving the cul-de-sac and the second being able to commence its pushback must be enforced between the cul-de-sac times which are applied to aircraft. Inequality 3 represents the effect upon the cul-de-sac time of delays due to the minimum cul-de-sac time separations.

Since the constraints upon the earliest pushback time can be modelled as constraints upon the cul-de-sac time (Inequality 2), and cul-de-sac contention can be modelled as sequence-dependent minimum separations between the cul-de-sac times for aircraft (Inequality 3), the take-off sequencing element of the problem has to consider only cul-de-sac times not pushback times. Once a cul-de-sac time is known for an aircraft, it becomes possible to consider the constraints upon the take-off time. Inequality (4) expresses the fact that an aircraft cannot take off before it can reach the runway, where the minimum runway hold (MRH_j for aircraft j) represents the minimum time that an aircraft will be ensured to have available to traverse the holding area and line up for take-off. This can vary between aircraft, but a value of one minute was used for all aircraft in the experiments described in this paper.

The second constraint upon the take-off time is that aircraft with CTOTs should not take off before the start of the fifteen minute slot implied by the allocated CTOTs. This constraint can be expressed by Inequality 5.

Finally, the minimum runway separations (defined in this paper as RS_{ij} for any ordered pair of aircraft i and j, when i takes off before j) must be obeyed for aircraft. These separation values depend upon the departure routes, speed groups and weight classes of the two aircraft and can vary depending upon the times at which the aircraft take off and the take-off runway that is used. Although the developed algorithm handles this variation, the datasets used for the tests are for a single runway and cover time periods for which the separation rules were fixed. For the experimental results presented in this paper, the value of RS_{ij} can be treated as a constant (independent of d_i) for any specific pair of aircraft i and j and the constraint upon the take-off time can thus be expressed by Inequality 6.

Given the above described constraints upon the take-off times, and the fact that the earliest take-off time for an aircraft depends upon both the take-off sequence (from Inequality 6) and the cul-de-sac sequence (from Inequalities 3 and 4) the aim of the take-off sequence prediction algorithm can be considered to be to determine the take-off sequence to use, and potentially to also have to consider the cul-de-sac sequencing that will be necessary.

3.4 The TSAT allocation stage

Once predicted take-off times are known for aircraft, an ideal cul-de-sac time can be determined by assuming that any delay beyond an ideal maximum runway hold value is absorbed as stand hold. The ideal runway hold (IRH_j for aircraft j) was assumed to be five minutes for all aircraft in the experiments described in this paper. The ideal cul-de-sac time (ict_j for aircraft j) can then be determined using Equation 8.

Given ideal cul-de-sac times for each aircraft, the cul-de-sac times to allocate are determined such that the cost of deviation between the ideal and allocated cul-de-sac times as expressed by Formula 9 is minimised, subject to the constraints upon the cul-de-sac times that are implied by Inequalities 2, 3 and 4. Formula 9 aims to find cul-de-sac times close to the ideal times, with a preference for pushing back early (increasing the delay at the runway hold, and the slack to allow for late arrivals at the holding area) rather than late, and with a non-linear penalty for deviations in order to prefer a more equitable allocation of any necessary deviations from ideal times.

Once a cul-de-sac time is known for each aircraft, a push-back time may be calculated using Equation 7, by assuming that it is always better to delay an aircraft at the stand prior to pushback than within the cul-de-sac once the engines have been started. The allocated pushback time is then reported as the TSAT to allocate, specifying the time at which the aircraft should commence its pushback prior to starting its engines.

4 SOLUTION APPROACH

The two problems are solved using different solution algorithms. The take-off sequencing stage is by far the harder of the two problems, and takes up the majority of the solution time, since it requires concurrent solution of both the take-off sequencing and cul-de-sac sequencing problems. The solution approaches for both stages are described below.

4.1 Solution method for the take-off sequencing problem

One major problem for the take-off sequencing problem is that individual aircraft may need to move a long way from the first-come-first-served sequence. The two main reasons that some aircraft may need to be delayed or prioritised in the departure sequence are in order to meet a take-off time-slot due to an allocated CTOT, or (to a lesser extent) due to the differing demands upon different departure routes, as discussed in Section 5. Sometimes the delay for CTOT compliance (or potentially, but less often, the advancement) can be a significant distance in the sequence, as discussed Section 5. These characteristics of the best schedules prevent the successful application of solution approaches which simplify the sequencing problem by limiting the maximum number of positions that a single aircraft can be moved into such as those proposed by Dear and Sherif (1989,1991) or Trivizas (1998). One of the aims of this paper is to provide greater insight into the trade-offs between the delay, CTOT compliance and positional shifts in the schedules. They also have to be overcome by the selected approach.

The basis of the solution approach for the take-off sequencing element is the application of a rolling window across the take-off sequence and the optimal sequencing of the aircraft within the window. Obviously, the application of a window applies a limitation based upon the window size to the positional movement of the aircraft in the schedule. The developed solution approach copes with the problem of aircraft having to move a long way in the schedule in three ways. Firstly, the window is rolled forwards from the first take-off to the last, with an overlap of all but one aircraft. Subsequent positions of the window can, therefore, delay aircraft further in the sequence when required. Secondly, multiple passes of the algorithm are applied, where each pass involves starting the window at the first aircraft in the sequence which resulted from the previous pass and rolling the window to the end. Subsequent passes can, therefore, be used to advance aircraft earlier in the take-off sequence in cases where they need to be advanced by more places than would be possible in a single window. Finally, a heuristic generation method is used to create an initial sequence which allows for potential advancements and delays to account for CTOTs.

The initial take-off sequence, which will be improved using the rolling window approach, is produced by generating an estimated take-off time for each aircraft, assuming a five minute delay beyond the earliest take-off time for the aircraft. If the estimation is later than the end of the take-off time-slot for the aircraft then the end of the take-off time-slot is instead, unless this is unachievable, in which case the earliest take-off time for the aircraft is used, thus advancing aircraft which have tight CTOTs. If the estimation is earlier than the earliest take-off time for the allocated CTOT, then the earliest time for the CTOT is used instead, delaying aircraft which will need to delay due to a CTOT. These estimated take-off times are then used to build an initial take-off sequence by ordering aircraft by non-decreasing estimated take-off time.

4.1.1 The rolling window branch-and-bound algorithm

The main take-off sequencing algorithm progressively improves the existing take-off sequence. A window size of nine aircraft was used for the experimental results presented in this paper, although the window size is a parameter of the selected algorithm. The algorithm optimally sequences the nine aircraft within the current window, ignoring any later take-offs but considering the take-off sequence and predicted take-off times for aircraft which take off earlier than the current window to be inviolate. The algorithm starts at the first nine aircraft, sequences these then fixes the position and take-off time of the first aircraft in the window. The window is then moved forward by one aircraft and the second to tenth aircraft are sequenced, after which the first aircraft in the window is again fixed (as the second aircraft in the overall sequence). This process continues, advancing the window by one aircraft at a time. When the last window position is reached the take-off sequence and times are adopted for all aircraft in the window and, an overall take-off sequence will have been produced. Four passes of the algorithm were utilised for the results presented in this paper, in order to attempt to improve the sequence further on each pass. Once the final pass has been completed, the resulting take-off sequence is passed to the TSAT allocation system, along with the predicted take-off times, so that TSATs can be determined for each aircraft.

Optimising the sequence of aircraft within the window involves determining a take-off time for each aircraft. Due to the separation rules, take-off times can only be predicted if the take-off times of earlier take-offs are known. This is the reason that aircraft which take-off prior to the current window have to be considered and why the window has to roll from first to last aircraft in the sequence. A branch-and-bound algorithm is used to optimally sequence the aircraft within the current window. All of the aircraft are taken out of the sequence then added back in one at a time, in potential take-off order, so that, at each stage, the take-off times for all previous aircraft are known. As each aircraft is added, the take-off time is predicted and the cost of the new aircraft in the sequence is determined using Formula 1. Lower bounds are then determined for the take-off times for the remaining aircraft. Using these take-off time lower bounds, lower bounds for the three components of the objective function are determined independently by ordering the remaining aircraft optimally for each objective in turn (for instance into first-come-first-served order for the equity objective), assigning the lower bound take-off times to the aircraft and determining the consequent cost. Summing the three components of the cost gives a lower bound for the cost of all sequences which start with the current sub-sequence, allowing pruning based on the sub-sequence if the lower bound of the cost exceeds the cost of a known full take-off sequence. At the end of this process, the take-off sequence with the lowest cost is adopted.

4.1.2 Take-off time prediction

As discussed earlier, take-off time prediction involves knowing a cul-de-sac sequence as well as a take-off sequence. However, the number of different cul-de-sacs at Heathrow is large enough that aircraft which are close together in the take-off sequence are often allocated to stands which are sufficiently far apart that cul-de-sac contention is not a problem. It only becomes a problem when the delay for an aircraft (which can be considered as the slack between the earliest take-off time and actual take-off time) is insufficient to absorb the additional delay at the stand which results from the cul-de-sac contention. Unfortunately, cul-de-sac contention is a problem often enough that it cannot be ignored. Experimentation revealed that there were cases where the best schedule found (i.e. that which would have otherwise been adopted) when cul-de-sac contention was ignored could not actually be achieved when cul-de-sac contention was considered. i.e. on occasions, the delays at the cul-de-sac were such that the predicted take-off times could no longer be achieved once cul-de-sac contention was considered, leaving an infeasible problem to be solved at the second (TSAT allocation) stage.

Although cul-de-sac contention cannot be ignored, the developed solution approach takes advantage of the relative infrequency of cul-de-sac contention delaying take-off times by first checking whether the cul-de-sac sequencing problem actually has to be solved before attempting to solve it. A potential cul-de-sac time is determined for each aircraft as it is added to the take-off sequence, chosen to be the earliest time which will meet Inequalities 2 and 3, making the initial assumption that all aircraft which take-off before the new aircraft also leave the cul-de-sacs earlier (i.e. are earlier in the cul-de-sac sequence). The earliest take-off time which will meet Inequalities 5 and 6 is then determined and compared against the earliest time which will meet Inequality 4. If Inequality 4 gives a time which is no later than that from Inequalities 5 and 6 (i.e. Inequalities 5 or 6 are the binding constraints upon the take-off time) then the potential cul-de-sac time is adopted and recorded, along with the take-off time which was determined from Inequalities 5 and 6. The algorithm then progresses to calculating the cost for the aircraft and adding the next aircraft to the sequence without attempting to optimise the cul-de-sac times. However, if Inequality 4 gives a later take-off time (so that Inequality 4 is the binding constraint), then an attempt is made to find an earlier cul-de-sac time. To do this, a lower bound for the cul-de-sac time is determined, by considering that it must leave the cul-de-sac after an aircraft which takes off earlier if not doing so would delay the take-off time of the earlier aircraft. If this lower bound is achieved by the current potential cul-de-sac time then the cul-de-sac time and (consequent) take-off time are both adopted without further work. If not then the cul-de-sac sequencing problem has to be solved. This involves finding the cul-de-sac sequence which will minimise the take-off time of this aircraft without delaying the take-off times of earlier take-offs (the potential cul-de-sac times for earlier aircraft can be modified, however). This problem is a simplified version of the problem considered in Section 4.2 since the same constraints apply, however the search can cease as soon as any cul-de-sac sequence is found which obtains the earliest take-off time which is no later than the earliest time which will meet Inequalities 5 and 6. Experimental results showed that the cul-de-sac sequencing sub-problem had to be solved very infrequently and that in the majority of cases the problems could be solved without enforcing a delay upon the take-off time (i.e. re-sequencing enabled the earliest time from Inequalities 5 and 6 to be achieved). Furthermore, the problem sizes that needed to be considered were usually very small, since only aircraft with overlapping time windows and required cul-de-sac separations need to be considered in the problem, as discussed in Section 4.2.

4.2 Solution method for the TSAT allocation problem

Once a take-off sequence has been determined, the take-off times and potential cul-de-sac times are passed to the TSAT allocation algorithm. This algorithm attempts to minimise the value of Formula 9, subject to the constraints upon the cul-de-sac time implied by Inequalities 2, 3 and 4. Allocated TSATs must be on minute boundaries, so there are usually very few options for the potential cul-de-sac times for each aircraft. Furthermore, the full problem can be decomposed into a number of relatively small problems (up to nine aircraft in the worst case on the test datasets), each consisting of those aircraft which have cul-de-sac contention with each other and which have time-slots for the cul-de-sac times which are close enough together to be able to influence each other.

Given a sub-problem consisting of aircraft which are in contention at the stands, the cul-de-sac sequencing problem is again solved using a branch-and-bound algorithm. A rolling window approach is again utilised if the problem size grows beyond a given limit. In that case an initial sequence is built using the predicted potential (and known to be feasible) cul-de-sac times from the take-off sequencing stage and passing a window over the sequence, optimising the sequence of the aircraft within the window while assuming that the sequence of the other aircraft is fixed.

5 RESULTS

Experiments were performed using ten datasets containing historic data covering 110 aircraft in each dataset. 110 aircraft was sufficient to cover at least two hours of take-offs, thus providing an opportunity to observe the (relatively) long term effects of earlier sequencing. The datasets contained information about the historic pushback times, taxi durations and take-off times, as well as the necessary details to determine any cul-de-sac contention and take-off separations.

The experimental results which are presented in this paper evaluate schedules according to three values, one of which measures each of the objectives: CTOT compliance, delay and equity. Firstly, the number of CTOTs which were missed (i.e. for which extensions were required or where even extensions were missed) is reported. Secondly, the total delay, summed across all 110 aircraft, is reported. Although all times are maintained in seconds, the results are presented in minutes to simplify the values. Delay is measured as the difference between the predicted take-off time and earliest time at which the aircraft could have reached the holding area, in the absence of any stand hold or contention with other aircraft, i.e. (d_j - ec_j) for each aircraft j. The third value which is reported is the sum of the squares of the positional movement of each aircraft in the take-off sequence (i.e. $(as_j - tsc_j)^2$ for each aircraft j). This is used as an approximation of the inequity in the take-off sequence since it provides a weighted measure of the deviation between the take-off sequence and first-come-first-served sequence. The shortcomings of this measure for this problem are discussed below.

Experiments were performed by varying the different weights W_1, W_2 or W_3. Values of W_1=1, W_2=100 and W_3=100 were found to give balanced weights to each of the objectives, such that increasing one of these weights while the other two were fixed had a perceptible effect upon the schedules which were produced. Further experiments were then performed to compare the pair-wise effects of the three objectives. W_1 was varied from 0 to 99 in steps of 1, and W_2 and W_3 were varied from 0 to 9900 in steps of 100 due to their proportionately lower effect upon the sequencing.

The cost for inequity was never set to zero in these experiments (i.e. $W_3 \geq 1$). This was extremely useful for removing symmetries in the problem. With 110 aircraft in a dataset, there will be many aircraft with the same weight class, speed group and departure route as other aircraft. When two aircraft leave the stands at similar times and a linear cost for delay is applied, reversing the take-off sequence for these two aircraft will result in an identically costing sequence, assuming that both aircraft can achieve both take-off times. There is nothing within the branch and bound method to take advantage of this symmetry so the solution space is greatly increased and solution pruning can be greatly impaired. By ensuring that a slight penalty for inequity is applied, these otherwise identical schedules have different costs and the solution speed was noticeably increased. Unless otherwise stated, removing consideration of one element of the objective function meant setting W_1 to 0, W_2 to 0 or W_3 to 1. Comparison of results with and without the equity cost (W_3=0 or W_3=1) showed that its effect was small enough to be overwhelmed by values of W_1=1 or W_2=100, and to have no effect upon the delay and CTOT compliance results in these experiments. However, the breaking of the symmetries in the problem reduced the solution time considerably and ensured that consistent and comparable schedules were produced. Without this symmetry breaking it was extremely difficult to determine whether the schedules produced by two sets of weights were actually equivalent. Unfortunately, despite the similarity between the characteristics of some aircraft, there are too many different combinations of aircraft characteristics to allow a grouping approach such as that proposed by Psaraftis (1980) to be useful.

Experiments were performed to investigate the combinations of weights for different objectives. Since it is only the relative values of the weights which matter, selected results are presented graphically for each pair of weights in Figures 1 to 5, showing the effect of increasing first one then the other weight rather than reporting full results for each pair of potential weights. This is equivalent to showing the first row and first column of a table cross-referencing the values of the weights. The centre-point of each graph, labelled ‘balanced costs' is where the weights are of the relatively balanced values (W_1=1,W_2=100 or W_3=100). Moving to the left of this point, one of the weights will gradually gain value, while moving to the right the other will gain value. The far left-most point has zero value for the one of the weights, and the far right-most point has zero value for the other weight. The x-axis of each of the graph can be considered to measure the ratios of the weights, starting with only a weight for the first objective, then having a ratio of 99:1, through a set of 98 further results with the ratio dropping gradually to 1:1, then rising gradually to 1:99 over the next 98 results, before the final point has only a weight for the second objective.

5.1 Delay reduction vs CTOT compliance

The effects of varying the weights for the CTOT compliance (W_1) and delay (W_2) can be observed in Figure 1. This is perhaps the most important trade-off to be aware of in these results, since it was the emphasis upon CTOT compliance that was previously identified as being a potential deviation between the selected objective function and the real controller behaviour. Two sets of points are plotted on the graph as lines. The top line on the graph shows the number of CTOTs which were missed and the bottom line the total delay. On the far left of the graph in Figure 1 is the solution when no CTOT miss cost is applied at all (W_1=0). The next point has a large delay cost and a moderate CTOT cost (W_1=1,W_2=9900), then the delay cost is gradually reduced in steps of 100 until the centre point (labelled balanced costs, where W_1=1,W_2=100) is reached. From the centre point, the CTOT cost was gradually increased, leaving the delay cost at a moderate value (W_2=100, W_1 increased in steps of 1), so that the penultimate point on the right has a value of W_1=99, W_2=100). The point on the far right had no cost for delay (W_1=1,W_2=0).

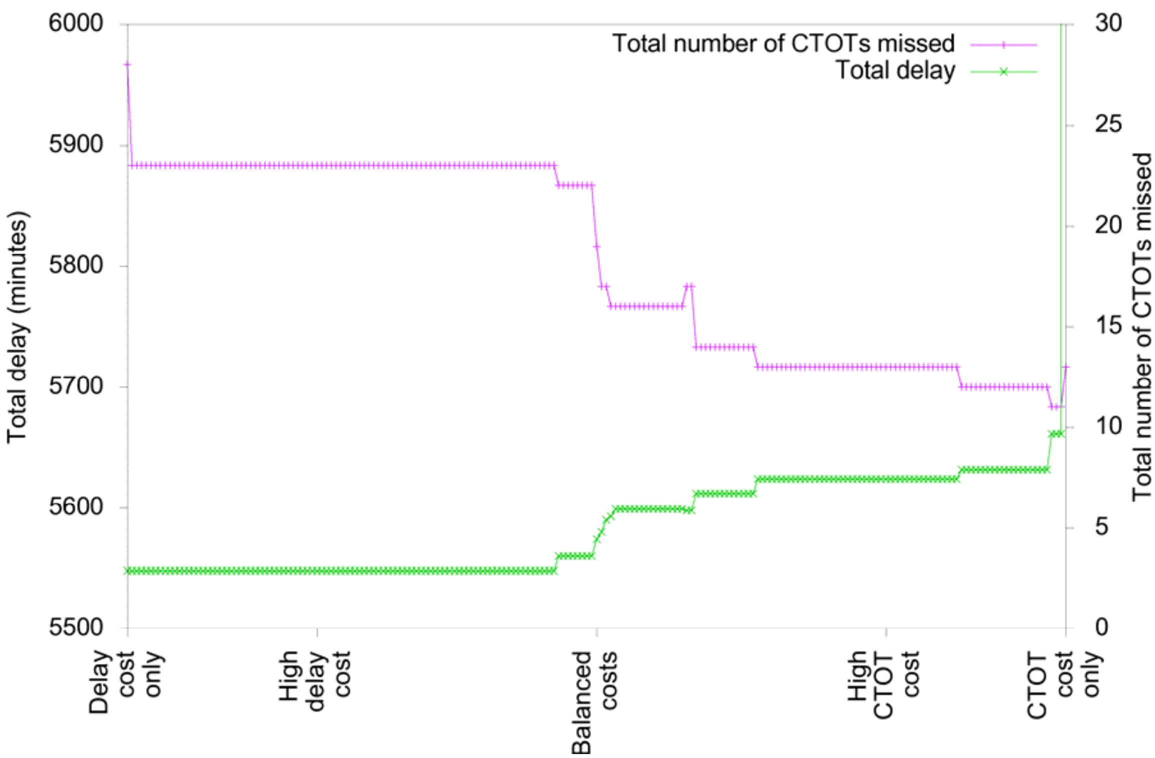

Figure 1: Graph of total delay and number of CTOTs missed as W_1 and W_2 are varied.

Figure 1 provides comparative results of the total delay for all aircraft (in minutes) and the total number of CTOTs which were missed. Since there are 110 aircraft per dataset and ten datasets, a total delay value of 1100 represents a mean delay of one minute per aircraft. The mean delay per aircraft for the results in Figure 1 is thus between 5 minutes 2 seconds and 5 minutes 9 seconds.

The results in Figure 1 show an obvious trade-off between these two objectives in the best solutions found, so improvements for one objective correspond to losses for the other. They also show that there are only a limited number of different solutions which are adopted, as can be seen from the many horizontal regions of the graph, determined by the limited number of possible CTOTs that can be missed. Finally, the extreme points (where one or other of the weights is zero) do not follow the trend of the graph. For example, the far right-most point (where W_1=1,W_2=0) has a higher number of CTOTs missed than the point immediately to the left of it (W_1=99,W_2=100).

In order to gain further insight into the trade-off the values for the three objectives at the extreme points are shown in Table 1 for each of the ten datasets. The column labelled "D.S." specifies the dataset number. The columns labelled "#C" specify the number of CTOTs which were missed. The columns labelled "Delay" specify the total delay for all aircraft, in minutes. The columns labelled "S.P.D." show the sum of the Squares of the Positional Deviation from the first-come-first-served sequence. So a schedule with "S.P.D." of 0 is the first-come-first-served sequence. Results are given for four experimental cases, labelled ECD1 to ECD4 (for Experiment, CTOT vs Delay), for different weights of the objective function.

The results in Table 1 provide more insight into what is actually happening. For example, comparison of the results for ECD1 and ECD2 shows that, in eight of the ten datasets, adding a small penalty for missing CTOTs did not increase the delay, although in two of these cases (datasets 7 and 10) the addition of the penalty meant that the number of CTOTs missed was decreased even though the delay was not increased. In both of these cases (and in datasets 4, 5, 6 and 9), the measure of inequity in the sequence was also increased by the addition of a penalty for CTOT compliance. Datasets 4 and 9 are examples where the addition of the penalty for CTOT compliance actually altered the schedule produced (the S.P.D. values changed) but schedules with the same delay were found. These both illustrate the presence of symmetries within the problem and foreshadow the trade-off between CTOT compliance and equity which will be observed later.

Table 1: Results showing number of CTOTs missed (#C), total minutes delay and inequity (squared positional deviation, S.P.D.) for extreme cases of CTOT compliance vs delay, with a very low cost for positional inequity (W_3=1)

	ECD1: W_1=0,W_2=100			ECD2: W_1=1,W_2=9900			ECD3: W_1=99,W_2=100			ECD4: W_1=1,W_2=0		
D.S.	#C	Delay	S.P.D.	#C	Delay	S.P.D.	#C	Delay	S.P.D.	#C	Delay	S.P.D.
1	1	29784	400	1	29784	400	0	30444	360	0	106320	86
2	0	27896	328	0	27896	328	0	27896	328	0	89523	160
3	1	30230	530	1	30230	530	0	30410	544	0	98425	90
4	1	25264	186	1	25264	204	1	25264	204	1	48424	30
5	4	37710	1270	2	37830	1508	0	38070	1556	1	108460	1306
6	2	35084	550	1	34974	570	0	35694	568	0	85480	146
7	3	30120	562	2	30120	628	1	30600	646	1	71220	376
8	2	29558	336	2	29558	336	1	29738	342	1	67898	112
9	7	48987	1904	7	48987	1948	4	50598	2118	4	105639	640
10	7	38217	600	6	38217	768	4	40917	874	5	80064	1226

Comparison of the results for ECD2 and ECD3 reveals the way in which the ratio of the weights for CTOT compliance and delay affect the schedules which are produced. Datasets 2 and 4 gave identical results for ECD2 and ECD3 regardless of the weights that were applied. This shows that the normal trade-off between CTOT compliance and delay in the good schedules (as seen by the trend in Figure 1 for the sum over all datasets) was not present in these datasets. Thus, at some times of the day it is possible to obtain the best CTOT compliance even without an increase in delay. This kind of information can be hidden in total results over all datasets. In the other cases, the difference between the best schedules with a high delay cost and the best schedules with a high CTOT miss cost is only a few CTOTs (1 in 5 cases, 2 in 2 cases and 3 in 1 case). The delay varies by less than 1% in 5 cases and by 1.5%, 2%, 2.2%, 3.2% and 6.6% in the other cases. The increase in delay for achieving the better CTOT compliance can, thus, be observed to vary greatly over the different datasets. The potential benefits of altering the weights can, therefore, vary greatly according to the time of the day. The effects of removing CTOT constraints when sequencing at the holding area were significantly greater than this (Atkin *et al*., 2006a), illustrating the greater flexibility to cope with delays and advancements for CTOTs when re-sequencing at the stands.

Comparison of the results for ECD3 and ECD4 reveals that the CTOT compliance in datasets 5 and 10 more CTOT slots are missed than when only CTOT compliance is penalised than when both CTOT compliance and delay are penalised. The problems raised by penalising CTOT compliance without a delay penalty are discussed in more detail later and are related to a lack of an incentive towards producing low delay schedules when CTOTs are not present (since W_3=1, the first-come-first-served sequence is preferred) leading to a large cumulative delay so that later aircraft with CTOTs cannot achieve them.

5.2 Equity of positional deviation vs CTOT compliance

Any positional movement of an aircraft within a sequence can be perceived to be unfair. For example, when the occupants of an aircraft with a positional delay observe other aircraft overtaking it in the queues in the holding area. Even when some of the overtaking is implicitly taking place at the stands (via stand holds), cost considerations mean that it is important not to penalise some airlines more than others. One way of limiting this kind of inequity in the take-off sequence is to penalise positional movement, or to apply a maximum position shift to aircraft, as discussed earlier. The results in this section will show that neither approach is appropriate for the Heathrow take-off problem.

Experiments were performed to determine the effects of varying the ratio between W_1 and W_3. The results are shown in Figures 2 and 3. In each case the total squared positional deviation across all ten datasets is shown along with the total number of CTOTs which were missed. Figures 2 and 3 differ in the value that was used for the delay cost weight W_2. W_2 was set to 0 for the results shown in Figure 2 and to 100 for the results shown in Figure 3. In both cases, the far left result is for a reasonable CTOT cost and a low equity cost (W_1=1,W_3=1), and the far right is for only an equity cost, not a CTOT cost (W_1=0,W_3=100). The second value from the left shows the results for a high cost for CTOT misses and only a moderate cost for positional inequity (W_1=99,W_3=100), then the CTOT miss cost decreases until the point labelled “balanced costs” (reducing W_1 in steps of 1 until the centre point has value W_1=1,W_3=100) after which point the equity cost is increased in steps of 100 until the penultimate point has the weights W_1=1,W_3=9900.

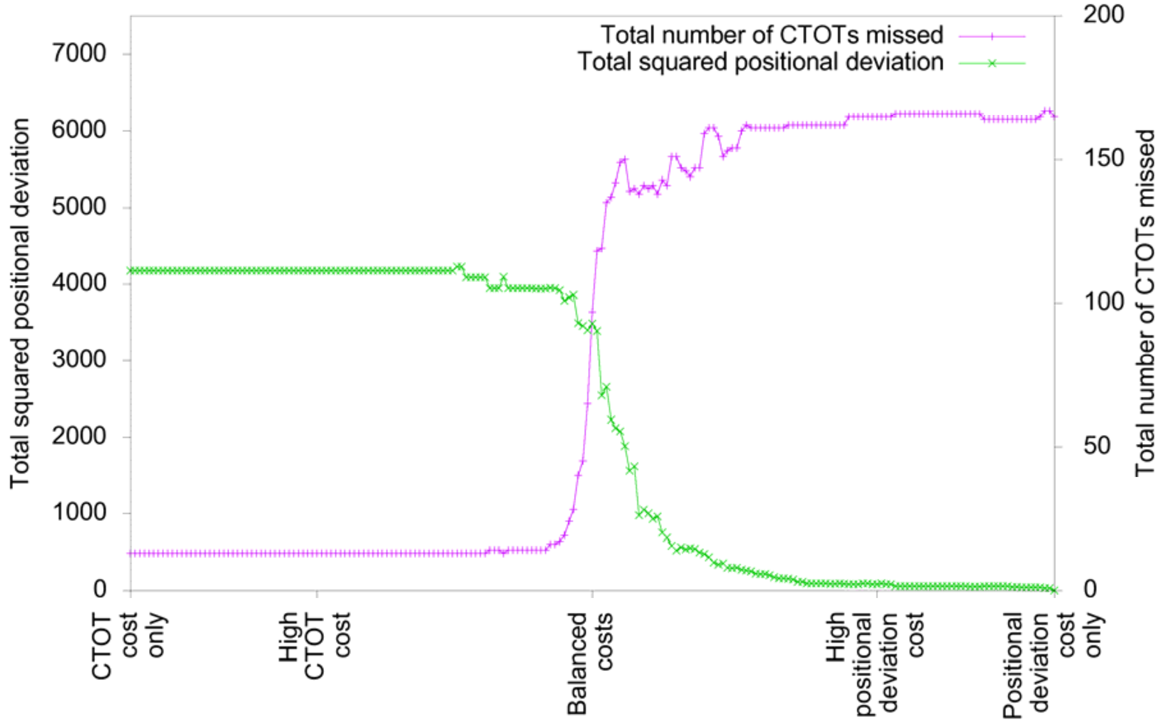

Figure 2: Total squared positional delay and number of CTOTs missed, with no penalty for delay (W_2=0), for varying values of W_1 and W_3 in the objective function

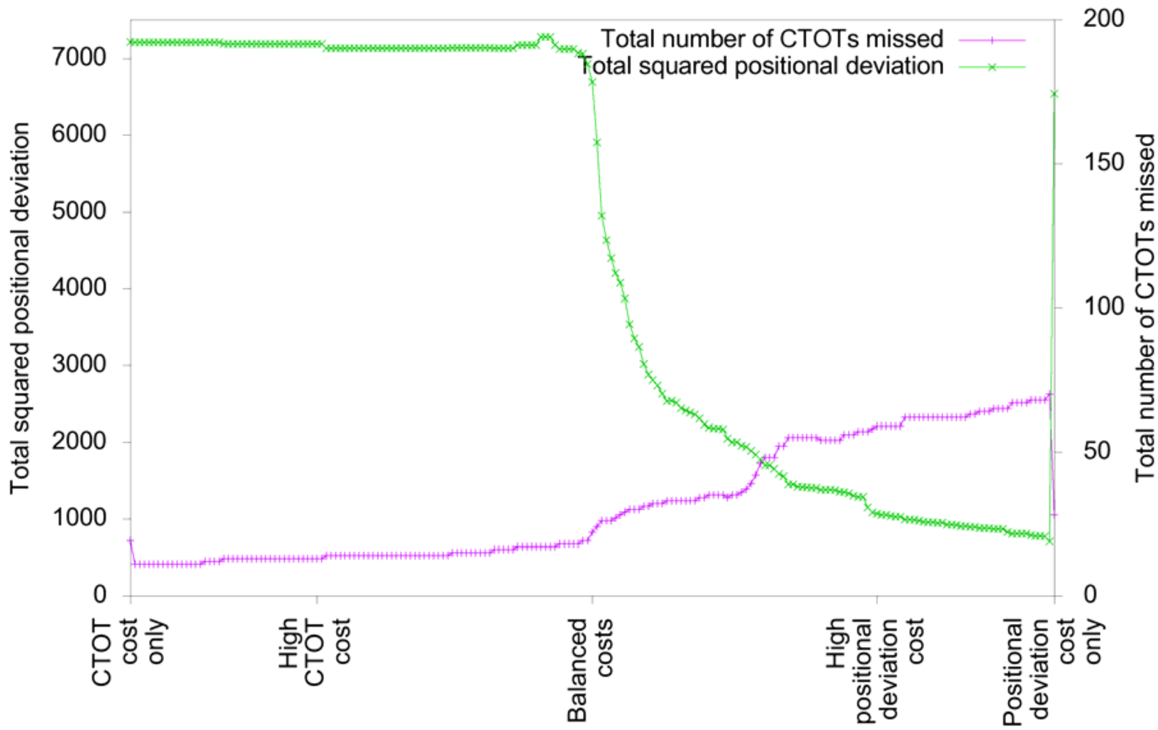

Figure 3: Total squared positional delay and number of CTOTs missed, with a penalty for delay (W_2=100), for varying values of W_1 and W_3 in the objective function

In both Figure 2 and Figure 3, a trade-off between the number of CTOT misses and the positional inequity is apparent, although the graph in Figure 2 is far from monotonically increasing (for CTOT misses) or decreasing (for equity) as the penalty for missing CTOTs is decreased or the equity penalty is increased. As discussed earlier, the reason for this is the conflict between CTOT compliance and positional delay within the optimisation windows. Large delays can be accumulated when aircraft with CTOTs are not present. The cost of positional deviation compared with the relatively small cost for changing the amount by which a CTOT is missed will often mean that aircraft within the current window which cannot immediately achieve their CTOTs will be sequenced into first-come-first-served sequence. What happens within an optimisation window is, therefore, very sensitive to the selection of aircraft with CTOTs which are present. This sensitivity to particular aircraft makes the schedule less stable since the value of the local solution within the window is far less related to the value of the overall schedule. As the cost for equity of positional deviation grows, the positional inequity can drop to very low values and the number of CTOTs missed can be very large. The schedules thus converge very quickly to the first-come-first-served sequence as the cost of positional deviations is increased, since any aircraft which cannot achieve a CTOT within a single window is forced into the first-come-first-served sequence. The addition of a delay penalty helps to reduce this sensitivity by improving the relationship between the value of the partial sequence within the window and the value of the full take-off sequence, reducing the accumulated delay and making it more likely that aircraft can achieve the CTOTs within a single optimisation window. This has the effect of reducing the number of CTOTs which are missed but greatly increasing the perceived inequity in the schedule.

The values for the extreme points of Figures 2 and 3 are shown in Tables 2 and 3 for each of the datasets individually. Table 2 shows the results with W_2=0, and Table 3 with W_2=100, to match the results in Figures 2 and 3, respectively. As expected, since delay is not penalised in ECE1 to ECE4, the delay in these schedules is much higher than in the schedules produced by ECE5 to ECE8, but the number of CTOTs missed is greatly reduced and the positional inequity is greatly increased. Even with the delay penalty some datasets (5, 9 and 10) have very large numbers of CTOTs missed, but without the delay penalty far more are missed in almost all of the datasets.

An interesting relationship between delay, CTOT compliance and positional equity is shown in these results. It can be illustrated by considering an aircraft *A* which has to wait for the start of a CTOT slot, as enforced by Inequality 5. There are two options from which a controller has to decide upon for aircraft which arrive at the holding area after this one. Either these aircraft can take off before *A*, utilising the otherwise wasted runway capacity, or they can take off after *A*. If an aircraft takes off before *A* it will be overtaking, so the take-off sequence will deviate more from the first-come-first-served sequence, and the sequence with the overtaking could be perceived as being less equitable than the sequence where the aircraft waits to take off until after *A*. However, in this case, the overtaking actually has no adverse effect upon aircraft *A*. This is the major reason why a high penalty for positional equity is counter productive for take-off sequencing. When a penalty is applied for delay here, the objective to reduce delay encourages other aircraft to take-off before *A* since doing so will reduce the overall delay, and increases the measured inequity in the sequence.

In fact a similar thing also occurs when there are a lot of aircraft queued for a busy departure route. Due to the departure rate restrictions along some routes a large delay can sometimes accumulate for the later aircraft. This is unavoidable. In this case it is desirable to allow aircraft for other routes to overtake these aircraft since they can do so without delaying the take-off times of the overtaken aircraft. Penalising positional movement will prevent this from happening.

Investigation of the sequences which were obtained for dataset 5 illustrates the interaction of a strong emphasis on equity and aircraft with CTOTs or which are queued for a busy departure route. In ECE6, two aircraft in dataset 5 were assigned a positional delay of fourteen places and one a delay of twelve places. In contrast, in ECE7 the maximum positional movement in dataset 5 was two places (three aircraft were advanced two places, 25 were advanced one place, 25 were delayed one place and 3 were delayed two places). Consideration of the actual delay that aircraft suffer identifies the three problematic aircraft immediately for ECE6 since they are the only ones with high delay values, but also reveals that, for ECE7, all of the other aircraft around these also suffer a large delay since they are not permitted to overtake to take advantage of the empty runway at that time. Obviously, this would never be allowed to happen in practice.

Table 2: Results showing number of CTOTs missed (#C), total minutes delay and inequity (squared positional deviation, S.P.D.) for extreme cases of CTOT compliance vs equity, when no cost is applied to penalise delay

	ECE1: W_1=1,W_3=1			ECE2: W_1=99,W_3=100			ECE3:W_1=1,W_3=9900			ECE4: W_1=0,W_3=100		
D.S.	#C	Delay	S.P.D.	#C	Delay	S.P.D.	#C	Delay	S.P.D.	#C	Delay	S.P.D.
1	0	106320	86	0	106320	86	7	136860	2	7	140520	0
2	0	89523	160	0	89523	160	5	99003	0	5	99003	0
3	0	98425	90	0	98425	90	1	127645	0	1	127645	0
4	1	48424	30	1	48424	30	4	55804	4	4	53464	0
5	1	108460	1306	1	108460	1306	28	200860	2	28	200860	0
6	0	85480	146	0	85480	146	14	122440	2	14	122380	0
7	1	71220	376	1	71220	376	27	145080	8	28	154080	0
8	1	67898	112	1	67898	112	8	79178	0	8	79178	0
9	4	105639	640	4	105639	640	29	176868	4	29	177888	0
10	5	80064	1226	5	80064	1226	44	153144	10	41	144684	0

Table 3: Results showing number of CTOTs missed (#C), total minutes delay and inequity (squared positional deviation,S.P.D.) for extreme cases of CTOT compliance vs equity, when a cost is applied to penalise delay (W_2=100)

	ECE5: W_1=1,W_3=1			ECE6: W_1=99,W_3=100			ECE7: W_1=1,W_3=9900			ECE8: W_1=0,W_3=100		
D.S	#C	Delay	S.P.D.	#C	Delay	S.P.D.	#C	Delay	S.P.D.	#C	Delay	S.P.D.
1	1	29784	400	0	30444	360	1	43644	62	1	29784	400
2	0	27896	328	0	27896	328	1	42716	74	0	27896	328
3	1	30230	530	0	30410	544	1	46790	40	1	30230	530
4	1	25264	204	1	25264	204	1	28264	52	1	25264	186
5	0	38070	1376	0	38370	1272	14	87450	74	4	37650	1144
6	1	35154	544	0	35694	568	4	60682	52	2	35084	550
7	2	30120	628	1	30600	646	8	75060	124	3	30120	562
8	2	29558	336	1	29738	342	2	39038	52	2	29558	336
9	6	49707	2062	4	50598	2118	18	115785	108	7	48987	1904
10	5	38637	800	4	41157	824	20	68637	66	7	38217	600

The ECE1 and ECE2 results show the values of the schedules which were obtained with a high penalty for CTOT compliance and a low penalty for inequity. The fact that the same schedules were found for both is hardly surprising since the ratio of W_1 to W_3 is very similar in both cases. In contrast, the results for ECE5 and ECE6 show the effect of the presence of the delay cost when the weights for CTOT compliance and equity are raised from W_1=1 and W_3=1 to W_1=99 and W_3=100 respectively. ECE5 has a comparatively high weight for delay compared with ECE6, and the improved CTOT compliance for ECE6 in many cases is only to be expected.

The results for datasets 1, 2, 3 and 8 are identical for ECE5 and ECE8. They were also identical for ECD1 and ECD2. This indicates that, in these cases, the CTOT compliance is obtained anyway for the low delay schedule. Considering the delay values in ECD1 or ECE1, it is obvious that these are also some of the lowest delay schedules (dataset 4 is the other low one) and were from the quieter times of the day (the intuitive conclusion that delay is related to the number of aircraft queuing for the runway is supported by the observations which were made by Idris *et al.*, 2002). The number of possible sequences is highly related to the mean delay – since aircraft can only switch positions in the better schedules with other aircraft which can attain the same take-off times. In this case it is not so surprising that the number of good schedules is limited and that the better CTOT compliance schedules also have low delays.

5.3 Total delay vs equity of positional delay

Finally, the relationship between total delay and equity of positional deviation has also been considered. The trade-off is illustrated in Figure 4. It is possible to observe from Figure 4 that these two objectives are in conflict in the better take-off sequences, as expected given the earlier observations. As the cost for delay is increased the positional inequity increases (although previous observations should have clarified that this may not be a problem) until it reaches a peak and levels off.

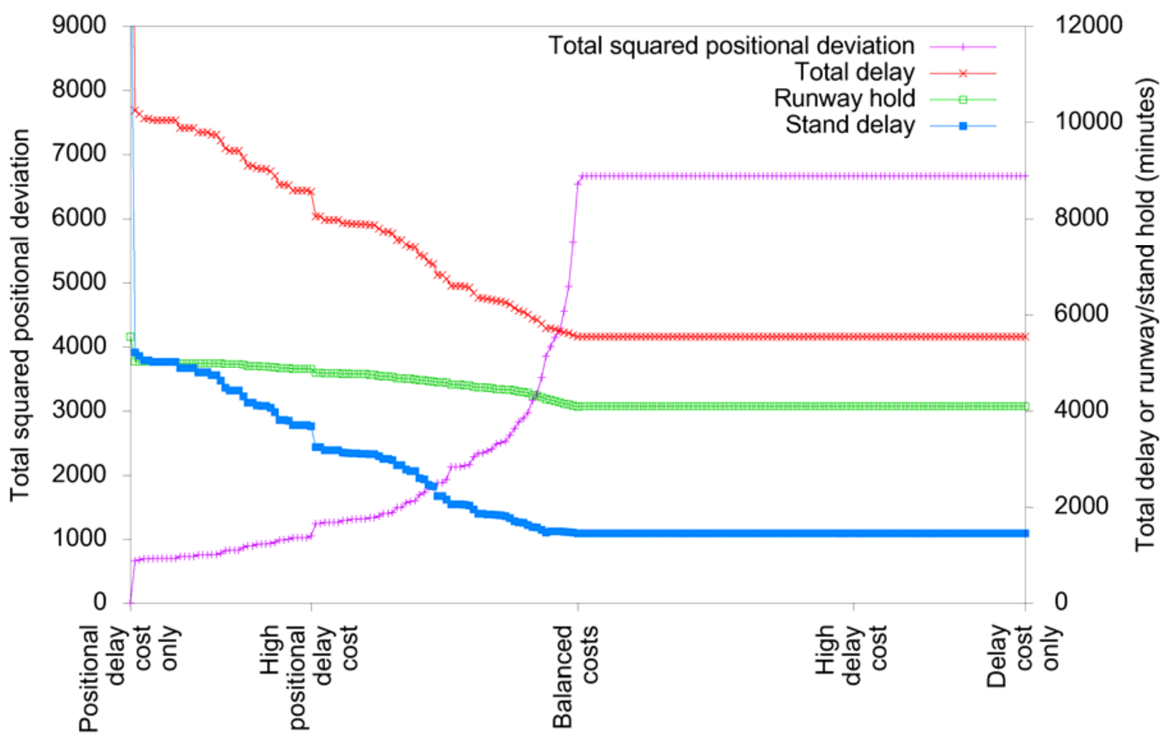

Figure 4: Total squared positional delay and total delay, with no penalty for CTOT compliance (W_1=0), for varying values of W_2 and W_3 in the objective function

The purpose of the TSAT allocation system is, obviously, to allocate TSATs, thus it is possible to also observe the breakdown of the predicted delay between stand delay (without the engines running) and runway hold delay (with the engines running), and to observe the performance of the TSAT allocation system in terms of the amount of the time that the system allocates to stand hold. The stand delay and runway hold delay are also plotted in Figure 4. Obviously, the sum of these two delays is the total delay. One encouraging feature of Figure 4 is that, as the overall system delay increases, the majority of the additional delay is absorbed as stand hold rather than as runway hold (when the engines would be running and fuel usage would increase). A slight increase in total runway hold should be expected, however, when the overall delay increases since only hold beyond the threshold value (five minutes in these experiments) is absorbed as stand hold. In the good schedules, many aircraft will have lower than the five minute threshold runway hold. As the delay in the schedule increases the delay for these aircraft will increase and, until it reaches the five minute threshold, this will be observed as increased total runway hold not stand hold. Figure 4 indicates that the TSAT system is correctly performing its desired function.

Table 4 shows the results for the extreme values of the weights for each of the individual datasets. For EDE1, the cost for positional inequity is high and there is no cost for delay or CTOT misses so the first-come-first-served sequence is, obviously, optimal and these are the sequences which the algorithm adopts (shown by values of S.P.D. of 0).

Table 4: Results showing number of CTOTs missed (#C), total minutes delay and inequity (squared positional deviation, S.P.D.) for extreme cases of equity vs delay, with no cost for CTOT misses (W_1=0)

	EDE1: W_2=0, W_3=100			EDE2: W_2=100, W_3=9900			EDE3: W_2=9900, W_3=100			EDE4: W_2=100, W_3=1		
D.S	#C	Delay	S.P.D.	#C	Delay	S.P.D.	#C	Delay	S.P.D.	#C	Delay	S.P.D.
1	7	140520	0	1	43644	62	1	29784	400	1	29784	400
2	5	99003	0	1	42716	74	0	27896	328	0	27896	328
3	1	127645	0	1	46790	40	1	30230	530	1	30230	530
4	4	53464	0	1	28264	52	1	25264	186	1	25264	186
5	28	200860	0	14	84810	72	4	37710	1270	4	37710	1270
6	14	122380	0	4	60682	52	2	35084	550	2	35084	550
7	28	154080	0	9	77100	110	3	30120	562	3	30120	562
8	8	79178	0	2	39038	52	2	29558	336	2	29558	336
9	29	177888	0	20	123345	82	7	48987	1904	7	48987	1904
10	41	144684	0	20	68637	66	7	38217	600	7	38217	600

Figure 5 shows how the results vary when a penalty is also applied for missing CTOTs (W_1=1 rather than W_1=0) and Table 5 shows the results for the extreme values in Figure 5. The trade-off between delay and equity of positional delay is again obvious. Adding the penalty for missing CTOTs increases the delay slightly, as expected, but not for all datasets when the penalty for delay (W_2) is high. As expected, the CTOT compliance is improved when a penalty for CTOT misses is added, and the positional inequity is also increased due to the movement required to meet CTOTs in the presence of a penalty for delay.

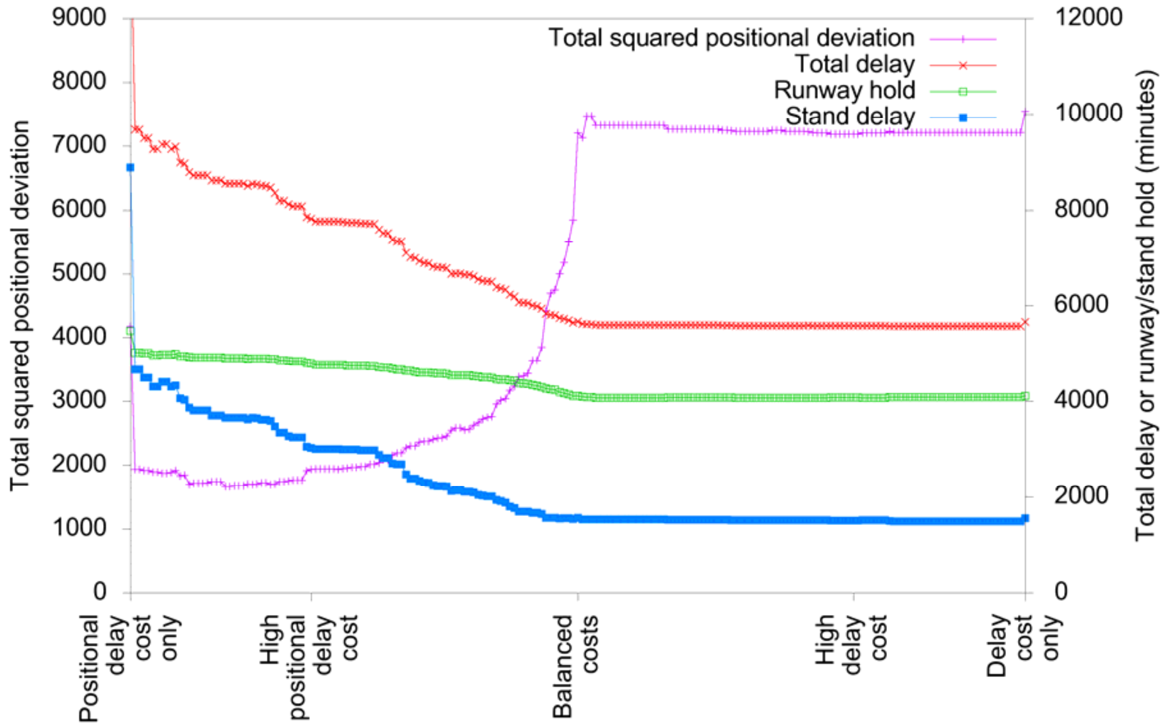

Figure 5: Total squared positional delay and total delay for schedules with a penalty for missing CTOTs (W_1=1), for varying values of W_2 and W_3 in the objective function

Table 5: Results showing number of CTOTs missed (#C), total minutes delay and inequity (squared positional deviation, S.P.D.) for extreme cases of equity vs delay, with a cost for CTOT misses (W_1=1)

	EDE5: W_2=0, W_3=100			EDE6:W_2=100,W_3=9900			EDE7:W_2=9900,W_3=100			EDE8: W_2=100,W_3=1		
DS	#C	Delay	S.P.D.	#C	Delay	S.P.D.	#C	Delay	PD	#C	Delay	S.P.D.
1	0	106320	86	1	43644	62	1	29784	400	0	30444	360
2	0	89523	160	1	42836	78	0	27896	328	0	27896	328
3	0	98425	90	1	46910	58	1	30230	530	0	30410	544
4	1	48424	30	1	28384	58	1	25264	204	1	25264	204
5	1	108460	1306	18	100410	344	0	38070	1376	0	38070	1556
6	0	85480	146	2	61342	80	1	35154	544	0	35694	568
7	1	71220	376	3	54180	322	2	30120	628	1	30600	646
8	1	67898	112	1	38798	68	2	29558	336	1	29738	342
9	4	105639	640	13	97239	390	6	49707	2062	4	50598	2118
10	5	80064	1226	10	67224	472	5	38637	800	4	40917	874

5.4 The problem of applying a high weight for positional equity

The problems (that were discussed earlier) with heavily penalising positional inequity in the sequence mean that a high penalty for positional inequity cannot be applied without adverse side-effects on those occasions when aircraft require a long delay for the start of CTOT or where certain departure routes are busier than others. Both of these occurrences are common. At these times there would be a choice of whether to accept a large positional delay (allowing aircraft which do not need to wait to overtake the delayed aircraft) or to accept unnecessary delays for other aircraft merely to avoid perceived positional inequity.

Penalising positional movement at first appears to be both a good way for measuring inequity and for penalising inequity, since it applies a penalty based upon the number of aircraft which an airline/pilot will observe to overtake their aircraft. However, the above cases show that this is not actually the case. A low (but non-zero) cost for positional inequity can be a good thing for breaking the symmetries in the problem and preferring more equitable schedules, but too great a cost will prevent the system from accurately modelling what would happen in real life.

In fact, a low cost for positional inequity was used by Atkin (2008) and Atkin *et al.* (2007,2008b) with great success when scheduling within the holding area, but this was only successful because the holding area structure itself limited the schedule sufficiently to avoid unnecessary long delays for aircraft (Atkin, 2008). Scheduling at the stands is far more flexible but some degree of equity must still be ensured. An alternative method is required.

5.5 Comparison with the real schedules

Now that a better understanding of the trade-offs between the objectives has been obtained, a comparison with the manual/real results is worthwhile. Three things must be clarified, however, when considering the manual schedules. Firstly, it is apparent that some CTOTs were renegotiated on the day and this is not reflected in the datasets. Unfortunately data about modified CTOTs was not available so the automated tests have assumed that the original CTOTs which applied to the aircraft were maintained. In particular, those aircraft which have a long wait for a CTOT will usually attempt to renegotiate an earlier slot if one becomes available, and this often happens. This means that the long delays (and perceived inequity) in the automated schedules while aircraft await the start of a CTOT slot do not occur to the same extent for the real schedules. One effect of this is a rise in the apparent number of CTOTs that were missed in the real schedule, since this count does not allow for the fact that some of these will have been renegotiated and hence not require extensions. The second thing to note is that the controllers were producing these schedules in real time, considering only a few aircraft at once, with imperfect knowledge of what would happen in future. The limited planning horizon that has to be utilised by controllers can reduce the potential for long term gains (Atkin *et al.*, 2006b) as can the lack of exact knowledge about aircraft taxi times (Atkin *et al.*, 2008b). The third thing to clarify is that the controllers were working within the restrictions of the holding area structures at Heathrow. The holding area structure restricts the possible re-sequencing (Atkin *et al.*, 2006a), giving the controllers a harder problem to solve. Controllers sometimes have to utilise additional extensions in order to avoid significant losses of runway capacity or unfairly long delays for other aircraft, which is, of course, why these discretionary extensions exist. This can be exacerbated by the holding area structure since it may be impossible for aircraft with CTOTs to overtake other aircraft in the holding area in order to achieve take-off within a CTOT time-slot. One side-effect of TSAT allocation should be a reduction in holding area congestion, enabling controllers to avoid that kind of situation.

With the above considerations in mind, Table 6 presents comparative results for the real schedules and some automatically generated ones summed across the ten test datasets. The first row presents the real results that the controllers attained, showing both the number of CTOTs missed (with the note that some of these will have been renegotiated so would not actually have been missed), the total minutes delay (measured as the total time from holding area arrival to take-off over all 1100 aircraft in the 10 datasets), and the total squared positional deviation from the first-come-first-served sequence based upon holding area arrival time. The positional deviation in the real schedules is very low, but the number of CTOT misses is high. In fact consideration of the real schedules shows that the majority of misses were in datasets 8, 9 and 10. These datasets were problematic due to the high number of CTOTs allocated at a time of high runway demand. Accumulated delay meant that a number of CTOTs had to be renegotiated or missed. The TSAT allocation system also required significantly more CTOT extensions for datasets 9 and 10 than for the other datasets because of this problem. The real take-off times imply that a significant number of these CTOTs had obviously been renegotiated, so probably did not require the use of extensions.

The remaining rows of Table 6 show the results from the TSAT allocation system for various objective function weights and include a breakdown of the delay into stand hold and runway hold since this is the primary expected gain of the system. The second set of results is for the first-come-first-served sequence (W_1=0,W_2=0,W_3=100) and the third is for a concentration upon CTOT compliance without a delay penalty (W_1=1,W_2=0,W_3=1). The remaining results are for W_1=1, W_3=1 and varying weights and powers for the delay. The fourth set of results has a moderate weight for delay while the fifth has a high weight for delay. The sixth and seventh are similar but the delay (in seconds) is raised to the power of 1.5 prior to applying the weight in the objective function, i.e. α=1.5 in function $D(d_j, bt_j)$. In the eighth and ninth the seconds delay is squared prior to applying the weight in the objective function, i.e. α=2.0 in function $D(d_j, bt_j)$.

Table 6: Comparison of automated vs real results, utilising a linear and non-linear cost for delay in order to promote equity of delay

Schedule	Total number of CTOTs missed	Total minutes delay	Total minutes runway hold	Total minutes stand hold	Total squared positional deviation
Manual/real	35	7629	7626	0	4430
First-come-first-served	165	21662	5546	16116	0
No delay cost, W_2=0	13	14358	5473	8885	4172
Linear, W_2=100	19	5574	4086	1488	7208
Linear, W_2=9900	23	5548	4074	1474	7220
Power 1.5, W_2=100	23	5585	4153	1432	5296
Power 1.5, W_2=9900	25	5595	4165	1430	5118
Squared, W_2=100	28	5707	4248	1458	4408
Squared, W_2=9900	28	5707	4248	1458	4408

The manual results reveal that the schedules that the controllers generate are extremely equitable in terms of positional delay. Part of the reason for this is the strong preference for equity, but another reason is the limited planning horizon over which their workload forces them to work. There are limits to how far ahead they can look, since there is so little thinking time available to them. It should also be noted that the controllers perform exceptionally well in the circumstances in which they work (Atkin *et al.*, 2007) and the reduction in delay between the first-come-first-served sequences and the controller produced sequences shows just how valuable their work is in keeping the throughput of the runway high.

The results in Table 6 for the TSAT allocation system show that the system can find schedules which are predicted to have lower overall delays that those which the controllers actually adopted, and furthermore it is able to find schedules which are as equitable as the controller-produced ones when the measure of delay is squared in the objective function. Moreover, the penalty paid in terms of the increase in overall delay from favouring schedules with minimal squared delay rather than minimal delay is relatively low (around 150 minutes) compared with the potential predicted overall benefits (around 1900 minutes).

Although these results indicate that perhaps the most appropriate cost for the delay involves squaring it in the objective function, further investigation needs to be performed into actual controller preferences, to determine whether a slightly lower power would be appropriate or not. A power of 2 appears to give results which best model the positional equity of controller produced schedules, but the question now is whether these were produced by preference or necessity (for instance from a limited planning horizon or congested holding area), and only tuning against feedback from real controllers (rather than historic data) will ever really capture these preferences. What these results do provide is a starting point as to how to model the current controller preferences and an indication that currently produced schedules show a very strong preference for equity of delay.

Of course the TSAT allocation system does not attempt to perform take-off sequencing for runway controllers. The main benefit of a TSAT allocation system can actually be observed in the contents of the fifth (rather than the third) column of Table 6. Namely, that it allows some of the hold to be absorbed at the stand rather than the runway. These results predict that significant amounts of the delay can be absorbed at the stand rather than with the engines running, with consequent environmental and financial benefits.

If the main reason for controllers not achieving the same kind of delay as the automated system is the lack of knowledge of aircraft sufficiently early enough to be able to utilise that knowledge in the sequencing, or from congestion in the holding areas, then a system which holds aircraft appropriately, removing from their consideration aircraft which are obviously unsuitable for take-off at that time, and thus ensuring a more appropriate selection of aircraft at the holding area should simplify the problem for the controllers and enable them to perform even better than they currently do. It is possible, therefore, that a TSAT allocation system could enable the controllers to bridge the gap between the delay predicted by these results and the actual delay that they achieve, and to do so without the TSAT allocation system ever giving any advice to controllers about which sequences to adopt. Furthermore, since the take-off time prediction method is known to be pessimistic (Atkin *et al.*, 2007) because controllers can safely reduce some of the separations at their discretion, but the TSAT allocation system does not assume that this is done, it is entirely possible that controllers could produce better schedules that those which are predicted here if only their workload could be appropriately reduced (for example by removing aircraft that they do not need to consider at that time, and by reducing the congestion at the holding area).

6 CONCLUSIONS

This is a very practical public transport scheduling problem and the performance of the TSAT system will affect around 68 million passengers per year (based on the 2007 passenger figures, BAA, 2007). Holding the correct aircraft at the stands for the correct amount of time is important in order to reduce fuel cost and pollution but not to reduce runway throughput; in fact it may increase runway throughput, by reducing congestion at the runway holding point. The aim of the TSAT allocation system is not to generate a required take-off sequence, but to predict the sequences that controllers would like, and ensure that they are not prohibited, while also reducing the number of inappropriate aircraft at the holding area. It is, therefore, imperative to adequately understand the trade-off between the various objectives which are used to make the hold decisions. The results which were presented in this paper have shown that all three objectives (delay, equity and CTOT compliance) are mutually in conflict in the better take-off sequences. However, they have also shown that occasionally the schedules with the best delay also have the best CTOT compliance, so reducing delay does not always adversely affect the CTOT compliance. These results show the importance of tuning a TSAT allocation system over a variety of different time periods, since the problem can change depending upon the delay for aircraft in the system at the time.

Importantly, the results in this paper have shown that, although it may seem like a good idea for favouring more equitable take-off sequences, applying a high penalty to positional inequity is actually an extremely bad idea in some circumstances and the resulting sequences would not match operational practice. On the other hand, the results have also shown that the schedules which the controllers produce have a very low total squared positional deviation from the first-come-first-served sequence. This implies that they are very equitable in terms of positional deviation. It has been observed that a non-linear objective function for the delay (for example, aiming to minimise squared delay rather than delay itself) generates sequences which have positional deviations closer to those in the manually produced schedules. Care must be taken, however, since the automatically produced sequences should have had a greater positional inequity than the manually generated ones due to the presence of long CTOT delays in the datasets, which the real take-off times showed had obviously been renegotiated in the real sequences. It is possible, therefore, that the weighting for equity was too high for other parts of the sequences, thus making up for the positional inequity where there were long delays. This implies that the squaring of the delay may actually be too much.

These results predict that, even with a heavy penalty applied for inequity of delay, it is possible to find schedules with a predicted delay significantly lower than the delay for the real schedules. Of course, for these experiments the system had perfect knowledge and a predictable environment, neither of which the controllers have in practice. Previous work with take-off sequencing (for example, Atkin *et al.* 2007) has shown that small schedule improvements, which were only possible from looking ahead to foresee problems, could have large delay benefits due to the cumulative effects of separations upon delay. Perhaps more importantly, predictions indicate that it should be possible to absorb significant amounts of the necessary delay at the stand rather than at the holding area. The hope is that, with increased knowledge at the stand, and the time to consider more aircraft in the sequencing, a more appropriate selection of aircraft can be provided to the controller at the holding area, reducing the complexity of the problem and enabling the controllers to outperform even the predictions from these results.

The results in this paper should provide the necessary insights to allow the tuning of the TSAT system towards controller preferences. What remains to be discovered is what these preferences actually are and whether historical results actually show the desired sequences, or those which were enforced upon controllers by the limited planning horizon or holding area structure. Since the TSAT allocation system that has been described in this paper is currently being integrated into a NATS system for CDM at Heathrow, a tuning phase will follow prior to going live, utilising the trade-off results which were described in this paper in order to more accurately reflect the desirable take-off sequences at Heathrow. Of course, the minimum and ideal runway hold values that are used ensure that a pool of aircraft will be available to a controller, and thus that any take-off sequence priorities in the algorithm need not be utilised by the controller, but it is important to ensure that the algorithm is tuned to controller behaviour in order to ensure that the pool of aircraft available includes the appropriate aircraft for the controller.

REFERENCES

Atkin, J.A.D. (2008). On-line decision support for take-off runway scheduling at London Heathrow airport. Ph.D. thesis, The University of Nottingham.

Atkin, J.A.D., Burke, E.K., Greenwood, J.S., Reeson, D. (2006a). An examination of take-off scheduling constraints at London Heathrow airport. *Electronic proceedings of the 10th International Conference on Scheduling of Public Transport (CASPT),* Leeds, UK.

Atkin, J.A.D., Burke, E.K., Greenwood, J.S., Reeson, D. (2006b). The effect of the planning horizon and freezing time on take-off sequencing. *Proceedings of the 2nd International Conference on Research in Air Transportation (ICRAT),* Belgrade, Serbia and Montenegro.

Atkin, J.A.D., Burke, E.K., Greenwood, J.S., Reeson, D. (2007) Hybrid meta-heuristics to aid runway scheduling at London Heathrow airport. *Transportation Science* 41(1), 90-106.

Atkin, J.A.D., Burke, E.K., Greenwood, J.S., Reeson, D. (2008a) A meta-heuristic approach to aircraft departure scheduling at London Heathrow airport. M Hickman, P Mirchandani, S Voss, eds., *Computer Aided Systems of Public Transport. Lecture Notes in Economics and Mathematical Systems*, Springer, Berlin.

Atkin, J.A.D., Burke, E.K., Greenwood, J.S., Reeson, D. (2008b) On-line decision support for take-off runway scheduling with uncertain taxi times at London Heathrow airport. *The Journal of Scheduling* 11(5), 323-346.

BAA. (2007) BAA Limited Annual Report, 2007. Available at : http://www.baa.com/assets/B2CPortal/Static%20Files/AReport_final07.pdf

Beasley, J.E., Krishnamoorthy, M., Sharaiha, Y.M., Abramson, D. (2000). Scheduling aircraft landings - the static case. *Transportation Science* 34, 180-197.

Beasley, J.E., Sonander, J., Havelock, P. (2001) Scheduling aircraft landings at London Heathrow using a population heuristic. *Journal of the Operational Research Society* 52, 483-493.

Beasley, J.E., Krishnamoorthy, M., Sharaiha, Y.M., and Abramson, D. (2004). “Displacement problem and dynamically scheduling aircraft landings." *Journal of the Operational Research Society*, 55(1), 54-64

Dear, R.G. and Sherif, Y.S. (1989). The dynamic scheduling of aircraft in high density terminal areas. *Microelectronics and Reliability* 29(5), 743-749.

Dear, R.G. and Sherif, Y.S. (1991). An algorithm for computer assisted sequencing and scheduling of terminal area operations. *Transportation Research Part A, Policy and Practice* 25, 129-139.

Ernst, A.T., Krishnamoorthy, M., Storer, R.H. (1999) Heuristic and exact algorithms for scheduling aircraft landings. *Networks* 34, 229-241.

Eurocontrol Experimental Centre. (2005) London Heathrow CDM WP1, EEC Note No. 03/05.

Idris, H.R., Clarke, J.-P., Bhuva, R., Kang, L. (2002) Queuing model for taxi-out time estimation. *Air Traffic Control Quarterly*, ATCA Publications 10(1), 1-22.

Psaraftis H.N. (1980) A dynamic programming approach for sequencing groups of identical jobs. *Operations Research* 28(6), 1347– 1359

Trivizas, D.A. (1998) Optimal scheduling with maximum position shift (MPS) constraints: A runway scheduling application. *Journal of Navigation* 51, 250-266