

The trade-off between taxi time and fuel consumption in airport ground movement

Stefan Ravizza · Jun Chen ·
Jason A. D. Atkin · Edmund K. Burke ·
Paul Stewart

Abstract Environmental impact is a very important agenda item in many sectors nowadays, which the air transportation sector is also trying to reduce as much as possible. One area which has remained relatively unexplored in this context is the ground movement problem for aircraft on the airport's surface. Aircraft have to be routed from a gate to a runway and vice versa and it is still unknown whether fuel burn and environmental impact reductions will best result from purely minimising the taxi times or whether it is also important to avoid multiple acceleration phases. This paper presents a newly developed multi-objective approach for analysing the trade-off between taxi time and fuel consumption during taxiing. The approach consists of a combination of a graph-based routing algorithm and a population adaptive immune algorithm to discover different speed profiles of aircraft. Analysis with data from a European hub airport has highlighted the impressive performance of the new approach. Furthermore, it is shown that the trade-off between taxi time and fuel consumption is very sensitive to the fuel-related objective function which is used.

Keywords Airport operations · Environmental impact · Graph-based approach · Ground movement · Multi-objective routing

Stefan Ravizza and Jason A. D. Atkin
School of Computer Science, University of Nottingham
Jubilee Campus, Nottingham, NG8 1BB, UK
E-mail: smr@cs.nott.ac.uk, jaa@cs.nott.ac.uk

Jun Chen and Paul Stewart
School of Engineering, University of Lincoln
Brayford Pool, Lincoln, LN6 7TS, UK
E-mail: juchen@lincoln.ac.uk, pstewart@lincoln.ac.uk

Edmund K. Burke
University of Stirling
Cottrell Building, Stirling, FK9 4LA, UK
E-mail: e.k.burke@stir.ac.uk

1 Introduction

Air transportation is a growing sector and this trend is predicted to continue in future. At the same time, the stakeholders are caring more and more about the environmental impact of the sector. Aircraft ground movement is an operation which is affected a lot by these two conflicting trends. With the increase in aircraft movements, it is likely that hub airports, especially, will form bottlenecks for air transportation. The ground movement problem is thereby key to reducing delays, through increasingly efficient airport operations, linking the runway sequencing, stand holding and gate assignment problems. However, lower accelerations may sometimes be more fuel efficient, even though movement times would be increased. An ambitious goal stated in the report of the High Level Group on Aviation Research for the European Commission attempts to have emission-free aircraft movements when taxiing in the year 2050 (European Commission 2011).

The details of the ground movement problem vary depending upon the aims of the airport but it can be summarised as the problem of producing conflict-free routings for aircraft on the airport's surface, usually from gates/stands to runways and vice versa. A variety of different constraints and objective functions have been used in the past. Previous research on ground movement often focused on minimising the total taxi time (Atkin et al. 2011b; Marín 2006; Pesic et al. 2001; Roling and Visser 2008) or minimising the makespan (the duration from first to last aircraft movement) (García et al. 2005; Herrero et al. 2005). Multi-objective approaches have also been used. In addition to minimising the total taxi time, penalising deviations from a scheduled time of departure/arrival was also considered (Balakrishnan and Jung 2007; Deau et al. 2009; Smeltink et al. 2004) and Gotteland et al. (2003) considered penalising deviations from a departure time interval. Other research has used a weighted linear objective function to simultaneously consider the total routing time, the delays for arrivals and departures, the number of arrivals and take-offs, the worst routing time and the number of controller interventions (Marín and Codina 2008). Although multi-objective approaches were used, we have not found any research focusing on the integration of objectives related to the environmental impact. A literature review was recently published by Atkin et al. (2010b), which also highlighted the differences between the existing approaches.

There is little coverage of environmental considerations in taxiing within the previous research. The main focus has been upon stand holding in order to reduce fuel burn (Atkin et al. 2010a, 2011a; Burgain et al. 2009). The assumption made by them was that by reducing the total taxi time, one can simultaneously improve the efficiency of airport operations and reduce the fuel consumption. However, as indicated in Chen and Stewart (2011), this may not be true for all cases or airports, since the detailed relationship between fuel consumption and the corresponding speed profile was not investigated in previous research. Atkin et al. (2010b) suggested the value of considering speed profiles when routing aircraft to avoid unnecessary fuel burn due to

acceleration and deceleration. Lesire (2010) applied a postprocessing stage in his routing approach to smoothen the speed profiles. Finally, Chen and Stewart (2011) presented an approach to analyse the trade-off between taxi time and fuel consumption for a single trajectory of an unimpeded aircraft.

In this research, we analyse the trade-off between the total taxi time and the fuel consumption for the conflict-free routing problem for aircraft on an airport's surface. In contrast to the approach of Chen and Stewart (2011), the interactions between multiple aircraft are considered. These interactions affect the speed profiles of the aircraft involved and massively increase the solution space of the routing approach, hence, a sophisticated new procedure had to be developed to make such an analysis possible.

Section 2 presents the case which was analysed, then the newly developed multi-objective approach for analysing the trade-off between taxi time and fuel consumption is detailed in Section 3. The results of the application of the algorithm to the dataset are then shown in Section 4; before the paper ends with some conclusions in Section 5.

2 Problem details

Different approaches for fuel burn estimation are considered first in this section, together with the settings which are used for maximal speeds and acceleration, then the dataset for Zurich Airport which was utilised is presented. The section ends with an explanation of the categorisation which is used for aircraft.

2.1 Fuel consumption, taxi speed and acceleration

As is common practice, the International Civil Aviation Organization engine emissions database (ICAO 2008) has been used for estimating the fuel consumption of aircraft. It states that the engine power setting for taxi/ground idle is 7% of full rated power but does not distinguish between the different phases of taxiing. Morris (2005) showed that levels of around 5% to 6% are more realistic for most engine types. A newer approach by Nikoleris et al. (2011) used a set of four different values for different taxi operation phases: 4% for idle thrust, 5% for taxiing at a constant speed or brake thrust, 7% for perpendicular turn thrust and 9% for breakaway thrust. In their study about air quality and public health impacts of UK airports, Stettler et al. (2011) used a setting of 4-7% (a uniform random distribution with a mean of 5.5%) for taxiing (for maintaining a constant speed, decelerating, or holding) and a setting of 7-17% (a triangular distribution with mode of 10%) for taxiway acceleration. Based on the results by Wey et al. (2006), they were stating that fuel flow of the engines is approximately proportional to the engine thrust setting. Khadilkar and Balakrishnan (2011) presented an approach to estimate fuel burn using linear regression. They concluded that the total taxi time is the

main factor, although the number of acceleration events was also a significant factor. Our analysis has approached the problem using a physics-based model which is introduced later, in Section 3.1.

Different researchers were working with different taxi speed settings. Rappaport et al. (2009) showed, using quantitative analysis, that the average speed on straight taxiways (15.9 knots) was higher than the average speed during turns (12.5 knots) at Detroit Metropolitan Wayne County Airport (DTW) in Michigan, USA. Cassell and Evers (1998) reported that 95% of aircraft taxi at less than 30 knots and the average speed was found to be 10 knots during turns. This setting was also used in the work by Chen and Stewart (2011), where the maximal speed during taxiing was set to 30 knots and the speed during turns to 10 knots. The same setting has been applied in this research, where a turn is considered to be when an aircraft has to make a change of direction of more than 30 degrees on a part of a taxiway. The maximal acceleration and deceleration is set to $0.1 \cdot g$, to ensure passenger comfort, as in the latter reference, where $g = 9.81 \text{ m/s}^2$ is the acceleration due to gravity.

It is assumed in this analysis that the airport has no significant taxiway slopes, no heavy wind occurred which affected the fuel burn of aircraft, and that no drag is needed in the model for estimating fuel consumption.

2.2 Dataset: Zurich Airport

Zurich Airport (ZRH) is the largest airport in Switzerland and a hub airport for Swiss International Air Lines AG. The airport has three runways, named 10/28, 14/32 and 16/34 according to their direction of operation, with the first and the last runways intersecting each other. We had access to data for an entire week's operations between the 27th of June and the 3rd of July 2011 for this analysis. There were 2806 arrivals and 2807 departures in total and no extraordinary occurrences happened. The considered data included information about the airport layout, the positions of stands and runway entrance and exit points and the layouts of all of the taxiways. This information was used to represent the entire airport layout as a directed graph, where the edges represent the taxiways and the vertices represent the junctions or intermediate points (Figure 1 illustrates a part of this graph). The data also included the real timings for the aircraft using the airport during each day.

2.3 Aircraft categorisation

Aircraft were classified into different groups and for each group the settings for a representative aircraft type were used for the calculation. This procedure was necessary due to the lack of detailed data in the provided dataset. Aircraft were distinguished by their wake vortex separation group. The group 'light' was represented by the settings for a Cessna 172 Skyhawk. The settings for an Airbus A320 were used for the wake vortex group 'medium'. Finally, the

group ‘heavy’ was represented by the settings for an Airbus A333. All these aircraft were the most common aircraft type in their category. The technical details of the aircraft and their engines can be found in Table 1.

Table 1 Specifications of aircraft and engines

	Cessna 172 Skyhawk	Airbus A320	Airbus A333
Maximum take-off weight	1100 kg	78000 kg	230000 kg
Rolling resistance	162 N	11.48 kN	33.85 kN
Engine	O-320	CFM56-5A1	PW4168
Number of engines	1	2	2
Maximal fuel flow	1×0.0112 kg/s	2×1.051 kg/s	2×2.884 kg/s
Rated output	unknown	2×112 kN	2×302.5 kN

The same approach was used to calculate the total rolling resistance as in Chen and Stewart (2011) and is defined as follows:

$$F_r = \mu \cdot m \cdot g, \quad (1)$$

where $\mu = 0.015$ is the rolling resistance coefficient on a concrete surface, m is the maximum take-off weight of the aircraft and $g = 9.81$ m/s² is again the acceleration due to gravity. The maximal fuel flow and the rated output values are based on the International Civil Aviation Organization engine emissions database (ICAO 2008) and the research by Stettler et al. (2011) and have to be multiplied by the number of engines specified in each setting.

3 Methodology

The focus of this research is entirely on the ground movement part of the airport operations of aircraft. In doing this, the pushback/landing time of aircraft are as specified by the dataset and assumed to be fixed.

This section first introduces the objective functions, before an overview of the developed integrated procedure is given. Two key elements of the procedure are given in separate sections afterwards.

3.1 Objective functions

This research analyses the trade-off between taxi time and fuel consumption in airport ground movement. The first objective function aims to minimise the total taxi time (including waiting times during taxiing) combined with moving possible waiting times to the gate where the engines are not running. The second objective function aims to minimise fuel burn. Similarly to the research by Chen and Stewart (2011), a fuel consumption index is used. This penalises high acceleration rates during taxiing and uses a physics-based model. Basically, the formula for the force of acceleration is assumed to be given by

$F_a = m \cdot a_p$, (Newton’s second law of motion) where a_p is the acceleration of an aircraft during a phase p and m its weight. The rolling resistance F_r is then also taken into account (see Section 2.3 for the formula and values). The fuel consumption index is defined as the sum of the force of acceleration plus the rolling resistance, multiplied by the time for which it was applied. If the sum of the force of acceleration plus the rolling resistance is negative, due to deceleration in a phase, the sum is set to zero, since aircraft need fuel to accelerate or taxi with constant speed but cannot recover fuel while decelerating. The trade-off between the two objective functions for an example taxi route is shown in Figure 2.

3.2 Integrated procedure

A routing approach was developed, based upon the algorithm presented in Atkin et al. (2011b), utilising the trade-off information gleaned from the algorithm proposed in Chen and Stewart (2011). It is a sequential, vertex-based, label-setting algorithm working on a graph representing the airport’s surface. Since two conflicting objective functions are considered, the approach has to be enhanced by using an adapted version of the algorithm, in a sophisticated integrated procedure.

The general idea of this procedure was proposed by Climaco and Martins (1982), whose aim was to develop a shortest path algorithm for finding the pareto-front of optimal paths for two criteria. The objective functions which they used were minimising the total time and minimising the cost of the path, where each edge had two values assigned to it. Their method generates a sequence of k shortest paths with respect to the first objective function, until the path with the minimal value with respect to the second objective function is obtained, leading to a pareto-front of all optimal paths.

Our problem differs from the problem which Climaco and Martins (1982) were facing in two main points. Firstly, not all edges are available at all times since other aircraft are also travelling at the airport and will block some parts of the taxiways at certain times. Secondly, the second objective function cannot be evaluated with a simple Dijkstra’s algorithm for finding the shortest path (Dijkstra 1959) in this situation, but needs a more elaborate method due to its non-additivity.

Algorithm 1 shows the proposed integrated procedure at a glance. The approximation of the global pareto-front is generated in a discretised way due to the complexity of the problem. The parameter l defines the number of generated points on the global pareto-front approximation. In each iteration of the outer loop (lines 2-10), the objective values are generated for both objectives, starting with the most time-efficient solution then incrementally changing to the most fuel-efficient solution. For each outer loop, the entire set of aircraft has to be scheduled. The algorithm routes and schedules the flights sequentially and is based on an initial sequencing (line 1) by pushback/landing

Algorithm 1: Integrated procedure for trade-off analysis

```

1 Sort all flights by pushback/landing time
2 foreach objective function discretisation  $i \leftarrow 1$  to  $l$  do
3   foreach aircraft  $a$  do
4     Find the best  $k$  routes w.r.t. minimal taxi times using the  $k$ -QPPTW
       algorithm
5     foreach route  $k$  of aircraft  $a$  do
6       Approximate the pareto-front of both objectives, using the population
       adaptive immune algorithm (PATT-PAIA)
7     Generate the combined pareto-front for the source-destination pair of
       aircraft  $a$ 
8     Discretise this pareto-front into  $l$  roughly equally spaced points
9     Select the  $i$ th point and reserve the relevant route for aircraft  $a$ 
10  Save the accumulated values for all aircraft for both objective functions for the
      global pareto-front
11 Output: Global discretised pareto-front

```

times of all aircraft. Different (adaptive) sequencing methods could be used, as was done by Ravizza and Atkin (2011), but this was not investigated here.

The first subroutine (line 4) finds the best k routes for aircraft a related to the total taxi time. In doing so, reservations of already routed aircraft have to be taken into account. The k -Quickest Path Problem with Time Windows (k -QPPTW) was developed for this purpose and is explained in more detail in Section 3.3. A possible set of generated routes can be seen in Figure 1.

The second subroutine (line 6) analyses each of the k routes independently. A population adaptive immune algorithm (PATT-PAIA) approximates the pareto-front of different speed trajectories for aircraft a on a particular route, complying with the unblocked time-windows for each edge and the detailed speed behaviours of this aircraft. A more detailed description of this subroutine is given in Section 3.4 and an example of the output can be seen in Figure 2.

The subroutine in line 7 combines the k different pareto-fronts and selects, with the same dominance rules as in the PATT-PAIA, the global pareto-front for a given source-destination pair of aircraft a (see Figure 3). The resulting pareto-front is discretised into l points, as equally spaced as possible (line 8). The approach aims to split the border of the pareto-front between the most time-efficient and most fuel-efficient solutions into equally spaced segments and always selects the closest non-dominated point to each of the ends of these segments. Line 9 selects the i th point (according to the outer loop of the algorithm) out of the l ordered representative points. In addition, the detailed route associated with this point is fixed for this aircraft and the scenario is changed in such a way that upcoming aircraft cannot use the same parts of the taxiways at the same time. The inner loop (lines 3-9) is repeated until all

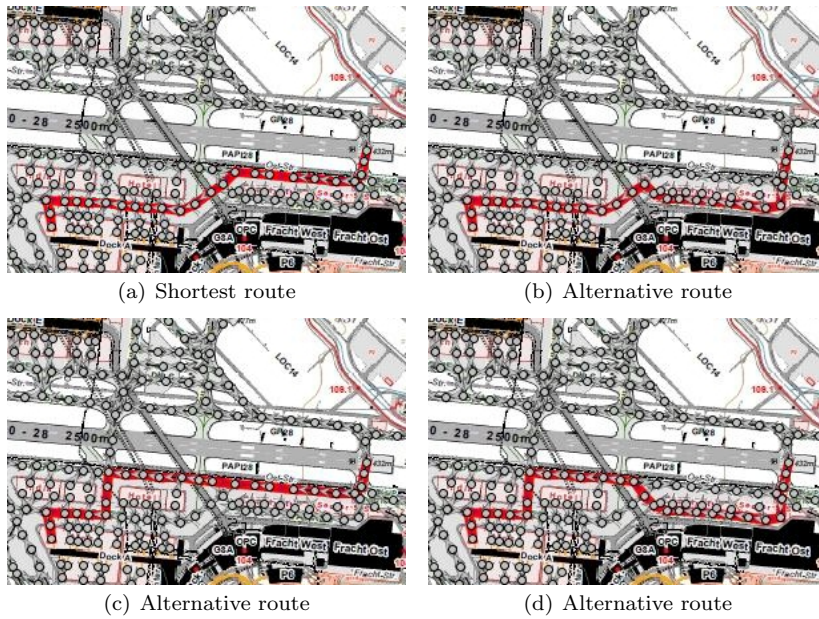


Fig. 1 Different routes from pier A to runway 28

of the aircraft from the dataset have been routed and the total taxi time and the total fuel consumption can be accumulated to generate a single point in the global pareto-front (line 10). Obviously, before repeating the outer loop (line 2) all of the reservations of the aircraft have to be reversed, since the scenario is then evaluated for a different objective function discretisation.

Since the subroutine on line 6 is comparably time consuming, the procedure could be parallelised for this stage and executed on a cluster of processors.

3.3 Sequential k -QPPTW

Schüpbach and Zenklusen (2011) recently showed that a simplified version of the conflict-free routing problem for a group of n vehicles is NP-hard, even when the underlying graph is a path, using a reduction from the 3-partitioning problem. Hence, an approach was developed which is based upon a sequential routing of the aircraft. Atkin et al. (2011b) proposed the Quickest Path Problem with Time Windows (QPPTW) algorithm for solving the airport ground movement problem more realistically than other proposed approaches. This vertex-based, label-setting algorithm works on a graph representing the airport's surface, does not need any time discretisation, respects reservations of parts of the taxiways for previously routed aircraft and can compute a route very quickly. Details of the approach can be found in Atkin et al. (2011b).

The k -QPPTW algorithm was extended by adapting it to not only generate the “best” (in this setting, fastest) route, but a set of the k best solutions. This




Fig. 2 Pareto-front of unimpeded taxi trajectories




Fig. 3 Combined Pareto-front from four different routes which are shown in Figure 1 having associated time-windows

extension is based upon the ideas of Yen (1971) and Lawler (1972). Yen (1971) introduced an algorithm to find the k shortest loopless paths in a network, where the computational upper bound of the algorithm only increases linearly with the value k . The main idea behind the approach is that the $(j + 1)$ th path can only deviate from the root of one of the best j paths in one vertex. Hence, it is only necessary to look for all shortest deviations from the best j

paths and then select the deviation which has the best objective value for the entire path.

The QPPTW algorithm was similarly adapted and generates the best route in the conventional way. It then iterates until it has found the best k routes. In each iteration j , it generates all deviations from the $(j - 1)$ th best routes which are different from routes which have already been found. The j th best route is then the best one of all of these routes which has not already been identified as one of the $(j - 1)$ best routes. To speed up the entire algorithm and to minimise storage space, only subroutes need to be stored along with the information about which route it is deviating from, instead of storing the entire route.

3.4 Planning Aircraft Taxiing Trajectories via a Population Adaptive Immune Algorithm (PATT-PAIA)

Chen and Stewart (2011) proposed an immune inspired multi-objective optimisation algorithm which utilised a physics-based aircraft dynamic model to search for different taxiing trajectories for a given route. Each of these trajectories represents a different trade-off between taxi time and fuel consumption. This algorithm has been extended in this research to incorporate time-window constraints. In the following, the PATT-PAIA is briefly discussed. Interested readers are referred to Chen and Mahfouf (2006) and Chen and Stewart (2011) for more details.

Algorithm 2 shows the proposed PATT-PAIA with time-window constraints at a glance. As discussed in Section 2.2, the entire airport layout is represented as a directed graph. Time-windows, corresponding to edges between the vertices, represent when a part of a taxiway is not used by any other aircraft. Unlike the vertices shown in Figure 1, PATT-PAIA only considers junctions and divides the entire taxi route of an aircraft into segments. Each of these segments may contain several intermediate vertices. There are two types of segments, namely straight segments and turning segments. The maximum speed for a straight segment is 30 knots and the speed during turns is fixed to 10 knots. For the straight segment, there are four consecutive transitional phases for an aircraft: a) acceleration phase, b) constant speed phase, c) deceleration phase and d) fast deceleration phase. By adjusting the acceleration and deceleration rates and the switching points between the phases, one can obtain different speed profiles and their corresponding fuel consumption indices using the aircraft categorisation (see Section 2.3) and the second objective function (see Section 3.1).

To obtain a good approximation of the pareto-front of both objectives, PATT-PAIA is devised as follows. First, an initial population pool is randomly generated around a feasible solution (trajectory) which fulfils all of the time-window constraints (lines 1 and 2). This feasible solution is generated using a heuristic to find the most time-efficient trajectory which takes into account a more realistic speed profile. Then, non-dominated sorting will be

utilised to distinguish between dominated and non-dominated solutions (line 3). The loop (lines 4-9) iterates Gen times to improve the current pareto-front and aims to have the pareto-front equally spread. One of the good solutions will be randomly selected in order to calculate its distance from the rest of the solutions, which defines the fitness of each solution (line 5). Based upon the fitness values, good solutions will be selected to be cloned with a higher probability (line 6). The cloned solutions will be mutated with small variation steps, to locally search the neighbourhood (line 7). On the other hand, bad solutions will only be cloned once and will be subjected to greater mutation in order to explore more of the search space. Constraint handling is used each time, immediately after new solutions are generated, to check whether the mutated solutions are still within the feasible bounds, as discussed by Chen and Stewart (2011), and also to calculate the arrival time at each vertex to see whether the current solution still complies with the given time-windows (line 8). The mutated solutions and the previous solutions will be combined and passed to a reselection stage, so that the best solutions survive into the next generation (line 9). The output of the algorithm is an approximation of the pareto-front for a given route of an aircraft fulfilling the given time-windows (line 10).

Algorithm 2: PATT-PAIA

- 1 Approximate the most time-efficient speed profile for each segment satisfying all given time-windows
 - 2 Randomly generate additional initial solutions in its neighborhood
 - 3 Objective evaluation and non-dominated sorting
 - 4 **foreach** of the Gen iterations **do**
 - 5 Fitness evaluation: evaluate the fitness of the candidate solutions
 - 6 Selection and cloning: selection of good solutions based on the fitness values; selected solutions are cloned
 - 7 Mutation: variation of clones
 - 8 Constraint handling: check if the mutated solutions meet all time-windows and feasible bounds
 - 9 Objective evaluation and reselection: reselect good solutions from the combined solutions based on non-dominated sorting
 - 10 **Output:** Approximation of pareto-front of taxi trajectories
-

The feasible solution which is generated using the heuristic in line 1 not only speeds up the search of the PATT-PAIA, but also guarantees at least one feasible solution at the end. To spread the solutions more equally between the most time-efficient route and the most fuel-efficient route, a possible improvement would be to also generate the most fuel-efficient trajectory.

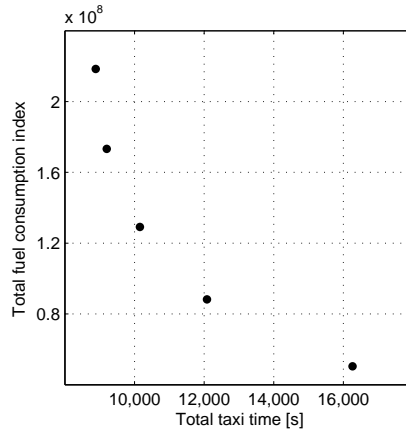


Fig. 4 Global discretised pareto-front of 57 aircraft

4 Results and discussions

The following section is split into three parts. It starts with the visualisation of the global pareto-front, continues with an analysis across an entire week's operations and finishes with a sensitivity analysis related to a different objective function.

4.1 Global pareto-front

Analysis showed that, by using the three best routes ($k = 3$, in Algorithm 1), the procedure can find very good approximations to the global pareto-front. Similarly, the number of iterations which were needed by the PATT-PAIA to find a good approximation was also tested, and this value was then fixed to $Gen = 40$. The execution time of Algorithm 1 on a personal computer (Intel Core 2 Duo, 3GHz, 2GB RAM) with these settings is around 100 minutes for one data point on the global pareto-front (inner loop of the algorithm), for a dataset of 57 aircraft. Due to the long execution time, the focus of this research was restricted to analysing the busiest time of the day, which was the hour between 11am and noon. Figure 4 shows the global pareto-front with five discretised values. The point at the top indicates the analysis where each of the 57 aircraft was taxiing as time-efficiently as possible, whereas the furthest point to the right indicates the analysis where each aircraft was taxiing as fuel-efficiently as possible.

4.2 Analysis over a week's operations

An analysis was performed to see how consistent the results were over a week's operations of the hour between 11am and noon. For this purpose, only the two

extreme cases were studied instead of analysing the entire global pareto-front. Table 2 shows the results for the analysed week. The values are reported as the average values per aircraft, since the number of aircraft in the dataset varied between 46 and 63 over the week. The first column restates the extreme values from Figure 4. The last row of each block highlights the growth from the best solution to the other extreme value. The values seem consistent and the dataset for Monday seemed to be a good representation of an average day.

Table 2 Analysis over a week’s operations of the hour between 11am and noon with the focus upon the extreme values

	Mon	Tue	Wed	Thu	Fri	Sat	Sun
∅ Taxi time							
Time-efficient [s]	156	157	128	174	152	165	154
Fuel-efficient [s]	285	293	214	320	292	316	295
Growth	83%	87%	67%	84%	92%	91%	91%
∅ Fuel cons. index							
Time-efficient [$\times 10^3$]	3832	3291	3492	4002	3173	3754	3718
Fuel-efficient [$\times 10^3$]	884	762	742	922	700	849	823
Growth	334%	332%	371%	334%	353%	342%	352%
Number of aircraft	57	58	46	58	56	63	52

4.3 Different objective function

A further experiment was run to see how sensitive the algorithm was to the fuel-related objective function. For this purpose, the setting from Stettler et al. (2011) was used as a replacement for the second objective function. As stated in Section 2.1, two different phases were used, one for acceleration and one for taxiing with constant speed, deceleration or holding. A stepwise function was utilised to measure the fuel used (in kg) during taxiing, based on the fuel flow coefficient. The acceleration phase was set so that an aircraft burns 10% of the maximal fuel flow and 5.5% in the other phase. With such a setting, the PATT-PAIA is forced to always accelerate with the maximal acceleration rate and mainly controls the length of these acceleration phases. Table 3 shows the results for the Monday dataset and is structured in the same way as Table 2, with the only difference being that the other fuel-related objective function is used instead.

It can be seen that the trade-off analysis very much depends upon the second objective function. With the approach based on the research by Stettler et al. (2011), there seems to be hardly any difference between optimising an aircraft’s trajectory for time-efficiency or fuel-efficiency. The actual values cannot be directly compared to those in Table 2, but, due to the fact that fuel flow of the engines is approximately proportional to the engine thrust setting (Wey et al. 2006, p. 7), the growth between the most extreme solutions indicates the potential for using a trade-off analysis.

Table 3 Analysis with the focus upon the extreme values where the fuel-related objective function was replaced in reference to the research by Stettler et al. (2011)

	Different objective function
∅ Taxi time	
Time-efficient [s]	155.5
Fuel-efficient [s]	156.7
Growth	0.8%
∅ Fuel flow	
Time-efficient [kg]	23.8
Fuel-efficient [kg]	23.5
Growth	1.2%

Calculations highlight why the trade-off is sensitive. On the example of the Airbus A320, the rolling resistance is 11.48 kN, which is around 5.1% of the total rated output. Hence, the two fuel-related objective functions behave very similarly in phases of constant speed, deceleration or during a hold. However, during acceleration with maximal acceleration rate, the physics-based objective function adds the rolling resistance and the force of acceleration which is 88.0 kN in total. This is around 39.3% of the total rated output, which is considerably more than the 10% of the other function. More research is needed to better understand in detail the fuel burn during taxiing, before the question can be answered of whether there is actually a trade-off between time-efficient and fuel-efficient taxiing, and to quantify any potential trade-off.

5 Conclusions and future research directions

A new model was developed to analyse the trade-off between two different, potentially conflicting, objective functions for the ground movement problem at airports. A sophisticated combination of two algorithms has enabled the development of a framework to run simulations for different datasets, and to perform sensitivity analysis. The first utilised algorithm finds the best possible routes for an aircraft at an airport and the second algorithm finds an approximation of the pareto-front for different speed profiles for each of these routes, in relation to the given objective functions.

Historic data from Zurich Airport was utilised for the analysis. The objective functions consisted of the taxi time, which is a commonly used measure, and a physics-based function related to the force needed from the aircraft engines during taxiing. The results show that the integrated procedure is able to tackle this hard problem in a comparatively efficient way. Furthermore, results seem to be consistent over several days. Sensitivity analysis has highlighted that the potential trade-off between the two objectives depends very much upon the actual modelling of the fuel-based objective function, which appears not to be well understood at the moment. Future research is mainly needed in better understanding the details of the fuel usage during taxiing and the standard practices of pilots during taxiing. Such insights could then be used

to more clearly show the influence of operational and environmental targets during the taxiing process.

Acknowledgements The authors wish to thank the Engineering and Physical Sciences Research Council (EPSRC) for providing the funding which made this research possible. We would also like to thank Flughafen Zürich AG who provided the real dataset and the trained commercial pilots who provided us with helpful insights into taxi behaviours at large airports, especially those who work for Swiss International Air Lines AG.

References

- Atkin JAD, Burke EK, Greenwood JS (2010a) TSAT allocation at London Heathrow: The relationship between slot compliance, throughput and equity. *Public Transport* 2(3):173–198, DOI 10.1007/s12469-010-0029-2
- Atkin JAD, Burke EK, Ravizza S (2010b) The airport ground movement problem: Past and current research and future directions. In: *Proceedings of the 4th International Conference on Research in Air Transportation (ICRAT 2010)*, Budapest, Hungary, pp 131–138
- Atkin JAD, Burke EK, Greenwood JS (2011a) A comparison of two methods for reducing take-off delay at London Heathrow Airport. *Journal of Scheduling* 14(5):409–421, DOI 10.1007/s10951-011-0228-y
- Atkin JAD, Burke EK, Ravizza S (2011b) A more realistic approach for airport ground movement optimisation with stand holding. In: *Proceedings of the 5th Multidisciplinary International Scheduling Conference (MISTA 2011)*, Phoenix, Arizona, USA
- Balakrishnan H, Jung Y (2007) A framework for coordinated surface operations planning at Dallas-Fort Worth International Airport. In: *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, Hilton Head, SC, USA
- Burgain P, Feron E, Clarke JP (2009) Collaborative virtual queue: Benefit analysis of a collaborative decision making concept applied to congested airport departure operations. *Air Traffic Control Quarterly* 17(2):195–222
- Cassell R, Evers C (1998) Development of airport surface surveillance performance requirements. In: *Proceedings of the AIAA/IEEE/SAE 17th Digital Avionics Systems Conference (DASC)*, vol 2, DOI 10.1109/DASC.1998.739821
- Chen J, Mahfouf M (2006) A population adaptive based immune algorithm for solving multi-objective optimization problems. In: *Artificial Immune Systems, Lecture Notes in Computer Science*, vol 4163, Springer Berlin, pp 280–293
- Chen J, Stewart P (2011) Planning aircraft taxiing trajectories via a multi-objective immune optimisation. In: *Proceedings of the 7th International Conference on Natural Computation (ICNC 2011)*, Shanghai, China, vol 4, pp 2235–2240, DOI 10.1109/ICNC.2011.6022587

- Climaco JCN, Martins EQV (1982) A bicriterion shortest path algorithm. *European Journal of Operational Research* 11(4):399–404, DOI 10.1016/0377-2217(82)90205-3
- Deau R, Gotteland JB, Durand N (2009) Airport surface management and runways scheduling. In: *Proceedings of the 8th USA/Europe Air Traffic Management Research and Development Seminar*, Napa, CA, USA
- Dijkstra EW (1959) A note on two problems in connexion with graphs. *Numerische Mathematik* 1:269–271, DOI 10.1007/BF01386390
- European Commission (2011) *Flightpath 2050, Europe’s vision for aviation. Report of the High Level Group on Aviation Research*, DOI 10.2777/50266
- García J, Berlanga A, Molina JM, Casar JR (2005) Optimization of airport ground operations integrating genetic and dynamic flow management algorithms. *AI Communications* 18(2):143–164
- Gotteland JB, Durand N, Alliot JM (2003) Handling CFMU slots in busy airports. In: *Proceedings of the 5th USA/Europe Air Traffic Management Research and Development Seminar*, Budapest, Hungary
- Herrero JG, Berlanga A, Molina JM, Casar JR (2005) Methods for operations planning in airport decision support systems. *Applied Intelligence* 22(3):183–206, DOI 10.1007/s10791-005-6618-z
- ICAO (2008) *International standards and recommended practices, annex 16, environmental protection: Aircraft engine emissions*. International Civil Aviation Organization
- Khadilkar H, Balakrishnan H (2011) Estimation of aircraft taxi-out fuel burn using flight data recorder archives. In: *Proceedings of the AIAA Guidance, Navigation, and Control Conference*
- Lawler EL (1972) A procedure for computing the k best solutions to discrete optimization problems and its application to the shortest path problem. *Management Science* 18(7):401–405
- Lesire C (2010) Iterative planning of airport ground movements. In: *Proceedings of the 4th International Conference on Research in Air Transportation (ICRAT 2010)*, Budapest, Hungary, pp 147–154
- Marín Á (2006) Airport management: Taxi planning. *Annals of Operations Research* 143(1):191–202, DOI 10.1007/s10479-006-7381-2
- Marín Á, Codina E (2008) Network design: Taxi planning. *Annals of Operations Research* 157(1):135–151, DOI 10.1007/s10479-007-0194-0
- Morris KM (2005) Results from a number of surveys of power settings used during taxi operations. Tech. rep., ENV/KMM/1126/14.8, British Airways
- Nikoleris T, Gupta G, Kistler M (2011) Detailed estimation of fuel consumption and emissions during aircraft taxi operations at Dallas/Fort Worth International Airport. *Transportation Research Part D: Transport and Environment* 16(4):302–308, DOI 10.1016/j.trd.2011.01.007
- Pesic B, Durand N, Alliot JM (2001) Aircraft ground traffic optimisation using a genetic algorithm. In: *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, San Francisco, California, USA, pp 1397–1404

- Rappaport DB, Yu P, Griffin K, Daviau C (2009) Quantitative analysis of uncertainty in airport surface operations. In: Proceedings of the AIAA Aviation Technology, Integration, and Operations Conference
- Ravizza S, Atkin JAD (2011) Exploration of the ordering for a sequential airport ground movement algorithm. Tech. Rep., 1543, University of Nottingham
- Roling PC, Visser HG (2008) Optimal airport surface traffic planning using mixed-integer linear programming. *International Journal of Aerospace Engineering* 2008(1):1–11, DOI 10.1155/2008/732828
- Schüpbach K, Zenklusen R (2011) Approximation algorithms for conflict-free vehicle routing. In: Proceedings of the 19th Annual European Symposium on Algorithms (ESA), DOI 10.1007/978-3-642-23719-5_54
- Smeltink JW, Soomer MJ, de Waal PR, van der Mei RD (2004) An optimisation model for airport taxi scheduling. In: Proceedings of the INFORMS Annual Meeting, Denver, Colorado, USA
- Stettler MEJ, Eastham S, Barrett SRH (2011) Air quality and public health impacts of UK airports. Part I: Emissions. *Atmospheric Environment* 45(31):5415–5424, DOI 10.1016/j.atmosenv.2011.07.012
- Wey CC, Anderson BE, Hudgins C, Wey C, Li-Jones X, Winstead E, Thornhill LK, Lobo P, Hagen D, Whitefield P, Yelvington PE, Herndon SC, Onasch TB, Miake-Lye RC, Wormhoudt J, Knighton WB, Howard R, Bryant D, Corporan E, Moses C, Holve D, Dodds W (2006) Aircraft Particle Emissions eXperiment (APEX). Tech. rep., NASA
- Yen JY (1971) Finding the k shortest loopless paths in a network. *Management Science* 17(11):712–716