When the Model Hits the Runway: The DOZE Algorithm for optimal dispatching of escorts to mobility impaired airline passengers

Control Team 22

February 6, 2006

Abstract

We model and algorithmically optimize the problem of assigning escorts to accompany Non-ambulatory Airport Passengers (NAPs). We accomplished the modeling by first constructing a scalable computer simulation of an airport, including flight schedules, concourse layouts, and variable congestion. Then we used census and survey data to obtain a rough estimate for how many passengers would require mobility assistance in order to make flight connections. We gathered travel agency data to model distributions of layover times and congestion, and used information from the Bureau of Transportation Statistics to model flight delays. Utilizing estimates on different costs associated with escorting NAPs, we were able to prioritize and codify all the operations taking place during NAP connections.

The algorithm employed uses two kinds of escorts, whom we called Daily-Ordered Escorts and Zoned Escorts. The Daily-Ordered escorts are responsible for the majority of scheduled NAP connections, whereas the Zoned Escorts redistribute wheelchairs throughout and among fabricated Zones. We develop a system of rules which allows for all NAP connections to be handled, while leaving allowance for unexpected delays and short-notice connection preparation. If time permitted we planned to use our algorithm, together with our airport simulation, to model how effectively the algorithm can deal with varying levels of congestion, NAP frequency, and airport size.

Contents

1 Introduction **3**
1.1 Responsibilities of Epison Airlines toward Non-ambulatory Passengers . . . 3
1.2 Goals of Analysis . . . 4
1.3 Method of Analysis . . . 5

2 Modeling Airport Activity **5**
2.1 Airport Layout . . . 5
2.1.1 Large Airport . . . 6
2.1.2 Medium Airport . . . 6
2.1.3 Small Airport . . . 6
2.2 Flight Schedule . . . 7
2.3 Airport Congestion . . . 7
2.4 Incoming Flight Delay . . . 7

3 Modeling Non-ambulatory Passengers, Escorts and Wheelchairs. **8**
3.1 Identifying NAP itineraries. . . . 9
3.1.1 The probability of being a NAP . . . 9
3.1.2 Layover Distribution . . . 10
3.1.3 Selecting Departing Flights . . . 10
3.2 Escort Behavior . . . 10
3.2.1 Escort Travel Time . . . 12
3.3 Wheelchair States . . . 12
3.3.1 *Occupied* and *Unoccupied* . . . 12
3.3.2 *In-Transit* and *Stationary* . . . 12

4 Cost Analysis **13**
4.1 Wheelchair Maintenance . . . 13
4.2 Escort Wages and Compensation . . . 13
4.3 Cost of Delaying NAP Departure . . . 13
4.4 Cost of Delaying NAP Pickup . . . 14
4.5 Constraining Wheelchair Storage . . . 14

5 The DOZE Algorithm **14**
5.1 Daily-Ordered Escort Scheduling . . . 15
5.2 Zone Assignment and Zoned Escorts . . . 16
5.3 Zoned Escorts and Wheelchair Distribution . . . 16
5.4 Escort Notification . . . 17

6 Simulating an Average Day at the Airport **18**

7 Conclusions **19**

1 Introduction

According to the 2000 census, 49.5 million Americans live with disabilities [7]. An estimated 23% of these need some sort of specialized assistance just to leave home, and 55% report having experienced added difficulty at airports. The US Department of Transportation requires that airlines provide wheelchair assistance for mobility impaired passengers to reach connecting flights while other mobility aids are checked and unavailable [9].

However, today in the age of budget air carriers and widespread airline bankruptcy, Epsilon Airlines must spend every penny carefully. This includes costs associated with escorting non-ambulatory passengers from arrival gate to departure gate during layover connections in fulfillment of the responsibilities set forth by the US Department of Transportation (USDOT). We are happy to have been chosen to consult with Epsilon Airlines in choosing a system that will most efficiently transport non-ambulatory passengers while incurring minimal direct and liability costs to the airline.

1.1 Responsibilities of Epison Airlines toward Non-ambulatory Passengers

Non-ambulatory passengers (NAPs) make up a small yet important part of air travelers. The legal and ethical responsibilities of Epsilon Airlines with regard to NAPs that are connecting between flights are:

- to provide a **wheelchair at the arrival gate** at or prior to the arrival time of the non-ambulatory passenger's incoming flight.
- to provide an **escort to transport the non-ambulatory passenger** from the arrival gate to the departure gate with sufficient time to board the departure flight. Standard airline policy is that passengers must be boarded 10 minutes before the scheduled departure [5].
- to allow the NAP **usage of a wheelchair for the duration of his or her layover.** (Note that the escort may leave the NAP after delivering the NAP to the departure gate. Gate attendants are trained to aid NAPs in boarding flights.)

Note that we are concerned only with moving NAPs between connecting flights. Per USDOT policy, Epsilon Airlines is *not responsible* for moving NAPs from check-in to initial boarding, or from final departure to baggage claim [7].

In constructing a plan for Epsilon Airlines perform these responsibilities, we will seek to minimize several direct and indirect costs and constraints:

- Costs associated with **wheelchair maintenance**. (*Direct Cost*)
- Costs of **compensating escorts**, including wages, benefits, and taxes. (*Direct Cost*)

- Cost associated with having **no available wheelchair** at the gate upon a non-ambulatory passenger arrival. These costs include decreased satisfaction (as the NAP grows impatient waiting for assistance leaving the plane) and a possible delay in future departing flights (*Indirect Cost*)
- Cost of **delaying departing flight** by not delivering a non-ambulatory passenger on time. (*Indirect Cost*)
- Liability hazard cost of leaving **wheelchairs in concourse unattended.** (*Liability Cost*)
- Constraint of **limited wheelchair storage** space at airports. (*Constraint*)

In this list of costs we have *not* included a purchase price for a fleet of wheelchairs. We assume that, in compliance with USDOT standards, Epsilon Airlines has a currently operating fleet of wheelchairs which will continue to be in service. Furthermore if our analysis finds that additional wheelchair purchases are necessary, this is a relatively insignificant one time cost. Any cost incurred through replacing worn or non-reparable wheelchairs should be included in the per-wheelchair annual maintenance cost.

1.2 Goals of Analysis

In our bid we will develop an algorithm to schedule the movement of wheelchairs and escorts throughout the airport. Our analysis will provide Epsilon Airlines with

- an estimate of the optimal **number of wheelchairs** needed to serve an airport based on the airport size and flight schedule.
- an estimate of the optimal **number of escorts** needed at an airport to adequately serve non-ambulatory passengers.
- a optimal **shift schedule** for escorts based on the size and flight schedule of the airport.
- an estimate and error analysis of **budget costs** associated with NAP escort service for a given airport.
- an algorithm for **tracking and dispatching escorts and wheelchairs** throughout the airport.
- potential **sources of bottlenecking** or high costs in the system.

1.3 Method of Analysis

In order to obtain our analysis we will follow the following steps:

1. Construct a scalable **computer simulation model** of Epsilon Airlines activities in an airport, including airport layout, incoming and outgoing flights, and congestion.

2. Obtain an estimate for NAP prevalence and determines simulation rules for NAP, escort, and wheelchair behavior.

3. Use industry data to estimate each of the previously identified costs.

4. Develop and test an algorithm for escort scheduling using the airport simulation model and calculate total associated costs.

5. Refine the algorithm based on primary sources of inefficiency.

6. Demonstrate present and future scalability of optimal algorithm by testing the algorithm with different size airports and varied NAP and congestion parameters.

2 Modeling Airport Activity

The first important aspect for testing an algorithm is to accurately simulate the activities of an airport.

- Airport Layout
- Flight Schedule
- Airport Congestion
- Flight Delays

2.1 Airport Layout

Modern airports are complex structures with many components: terminals, baggage claim, concourses, runways, gates, parking garages, food courts, security checkpoints, etc. However, because all that we are really concerned with is transferring passengers from one flight to another connecting flight, the only aspect of the airport which we need to model is the trip a passenger takes from his or her arrival gate to the departure gate for the next flight. Typically at this point the passenger does not need to check in baggage, has cleared security checkpoint lines, and may travel freely from gate to gate. Furthermore, passengers typically fly the same airline or a code-sharing airline through their whole itinerary, so we assume that all passengers arriving on Epsilon Airlines will be connecting to another flight on Epsilon Airlines. Thus in the model, airports consist only of **concourses and gates which are serviced by Epsilon Airlines,** where the **distances between gates are measured by walking time under low congestion.**

Each concourse in the airport is parameterized as a symmetric matrix. For example, Concourse A with n gates serviced by Epsilon Airlines, is given by an $n+1 \times n+1$ matrix A where the rows and columns $\{1, \ldots, n\}$ correspond to each gate and the $n+1$ row and column correspond to the entry way of the concourse. The a_{ij} entry is the time in minutes for an unimpeded passenger to walk from gate i to gate j.

For the case that a passenger need to change concourses, there another matrix fixes travel times between each concourse serviced by Epsilon Airline. These travel times may account for time other than simply walking time. For example, one reason that a passenger may need to change concourses is to transfer from an international flight in Concourse E to a domestic flight in Concourse A. In this case the passenger would need to go through customs, incurring an added travel time dependent on airport congestion.

2.1.1 Large Airport

Epsilon Airlines' operations in a large airport are modeled after the operations of Delta Airlines at Hartsfield-Jackson Atlanta International Airport. Delta operates 102 gates in the Atlanta Airport: 34 in Concourse A, 36 in Concourse B, 15 in Concourse D, and 17 in Concourse E [6].

Each concourse is straight with gates on either side, consecutive gate numbers are across from one another. The Manhattan metric is applied to calculate travel times between gates within the same concourse where the base travel time under low congestion between adjacent gates is one minute.

Under low congestion, travel between adjacent concourses takes 5 minutes. The concourses are aligned in linearly, so for example to travel from Concourse A to Concourse D would take 15 minutes. The exception is to travel from the international Concourse E to domestic Concourse D, for which the base travel time is 15 minutes to allow for US customs.

2.1.2 Medium Airport

Epsilon Airlines' operations at a medium sized airport are modeled after the Delta Airlines operations at Boston Logan Airport. In Boston, Delta operates 24 gates in two Concourses. Delta gates are only located on one side of the concourse and the base travel time between adjacent gates is 1 minute. The base travel time between the concourses is ten minutes.

2.1.3 Small Airport

Epsilon Airlines' operations in a small airport are modeled after Alaska Airlines' operations in Anchorage, Alaska. Alaska Airlines operates 11 gates in Concourse C of the airport. The gates are all adjacent and the based travel time is assessed using the Manhattan metric with one minute of travel time between adjacent gates.

2.2 Flight Schedule

The simulator generates an arrival and departure time for each Epsilon Airlines flight in an airport on a day using a fixed number of arriving and departing flights and a probability distribution of flights over the day. After obtaining times for each flight, the flights are distributed randomly to the gates served by Epsilon Airlines such that no two flights arrive or depart at the same gate within 45 minutes (the minimum time to deplane and prepare another flight for takeoff).

According to Delta Airlines [1], flights are distributed as evenly through the day as possible throughout the day between the hours of 7 AM and 11 PM. The Delta flight schedule in Atlanta and Boston and the Alaska Air flight schedule in Anchorage support this distribution [12][5]. Thus we have assigned a uniform distribution between the hours of 7 AM to 11 PM to the arrival and departure schedules for each airport.

The number of Epsilon Airlines flights each day is 800 arrivals and departures each day in the large airport, 200 arrivals and departures each day in the medium airport, and 80 arrivals and departures each day in the small airport [12][5].

2.3 Airport Congestion

As any traveler knows, airport congestion can severely affect the amount of time that it takes to move around an airport. At peak times it can take up to three times as long to walk from one place to another [4]. Although the flight distributions are uniform throughout the day, congestion in concourses tends to worsen during the midday due to overlapping layovers and the nature of traffic flow. To account for this, we multiply the base travel time between any two gates by a 'congestion factor', which depends on the time of day. The congestion factor is 1 at 7:00 AM, increases linearly to 3 (indicating that it takes three times as long to travel from location to location) at 3:00 PM, and decreases linearly back to 1 at 11:00 PM. Using the congestion factor, the travel time for an ordinary passenger to walk between two locations in the airport can be expressed by:

$$\text{travel time} = \text{base travel time} \times \text{congestion factor}$$

2.4 Incoming Flight Delay

Delayed arriving flights severely impact standard airport operations. They cut into passengers' layover time and disrupt other scheduled airline events such as escorts providing wheelchair assistance. In order to accurately model flight delay, we analyzed data raw data from the Bureau of Transportation Statistics about flight delay in December 2005 [3]. The data shows that 68.67% of flights

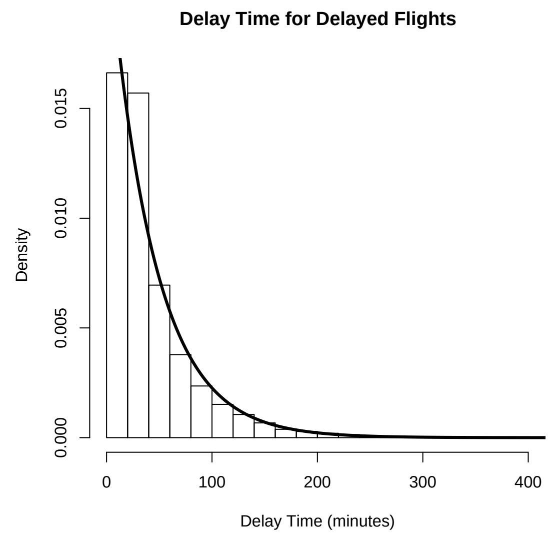

Figure 1: Exponential distribution of flight delay minutes for flights with non-zero delay. Histogram represents empirical delay times for all air carriers during December 2005.

have zero minutes of delay. For the roughly 30% of flights that are delayed, the amount of delay in minutes follows an exponential distribution with parameter $\lambda = 1/42.90$ as seen in Figure 1. This translates to a median delay time of 29.73 minutes.

An unpredictable delay time is assessed to each incoming flight using a random number between zero and one. If the random number is less than .6867 the flight is assessed no delay. If the random number is greater than .6867, another random number is chosen between zero and one and is compared to the exponential distribution to obtain the number of delay minutes added to the scheduled arrival time of the incoming flight to obtain the actual arrival time.

Even though sometimes delays in airports can be anticipated, often they cannot. We assume that delays are unexpected. When scheduling wheelchair assistance and escorts, we may only use scheduled arrival times, not delayed arrival times.

3 Modeling Non-ambulatory Passengers, Escorts and Wheelchairs.

There are essentially three types of non-ambulatory passengers that we must account for:

1. NAPs who request wheelchair assistance at the time of booking their ticket. ("Long-notice NAPs")

2. NAPs who request wheelchair assistance at their flight check-in time. ("Medium-notice NAPs")

3. NAPs who request wheelchair assistance only just before arrival. ("Short-notice NAPs")

The medium-notice and short-notice NAPs are collectively known as "unscheduled NAPs."

We conservatively assume that 70% of the non-ambulatory passengers fall into the first category. The itinerary of these long-notice NAPs (in particular, their scheduled arrival and departure times and gates) are known at the beginning of the day. These passengers provide the airline with enough notice that it will be possible to plan ahead for them.

The second category comprises 25% of the non-amulatory passengers. These passengers provide about 4-5 hours notice to the airline that they will need assistance. Thus the airline must have some flexibility in its scheduling. It cannot simply draw up a schedule at the beginning of the day and assume that there will be no changes.

The remaining 5% of non-ambulatory passengers are short-notice NAPs. The airline has only 15-20 minutes to meet one of these passengers at the gate. It is very unlikely that an airline would be able to provide assistance in a timely fashion to one of these passengers unless it had someone in close proximity to each gate at all times.

3.1 Identifying NAP itineraries.

After generating a daily flight schedule, the simulator uses a (very small) individual probability that a passenger on an incoming flight that day requires wheelchair assistance to stochastically generate a set of NAPs who will be traveling through the airport that day. Then for each NAP a layover time is generated randomly from the layover time distribution, and the departure gate with the next departure following the

3.1.1 The probability of being a NAP

Every passenger on an incoming flight has a certain probability of being non-ambulatory. To calculate this probability we use data from [11]. The rough proportion of NAPs to all passengers is expressed as the conditional probability that a person is mobility-disabled and requires special assistance, given that s/he travels long distances by air. The fraction of disabled citizens in the U.S. is 19%, of whom approximately 76.2% are able to travel long distances, making 14.5% of the population disabled long-distance travelers. Of these, only 18.8% travel by air, making the probability that a given citizen is able to travel by air

2.73% (the disabled portion) + 81% (the non-disabled part of the U.S.), or an 83.7% probability that a given person is able to travel large distances by air.

Now of those disabled who are able to travel long distances, 3.1% require special assistance (such as a wheelchair). Of these, 87.6% are mobility disabled. Lastly, the proportion of mobility-disabled long-distance travelers who travel by air is 15.5%. Thus we can estimate the probability that a given person is mobility-disabled, but able to travel long distances by air when given special assistance, as 0.045%. Then the conditional probability that a long-distance air traveler is also mobility disabled and needs assistance is 0.045%/83.7% = 0.054%, corresponding to roughly one in 1850 travelers.

3.1.2 Layover Distribution

After a NAP's arrival is stochastically generated using the probability in the previous section, the duration of the NAP's layover between connecting flights is chosen randomly from a distribution of layover times. The layover time distribution is a *Beta distribution* with parameters $a = 2.651$ and $b = 12.409$. Since a Beta distribution must be defined over a compact interval, we allow zero to be the minimum layover time and set the maximum layover time to be 600 minutes (10 hours). This greatly exceeds any observed layover times. This distribution yields a mean layover time of 105.615 minutes and standard deviation 57.019 minutes. The shape of the Beta distribution is seen in Figure 2.

In order to obtain this layover distribution we collected a random sample of layover times from actual flight itineraries provided by Orbitz.com and Expedia.com. The sample included a variety of airlines and airports of various sizes. The size of the sample is $n = 286$ and the distribution of the collected data can be seen in the histogram in Figure 2. This illustrates that the shape of the Beta distribution fits the observed data quite well.

3.1.3 Selecting Departing Flights

After each NAP has been generated on an incoming flight and given a layover time, the next scheduled Epsilon Airlines departure flight after the NAP's layover is chosen and the designated gate selected as the NAP's departure gate to which the NAP will be escorted.

3.2 Escort Behavior

Escorts are hired by Epsilon Airlines for the purpose of moving non-ambulatory passengers between gates. Each escort is hired for a four hour shift during which we assume he or she is accessible via radio control to be dispatched to or from any place in the airport. Escorts are sent to the arrival gate at the expected arrival time of the NAP's incoming flight. If the central dispatch unit believes that there is an available wheelchair at the gate, the escort will be sent directly to the gate, otherwise the central dispatch unit will send an escort to a location to obtain a wheelchair prior to retrieving the NAP.

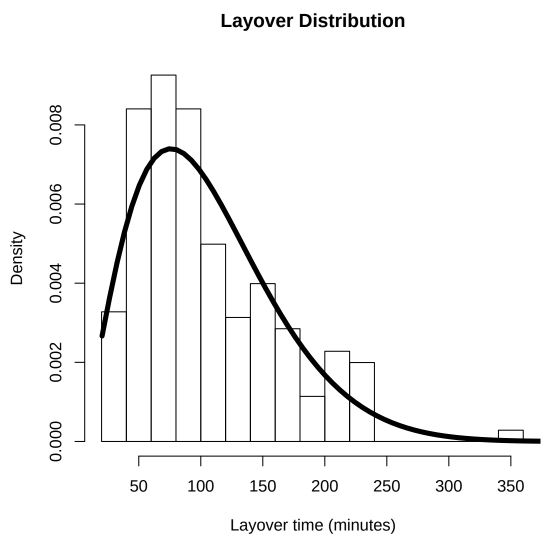

Figure 2: A Beta distribution with parameters $a = 2.651$ and $b = 12.409$ normalized to the time interval $(0, 600)$ from which we sample layover times. The underlying histogram depicts observed layover times from a random sample of $n = 286$.

After retrieving the NAP, the escort will proceed directly to the NAP's departure gate. Upon arrival at the departure gate, if the NAP's departure flight has not already boarded, the escort will leave the NAP and the wheelchair at the departure gate and again become free for assignment.

When not occupied moving a passenger, escorts may also be dispatched to balance the distribution of escorts or wheelchairs.

3.2.1 Escort Travel Time

Since pushing an unoccupied or occupied wheelchair affects an individuals walking speed, the time which it takes escorts to travel from gate to gate should vary depending on whether they are traveling freely, traveling with an unoccupied wheelchair, or pushing a NAP. In order to obtain information about the variable rates of these three modes of travel, we constructed an experiment in which we pushed team members in a rolling chair through open areas and the campus food court. These experiments showed that in both congested and open areas, pushing an empty chair makes travel time about 1.5 times longer and pushing an occupied chair makes travel time approximately twice as long.

Thus escort travel time calculated by multiplying the ideal travel time between two gates by the congestion factor obtained in section 2.3, times a 'travel load' coefficient of 1, 1.5, or 2 depending on if the escort travels alone, with an empty wheel chair, or an occupied wheelchair respectively:

$$\text{travel time} = \text{base time} \times \text{congestion} \times \text{travel load}$$

3.3 Wheelchair States

3.3.1 *Occupied* and *Unoccupied*

Wheelchairs may take two basic states; they may be *occupied* or *unoccupied.* A wheelchair changes between these states only when a NAP deplanes from an arriving flight and hence occupies and unoccupied wheelchair, or enplanes on a departure flight and hence vacates an occupied wheelchair. Note that an escort does not need to be present for either of these events to occur, as Epsilon Airlines gate attendants are trained in assisting disabled passengers between wheelchair and the aircraft. However, if there is no unoccupied wheelchair at the gate upon arrival of a NAP, an escort must bring a wheelchair to the gate.

3.3.2 *In-Transit* and *Stationary*

Furthermore a wheelchair (occupied or unoccupied) may be *in-transit* or *stationary.* An in-transit wheelchair must be accompanied by an escort. A wheelchair may only be 'stationary' if it is at a gate (it may be occupied or unoccupied). Stationary and unoccupied wheelchairs may be either *in-storage* or *not-stored* where an in-storage wheelchair is in a designated storage area at the gate. If there is storage available at the gate where the wheelchair is located, an unoccupied stationary wheelchair will automatically be stored by gate attendants.

A not-stored unoccupied wheelchair will be presumably sitting in an open area around the gate and could be a safety hazard to other airline passengers. To move a not-stored wheelchair to storage at another gate, an escort can come to the gate to get the wheelchair and move it.

4 Cost Analysis

We need to analyze and balance all of the following respective constraints:

- The hard cost of maintaining the fleet of wheelchairs.
- The hard cost of hiring and compensating escorts.
- The internal cost of delaying NAP departure.
- The internal cost of delaying NAP pickup.
- The constraint of storing unoccupied wheelchairs.

The first two are costs which the airline must pay directly. The next two are costs which the airline pays indirectly, through decreased profit. Last is our interpretation of the liability hazard of unattended, unoccupied wheelchairs which are not stored at gates.

4.1 Wheelchair Maintenance

As stated in Section 1, we count the cost of a wheelchair not by the one-time purchase price, but by the average cost of maintenance. It is estimated that for each wheelchair, the airline incurs a maintaining cost of $1,000 over a five-year period, so about $200 per wheelchair per year. We will see that this cost is nearly negligible when compared to the cost of hiring escorts or delaying flights.

4.2 Escort Wages and Compensation

We need an accurate estimation of how much it costs to hire each escort. From [2], we know that the annual compensation of airline employees ranges from $10,000 to $200,000, including wages and fringe benefits, with an average of $47,000. We chose to use this last figure to estimate the annual cost of hiring each escort. Assuming that all the escorts work full-time (40 hours per week) for 50 weeks out the year, this gives us an hourly cost of $23.50 per escort.

4.3 Cost of Delaying NAP Departure

We needed to obtain a rough estimate for how much flight delays cost by the minute. A first-order approximation could include the hourly operating cost of the airplanes, but then we would have to consider the domino effect on other flights, the grievances of aggravated passengers, etc., all of which are difficult to directly estimate. As a result, it would be more accurate simply to take the econometric estimate of total delay costs, and average over the total number of delay hours. We find from [8] that the economic cost of delays to the airlines in the year 2000 was approximately $3.13 Billion while the number of delay hours was 1.69 Million. This gives us a delay cost per minute of about $30.88.

Since flights generally issue a final warning to have their passengers board at least 10 minutes before departure, we decided that we would incorporate flight delay cost as follows:

- If a NAP arrives at the departure gate at least 10 minutes before the departure time, no delay cost is incurred.
- If a NAP arrives after the 10-minute warning, $31 is charged for each delayed minute.

For example, if a NAP is scheduled to board a plane departing at 3:00 and does not arrive until 2:54, the associated cost is $4 \cdot \$31 = \$124$.

4.4 Cost of Delaying NAP Pickup

There is a similar cost for delaying the pickup time of an arriving NAP. We now allow a 10-minute grace period for pickup after the scheduled arrival time. After this grace period ends, we are delaying the next outgoing flight from that gate, and so there is a $31 charge per minute again. However, there is another (smaller) charge to consider: that of delaying the NAP. The value of passenger time is currently estimated at approximately $58 per hour, or just under $1 each minute. [10] This cost is incurred whenever an empty wheelchair is not present at the time of NAP arrival. For example, if the scheduled arrival time of a flight is 3:30, but the NAP is not seated until 3:41, the total cost incurred is $11 \cdot \$1$ for delaying the passenger plus $1 \cdot \$31$ for delaying the plane, totaling a $42 charge.

4.5 Constraining Wheelchair Storage

We decided to avoid directly using the indeterminate liability of leaving an unoccupied wheelchair unattended, and instead build wheelchair capacity directly into our model of an airport. From anecdotal evidence, we realized that at most two unoccupied wheelchairs may be located at any gate and time, before the collection of wheelchairs becomes hazardous. (The gate capacity may be changed for use in different airports.) As a result, in our optimization algorithm we made it a necessity for escorts to continually redistribute wheelchairs so that the chair distribution would avoid piling up at a single gate, see Section 5.3

5 The DOZE Algorithm

The Daily-Ordered and Zoned Escorts (DOZE, for short) Algorithm is the key ingredient to optimizing the escort schedule. DOZE uses all known information about the airport, including current escort and wheelchair locations, certain pending NAP arrival times, and certain flight delays. Our algorithm then follows the following rough outline:

- **DOZE compiles a schedule for all known incoming NAPs to have escorts within their allotted layover time.** This is accomplished by choosing a collection of escort movement instructions so that each scheduled escort accompanies as many successive NAP connections as possible within a four-hour shift. The escorts required by this phase of the algorithm are called **Daily-Ordered Escorts**, since the algorithm projects escort assignments 24 hours into the future.

- **DOZE sorts the gates in each concourse into Zones, and assigns an extra "Zoned Escort" to govern each zone.** Since the number of Daily-Ordered Escorts was chosen minimally to satisfy the long-notice NAP demands, it is possible that an unscheduled NAP may arrive at a time when no one can escort him or her to the departing flight. The zones are therefore chosen so that each Zoned Escort can reach any gate in the zone, in time to accompany any unscheduled NAPs to their destinations.

- **DOZE compiles a schedule for all Daily-Ordered or Zoned Escorts to distribute wheelchairs as necessary for NAP arrivals and layovers.** While they are not rescuing unexpected NAPs, Zoned Escorts are to redistribute unoccupied wheelchairs according to rules which optimize NAP connections. DOZE also decides whether a Daily Ordered Escort should distribute wheelchairs *en route* to retrieving NAP arrivals.

- **DOZE notifies all escorts of their next assignments.** Even though DOZE predicts escort schedules well into the future, it need only notify escorts of their next intended action. This is due in part to avoid confusion on the part of the escort, and in part because DOZE will be continually recalculating its optimized schedule as it receives updated information.

5.1 Daily-Ordered Escort Scheduling

The first step in optimizing the escort schedule is to find out how many Daily-Ordered Escorts are needed to satisfy the needs of all the expected NAP connections. We accomplish this by assigning NAP connections to escorts as follows:

1. Select the first pending NAP arrival that still needs an escort, and assign a new escort to it.

2. Block in a 95% confidence travel time estimate for the escort to deliver the NAP to the departure gate. The escort is occupied during this time.

3. Select the first NAP arrival which the escort can reach by the arrival time after dropping off the previous NAP, still using a 95% confidence level.

4. Repeat until the total amount of time the Escort has been working exceeds four hours; the escort's shift is over.

5. Repeat from step 1 until all NAPs have been assigned to escorts.

This will allow all known NAP connections to be covered by the collection of Daily-Ordered Escorts.

5.2 Zone Assignment and Zoned Escorts

Creating the Zones is a relatively simple task for DOZE. We suppose that 95% of all short-notice NAP connections give at least a certain warning, (say) of 15 minutes. Then within each concourse, begin with the first gate and place all gates within 15 minutes walking time into the first zone. Then take the next gate not contained in the first zone and make the second zone similarly; in this way each concourse is partitioned into zones within which travel time is less than nearly all warning times for short-notice NAP arrivals.

Now DOZE assigns a Zoned Escort to each zone. We can see that the Daily-Ordered and Zoned Escorts are both necessary and sufficient to escort all the NAPs through their connections (necessary in the sense that without any particular escort it is possible to have an arriving NAP not retrieved by an escort, and sufficient in that given all of these escorts it should be possible to handle nearly every collection of NAP layovers).

5.3 Zoned Escorts and Wheelchair Distribution

The behavior of the Daily-Ordered Escorts is a simple greedy algorithm designed to cover all expected NAP connections. The Zoned Escort behavior is inherently more complex, because they must deal with a more probabilistically oriented NAP expectation. We settled on the following three kinds of actions which Zoned Escorts may perform:

- **Whenever Zoned Escorts move wheelchairs, they choose the wheelchair to balance the overall distribution against NAP needs.** The Zoned Escort chooses a wheelchair at either (a) a gate with more empty wheelchairs than can be safely stowed, or (b) a gate with the most time until its next arriving NAP. The Zoned Escort takes this wheelchair either to (a) a gate with a NAP waiting for a chair, or (b) a gate with the least time until a NAP arrival.

- **Zoned Escorts accompany NAPs which the Daily-Ordered Escorts cannot handle due to unexpected circumstances.** Suppose that a NAP arrives in Zone A, and needs to connect to Zone B. If Zoned Escort A is not currently accompanying a NAP, then Zoned Escort A will arrive as soon as possible and begin the travel to the destination in Zone B. At this point, Zoned Escort A becomes the assigned Zone Escort of Zone B, and Zoned Escort B relocates to Zone A as soon as possible. In this way, the NAP is successfully escorted and the Zoned Escorts remain efficiently distributed.

If Zoned Escort A is already accompanying a NAP at the time that another NAP arrives unexpectedly in Zone A, the nearest unoccupied Zoned Escort (of, say, Zone C) takes the job of connecting the new NAP. Then Zone Escort C is relabeled Zone Escort B and vice versa.

- **Zoned Escorts distribute wheelchairs among Zones.** If Zone A has no unoccupied chairs, Zoned Escort A identifies all zones with available Zoned Escorts (that is, those who are not currently escorting NAPs). Of these zones, Zoned Escort A chooses the closest zone (say, Zone B) with at least two wheelchairs. Then Zoned Escort A travels to Zone B and becomes its Zoned Escort, and Zoned Escort B takes a wheelchair to Zone A, becoming Zone A's Zoned Escort. If there is no such Zone B to choose from, then Zoned Escort A simply retrieves the nearest unoccupied chair and returns.

  Conversely, if every gate in Zone A has at least the maximum number of chairs, then Zoned Escort A chooses the closest zone (Zone B) with at least one space for a free wheelchair, and takes one of its extra chairs there. Once again, Zoned Escort A trades places with Zoned Escort B if possible, but otherwise simply returns to Zone A.

- **Zoned Escorts distribute wheelchairs within Zones.** In Zone A, Zoned Escort A will always place a wheelchair at a gate, with priority for those NAPs arriving first. The Zoned Escort continues until there are no more unassigned wheelchairs in the Zone, or a gate has too many empty wheelchairs, at which the Zoned Escort redistributes the chairs according to the rules in part 1.

Altogether, the order of priority is as follows: If there is a NAP connection which the Daily-Ordered Escorts cannot handle, the Zoned Escort takes care of the situation. If no such action is required, but there are no unoccupied wheelchairs (or too many) in a Zone, the Zoned Escort redistributes the wheelchairs among Zones. If there are neither too few nor too many wheelchairs in the Zone, then the Zoned Escort distributes a wheelchair to each pending NAP's incoming gate. In this way, the Zoned Escorts are able to distribute the wheelchairs so that (a) the connecting NAPs all have wheelchairs, and (b) the overall wheelchair distribution never becomes too unbalanced.

5.4 Escort Notification

The DOZE algorithm finally notifies the escorts of their next assignments. Each Zoned Escort receives the next instruction based on the above rules for wheelchair distribution and NAP pickup. The Daily-Ordered Escorts also receive their next NAP assignments, but with a slight addition. Whenever a Daily-Ordered Escort walks from one Zone to another without escorting a NAP, DOZE checks

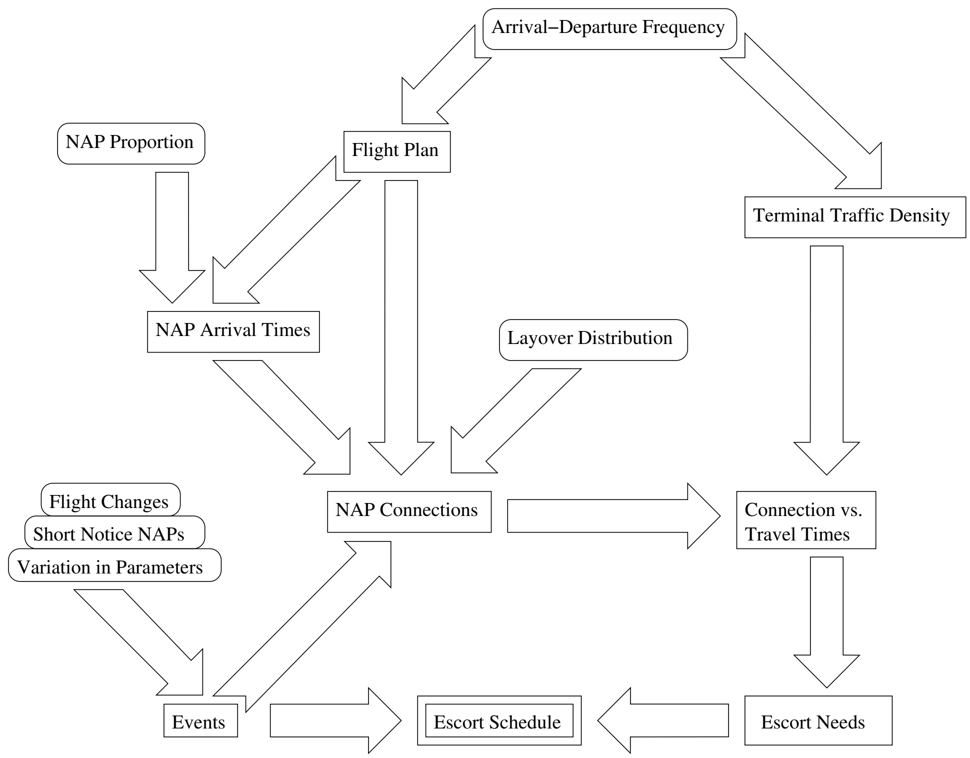

Figure 3: A schematic showing how the airport model information, NAP parameters, and simulated events combine to give the Escort Schedule. Each arrow shows a piece of our simulation.

the relative densities (average wheelchairs per gate) of the two Zones, and tells the Daily-Ordered Escort to bring a wheelchair along if it would make the densities more equal. Then DOZE releases the combined NAP assignments and wheelchair movement orders to the Daily-Ordered Escorts.

6 Simulating an Average Day at the Airport

The diagram in Figure 3 may be understood as follows:

1. Begin with the arrival and departure frequencies.
2. As discussed in Section 2.2, we use these distributions to generate a flight plan for the simulation period of one week, and a varying traffic density/congestion profile used to estimate gate-wise travel times.
3. From the proportion of NAPs among our passengers, we use the flight plan to randomly generate a list of NAP arrival times.

4. We incorporate the layover distribution and NAP arrival times, and obtain from the flight schedule a complete list of NAP connection gates and times.

5. The DOZE algorithm combines the connections and traffic density to produce its estimation of where and when all the escorts will be needed.

6. From these needs, DOZE produces the first iteration of the complete escort schedule, and begins directing where the escorts should go.

7. As various unexpected airline events occur, such as flight changes, short- and medium-notice NAPs, and variation in actual travel times, the predicted escort connections will slightly change, as well as the actual locations of escorts and wheelchairs.

8. The new connections feed back into DOZE, which produces the next iteration of the Escort Schedule.

9. This process repeats indefinitely, so that the Escort Schedule is continually refined to match current events.

To simulate this process, we divided the process into two pieces:

- Event Planning
- DOZE Optimization

During the event planning, all of the NAP connections (including, short-, medium-, and long-notice NAPs), flight schedules, and flight delays are preplanned. Every unexpected NAP arrival and flight delay has an associated warning time: long-notice NAPs are known well in advance, medium-notice typically allows four hours, and short-notice gives as little as 15 minutes of notice; flight delays allow for approximately five hours' notice.

During DOZE Optimization, our algorithm attempts to give the best possible escort schedule and update it as events unfold. Every minute, the simulation feeds all known information to the Doze algorithm, including certain flight delays and NAP connections, current escort and wheelchair locations, etc. DOZE runs on this data and supplies instructions to the escorts, who begin to act on these instructions. Their traveling generates new events (such as delays in traveling from gate to gate), which then feed back into the Escort Schedule for the next round of DOZE optimization.

7 Conclusions

Unfortunately, times are busy at Control Team 22 Consulting, and work on this bid could not begin until a mere 96 hours before the bid was due. As developing, programming, and implementing stochastic models is a very time consuming process, we have not been able to obtain the simulation results which we intended. In particular, while simulation would upon completion provide a detailed breakdown of costs associated with the DOZE algorithm, at this point we are not in a position to make specific budget recommendations for Epsilon Airlines.

However, based on the merits and characteristics of the DOZE algorithm, we are able to make some predictions about its scalability and effectiveness in different airports. The zoned escorts are necessary for adequate wheelchair redistribution and short-notice coverage, but in smaller airports where non-ambulatory passengers are less prevalent, incidents occur less often. Thus hiring Zoned escorts and keeping extra wheelchairs available amounts to an extra cost, and increases the overall cost per passenger.

However, in larger airports, where there are many more NAPs and the Daily-Ordered Escorts are more closely scheduled, the presence of Zoned escorts ensures that all NAPs are taken care of even in the case of airport delays or short notice. In this way, the cost of the Zoned Escorts is minimal compared to the averted costs of delaying flights or failing to pick up a passenger needing assistance. Thus we can expect our algorithm to be more cost-effective in Epsilon Airline's larger airports.

Furthermore, in the future as the fraction of passengers who require mobility assistance increases, the presence of Zoned Escorts allows Epsilon Airlines to hire only a minimal number of new Daily-Ordered Escorts, thereby keeping costs low. Thus we expect that the DOZE algorithm will only increase in cost-efficiency in the future.

References

[1] Adams, Marilyn. (2005) "Day 1 of Delta's Atlanta overhaul." *USA TODAY* http://www.usatoday.com/money/biztravel/2005-02-01-clockwork-usat_x.htm, Accessed on February 5, 2006.

[2] *Airline Careers: The people of the airlines.* Chronicle Guidance Publications, CGP Reprint 514

[3] *Airline On-Time Statistics and Delay Causes.* (2005) Bureau of Transportation Statistics. Raw data: http://www.transtats.bts.gov/OT_Delay/ot_delaycause1.asp?display=download &pn=0&month=12&year=2005, Accessed February 4, 2006.

[4] *Airport Tips.* Delta Airlines Webpage. http://www.delta.com/traveling_checkin/travel_tips/airport/index.jsp, Accessed February 6, 2006.

[5] *System Timetable* (2006) Alaska Airlines. http://www.alaskatimetable.com/AS.pdf, Accessed February 4, 2006.

[6] Atlanta Airport Terminal Map. http://www.atlanta-airport.com/default.asp?url=http://www.atlanta-airport.com/sublevels/terminal/termmap.htm, Accessed on February 4, 2006.

[7] Bureau of Transportation Statistics (2003). *Freedom to Travel.* Washington, DC: US Dept. of Transportation, http://www.bts.gov/publications/freedom_to_travel/pdf/entire.pdf, Accessed February 4, 2006.

[8] *The National Economic Impact of Civil Aviation.* (July 2002) DRI•Wefa, Inc.

[9] *New Horizons.* (2004) US Department of Transportation. http://airconsumer.ost.dot.gov/publications/HorizonsPrintable.doc, Accessed February 4, 2006

[10] *Total Cost for Air Carrier Delay* (1994) APO-130, Economic Analysis Branch

[11] *Transportation and Disability in Canada An Overview.* (1995) Goss-Gilroy Inc. TP 12545E

[12] *Worldwide Timetable.* (2006) Delta Airlines. http://images.delta.com/delta/pdfs/flight_schedules.pdf, Accessed on February 3, 2006