

Airline Planning Benchmark Problems

Part II: Passenger Groups, Utility and Demand Allocation

Kerem Akartunali^a, Natasha Boland^b, Ian Evans^c, Mark Wallace^d, Hamish Waterer^{b,d,*}

^aManagement Science Department, University of Strathclyde, Glasgow G1 1QE, Scotland, United Kingdom

^bSchool of Mathematical and Physical Sciences, University of Newcastle, Callaghan NSW 2308, Australia

^cConstraint Technologies International, Level 7 224 Queen St, Melbourne VIC 3000, Australia

^dFaculty of Information Technology, Monash University, Caulfield VIC 3145, Australia

Abstract

This paper is the second of two papers entitled “Airline Planning Benchmark Problems”, aimed at developing benchmark data that can be used to stimulate innovation in airline planning, in particular, in flight schedule design and fleet assignment. The former has, to date, been under-represented in the optimization literature, due in part to the difficulty of obtaining data that adequately reflects passenger choice, and hence schedule revenue. Revenue models in airline planning optimization only roughly approximate the passenger decision process. However there is a growing body of literature giving empirical insights into airline passenger choice. Here we propose a new paradigm for passenger modelling, that enriches our representation of passenger revenue, in a form designed to be useful for optimization. We divide the market demand into market segments, or *passenger groups*, according to characteristics that differentiate behaviour in terms of airline product selection. Each passenger group has an origin, destination, size (number of passengers), departure time window, and departure time utility curve, indicating willingness to pay for departure in time sub-windows. Taking as input market demand for each origin-destination pair, we describe a process by which we construct realistic passenger group data, based on analysis of empirical airline data collected by our industry partner. We give the results of that analysis, and describe 33 benchmark instances produced.

Keywords: Airline planning, passenger choice models, revenue models, benchmark data

1. Introduction

This paper is the second of two papers entitled “Airline Planning Benchmark Problems”. Our primary goal in these papers is to stimulate and facilitate further research in airline planning, by developing airline planning benchmarks and, more importantly, a process for constructing benchmarks that reflect the requirements of different kinds of airlines and airline networks.

Such a network-specific benchmark makes it possible, for the first time, for an airline to set up a fair and objective comparison between different algorithms and approaches, without the need to hide commercially sensitive aspects of the data. Moreover by making it possible to generate such benchmarks quickly and cost-effectively, using relatively few parameters, we hope to encourage every airline, large or small, to use such benchmarks to choose the best algorithm for the job at hand for the good of the airline and its passengers. Finally the availability of such benchmarks will enable academic researchers to compare novel approaches and algorithms against the state-of-the-art and encourage the flow of ideas not only between academics but also across to industry.

In these two papers, we focus on the development of airline demand data, which is the starting point for generating benchmarks for every aspect of an airlines operation from schedule design, through all the stages of airline planning and scheduling, to the day of operation.

In this second paper, we take a set of origin-destination demands across a network, and show how to break down the demand between an origin and destination into passenger groups across the day, with different tradeoffs between cost and time. The relatively simple groupings proposed here provide a realistic basis for generating schedules which match passenger demand and generate a high revenue in comparison with operational cost. Our aim is to provide “raw” realistic demand data that has not yet been manipulated by schedule choices or revenue management strategies, so that the effects of such strategies can be tested and compared.

The groupings and their basis are presented and justified in this paper. Nevertheless for any particular airline and origin-destination (OD) pair, if detailed passenger demand preference data is available then the parameters underlying the groups can be easily adjusted to match the data.

In what follows, we describe in detail the methodology we used to construct passenger group data for our benchmark problems, taking as input a flight network, and OD-pair demand, together with other design parameters and template data that we discuss in detail as required. Unlike Part I of this paper, Akar-

*Corresponding author

Email addresses: kerem.akartunali@strath.ac.uk
(Kerem Akartunali), natashia.boland@newcastle.edu.au
(Natashia Boland), ian.evans@contecint.com.au (Ian Evans),
mark.wallace@monash.edu (Mark Wallace),
hamish.waterer@newcastle.edu.au (Hamish Waterer)

tunali, Boland, Evans, Wallace, and Waterer (2011), in which an optimization model was developed to solve the inverse problem of deriving OD-demand from observed passenger loads on flight network arcs, the methodology here does not require sophisticated mathematics. However it does require some analysis and modelling of airline data, and careful explanation, so that the resulting benchmarks can be properly understood.

In the following section, Section 2, we survey passenger demand models, the availability of data that supports them, and their exploitation in schedule design algorithms. Section 3 motivates and explains the abstraction of passengers into passenger groups whose members share similar choice behaviour. Next, in Section 4, we discuss our key assumptions, and define necessary terminology and notation. We also give our template for travel time utility curves. In Section 5, we explain how we determine the number of passenger groups to create for each OD-pair, and how we define values for the time of day, fare and utility curve for each group. Section 6 describes how we partition the given OD-pair demand between passenger groups, and, finally, Section 7 describes a set of benchmark instances generated by our method.

2. Passenger demand models and data

2.1. Passenger demand models

Modelling of passenger demand for scheduled transport services is required when it is desired to optimise schedules or optimise revenue for a given schedule using revenue management. Here we will use the term demand to mean the number of passengers who desire to travel from a given origin to a destination, with each passenger potentially having a different utility for each offered service or combination of services that can provide this travel.

Passenger utility can vary in several dimensions:

- Fare classes on a given service
- Ability to book close to the time of departure, or make late changes to bookings
- Time or date of departure
- Time taken to travel from origin to destination (and the related travel route)
- Quality of service provided by a given transport operator
- Mode of transport

The first four of the above can affect choices between services offered by the same transport operator, while all of these can also affect choices between competing transport operators. Apart from the booking time frame, all of these dimensions are shared for the problems of schedule optimisation and revenue management.

A large proportion of passenger demand and utility modelling for airlines has been in the context of revenue management, which traditionally has relied on the *independent demand* model which assumes independent demand for each product, but the limitations of these product-based models are becoming more apparent (Van Ryzin, 2005). More recently there have

been efforts made to move from product-based models to more realistic models describing customer behaviour — for example models that consider demand for overlapping sets of products (Bront, Mendez-Diaz, and Vulcano, 2009).

Passenger demand and utility modelling is also very important in the airline schedule design process, and again the traditional assumption of independent demand (referred to as *flight leg independence* in this literature) has been shown to miss an important aspect of the fleet assignment problem, with more realistic models being developed to take into consideration interrelationships between itineraries using techniques such as a *spill* and *demand recapture* model (Lohatepanont and Barnhart, 2004) and a passenger flow time-space network (Yan and Tseng, 2002; Yan, Tang, and Lee, 2007).

2.2. Inaccessibility of passenger demand data

The development of optimisation algorithms that can make use of more complex customer behaviour models leads to another problem — how to obtain data about the actual demand and utility of passengers for a given situation. The primary data that airlines have to work from is historical booking records, but it is well known that this demand is censored (or constrained) by the presence of booking and capacity limits on past demands (McGill and Van Ryzin, 1999; Lohatepanont and Barnhart, 2004). Techniques for uncensoring (or unconstraining) this demand have traditionally relied on the independent demand assumption (often referred to as single-class demand unconstraining methods), but more recent work has included multi-class and multi-flight spill methods (Ratliff, Rao, Narayan, and Yellepeddi, 2008). Each possible itinerary in a market receives a rating, based on an industry standard known as the *Quantitative Share Index (QSI)*¹, for measuring the “attractiveness” of an itinerary, taking into account time of day of departure, length of trip, and number of connections. These methods require customer choice models that include the concept of passenger utility that takes into account factors related to the QSI.

A second difficulty with estimating demand is the presence of a large number of rare itineraries in large hub-and-spoke networks. The number of passengers travelling on many itineraries is very low, but together these can form an important part of revenue (McGill and Van Ryzin, 1999). A similar problem is the lack of data on departure time preferences on sectors with low demand, as there will be few flights per day. These difficulties are typically resolved by some form of data aggregation, and here it is important to select aggregation groups so that important characteristics related to passenger utility are retained.

Lastly, there is the difficulty of taking competition into account when estimating demand. Current demand unconstraining approaches (see e.g. Ratliff et al., 2008) classify competition for demand using three categories:

- Travel on the host airline
- Travel on a competing airline or
- No travel on this date (or *do not fly* demand).

¹Also sometimes referred to as the Quality Service Index.

This is normally further simplified using the open competitor assumption that competing airlines have available capacity, so that it is assumed that the *do not fly* demand is negligible. This approach uses the market share to estimate the amount of demand on competing airlines as more detailed information is hard to obtain, and ignores effects from competing utility except between different products offered by the host airline. This approach suffers, however, when competition occurs from dissimilar competitors such as a full-service airline competing with a low-cost airline or with other low cost transport modes such as buses. Low-cost competitors compete for a different type of demand with a utility that is sensitive primarily by cost and less by other quality of service indicators such as departure time, trip duration and number of stops. In the leisure market, demand may even spill from one OD pair to another based on price if similar holiday experiences are offered at different destinations.

2.3. Passenger demand and schedule design

Yan and Tseng (2002) formulate schedule design as a kind of network design problem in a time-space network, in which the passengers are represented as explicit multicommodity flows in the network designed by the aircraft variables. There is a commodity for each OD pair, and the total commodity flow is bounded above by the market demand. They use a standard ticket price for each commodity on each flight leg, and apply a “holding cost” to time passengers spend on the ground at intermediate stops, dependent only on OD pair and time spent. The latter is the only way in which passenger utility with respect to alternate itineraries is considered. Yan et al. (2007) extend this work to embed a more sophisticated passenger choice model, in which the presence of competition in the market is modelled by assuming passengers will be unavailable to the airline if they have to wait too long to depart. The passenger loss parameters require a complicated calculation taking into account airline preferences, ticket prices, flight frequency and travel time, which is nonlinear in the decision variables, and seems to be decoupled from subsequent flights in a multi-flight itinerary (only loss at the port of origin is considered). Lohatepanont and Barnhart (2004) combine fleet assignment with incremental flight schedule design, and take into account passenger choice via a spill-and-recapture model based on earlier work of Kniker (1998). The QSI (introduced in Section 2.2 above) is converted into a probability that passengers will be recaptured by an itinerary if spilled from the preferred itinerary for the market. The measure is independent of fare category, and depends only on market. We note that itinerary-based market share indicators are still a subject of current research, see Coldren, Koppelman, Kasturirangan, and Mukherjee (2003), for example. However like the QSI, Coldren et al. (2003), treat the market (OD pair) as whole, e.g., only average fares for an airline are considered.

3. Introducing passenger groups

3.1. Grouping passengers by their choice behaviour

A common feature of passenger choice models in schedule design optimization is that passengers with the same OD

pair are treated identically. Furthermore the effects of passenger choice on revenue are modelled in a relatively simple way: (i) revenue penalty for long connections in multi-flight itineraries (Yan and Tseng, 2002), (ii) lost passengers due to long waits for departure (Yan et al., 2007) and (iii) lost passengers at itinerary-based rates with recapture options (Lohatepanont and Barnhart, 2004). However both common sense and emerging research indicate that the passenger market is segmented, with different groups of passengers with the same OD pair showing different passenger choice behaviours, and exhibiting different utility functions with respect to a range of factors. For example, Walker (2006) analyse empirical data to show that passengers with different time-of-day preferences show different utility functions with respect to schedule delay. Walker (2006) also indicates that whether the traveller is business or leisure, or whether the traveller is more sensitive to their arrival time rather than departure time, can affect their utility functions for schedule delay. Koppelman, Coldren, and Parker (2008) also finds differences in schedule delay utility for business versus leisure travellers, and furthermore observes differences depending on whether the trip is out-bound or in-bound. Walker (2006) finds that the business and leisure segments put different dollar figures on factors such as total time for the trip, number of stops, aircraft changes, and so on. This is confirmed by the work of Garrow, Jones, and Parker (2007), who also showed that departure versus arrival sensitivity can affect preferred travel time, with departure-sensitive travellers showing strong morning and evening peaks, with arrival-sensitive passengers having a mid-day peak.

We conclude that the airline passenger market is segmented, and propose that schedule design optimization should be based on revenue models that better reflect this market segmentation.

3.2. Identifying passenger groups

Designing an airline schedule around the airline product selection behaviour of each individual potential customer is obviously impractical, both because of the number of individuals and the unavailability of data at this level of detail. It is necessary, therefore, to abstract airline demand into *passenger groups*, according to characteristics that differentiate their choice behaviour.

In short, we make several suggestions for simplification to make passenger demand modelling practically tractable:

- passengers are grouped into those with similar utility functions
- these groups can be reasonably inferred for each OD pair from the OD demand and the distance and direction of the destination from the origin
- utility translates directly into a price passengers are willing to pay as a function of time
- the time-of-travel function take one of three forms:
 - asymmetric preferring before a certain time (increasing up to that time then dropping to zero)

- asymmetric preferring after a certain time (zero before the time and decreasing afterwards)
- symmetric about a certain preferred time
- the time-of-travel function is a step function with three steps, associated with the three forms itemised previously.

We argue that this can be naturally modelled in an optimization setting. For example, in multi-commodity flow models, such as those of Yan and Tseng (2002) and Yan et al. (2007), a commodity for each passenger group could be defined instead of for each OD-pair. In itinerary-based models such as those of Lohatepanont and Barnhart (2004), an extra subscript would be required for passenger-itinerary variables. This is likely to increase model size. However we propose that departure time window should be a defining characteristic of a passenger group, so the subnetwork for each group will be necessarily limited. This will to some extent mitigate increases in model size. Furthermore, dealing with larger models is a challenge that drives the optimization research community forward, so we hope this proposal, with its associated benefits for improved revenue modelling, will stimulate such research. Indeed, we have made a small start on such work: our companion paper Akartunalı, Boland, Evans, Wallace, Waterer, and Smith (2010) studies an integrated airline schedule design and fleet assignment problem, for which passenger groups provide the necessary detailed information.

We have taken a pragmatic approach, applying a mix of common sense and the insights from the empirical literature, together with the insights of our industry partner and their analysis of airline passenger data (Evans, 2010). We view this as a first step, that can be used to test the concept. Thus we focus on what we believe are the key features. Of course a passenger group must have an OD-pair, and represent a specified number of passengers. The business versus leisure divide features in most recent analysis; here we do not explicitly define those as characteristics, but instead associate a fare with each passenger group, representing what they are willing to pay to travel at their preferred time. Since preferred time of day has been identified as having an important impact on schedule delay utility, (see e.g. Walker, 2006), we also characterize passenger groups by their departure time window, and a travel time utility curve, which subtracts schedule delay disutility from the fare the passenger is willing to pay to travel at their preferred time.

Clearly, substantial new market research would be required to confirm that our method yields the “best” way of representing passenger groups.

3.3. Basis of our approach to grouping passengers

In order to investigate the performance and scaling of transport optimisation algorithms, it is desirable to be able to use realistic passenger demand and fare profiles. However, passenger demand and fare profile data, even when it exists, is usually confidential as it has significant commercial implications. Data derived from passenger surveys would be ideal for our purposes, but the literature in this area is quite sparse, (see e.g. Walker, 2006; Garrow et al., 2007).

The passenger groups introduced in this paper are based on the UK Department of Transport National Travel Survey 2006 (UK Department of Transport, 2006, Table 6.6b). This table shows the morning and afternoon peaks, and a single broad spread centred on the middle of the day. Time constraints are tighter for the first two types, and such passengers are willing to pay a higher fare to travel at their preferred time compared to the third type.

Garrow et al. (2007) supports the above findings, notes that most passengers are departure-time-sensitive, and also finds that the minority arrival-time-sensitive passengers are generally mid-day travellers, speculating that hotel check-in times are the cause.

Fare data is accessible for forthcoming flights, and passenger load data is also publically available in some cases after the flight. However, both types of data can only represent preferences for the offered itineraries.

For economic reasons, airlines must keep their high capital-cost assets working continuously, and so must fly during the middle of day, when passenger survey data suggests demand is low. The airlines handle this via revenue management, seeking to attract budget-conscious *time-insensitive* passengers onto flights in the middle of the day, and by attempting to locate middle-of-the-day capacity on flights between port pairs for which there is low frequency of service, so passengers have little choice. We note that this makes it more difficult to infer passenger preferences from the above offered itineraries.

Cognizant of this limitation, one of the authors analysed publically available airline schedules and fares, to infer passenger demand (Evans, 2010). This paper analysed schedules operating at the end of 2007 for 44 carriers flying to 994 different airports worldwide. The data used only included scheduled services, and did not include actual passenger data, so only the available capacity of each leg was used in the analysis. This analysis, described in more detail in Section 6.2 below, reveals the start of the morning peak to be based on a departure time of around 06:00. Interestingly, the evening peak is split so that any legs due to arrive after about midnight will instead be deferred so that they arrive at about 05:00 the following day. The paper also reports an analysis of fare data on two typical busy mid-week days on a single short-haul route with a high number of flights.

Evans (2010) concludes that the data satisfies the following hypotheses:

- Three types of time-sensitive passenger groups exist, with preferred flying times consisting of a morning peak, an evening peak and a broad spread around the middle of the day.
- The preferred travel times for these passengers are based mainly on local departure time, but there is an effect related to arrival time for long west-to-east legs for evening departures.

Based on independent studies, (see e.g. Commission of the European Communities, 2007; Civil Aviation Authority, 2008), we assume morning and evening passenger groups consist of about 25% business travellers and 75% leisure travellers, and

that the proportion of business travellers in any midday group is negligible.

4. Introducing utility curve templates

This analysis is restricted to short-haul air travel where most flying occurs during the day and evening, and there is similar flying from one day to the next. In this context, we have noted that most time-sensitive passengers fall into one of three types:

- Those wishing to travel early in the morning, some of whom are business travellers who are to work at the destination for a day and return in the evening;
- Those wishing to travel early in the evening, some of whom are business travellers returning from a day trip or business travellers travelling the night before to be ready for an early morning meeting; and
- Those preferring to travel at a time that does not involve early rising or late dining, and thus with a preference to travel around the middle of the day.

We conjecture a fourth type of budget-conscious passenger: those who are unconcerned with the time of flying (time-insensitive), as long as the fare is low. The presence of this type of passenger is inferred from the revenue management strategies mentioned in Section 3.3 above that are in place at most airlines around the world.

We refer to these four types of passenger as *morning*, *evening*, *midday* and *time-insensitive*, respectively.

It would be interesting to try and map these passenger groups onto real booking classes. However, the number of booking classes for a flight is quite large, and the level of granularity of the groups is quite coarse so that passengers in the same group would have to be mapped to multiple (possibly up to a dozen) booking classes. On the other hand a broad-brush mapping (to business/economy/budget) is more or less equivalent to published estimates such as those of Commission of the European Communities (2007) and Civil Aviation Authority (2008).

In the first three cases, for the sake of simplicity, we assume passengers are departure-time sensitive. In the context of short-haul operations and midday travel, when time zone impacts are relatively minor, we believe preferred arrival time can reasonably be “mapped back” to preferred departure time, and so restricting our attention to departure-time sensitive groups is a reasonable approximation (see also Evans, 2010).

For each of the four passenger types, we propose a different shape of travel time utility curve. The shapes for the first three are shown in Figure 1. Moving from left to right across the time axis, labelled t , we see the morning, midday and evening curves respectively. In all cases, we use a piecewise constant function with three pieces, defined over a time window. For morning passengers, since we assume these are travelling for some scheduled activity at their destination, we assume their preference is to travel at the right time to make the activity, but could be willing to travel earlier if the price was right. For




Figure 1: Travel time utility curves for morning, afternoon and evening passenger groups

evening passengers, we assume they need to complete some activity at their origin prior to travel (e.g. completing the activity they travelled for before returning home, or completing a day at their home office before travelling for an early start the following day), and so reverse the shape of the curve for morning passengers. Midday passengers simply have a preferred time of travel, and are willing to travel either earlier or later, for the right price. The curve used for time-insensitive passengers is simply a flat line, indicating that they will only fly if the price is right, and then would be willing to fly at any time for that price.

Before discussing our assumptions about these curves further, we first define some notation and terminology. Since we use a daily setting, we let \mathcal{T} denote the length of a day in time units, e.g. minutes. So time parameters for passenger groups are chosen from the interval $[0, \mathcal{T}]$. G denotes the index set of all passenger groups, each $g \in G$ having an origin o_g and destination d_g . There are η_g passengers in group $g \in G$. Each group $g \in G$ also has an earliest departure time, t_g^s , and latest departure time, t_g^e , defining the start and end of the time window $[t_g^s, t_g^e]$ in which the passengers from this group are willing to travel. Each group $g \in G$ has a preferred interval for travel, $[\tilde{t}_g^s, \tilde{t}_g^e] \subseteq [t_g^s, t_g^e]$, and a fare they are willing to pay to travel in that interval, ρ_g . We call these the *peak interval* and *peak fare* respectively; these correspond to the peaks in the utility curves shown in Figure 1.

Given the time window, peak interval and peak fare for a passenger group, the rest of the utility curve is specified by a fixed formula, for each passenger type. We take the value of time to be $v^B = \$68.97$ and $v^L = \$19.64$ per hour for business and leisure travellers respectively, based on the study of Garrow et al. (2007), and thus the value of time for morning and evening groups to be on average $v = 0.25v^B + 0.75v^L$, and v^L for midday groups. We assume each utility curve has identical width $u_g = \tilde{t}_g^e - \tilde{t}_g^s$ for its three pieces, and so can compute the formulae

$$P_g(t) = \begin{cases} \rho_g, & \tilde{t}_g^s \leq t \leq \tilde{t}_g^e \\ \rho_g - vu_g, & \tilde{t}_g^s - u_g \leq t \leq \tilde{t}_g^s \\ \rho_g - 2vu_g, & \tilde{t}_g^s - 2u_g \leq t \leq \tilde{t}_g^s - u_g \end{cases}$$

for morning groups g , a similar formulae for evening groups, while for midday groups g the formula is simply

$$P_g(t) = \begin{cases} \rho_g, & \bar{t}_g^s \leq t \leq \bar{t}_g^e \\ \rho_g - v^L u_g, & t \in [\bar{t}_g^s - u_g, \bar{t}_g^s] \cup [\bar{t}_g^e, \bar{t}_g^e + u_g]. \end{cases}$$

Of course time-insensitive passengers simply have utility function $P_g(t) = \rho_g$ for all $t \in [0, \mathcal{T}]$.

Other key assumptions we make in generating passenger group data for our benchmark instances are itemized below. If the user has additional data to hand, it would not be hard to relax these assumptions and still apply the methodology. Indeed our methodology is completely flexible, and allows passenger groups to be defined with any peak interval, non-zero utility width, shape and height. On a particular airline network for which accurate and detailed passenger choice behaviour was available, the precise peak interval, utility width, shape and height could be supplied by the user.

- We assume that each OD-pair has a single time-insensitive passenger group and that the proportion of time-insensitive passengers may be dependent on the OD pair, i.e. there is a constant for each OD pair (o, d) , which we refer to as $\phi_{o,d} \in [0, 1]$, so that the proportion of passengers in an OD market allocated to the time-insensitive group is $\phi_{o,d}$. In the benchmarks generated to date, we have used the same value of ϕ for all OD pairs, except those with low demand.
- We assume that for any given OD pair, there is no overlap between peak intervals for the passenger groups with that OD pair (other than with the time-insensitive group, which has a single “peak” interval covering the whole time period). For passenger groups of the same type, two overlapping groups have the same impact on schedule design as a single larger group so this assumption is valid for practical purposes. An overlap between two groups of a different type would imply that two different sets of passengers with the same preferred departure time have different utility values, with different revenue implications for schedule design. This impact is approximated in our approach by setting the utility of that time point to a value that reflects the proportion of passengers of each type. As a consequence of the no-overlap assumption, the peak interval (preferred time of travel) for morning, midday and evening groups uniquely defines the group. Ideally, we would have a more complex utility curve to represent the case where passenger groups of different types overlap. However at present there is insufficient data to support this level of detail.
- For each OD pair (o, d) , there is a basic length of time, which we call the *unit width* for that pair, denoted by $\mu_{o,d}$, so that for all passenger groups $g \in G$ with $o_g = o$ and $d_g = d$, $u_g = \mu_{o,d}$ for g a morning or evening group, and $u_g = 2\mu_{o,d}$ for g a midday group. The difference for midday travellers is due to our assumption that their preferred travel time arises from a desire to avoid early or

late travel, rather than from a need to travel before or after a particular activity, and our assumption that they are predominantly leisure travellers. Both factors imply greater flexibility for midday groups. The use of a unit width for each OD pair reflects the idea that passengers’ willingness to be flexible in the time they travel is a function of the frequency of service on the OD-pair; clearly passengers in remote locations expect to wait longer for service, whereas passengers travelling between major cities will expect a flight close to their preferred travel time. In other words, the more itineraries available to passengers from an OD pair, the narrower their utility curve will be.

We discuss the latter point further in the next section.

5. Creating passenger groups

In this section, we describe how passenger groups are created, and how the time of day data, as well as the data defining the travel time utility curves, are determined for each group.

5.1. Calculating unit width

In order to generate specific passenger groups and data values for those groups, we first calculate the unit width value $\mu_{o,d}$ for each OD pair (o, d) . As mentioned above, we postulate that passengers willingness to be flexible about when they travel, (and hence the unit width), depends on the frequency of service they expect to see on that OD pair. At first sight, one might think this is a function of $D_{o,d}$, the given market demand for OD pair (o, d) (total number of passengers wishing to travel (daily) from o to d , generated as described in Akartunali et al. (2011)). However this figure could be misleading. To illustrate, consider an OD pair (o, d) where the flight network includes direct flights, but where there is also a reasonable one-stop service via port i . Then the frequency of the one-stop service depends on the expected frequency for OD pairs (o, i) and (i, d) . This could get quite difficult to estimate, since these in turn depend on expected frequency on indirect paths, so for simplicity, we assume that frequency of one-stop services (or other indirect services), such as (o, i, d) , can be well estimated from the OD demands $D_{o,i}$ and $D_{i,d}$. We assume whichever is the smaller of these creates a bottleneck, and the frequency of service is predicated on that value. We thus solve a maximum flow problem for each OD-pair, to maximize the flow from o to d on all reasonable paths from o to d in the flight network (our definition of “reasonable” is given in Akartunali et al. (2011)), with arc capacities given by the values $D_{i,j}$ for each arc (i, j) . We call the resulting maximum flow value the *surrogate demand*, denoted by $\bar{D}_{o,d}$ for each OD pair (o, d) . This is converted to approximate the number of services that might be available to the OD pair, $N_{o,d}^{serv}$, by using the average aircraft capacity, while accounting for average load factors. We also scale down by $1 - \phi_{o,d}$, since airlines typically do not put on services to meet the low-budget demand, instead using revenue management to manipulate that demand to “fill gaps”.

To determine unit width $\mu_{o,d}$ from expected service frequency $N_{o,d}^{serv}$, we assume that passengers are willing to be flexible enough

to wait for the interval of time between services, without reduction in their utility, and simply assume an even spread of service over the available time. Since most airports do not operate 24 hours per day, we define the *airport operating interval*, which we take to be 18 hours, (e.g. if the port operates from 5am to 11pm), denoted by \mathcal{T}_{oper} , and assume services are spread across this interval. Thus unit width can be calculated simply as

$$\mu_{od} = \mathcal{T}_{oper} / N_{o,d}^{serv}.$$

5.2. Creating passenger groups

Since unit width determines the width of the peak interval, u_g , for each passenger group g , and since we make the assumption that no two passenger groups (other than the time-insensitive group) can have overlapping peak intervals, we see that creating the passenger groups for each OD pair is simply a matter of stepping through time. (We also assume that each time is in a peak interval for *some* passenger group. If the demand profile doesn't support this, then we would expect the passenger group to be assigned zero passengers in Section 6, so the group could be ignored.) Defining morning to end at time T_m and evening to start at time T_e , we can simply create a passenger group for each interval of length $\mu_{o,d}$ from the start of the day until T_m , then create a group for each interval of length $2\mu_{o,d}$ until T_e , and finish by creating a group for each interval of length $\mu_{o,d}$ until the end of the day. (In addition, we must create the single time-insensitive group: this is explained in Section 4.) However depending on how nicely \mathcal{T} , T_m and $T_e - T_m$ divide by $\mu_{o,d}$, this could lead to a somewhat skewed collection of passenger groups. So instead we start at the centre of the day, and work outwards. Our algorithm is specified as Algorithm 1. It is helpful to note that if the peak interval $[\bar{t}_g^s, \bar{t}_g^e]$ falls so that its mid-point is at time T_m or earlier, it is deemed a morning group, at time T_e or later an evening group, and otherwise it is a midday group. The first while loop calculates the time characteristics for intervals before midday, taken to be $(T_m + T_e)/2$, and the second while loop does it for after midday intervals. We use A to denote the set of morning ("am") groups, M to denote the set of midday groups, and P to denote the evening ("pm") groups created for an OD pair. Note that we only explicitly describe calculation of the peak interval parameters \bar{t}_g^s and \bar{t}_g^e for each group g , since all other times, including u_g , can be calculated from these or from $\mu_{o,d}$ using the formulae given in Section 4 and the knowledge of which type of group it is.

As Algorithm 1 is stated, it creates passenger groups across the whole time interval $[0, \mathcal{T}]$. However, as mentioned earlier, some airports may have constrained operating hours, and this may affect passenger expectations for travel. In such cases, one could just take zero to be first time at which passengers may be willing to travel, and replace \mathcal{T} in Algorithm 1 with the length of the time interval over which passengers wish to travel.

We note that this algorithm and some of our assumptions might need to be adjusted for longer flights, particularly in a west-to-east direction, when changes in time zones would seem likely to affect passenger preferences. We discuss this point further in Section 6.

Data: Unit width $\mu_{o,d}$ for each OD pair (o, d)

Result: The set of passenger groups G , together with time characteristics of each group $g \in G$

Initialize $g := 0$;

forall OD pairs (o, d) **do**

Set $A := \emptyset$, $M := \emptyset$ and $P := \emptyset$;

Set $g := g + 1$, $o_g := o$ and $d_g := d$;

Create the time-insensitive group:

set $\bar{t}_g^s := 0$ and $\bar{t}_g^e := \mathcal{T}$;

if $\mu_{od} \leq \frac{\mathcal{T}}{2}$ **then**

Set $g := g + 1$, $o_g := o$ and $d_g := d$;

Create a first midday group:

calculate $t := \frac{T_m + T_e}{2}$;

set $\bar{t}_g^s := t - \mu_{od}$ and $\bar{t}_g^e := t + \mu_{od}$;

add g to M ;

Set $t := t - \mu_{od}$;

while $t > 0$ **do**

Set $g := g + 1$, $o_g := o$ and $d_g := d$;

if $t - \mu_{od} \geq T_m$ **then**

Create another midday group:

set $\bar{t}_g^s := t - 2\mu_{od}$ and $\bar{t}_g^e := t$;

add g to M ;

Set $t := t - 2\mu_{od}$;

else

Create a morning group:

set $\bar{t}_g^s := t - \mu_{od}$ and $\bar{t}_g^e := t$;

add g to A ;

Set $t := t - \mu_{od}$;

end

end

Set $t := \frac{T_m + T_e}{2} + \mu_{od}$;

while $t < \mathcal{T}$ **do**

Set $g := g + 1$, $o_g := o$ and $d_g := d$;

if $t + \mu_{od} \leq T_e$ **then**

Create another midday group:

set $\bar{t}_g^s := t$ and $\bar{t}_g^e := t + 2\mu_{od}$;

add g to M ;

Set $t := t + 2\mu_{od}$;

else

Create an evening group:

set $\bar{t}_g^s := t$ and $\bar{t}_g^e := t + \mu_{od}$;

add g to P ;

Set $t := t + \mu_{od}$;

end

end

end

end

Set $G := \{1, \dots, g\}$

Algorithm 1: Creating passenger groups and their time characteristics

To complete our specification of the travel time utility curve for each passenger group, we need to determine the peak fare ρ_g for each group g .

5.3. Calculating peak fares

For each passenger group, we need to determine the fare passengers in the group are willing to pay to fly at their preferred time. To do this, we adopt a standardized *fare profile*, which describes fare variation by time of day, irrespective of OD pair. The fare profile is thus a function $f(t)$ defined over $t \in [0, \mathcal{T}]$. We discuss this function and how we arrive at it in more detail below. For now, we note only that it is normalized, and applied to a specific OD-pair by scaling with OD-specific costs, for example, related to the shortest path in the flight network between the origin and destination. This gives a fare profile, $f_{o,d}(t) = \alpha_{o,d} f(t)$, specific to OD pair (o, d) , where $\alpha_{o,d}$ is the scale factor. For each passenger group $g \in G$, we then calculate the average fare over its peak interval to arrive at the peak fare, i.e. we set

$$\rho_g = \frac{\int_{t_g^s}^{t_g^e} f_{o,g,d_g}(t) dt}{t_g^e - t_g^s}.$$

The scale factor $\alpha_{o,d}$ is computed so that the minimum fare seen in f maps to an estimate of the lowest fare one would expect to see for travel from o to d . We take this to be $b + t_{o,d}^{\min} r$, where b and r are assumed to be industry standard values indicating the fixed cost and per unit flight time charge, per passenger, per trip respectively, and $t_{o,d}^{\min}$ indicates the least flight time a passenger could reasonably expect to accumulate in travelling from o to d in the given flight network. For these first two parameters, we use the current industry standards of $b = \$50.50$ and $r = \$0.60$ per minute flight time. Thus $\alpha_{o,d}$ can be calculated by

$$\alpha_{o,d} = (b + t_{o,d}^{\min} r) / \min_{t \in [0, \mathcal{T}]} f(t).$$

We note that for each time-insensitive group g , we simply set ρ_g to be the minimum fare that could be expected for its OD pair, i.e. $\rho_g = b + t_{o_g,d_g}^{\min} r$.

To determine peak fares for time-sensitive groups, we now return to discuss our standardized fare profile, $f(t)$, in more detail. To derive such a profile, our industry partner analysed fare data on two typical busy mid-week days on a single 700km short-haul route between two large cities with a high number of flights. This could be classed as a “shuttle” route, i.e. a relatively short route with a high frequency of service. The reason for analysing such a route is the hope that it would reveal willingness to pay, independent of offered capacity, and independent of strategies used by airlines to smooth demand. Such a route would also maximize the length of time over which data points could be observed.

The airline chosen for analysis was one with the following characteristics: (i) on the route selected, the interval between flights ranged from 15 minutes to one hour according to the expected passenger load, and (ii) its revenue management system includes multiple fare categories, with higher fares at times where there is more demand and where customers have proven

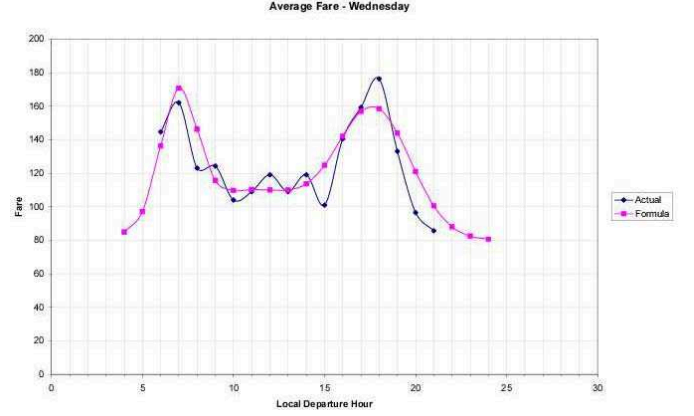


Figure 2: Actual and fitted fare profiles

Type (y)	τ_y	σ_y	w_y
Morning	7.0	1.0	75
Midday	11.0	3.5	30
Evening	17.75	2.0	75

Table 1: Fare profile function parameters

willing to pay higher prices to travel at those times. These characteristics enabled us to gather fare data from the airline’s web booking system. Data comprising the lowest available fare and the number of scheduled flights in each hour was averaged over flights in both directions on the route to minimise the effects of “waves” of flights (sometimes known as “banking”), and compared between two days to give an indication of variability.

Both days showed morning and evening peaks at the same time (around 7am and 6pm respectively, which appear to be very close to the peaks in Figures 2 and 4 of Koppelman et al. (2008)). Both also showed soft midday peaks. However there was significant variation between the two days in terms of relative fare values: Tuesday’s morning peak was much higher than its evening peak, whereas the two peaks for Wednesday (see Figure 2), were of similar height (see Evans (2010) for further details).

The data for these two days was fitted with a combination of 3 normal distribution functions, one for each of the morning, midday and evening peaks, i.e. we constructed

$$f(t) = h^a(t) + h^m(t) + h^e(t)$$

where each function took the form

$$h^y(t) = w_y \exp\left(-\frac{(t - \tau_y)^2}{2\sigma_y^2}\right) + 80$$

for each $y \in \{a, m, e\}$, where τ_y and σ_y are the two normal distribution function parameters (representing mean and standard deviation respectively), and w_y is a weighting factor. (\$80 was the “baseline” fare for this route.) The parameters found to best fit the two days of data are shown in Table 1, and resulting function f is shown in Figure 2.

Like the passenger group peak interval definitions, the fare profile might need to be adjusted for longer flights, particularly

in a west-to-east direction. We discuss this further in the next section.

6. Demand allocation

The final step of our approach is to partition market demand $D_{o,g}$ for each OD pair, (o, g) , (calculated in our first paper Akarunali et al. (2011)), amongst the passenger groups for that OD pair. We call this step *demand allocation*.

For each time-insensitive group g the number of passengers associated with this group, η_g , is simply

$$\eta_g = \phi_{o_g, d_g} D_{o_g, d_g}$$

where ϕ_{o_g, d_g} is the fixed proportion of time-insensitive passengers. In rare cases the surrogate demand for the OD-pair is very low, in which case $\mu_{o,d}$ exceeds $\mathcal{T}/2$. If this occurs, *all* the passengers for this OD-pair are time-insensitive and we have

$$\phi_{o_g, d_g} = 1.$$

As for the peak fare calculation, for time-sensitive groups we use a standardized *demand profile*: a function $\mathbf{d}(t)$ for each $t \in [0, \mathcal{T}]$ that indicates the proportion of demand seen at time t . We discuss below how we arrive at this function. The function is normalized so that the area under the curve is 1, and then the number of passengers η_g associated with passenger group $g \in G$ is calculated as follows:

$$\eta_g = (1 - \phi_{o_g, d_g}) D_{o_g, d_g} \left(\int_{t=\tau_g^-}^{\tau_g^+} \mathbf{d}(t) dt \right).$$

We note that for both the fare and demand profiles, our main motivation for using standardized profiles is the lack of adequate data to support alternatives. The data collected to date is not rich enough to allow us to differentiate demand profiles by OD pair characteristic (with the possible exception of travel direction, which we discuss further below). Since this seems to be emerging as a growing field of research, we hope in the future studies will be carried out that will permit greater differentiation of OD pairs. In the meantime, we adopt standardized profiles as a pragmatic approach that we believe yields realistic data useful for optimization benchmarks.

6.1. Finding the standardized demand profile

To determine the function $\mathbf{d}(t)$, our industry partner analysed the same shuttle route data as was used for deriving the fare profile (see Evans (2010) for more detail). In this case, the number of flights per hour was recorded across each of the two days investigated. The data for one of those days is plotted in Figure 3. Again, the data showed strong morning and evening peaks, with a soft midday peak. Figures 4 and 5 of Garrow et al. (2007) show similar results, with the morning and evening peaks occurring for one type of travellers, and the central peak for another. We thus again seek to fit a combination of 3 normal distributions to this data, i.e. we construct

$$\hat{\mathbf{d}}(t) = \beta^a(t) + \beta^m(t) + \beta^e(t)$$

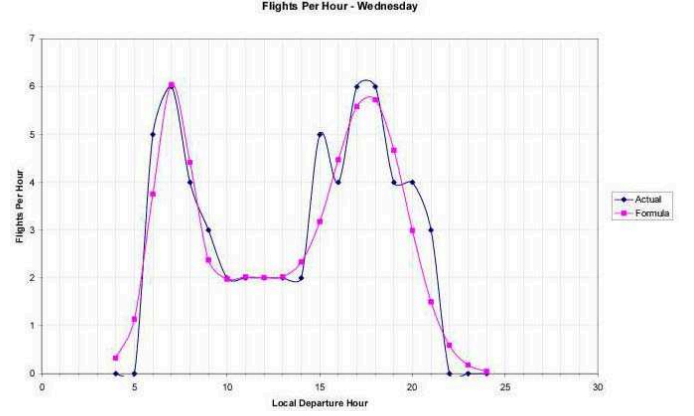


Figure 3: Actual and fitted demand profiles

Type (y)	μ_y	σ_y	w_y
Morning	7.0	1.0	5
Midday	11.0	3.5	2
Evening	17.75	2.0	5.5

Table 2: Demand profile function parameters

where each function took the form

$$\beta^y(t) = w_y \exp\left(-\frac{(t - \tau_y)^2}{2\sigma_y^2}\right)$$

for each $y \in \{a, m, e\}$, where τ_y and σ_y are the two normal distribution function parameters (representing mean and standard deviation respectively), and w_y is a weighting factor. ($\mathbf{d}(t)$ will be set to $\hat{\mathbf{d}}(t)$ multiplied by a normalizing factor.) The parameters found to best fit the two days of data are shown in Table 2, and the resulting function $\hat{\mathbf{d}}(t)$ is shown in Figure 3.

We note that Walker (2006) also suggests approximating the demand distribution by time of day with a combination of normal distributions.

6.2. Differentiating by direction of travel

As we have already mentioned, passengers' preferred travel time can be affected by the length of the trip, and its direction, in particular if a change in time zone is involved. Research presented by Walker (2006), Garrow et al. (2007) and Evans (2010) suggests that travel time preferences are similar for all directions of travel except for west-to-east, and the difference is more noticeable for travel which includes at least one time zone change, and takes more than 3 hours.

To illustrate, we give a sample of the results given in Evans (2010), obtained by analysing schedules operating at the end of 2007 for 44 carriers flying to 994 different airports in Europe, North America, Asia, Africa, South America and Australia, including all airlines in the Star alliance, all airlines in the oneworld alliance, Malaysia Airlines and Jetstar. The data used only includes scheduled services, and does not include actual passenger data, so only the available capacity of each leg is used in the analysis. Plots of total flight capacity by local hour

of departure are given for all non-west-to-east flights (with net west-to-east travel less than 750km), broken down by length of trip, in Figure 4. Similar plots for west-to-east travel are shown in Figure 5 (the case of trips less than 1000km is omitted, because it is very similar to the non-west-to-east case).

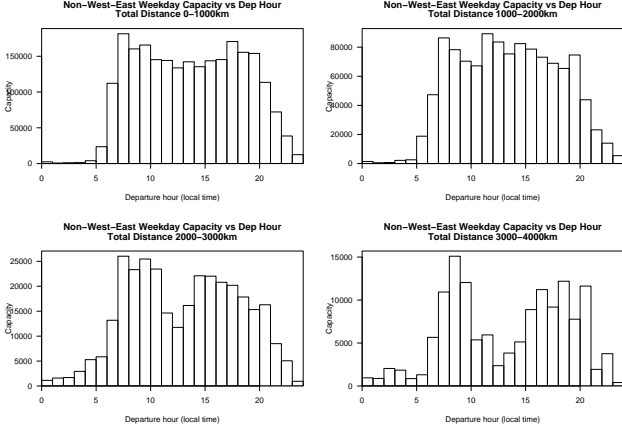


Figure 4: Non W-E Capacity versus local hour of departure

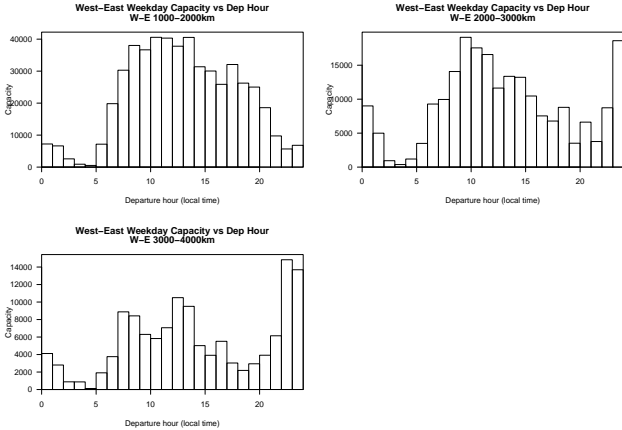


Figure 5: W-E Capacity versus local hour of departure

The differences are particularly apparent for trips of greater than 2000km. In these cases, it appears that a significant portion of the late afternoon/evening demand is shifted to arrive after 5am instead of before midnight. This may be a natural consequence of airport curfews, or may reflect actual passenger preferences.

Thus we adjust our demand profile function on OD pairs corresponding to long west-to-east trips, by shifting the portion of demand that would arrive between midnight and 5am to the next day. Mathematically, we construct an OD-specific demand profile

$$\mathbf{d}_{o,d}(t) = \begin{cases} \mathbf{d}(t) + \mathbf{d}(T_0^s + t - T^{5am}), & T^{5am} \leq t \leq T^{5am} + T_0^e - T_0^s \\ 0, & T_0^s \leq t \leq T_0^e \\ \mathbf{d}(t), & \text{otherwise,} \end{cases}$$

where $[T_0^s, T_0^e]$ is the interval of time during which trips departing o can be expected to arrive at d in the midnight to 5am interval, and T^{5am} is obviously 5am in commensurate time units.

This is clearly something of an ad hoc adjustment, and in future work, we plan to derive fare and demand profiles, and adjust the timing for types of groups, differentiated by direction and length of travel.

7. Benchmark problems

In this section we present summary statistics of the 33 benchmark instances created using the methodology described in this paper, taking as input the flight network and OD-pair demand data generated in our first paper Akartunali et al. (2011). The complete set of instances and supporting material is available at the URL www.infotech.monash.edu.au/~wallace/airline_benchmarks/ along with the reference Evans (2010).

The 33 instances consist of 30 with a single hub, and 3 with two hubs. The instances are grouped into threes. Each group of three networks were generated using the same parameters except the number of spokes which was varied to provide networks of different sizes. The instances show a variety of characteristics, from small to large in terms of the total passenger market, flying range, and size of the flight network. They also have differing degrees of asymmetry in the geometry of the flight network, as well as in the OD-pair demand, which also range over the proportion of passengers with an available direct service, versus those with at best a one-stop service. More details can be found in Akartunali et al. (2011).

Here in Table 3 and 4 we show for each instance the flight network size, total number of OD-pair demands, total number of passengers and summary statistics of the passenger load on arcs that was used to infer the OD-pair demand. We also give, in the column ‘‘Pax groups’’, the total number of passenger groups generated by the methodology we have described here, which ranges from just over 100 up to nearly 25,000 in the largest case. In the four columns headed ‘‘OD-pair pax groups’’ we give summary statistics of the number of passenger groups generated for each OD-pair. These range from 1, meaning that the expected frequency of service on that OD-pair was so low that only the time-insensitive group was generated, to a high of 29. In the final four columns, headed ‘‘OD-pair max revenue’’ we give summary statistics for the total revenue available in each OD-pair market. This is calculated by multiplying the number of passengers in each group by its peak fare, and is the revenue that would be collected if every passenger could be carried at their preferred time of travel. In these tables the hub name is preceded by either an ‘‘s’’, ‘‘m’’, or ‘‘l’’, indicating whether the instance is considered to be a short-, medium-, or long-haul network, respectively.

One caveat to these results is that our approach is tailored to short-haul networks. Long-haul networks have much higher ratios of single-leg passengers and involve rather more complicated patterns of demand versus time of day that are based on some combination of departure time, arrival time, and flying time. Extending our approach to these cases is the subject of future work.

Hub name	Spoke ports	Arcs	OD pairs	Pax count	Arc pax count				Pax groups	OD-pair pax groups				OD-pair max revenue			
					avg	stdev	min	max		avg	stdev	min	max	avg	stdev	min	max
sHAA	24	48	474	71410	1487.71	1340.70	124	4860	1849	3.89	4.09	1	29	28866.20	90282.50	298.99	821120.00
sHBA	72	144	3823	106474	739.40	821.65	28	4860	7272	1.90	1.86	1	29	5020.90	23758.40	195.95	536219.00
sHCA	120	240	10624	224898	937.08	911.61	68	4860	23436	2.21	1.91	1	29	3714.04	23672.30	191.26	737973.00
sHDA	24	48	446	71410	1487.71	1340.70	124	4860	1788	4.00	4.37	1	29	30256.00	86200.90	252.15	809780.00
sHEA	72	144	3840	106474	739.40	821.65	28	4860	7484	1.95	1.90	1	29	4997.48	23792.50	200.63	583023.00
sHFA	120	240	10877	224898	937.08	911.61	68	4860	24278	2.23	1.95	1	29	3630.56	22979.50	191.26	717117.00
sHGA	24	48	480	53230	1108.96	1189.57	83	4860	1452	3.02	3.51	1	29	22731.40	67088.00	247.47	623008.00
sHHA	72	144	3795	104278	724.15	808.63	28	4860	6629	1.75	1.72	1	29	5354.16	26765.40	219.37	666237.00
sHIA	120	240	9766	173132	721.38	789.65	36	4860	17063	1.75	1.64	1	29	3455.96	19337.30	195.95	566700.00
sHJA	24	48	492	53230	1108.96	1189.57	83	4860	1523	3.09	3.60	1	29	22262.50	63088.20	233.42	545797.00
sHKA	72	144	3833	104278	724.15	808.63	28	4860	6963	1.82	1.79	1	29	5293.79	24925.50	219.37	611522.00
sHLA	120	240	10175	173132	721.38	789.65	36	4860	18053	1.77	1.65	1	29	3335.20	19553.70	200.63	604141.00
mHMA	12	24	116	17564	731.83	810.76	50	2640	189	1.62	1.88	1	12	45979.80	132320.00	317.73	901327.00
mHNA	36	72	1051	56430	783.75	706.39	55	2640	1588	1.51	1.50	1	12	14412.20	51170.80	238.10	618896.00
mHOA	60	120	2352	51246	427.05	479.20	15	2640	2532	1.08	0.63	1	12	6570.01	39141.20	205.31	1046240.00
mHPA	12	24	128	17564	731.83	810.76	50	2640	208	1.62	1.92	1	12	41675.20	113428.00	256.84	790089.00
mHQA	36	72	1092	56430	783.75	706.39	55	2640	1676	1.53	1.53	1	12	13912.40	50064.00	214.68	673718.00
mHRA	60	120	2372	51246	427.05	479.20	15	2640	2588	1.09	0.64	1	12	6520.37	40317.60	205.31	1143550.00
IHSA	12	24	128	20062	835.92	403.54	253	1472	176	1.37	1.00	1	4	66593.60	102355.00	2640.20	611666.00
IHTA	36	72	976	48332	671.28	332.73	106	1472	1006	1.03	0.30	1	4	24902.50	71934.00	195.95	678646.00
IHUA	60	120	2610	56192	468.27	256.74	155	1472	2616	1.00	0.09	1	4	10321.80	37745.50	392.67	543435.00
IHVA	12	24	138	20062	835.92	403.54	253	1472	197	1.42	1.05	1	4	61312.50	99331.00	1759.61	561901.00
IHWA	36	72	1047	48332	671.28	332.73	106	1472	1077	1.03	0.29	1	4	23086.90	67151.40	364.57	633348.00
IHXA	60	120	2773	56192	468.27	256.74	155	1472	2779	1.00	0.08	1	4	9691.16	35657.60	369.25	520756.00
IHYA	12	24	100	14160	590.00	327.26	212	1472	106	1.05	0.43	1	4	88832.50	157798.00	907.91	970799.00
IHZA	36	72	715	19526	271.19	243.00	46	1472	721	1.01	0.16	1	4	16899.60	59693.20	434.83	851970.00
IH1A	60	120	1955	55908	465.90	258.45	155	1472	1961	1.00	0.10	1	4	17838.40	66370.60	697.13	833341.00
IH2A	12	24	106	14160	590.00	327.26	212	1472	112	1.05	0.42	1	4	83840.40	135377.00	5055.56	905859.00
IH3A	36	72	744	19526	271.19	243.00	46	1472	750	1.01	0.16	1	4	16220.10	51214.00	865.75	786835.00
IH4A	60	120	2131	55908	465.90	258.45	155	1472	2137	1.00	0.09	1	4	16301.30	53455.60	720.55	695456.00

Table 3: Summary statistics for single hub benchmark instances

Hub name	Spoke ports	Arcs	OD pairs	Pax count	Arc pax count				Pax groups	OD-pair pax groups				OD-pair max revenue			
					avg	stdev	min	max		avg	stdev	min	max	avg	stdev	min	max
sHAB	22	74	486	94290	1274.19	1030.61	167	4860	2258	4.61	3.79	1	29	37322.10	58811.00	238.10	580291.00
sHCB	70	226	3849	184406	815.96	888.14	25	4860	10857	2.82	2.48	1	29	8684.41	22447.30	224.05	445257.00
sHEB	118	378	10159	287234	759.88	825.92	35	4860	24261	2.39	2.16	1	29	5496.57	20888.10	214.68	667310.00

Table 4: Summary statistics for two hub benchmark instances

As can be seen from the tables, the instances display a range of features. They offer good opportunities for testing optimization approaches, with the ability to investigate performance against various characteristics, as well as to test scalability algorithms.

8. Conclusions and future work

In this paper, we have argued that market segmentation and the use of passenger groups represent a natural “next step” in passenger choice models for use in airline planning optimization. We have described one approach to defining such groups, which attempts to balance sometimes divergent empirical insights against what can practically be instantiated using currently available data. We have presented a methodology by which more realistic passenger group data can be generated from a given flight network and OD-pair demands, and produced 33 benchmark instances exhibiting a range of characteristics useful for testing optimization algorithms.

The data we have presented here can be used immediately for testing alternative passenger choice models, provided an airline schedule is also available. We intend in future work to generate schedules to accompany this data, indeed Akartunali et al. (2010) explores first steps in this direction. Clearly the generation of realistic schedules depends on the available airline resources. While much data on aircraft types and configurations, airline fleets, and operating costs is publicly available, in future work, we plan to extend our benchmark data to include aircraft, and so provide complete sets of input data for the flight schedule design process. We hope that in doing so we can stimulate further research in this hitherto less-studied area of airline planning.

Acknowledgements

The authors are very grateful to Ian Evans and Alan Dormer (CTI Pty Ltd) for their ongoing support and guidance on a variety of practical airline-related issues and for numerous technical suggestions and insightful feedback that improved the content and exposition of this work. This research is supported by the Australian Research Council, under Linkage Projects LP00668076 and LP0883855, and by CTI Pty Ltd.

- Akartunali, K., Boland, N., Evans, I., Wallace, M., Waterer, H., 2011. Airline planning benchmark problems. Part I: Characterizing networks and demand using limited data. Tech. rep., University of Newcastle.
- Akartunali, K., Boland, N., Evans, I., Wallace, M., Waterer, H., Smith, O., 2010. Airline schedule design: Network design optimization and heuristics. Working paper, The University of Melbourne.
- Bront, J. J. M., Mendez-Diaz, I., Vulcano, G., 2009. A column generation algorithm for choice-based network revenue management. *Operations Research* 57 (3), 769–784.
- Civil Aviation Authority, 2008. CAA passenger survey report. Tech. rep., Economic Regulation Group.
- Coldren, G., Koppelman, F., Kasturirangan, K., Mukherjee, A., 2003. Modeling aggregate air-travel itinerary shares: logit model development at a major US airline. *Journal of Air Transport Management* 9, 361–369.
- Commission of the European Communities, 2007. Ryanair/Aer Lingus merger procedure report. Office Journal of the European Union Case No COMP/M.4439, Commission of the European Communities.

- Evans, I., 2010. Typical airline demand and fare profiles. Tech. rep., Constraint Technologies International, available from http://www.infotech.monash.edu.au/~wallace/airline_benchmarks/.
- Garrow, L., Jones, S., Parker, R., 2007. How much airline customers are willing to pay: an analysis of price sensitivity in online distribution channels. *Journal of Revenue and Pricing Management* 5, 271–290.
- Kniker, T. S., 1998. Itinerary-based airline fleet assignment. Ph.D. thesis, Massachusetts Institute of Technology.
- Koppelman, F., Coldren, G., Parker, R., 2008. Schedule delay impacts on air-travel itinerary demand. *Transportation Research Part B* 42, 263–273.
- Lohatepanont, M., Barnhart, C., 2004. Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment. *Transportation Science* 38 (1), 19–32.
- McGill, J. I., Van Ryzin, G. J., 1999. Revenue management: Research overview and prospects. *Transportation Science* 33 (2), 233–256.
- Ratliff, R. M., Rao, B. V., Narayan, C. P., Yellepeddi, K., 2008. A multi-flight recapture heuristic for estimating unconstrained demand from airline bookings. *Journal of Revenue and Pricing Management* 7 (2), 153–171.
- UK Department of Transport, 2006. National travel survey: 2006. Tech. rep., <http://webarchive.nationalarchives.gov.uk/+http://www.dft.gov.uk/adobepdf/162469/221412/221531/223955/322743/NTS2006V3.pdf>.
- Van Ryzin, G. J., 2005. Models of demand. *Journal of Revenue and Pricing Management* 4 (2), 204–210.
- Walker, J., 2006. Time of day in airline passenger demand, presentation at the INFORMS Annual Meeting, Pittsburgh PA.
- Yan, S., Tang, C.-H., Lee, M.-C., 2007. A flight scheduling model for taiwan airlines under market competitions. *Omega* 35 (1), 61–74.
- Yan, S., Tseng, C.-H., 2002. A passenger demand model for airline flight scheduling and fleet routing. *Computers and Operations Research* 29, 1559–1581.