Improving Robustness through Buffer Placement in the Gate Assignment Problem: An Empirical Approach

Kevin Zimmer, Hamsa Balakrishnan, Amedeo Odoni
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, United States of America
bowserjr@mit.edu, hamsa@mit.edu, arodoni@mit.edu

***Abstract*—The allocation of flights to arrival and departure gates is one of the most important tasks performed daily by airport operators. The great majority of existing research develops models for optimizing gate assignments for various specified objective functions. A critical question in this respect is the robustness of these solutions in view of pervasive uncertainty regarding flight delays and gate occupancy times. In this paper, we apply an optimization algorithm to historical data to empirically examine the tradeoff between robustness (measured by the number of flights and passengers subjected to a gate change) and intra-terminal passenger walking times. We explore how different values of buffer time (i.e., the enforced time of vacancy between consecutive occupancies of a gate) affect this tradeoff as well as gate utilization metrics. Using data from a major hub airport, we show that a considerable decrease in the number of gate changes and of negatively affected passengers can be achieved at the expense of just a modest increase in passenger walking times. At the same time, setting buffer times at excessive levels may lead to infeasibility in developing a schedule of gate assignments that accommodate all flights.**

***Keywords*—airport management and operations; gate assignment problem; robustness-efficiency tradeoffs; buffer; slack**

I. Introduction

Gate assignments drive the flows of passengers through terminal buildings and of aircraft through the airport's network of taxiways, and shape passenger perceptions of level of service. Gates are also an expensive and highly-constrained resource: In 2023, an expansion of Boston Logan International Airport's Terminal E that cost nearly $650M was able to add just four additional international gates [1]. The efficient use of gate resources is therefore important if such costs are to be recovered within a reasonable number of years. However, while high demand and tight gate schedules with no slack will increase utilization, they will also result in cascading delays when disruptions occur.

The problem of gate assignment (also known as gate allocation) at an airport is a complex process that involves a variety of stakeholders, objective functions, and operational constraints. Stakeholders in the process include the passengers, the airlines, and the airport [2]. The different stakeholders also differ in their objectives, which are often at odds with each other. Passengers generally seek a minimum possible intra-terminal walking distance. Airlines may prefer gates with easy access to the runways for expedient arrivals and departures. And airports may want heavier utilization of gates that are close to shopping and dining amenities, so as to maximize these sources of revenue. Gate allocation is also subject to a number of operational constraints. For instance, at many major airports, international flights are typically limited to a specific set of gates connected to an airport's international zone. Gates also come in different sizes, and a given aircraft can only be assigned to a gate which can accommodate it. Constraints may also vary by airport based on local practice: for example, there has been a trend towards gate sharing among airlines at major hubs in Europe and Asia, while US airports largely still constrain sets of gates to a single carrier's operations. The recognition of the central role played by gate assignment in airport operations has led to several recent efforts, both in the research and practitioner communities, to address the underlying challenges [3], [4].

Gate allocation at a large airport is typically carried out in two steps. In Step 1, performed a day in advance or very early in the morning of the day-of-operations, an initial gate assignment plan is prepared for all the flights scheduled for that day based on the nominal arrival and departure flights plus, possibly, any other available information on the particular conditions expected on the day in question. In Step 2, the initial assignments are revised during the day as new (real-time) information about flight delays and other developments becomes available. Thus, a critical goal is to develop a "robust" initial assignment in Step 1 that will minimize the number of gate changes (and of passengers affected) during the day of operations (Step 2).

A. Prior work

The classical and most-studied criterion in gate assignment is the intra-terminal passenger walking distance [5]–[8]. More

This research is supported by the National Research Foundation (NRF), Prime Minister's Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) programme. Mens, Manus and Machina (M3S) is an interdisciplinary research group (IRG) of the Singapore MIT Alliance for Research and Technology (SMART) centre.

recently, there have been multiple studies which focus on the robustness of gate schedules, in recognition of the dynamic and uncertain nature of the airport environment [9]. Kim et al. [10] examined the tradeoffs between three optimality metrics, which were intra-terminal passenger transit time, aircraft taxi time, and the total duration of gate conflicts (a gate conflict occurs when an aircraft must wait for its gate because the previous aircraft has not yet left). Here, the gate conflict duration metric was used as a measure of the robustness of the gate assignment. Yu et al. proposed an adaptive large neighborhood search algorithm methodology to the gate assignment problem which factored in robustness [11]. Dell'Orco et al. demonstrated a Fuzzy Bee Colony Optimization metaheuristic algorithm to finding a robust solution to the gate assignment problem [12]. Diepen et al. introduced a novel integer linear programming model to come up with robust gate assignments for the flights at Amsterdam's Schiphol Airport [13]. Other efforts have explored stochastic, robust and chance-constrained optimization formulations of the gate assignment problem [14]–[16]. The gate assignment process as currently implemented at airports is quite laborious and time-consuming, and requires significant human inputs. This has motivated the investigation of various learning- and search-based heuristics to address the problem [17], [18].

B. This paper

We take a step back from the prior efforts described above, which have focused on developing robust formulations to the gate assignment problem. Instead, we explore what is, in many ways, the simplest way to improve the robustness of a schedule: adding slack or buffer time to mitigate the need to re-plan in the event of even small delays [19]–[21]. We investigate the potential of adopting this simple approach at a large airport. To answer this question, we consider a hypothetical airport loosely modeled on Dallas Forth Worth International Airport (DFW), but assuming the prevalence of common-use gates. We determine the initial gate assignments using an optimization formulation that takes the planned flight schedules and operational constraints as input, for different values of buffer time (i.e., the required slack between consecutive occupancies at any gate). We consider buffer times ranging from zero (i.e., no slack) to 30 min. The intra-terminal walking time is the objective of this initial optimization. We then consider what actually happened on a given day, in terms of the actual flight arrival and departure times. As the initial gate schedules may no longer be feasible, we reoptimize the gate assignments with the objective of minimizing the number of passengers who will experience a gate change from the initial schedule. Repeating this process for each day during a 3-month period, we empirically evaluate the tradeoffs between robustness (measured by the number of passengers and flights that are subjected to gate changes) and intra-terminal passenger walking time for different values of buffer time. We also evaluate the impact of buffer size on gate utilization metrics.

II. Airport setup

For our empirical study, we consider an airport that is based on Dallas Fort Worth International Airport (DFW), but assume that the gates are all shared among airlines (i.e., they are common-use), which is a departure from the actual operations at airports in the United States, including DFW. Fig. 1 illustrates the layout of the airport, including the approximate locations of gates. Using maps of the airport layout, we estimate the intra-terminal walking times between gates and various waypoints in the airport's terminals.

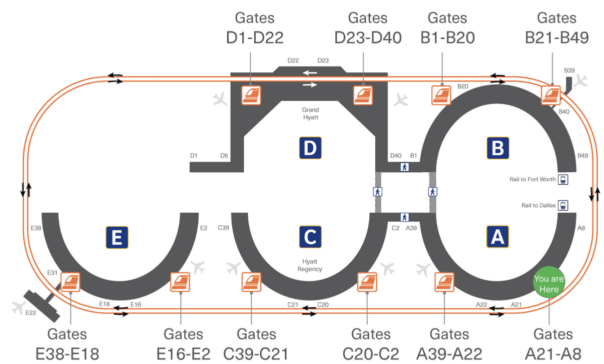

Figure 1. Layout of DFW terminals, showing the approximate locations of gates [22]. This paper focuses on Terminals A, B and C.

A. Gate attributes and constraints

We consider four characteristics for each gate: (1) the name, or unique identifier; (2) the size of aircraft that each gate can accommodate; (3) the type of flight (domestic or international) that a gate can accommodate; and (4) the estimated walking time from the nearest security screening checkpoint. We assume three aircraft sizes, corresponding to regional aircraft, narrow-body aircraft, and wide-body aircraft. An aircraft can be assigned to a gate of a matching or larger size (e.g., a regional aircraft can be assigned to any gate, but a wide-body aircraft cannot occupy a narrow-body gate). We assume that DFW is set up such that departing international passengers need to go through exit passport control and are thus segregated from departing domestic passengers, meaning that domestic flights cannot use international gates. This is not true in reality at DFW, as the United States does not enforce exit immigration controls, but is commonly the case outside of the United States. We note in Fig. 1 that Terminal E is detached from the other terminals. Terminal D serves the majority of International flights at DFW. For these reasons, we focus this initial investigation on only the gates in Terminals A, B and C. In other words, we consider 96 of the 171 gates at DFW. As a further simplifying assumption, we only consider departing passengers, all of whom start their trip at DFW. Consequently, the walking times from the gate to baggage claim or from one gate to another are not considered. Future work will consider scaling up the process to include operations in Terminals D and E, as well as arriving and transfer passengers.

B. Flight data archives

We consider flights operating at DFW in a 3-month period (May-July 2016). This period covers the busiest travel months of the year. The data set includes the airline code and flight number, the origin and destination airports (one of which is DFW), the unique registration number ("tail number") of the aircraft which operated the flight, the scheduled arrival and departure times, the actual arrival and departure times at the gate, as well as the gate that was actually used by each flight [23].

C. Pre-processing

The flight dataset is pre-processed prior to optimization. First, for the aircraft associated with each departing flight, we identify the corresponding arriving flight. This represents an aircraft turnaround at DFW. If the arrival time of that aircraft is within 4 hours of its departure time for the outbound leg, we assume that the aircraft remained at its gate from arrival until its subsequent departure, and record the arrival time of the inbound leg as the gate occupancy start time. Otherwise, we assume that the aircraft starts occupying the gate 90 minutes prior to the flight's scheduled departure time. In other words, we assume that in the case of turnaround times greater than 4 hours (e.g., overnight), the aircraft is moved to a remote hangar and brought back to the gate 90 min prior to its scheduled departure time. In addition to only considering departures, we also filter out cargo flights and duplicated code-share listings. For each flight, we infer an aircraft size (regional, narrow-body, or wide-body) as well as its domestic/international status using the origin/destination data and gate used. Future work could consult a fleet registry database for this purpose [24]. For simplicity, we assume a number of passengers based on the aircraft size: 75, 150, and 300 for regional, narrow-body, and wide-body aircraft, respectively. While we process data for all flights in the 3-month period, the analysis of this paper is limited to the flights that actually departed from Terminals A, B or C, as discussed in Section II-A.

III. Gate Assignment Optimization

For each day in the dataset, we use the flight schedules along with the scheduled gate occupancy start and end times as input to produce an initial gate schedule. In this initial schedule optimization, we seek a gate assignment plan which minimizes the average passenger walking time. Recognizing that actual operations will almost inevitably deviate from flight schedules, we produce initial gate assignment schedules that incorporate a buffer time between consecutive occupancies at each gate. This represents Step 1 of the process described in Section I.

To mimic what would need to happen on the day-of-operations once delays occur, we rerun the gate assignment optimization algorithm using the actual gate occupancy start and end times. The constraints are the same as before, but this time, our objective is to minimize the number of passengers subjected to a gate change relative to the initial gate schedule. We note that in reality, flight delay information (i.e., the actual arrival/departure times) is only known as the day unfolds and not all-at-once. However, the solution we obtain by rerunning the optimization helps us benchmark the potential benefits that could be achieved in Step 2 (Section I), and we show that, through use of a buffer, a significant reduction in passenger gate changes can be achieved at relatively little expense in walking time.

A. Notation

The following notation is used in our formulation:

$\mathcal{F}$	Set of all scheduled flights
$\mathcal{G}$	Set of all gates
$\texttt{arr}_{f,\text{sch}}$	Scheduled gate arrival time for a flight $f \in \mathcal{F}$
$\texttt{dep}_{f,\text{sch}}$	Scheduled time at which a flight $f \in F$ will have vacated its gate.
$\texttt{arr}_{f,\text{act}}$	Actual gate arrival time for a flight $f \in \mathcal{F}$
$\texttt{dep}_{f,\text{act}}$	Actual time at which a flight $f \in \mathcal{F}$ will have vacated its gate
M	Decision variable. $\mathcal{F} \times \mathcal{G}$ matrix where entry (i,j) is 1 if flight i is assigned to gate j, and 0 otherwise. This matrix represents the initial gate assignment plan.
$\texttt{buffer}$	Enforced buffer time between subsequent occupancies of the same gate in Step 1.
$\texttt{ac}_f$	Numerical indicator of aircraft type that is used to operate $f \in \mathcal{F}$. $\texttt{ac}_f \in \{1,2,3\}$ corresponding to wide-body, narrow-body or regional
$\texttt{gc}_g$	Numerical indicator of largest type of aircraft that gate $g \in \mathcal{G}$ can accommodate. $\texttt{gc}_g \in \{1,2,3\}$ corresponding to wide-body, narrow-body or regional
$\texttt{FT}_f$	Categorical variable for the type of destination that flight $f \in \mathcal{F}$ is going to, which can be *Domestic* or *International*
$\texttt{GT}_g$	Categorical variable for the type of flight that gate $g \in \mathcal{G}$ can serve, which can be *Domestic*, *International* or *Swing*
W_g	Numerical variable representing the walking time to reach gate $g \in \mathcal{G}$ from the closest security screening checkpoint
P_f	Numerical variable representing the passenger capacity of flight $f \in F$
U	Decision variable. Binary $\mathcal{F} \times \mathcal{G}$ matrix, where entry (i,j) is 1 if flight i is assigned to gate j, and 0 otherwise. This matrix represents the updated gate assignments, once the actual flight times (delays) are known.

B. Constraints

We incorporate the following constraints in the optimization formulation:

1) *Two flights that need to occupy a gate at the same point in time cannot be assigned to the same gate.* In other words, $\forall f_1, f_2 \in \mathcal{F}$ such that $(\texttt{arr}_{f_1,\text{sch}} < \texttt{dep}_{f_2,\text{sch}} + \texttt{buffer}) \wedge (\texttt{dep}_{f_1,\text{sch}} + \texttt{buffer} > \texttt{arr}_{f_1,\text{sch}})$, $M_{f_1,g} + M_{f_2,g} \leq 1$ $\forall g \in \mathcal{G}$.
2) *Every flight must be assigned to exactly one gate*, i.e., $\sum_{g\in\mathcal{G}} M_{f,g} = 1, \quad \forall f \in \mathcal{F}$.
3) *A gate cannot be used by an aircraft which is larger than it is able to accommodate,* i.e., $\forall f \in \mathcal{F}, g \in \mathcal{G}$ such that $\texttt{ac}_f < \texttt{gc}_g$, $M_{f,g} = 0$.
4) *An international gate cannot be used by a domestic flight*, i.e., $\forall f \in \mathcal{F}, g \in \mathcal{G}$ such that $(\texttt{FT}_f = Domestic) \wedge (\texttt{GT}_g = International)$, $M_{f,g} = 0$.
5) *A domestic gate cannot be used by an international flight,* i.e., $\forall f \in \mathcal{F}, g \in \mathcal{G}$ such that $(\texttt{FT}_f = International) \wedge (\texttt{GT}_g = Domestic)$, $M_{f,g} = 0$.

C. Objective Function

In the optimization to determine the initial gate assignments using the scheduled arrival and departure times (Step 1), the objective function is as follows:

$$\min \sum_{f\in\mathcal{F},g\in\mathcal{G}} W_g P_f M_{f,g} \tag{1}$$

This objective function minimizes the total passenger walking time.

D. Update Process

After the initial gate assignment plan is created, we re-plan the gate assignment using the actual gate occupancy start and end times and disregarding the buffer. The objective function used is as follows:

$$\min \alpha \sum_{f\in F,g\in G} |M_{f,g} - U_{f,g}| P_f + (1-\alpha) \sum_{f\in F,g\in G} W_g P_f U_{f,g} \tag{2}$$

The parameter α can range between 0 and 1, and represents the relative weights given to passengers who experience gate changes and the walking distance for all passengers. A high value of α will place a relatively high penalty on gate changes (one of the goals of this work) relative to the passenger walking distance, while a smaller value will result in a solution similar to that given by (1). We also note that the maximum difference in walking time between 2 gates in our layout is 8 minutes. For the airport layout that we consider in this paper, we find that the solutions do not change vary much for $\alpha \geq 0.7$. We use $\alpha = 0.999$ in the analyses that follow below, but recognize that the choice of weights depends on the airport layout as well as the priorities of stakeholders.

IV. Analysis of Results

A. Impact of buffer size on passenger walking times

As the buffer size increases, we would expect that a larger set of gates with potentially longer walking times would need to be used. Fig. 2 shows the average passenger walking time corresponding to the initial gate schedule, for different values of buffer size. As expected, we observe an increase in the average passenger walking time with increasing buffer times. However, the increase is certainly not drastic: In fact, the difference in average passenger walking time between no buffer (i.e., a buffer time of 0 min) and a 30 min buffer is only about 10 seconds!

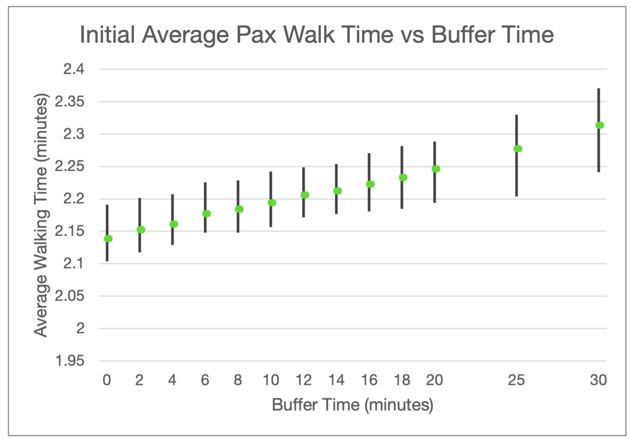

Figure 2. Average passenger walking time in the initial gate assignment plan for various buffer times. The green point is the mean, and the black bar shows the range from the 25th to 75th percentile.

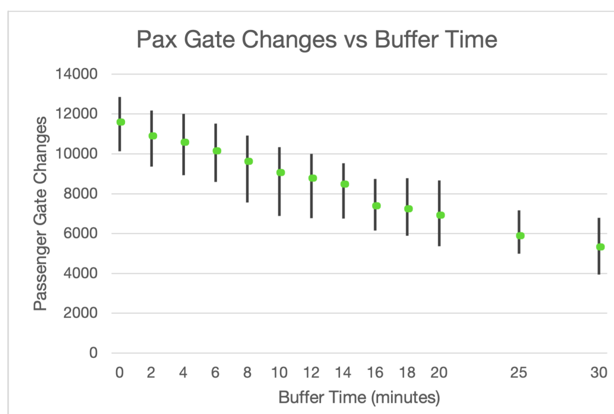

Figure 3. Number of passengers that experience a gate change when re-planning the initial gate assignment plan to account for delays. The green point is the mean, and the black bar shows the range from the 25th to 75th percentile.

B. Impact of buffer size on gate changes

The primary rationale behind adding a buffer was to improve schedule robustness, i.e., to decrease the number of passengers or flights that would be impacted by a gate change due to disruptions. We would therefore expect changes in buffer time to have a substantial impact on the number of passengers subjected to a gate change. Fig. 3 shows the number of passengers subject to a necessary gate change for different values of buffer time. With no buffer, we have an average of approximately 11,600 passengers who have to change gates, but with a 30 min buffer, this decreases to less than half of that amount, and we have on average only about 5,300 passengers needing to change gate!

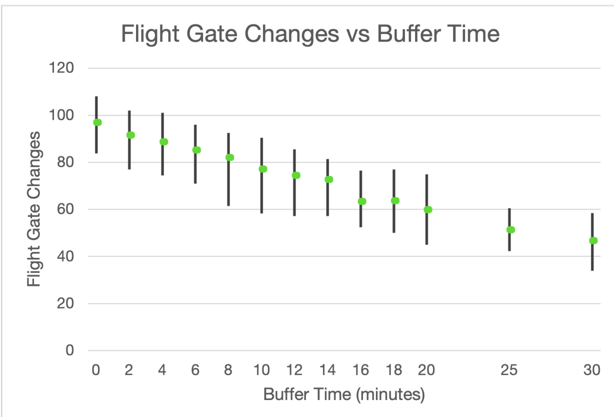

Figure 4. Number of flights that experience a gate change when re-planning the initial gate assignment plan to account for delays. The green point is the mean, and the black bar shows the range from the 25th to 75th percentile.

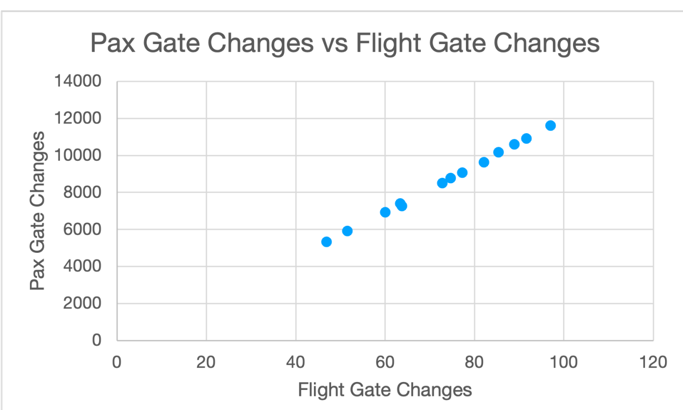

Figure 5. Scatterplot showing the number of flight gate changes and of passenger gate changes at the various buffer times.

Although the objective function (2) explicitly considers the number of passengers who experience gate changes, we also see a benefit in terms of the number of flights that require gate changes (Fig. 4). This result is not surprising when one considers the relationship between passenger gate changes and flight gate changes, as shown in Fig. 5. In this figure, we notice the apparent linear relationship between the number of flights requiring gate changes and the number of passengers similarly impacted. One possible reason for this behavior may be that with the many international (and therefore often wide-body) flights operating in Terminal D not being considered, the remaining fleet is quite homogeneous in terms of passenger capacity.

C. Tradeoffs between gate changes and walk times

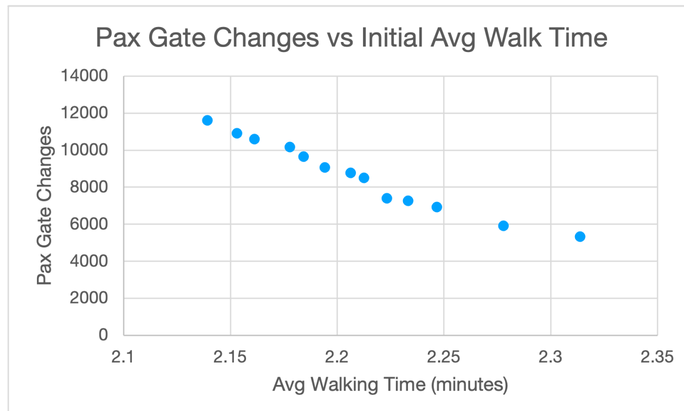

Figure 6. Scatterplot showing the initial average passenger walking time and the number of passenger gate changes at the various buffer times.

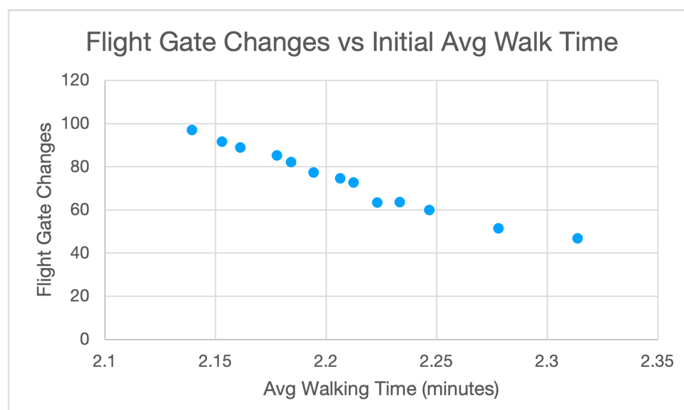

Figure 7. Scatterplot showing the initial average passenger walking time and the number of flight gate changes at the various buffer times.

Considering the previously discussed effects that buffer sizes have on both gate changes (i.e., decrease with increasing buffer size) and the initial passenger walk times (i.e., increase with increasing buffer size), we can evaluate the tradeoffs that exist between passenger/flight gate changes and passenger walk times. These tradeoffs are shown in Figs. 6 and 7. We notice that there are diminishing benefits to increasing buffer times beyond a certain point: while initially a significant improvement in passenger gate changes is seen with increasing buffer size (at the expense of a modest increase in passenger walk times), the slope of the tradeoff curve appears to change at higher buffer values, and there is not as much benefit in terms of decreasing necessary gate changes. For instance, from no buffer to a 10 minute buffer, passenger gate changes drop by 2532, then by a further 2150 when going to a 20 minute buffer, but only by 1590 more when increasing to a 30 minute buffer. Furthermore, the prevalence of infeasibility begins to rise as the buffer time increases beyond 30 minutes.

D. Impact of buffer size on gate utilization

As the buffer size increases, more time on a gate's schedule is set aside for a particular flight, in order to protect against potential flight delays. However, an increasing buffer will reduce the number of flights that can be scheduled at a particular gate. As gates are (at least in the United States) typically leased to airlines on an annual basis, these fixed costs are then spread across fewer flights, increasing the per-flight expense incurred for the use of the airport facilities. Furthermore, if gates are underutilized, the operations of the airline will be more spread out, requiring more personnel and increasing operating costs, as well as increasing the complexity of the operation.

There are several candidate metrics by which one could evaluate gate utilization. Here, we consider the ratio of the actual gate occupancy time of a flight to the time allocated to it in the initial gate assignment plan, averaged over all the flights. Using this definition and the notation introduced in Section III-A, the gate utilization of flight f is given by:

$$\texttt{util}_f = \frac{\texttt{dep}_{f,\text{act}} - \texttt{arr}_{f,\text{act}}}{\texttt{dep}_{f,\text{sch}} - \texttt{arr}_{f,\text{sch}} + \texttt{buffer}} \tag{3}$$

The average gate utilization is then calculated by averaging the ratio in (3) over all flights in the dataset. Fig. 8 shows the variation of gate utilization with buffer size. Interestingly, we notice that the average gate utilization is 110% for buffer size zero; in other words, the gates are actually being used for 10% longer than scheduled on average. In the absence of a buffer, such schedules are brittle and necessitate costly gate changes. Fig. 8 indicates that for a buffer size of 14 min, the average gate utilization decreases to a more manageable–yet high–value of 90%. Furthermore, our previous analyses indicate that this increase in robustness ($\approx$ 30% decrease in passenger gate changes) will be at only a modest increase (5% or 6 seconds) in the average passenger walking times.

An alternative definition of gate utilization considers the amount of gate time "set aside" or "reserved" for gate occupancy (including the buffer time) as a ratio of the total gate time available from the 96 gates at the airport. As the buffer size increases, so does this measure of gate utilization. However, we

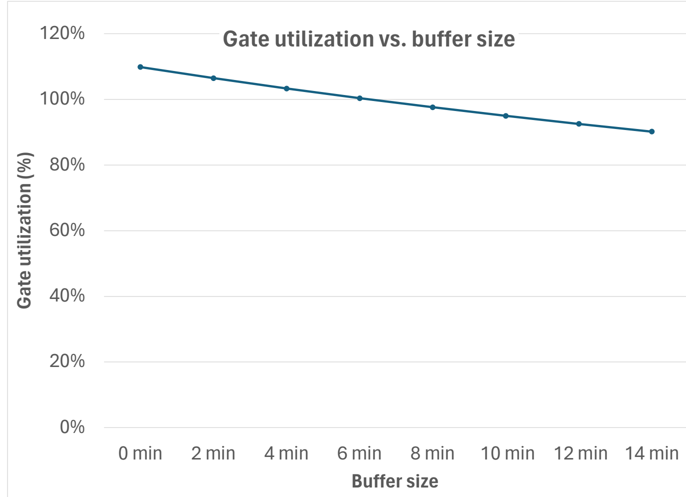

Figure 8. Average gate utilization (as given by (3)) as a function of the buffer size.

note that this quantity is inversely correlated with efficiency; an unnecessarily large amount of resource reserved for a flight may be overly conservative and even wasteful.

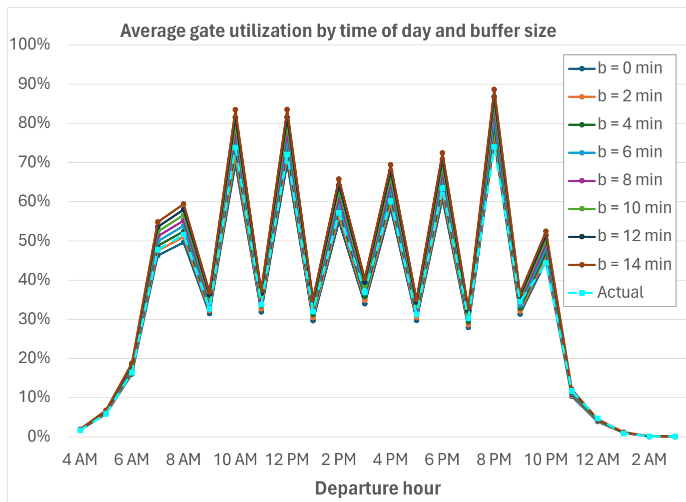

Figure 9. Gate utilization by time-of-day, as a function of the buffer size. This measure of gate utilization considers the fraction of available gate resources "set aside" to serve the given flight demand, and increases with buffer size.

Fig. 9 plots this measure of gate utilization as a function of the time-of-day. The peaks and valleys mirror the departure banks at DFW. We notice that during the periods of highest demand (e.g., for flights departing during the 8PM hour), nearly 90% of all available gate resources have been reserved for the scheduled traffic. This helps explain our observation that several days do not have a feasible solution to the gate assignment problem at buffer times of more than 30 minutes: there simply aren't sufficient gate resources to set aside a large buffer for each flight, especially during periods of peak demand.

In summary, the relatively modest increase in average walking time, together with the significant decrease in passenger gate changes achieved with an increased buffer time highlight the value of planning a gate assignment which is robust to delays. On the other hand, the amount of time during which gates are idle (i.e., are not occupied) increases as buffer times increase, i.e., the measure of gate utilization given by (3) decreases.

E. An example: Daily schedule at a gate

To illustrate how this works in practice, we consider the extract of the gate schedules on a particular day (01 May 2016) in our dataset. We focus on what happens at Gate C6, a gate which is very close to the security screening checkpoint and thus considered highly desirable. Table I shows the initial schedule when the required buffer time between consecutive gate occupancies is set equal to zero, as well as the actual arrival and departure times of these flights; it also shows (rightmost column) the gates that the scheduled flights actually left from, as a result of the delays that took place on that day.

It can be seen that Gate C6, in the zero-buffer case, was initially scheduled to serve 13 flights, i.e., to be very heavily utilized throughout the day. However, as a result of that day's delays, C6 ended up actually serving only 9, with the other 4 being forced to change gate. Clearly the initial "optimal" schedule was not robust in this case.

Table II next shows what happens if, instead, the initial schedule requires a buffer time of 20 minutes. In that case, only 11 flights are scheduled at the highly-desirable gate C6. But none of these flights needs to change gates after accounting for delays on that day and all 11 flights end up using C6. Although 2 fewer flights were scheduled to use C6, eventually 2 more flights actually used it than in the no-buffer case, and no passengers needed to change gates.

V. Conclusions and Future Work

This paper explored the use of a simple approach –buffer placement– to improve the robustness of gate schedules, using data based on a major airport, DFW. We found that buffer placement has the potential to decrease the number of passengers or flights impacted by disruptive gate changes by more than 50%, with only a modest increase of about 10 seconds in the average walk times. These promising findings motivate the further investigation of buffers as a way to improve robustness without significant sacrifice in efficiency.

In the work presented, we considered only the departing passengers, and we assumed that they all began their trips at DFW. In reality, 70% of DFW's passengers (and a large percentage at most other hub airports) are connecting passengers. For connecting passengers, the walking distance would correspond to the distance between the gates at which they disembark and embark at DFW. This exercise would then further extend, with some complications, to passengers connecting to, or from, international flights. We are currently in the process of incorporating information on passenger connectivity [25]. Our next task will be to account for connecting passengers in our gate assignment plan.

We also used a constant buffer between consecutive occupancies of any given gate, which assumes that every flight is equally likely to be delayed (or to arrive at the gate early). This is certainly not true in practice. Using historical data about the on-time performance of a given flight, we can compute each flight's delay time distribution, and can use this information to create a customized buffer between two flights assigned to the same gate specifically calibrated to keep the chance of a gate conflict below a given probability. Such variable buffer sizing will result in a more efficient use of gate resources by using large buffers only for flights that are particularly prone to delays.

Finally, we note that the Step 2 optimization, as presented in this paper, provides only an upper bound on the performance of the re-planning algorithms. An important challenge remains in the development of approaches for effective real-time re-planning of gate assignments to account for delays as they materialize.

References

[1] Boston Globe, "Logan's Terminal E architect explains the new 'Boston Red' building and fourth-generation airports," https://www.boston.com/news/the-boston-globe/2023/04/14/new-red-terminal-e-logan-airport/, 2023.

[2] G. S. Das, F. Gzarab, and T. Stützle, "A review on airport gate assignment problems: Single versus multi objective approaches," *Omega*, vol. 92, no. 10, April 2020.

[3] Z. Liu and Q. Xiang, "A branch-and-price algorithm for the airport gate assignment problem considering the trade-off between robustness and efficiency," *Transportation Research Part C: Emerging Technologies*, vol. 154, p. 104232, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0968090X23002218

[4] CNBC, "How airlines are shaving minutes off flight times to save millions," https://www.cnbc.com/2023/11/23/airlines-shave-minutes-off-flight-times.html, 2023.

[5] J. Braaksma, "Reducing Walking Distances at Existing Airports," *Airport Forum*, August 1977.

[6] O. Babić, D. Teodorović, and V. Tošić, "Aircraft stand assignment to minimize walking," *Journal of Transportation Engineering*, vol. 110, 1984.

[7] S. Wirasinghe and U. Vandebona, "Passenger walking distance distribution in single- and dual-concourse centralized airport terminals," in *Transportation Research Record*, 1988.

[8] R. S. Mangoubi and D. F. X. Mathaisel, "Optimizing gate assignments at airport terminals," *Transportation Science*, vol. 19, no. 2, pp. 173–188, 1985. [Online]. Available: http://www.jstor.org/stable/25768171

[9] M. E. Narciso and M. A. Piera, "Robust gate assignment procedures from an airport management perspective," *Omega*, vol. 50, pp. 82–95, 2015. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0305048314000760

[10] S. H. Kim, E. Feron, J.-P. Clarke, A. Marzuoli, and D. Delahaye, "Airport Gate Scheduling for Passengers, Aircraft, and Operations," *Journal of Air Transportation*, vol. 25, October 2017.

[11] C. Yu, D. Zhang, and H. Y. Lau, "An adaptive large neighborhood search heuristic for solving a robust gate assignment problem," *Expert Systems with Applications*, vol. 84, pp. 143–154, October 2017.

TABLE I
EXTRACT OF GATE SCHEDULES WHEN BUFFER EQUALS 0 MIN, FOR EXAMPLE CONSIDERED IN SECTION IV-E.

Flight No.	Sch. Occ. Start	Sch. Occ. End	Act. Occ. Start	Act. Occ. End	Initial Assigned Gate	Final Assigned Gate
AA2255	5:35AM	7:05AM	5:36AM	7:06AM	C6	C6
AA1035	7:44AM	9:10AM	7:58AM	9:05AM	C6	A35
AA1561	9:10AM	10:05AM	9:00AM	10:04AM	C6	C6
AA259	10:09AM	11:05AM	9:49AM	11:56AM	C6	C28
AA1560	11:26AM	12:20PM	11:18AM	12:21PM	C6	C6
AA1198	12:28PM	1:30PM	12:25PM	1:23PM	C6	C6
AA5829	1:59PM	2:40PM	1:35PM	2:41PM	C6	C6
AA2734	2:40PM	3:50PM	2:48PM	3:51PM	C6	C6
AA1072	3:55PM	5:25PM	3:49PM	5:29PM	C6	C6
AA2371	5:39PM	6:20PM	5:39PM	6:19PM	C6	C6
AA1075	6:21PM	7:15PM	6:03PM	7:11PM	C6	A34
AA551	7:24PM	8:35PM	7:24PM	9:01PM	C6	B10
AA1273	9:04PM	10:00PM	8:45PM	11:30PM	C6	C6

TABLE II
EXTRACT OF GATE SCHEDULES WHEN BUFFER EQUALS 20 MIN, FOR EXAMPLE CONSIDERED IN SECTION IV-E.

Flight No.	Sch. Occ. Start	Sch. Occ. End	Act. Occ. Start	Act. Occ. End	Initial Assigned Gate	Final Assigned Gate
AA2064	5:30AM	7:00AM	5:23AM	6:53AM	C6	C6
AA277	7:20AM	8:50AM	7:17AM	8:47AM	C6	C6
AA149	9:10AM	10:20AM	9:14AM	10:16AM	C6	C6
AA2564	10:55AM	12:15PM	10:32AM	12:13PM	C6	C6
AA1599	1:00PM	1:55PM	12:52PM	1:51PM	C6	C6
AA2302	2:20PM	3:20PM	2:08PM	3:16PM	C6	C6
AA301	3:48PM	4:50PM	3:21PM	4:49PM	C6	C6
AA2385	5:35PM	5:55PM	5:12PM	5:55PM	C6	C6
AA1075	6:21PM	7:15PM	6:03PM	7:11PM	C6	C6
AA1408	7:58PM	8:40PM	7:55PM	8:36PM	C6	C6
AA2367	9:00PM	9:55PM	8:50PM	9:52PM	C6	C6

[12] M. Dell'Orco, M. Marinelli, and M. G. Altieri, "Solving the gate assignment problem through the Fuzzy Bee Colony Optimization," *Transportation Research Part C: Emerging Technologies*, vol. 80, pp. 424–438, July 2017.

[13] G. Diepen, J. M. van den Akker, J. A. Hoogeveen, and J. W. Smeltink, "Finding a robust assignment of flights to gates at Amsterdam Airport Schiphol," *Journal of Scheduling*, vol. 15, pp. 703–715, October 2012.

[14] M. Seker and N. Noyan, "Stochastic optimization models for the airport gate assignment problem," *Transportation Research Part E: Logistics and Transportation Review*, vol. 48, no. 2, pp. 438–459, 2012. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1366554511001335

[15] L. Xu, C. Zhang, F. Xiao, and F. Wang, "A robust approach to airport gate assignment with a solution-dependent uncertainty budget," *Transportation Research Part B: Methodological*, vol. 105, pp. 458–478, 2017. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0191261517303648

[16] J. Kim, B. Goo, Y. Roh, C. Lee, and K. Lee, "A branch-and-price approach for airport gate assignment problem with chance constraints," *Transportation Research Part B: Methodological*, vol. 168, pp. 1–26, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0191261522002077

[17] M. Muhafiz Yıldız, U. Avcı, M. Örnek, and C. Öztürk, "Flight gate assignment problem with reinforcement learning," in *Intelligent and Fuzzy Systems. INFUS 2023*, 2023.

[18] M. Li, J.-K. Hao, and Q. Wu, "Learning-driven feasible and infeasible tabu search for airport gate assignment," *European Journal of Operational Research*, vol. 302, no. 1, pp. 172–186, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0377221721010456

[19] S. AhmadBeygi, A. Cohn, Y. Guan, and P. Belobaba, "Analysis of the potential for delay propagation in passenger airline networks," *Journal of Air Transport Management*, vol. 14, pp. 221–236, September 2008.

[20] V. Chiraphadhanakul and C. Barnhart, "Robust flight schedules through slack re-allocation," *EURO Journal on Transportation and Logistics*, vol. 2, no. 4, pp. 277–306, 2013. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2192437620301126

[21] A. J. Davenport, C. Gefflot, and J. C. Beck, "Slack-based techniques for robust schedules," in *Proceedings of the Sixth European Conference on Planning*, 2014.

[22] "Connect/Transfer at DFW Airport: Walk or Ride?" https://www.dfwairport.com/explore/plan/connect/walkride/, accessed: 2024-02-15.

[23] OAG, "Airport operations data," https://www.oag.com, 2016.

[24] Bucher and Company, "JP Airline Fleets International."

[25] C. Barnhart, D. Fearing, and V. Vaze, "Modeling passenger travel and delays in the national air transportation system," *Operations Research*, vol. 62, no. 3, pp. 580–601, 2014. [Online]. Available: https://doi.org/10.1287/opre.2014.1268