Journal of Air Transport Management 94 (2021) 102077

Contents lists available at ScienceDirect

Journal of Air Transport Management

journal homepage: www.elsevier.com/locate/jairtraman

Route optimization for energy efficient airport shuttle operations – A case study from Dallas Fort worth International Airport

Devon Sigler *, Qichao Wang, Zhaocai Liu, Venu Garikapati, Andrew Kotz, Kenneth J. Kelly, Monte Lunacek, Caleb Phillips

National Renewable Energy Laboratory, 15013 Denver West Parkway, Golden, CO 80401, USA

ARTICLE INFO

Keywords:
Air travel
Bus schedule optimization
Energy efficiency
Airport shuttles
Discrete event simulation

ABSTRACT

The objective of this research is to reduce energy consumption from intra airport shuttle operations by optimizing routes and schedules, without compromising on passenger travel experience. To achieve this objective, we propose an optimization model that generates optimal airport shuttle routes for a given set of constraints and a discrete-event simulator that evaluates the optimal shuttle routes in a stochastic environment to understand the tradeoffs between the amount of time passengers wait for shuttles, and shuttle energy consumption. The proposed optimization model and stochastic simulation are tested using shuttle route data provided by the Dallas Fort Worth International Airport. Results indicate that optimized routes can lead to a 20% energy reduction in shuttle operations with a modest 2-min increase in average shuttle wait times. The optimization model and simulator presented here are designed to be generalizable and can be adapted to optimize shuttle operations at any major airport.

1. Introduction

Airports, like Dallas Fort Worth International (DFW),[1] are considered ‘special generators’ since a considerable portion of any cities’ traffic originates or terminates at an airport. Understandably, air travel holds a major stake in long-distance travel, particularly for trips greater than 1000 miles from the point of origin (Bureau of Transportation, 2017). With air travel on the rise in pre-COVID times (Henao et al., 2018) and seeing a steady growth as travel is slowly returning to a new normal in light of the pandemic (Chokshi, 2020), it is reasonable to expect continued growth in passenger traffic to airports. Airport planning and operations focus on bridging the gap between a passenger reaching the airport and taking flight. Aspects of airport planning and operations span ensuring reliable ground access to airports (Malandri et al., 2017), ease of baggage handling (Budd et al., 2014), curb management, and infrastructure planning, not to mention the myriad airside operations. One such aspect of airport operations is transporting passengers to and from airport parking lots or rental car centers to the terminals. Within airport shuttle buses account for a considerable portion of the energy expenditure and emissions (that are within the control of the airport operators) at the airport. For example, the rental car shuttle fleet at DFW uses over 693,000 gasoline gallons equivalent (GGE) of compressed natural gas (CNG) and generates 5700 tons of CO_2 each year (Kotz et al., 2020). Therefore, optimizing within airport shuttle routes to minimize energy consumption, could increase the mobility and energy efficiency of airport ground transport operations.

A passenger’s experience in accessing an airport can be broken down into three mutually exclusive portions. The first part is airport ground access (labeled ‘travel to/from the airport’), i.e., accessing an airport from anywhere in the city. Passenger experience for this leg of the airport trip is influenced by factors such as reliability of ground access (Malandri et al., 2017), ease of baggage handling (Budd et al., 2014), cost of airport access (Jou et al., 2011), or level of safety and security in ground access modes (Budd et al., 2016). The second part is terminal

* Corresponding author.
E-mail addresses: devon.sigler@nrel.gov (D. Sigler), qichao.wang@nrel.gov (Q. Wang), zhaocai.liu@nrel.gov (Z. Liu), venu.garikapati@nrel.gov (V. Garikapati), andrew.kotz@nrel.gov (A. Kotz), kenneth.kelly@nrel.gov (K.J. Kelly), monte.lunacek@nrel.gov (M. Lunacek), caleb.phillips@nrel.gov (C. Phillips).

[1] Abbreviations: Dallas Fort Worth International Airport (DFW), gasoline gallons equivalent (GGE), compressed natural gas (CNG), vehicle routing problems (VRP), National Renewable Energy Laboratory (NREL), genetic algorithm (GA), depth-first search algorithm (DFS), Airport Shuttle Planning and Improved Routing Event-Driven Simulation (ASPIRES), Simulated Urban Mobility (SUMO), Controller Area Network (CAN), Spatial Positioning on Transit (SPOT), High Performance Computer (HPC), Gasoline Gallon Equivalents (GGEs), The maximum in-vehicle travel time (*ivtt*), headway (*hw*).

https://doi.org/10.1016/j.jairtraman.2021.102077
Received 25 September 2020; Received in revised form 16 April 2021; Accepted 21 April 2021
Available online 3 May 2021
0969-6997/

access (labeled ‘travel within the airport premises’), that is travel from rental car center or parking facilities to the terminal. However, passengers must wait for some amount of time, and travel in these shuttles for a few minutes to reach the terminal. So, time can be an influencing factor in passenger travel experience. The final part is gate access (labeled ‘travel inside the terminal). Aspects that affect passenger experience in this part of the airport trip are wait time at the ticketing counters, security wait time, and wait time at the gates. Aspects of ‘travel to the airport’ (such a transit fares), and ‘travel in the terminal’ (such as wait time at the ticketing counters) legs of the airport journey are often not in the control of airport authorities, making it difficult to address all the factors noted above using a unified optimization framework. This work intends to focus on the aspects that airport authorities do have control over, hence aspects of passenger ‘travel’ experience within the airport premises’ are chosen as the topic of research. In particular, the research effort aims to optimize airport shuttle operations with shuttle wait times and in-vehicle travel times as the indicators for passenger travel experience. We believe that the solutions proposed from this research are both easy to implement and are well within the purview of an airport operation team’s authority. A naïve solution to improve passenger travel experience (i.e., minimize shuttle wait times and ride times) would be to increase the number of airport shuttles; analogously, electrification of the airport’s shuttle fleet would reduce an airports energy footprint. However, increasing the number of shuttles to minimize passenger wait times is a sub-optimal and expensive solution, and electrifying an airport’s shuttle fleet requires large capital investments. By leveraging the power of data, modeling, and simulation, airport authorities can optimize existing shuttle operations for energy efficiency and passenger satisfaction.

Bus schedule and route optimization falls under the category of vehicle routing problems (VRP) where passengers with different origins and destinations are grouped to be served using a single vehicle. There is a wealth of literature on this topic both in the context of traditional transit route optimization (Zhao and Zeng, 2008; Yu et al., 2012), as well as taxi (Nunes et al., 2011), and ride hailing route optimization (Feng et al., 2014). However, the application of VRP in the context of airport shuttle route optimization is relatively sparse (Bao et al., 2018; Linqing et al., 2019), and research focusing on simultaneously tackling travel experience and energy/emission outcomes through vehicle routing is even sparser (Pei-Ying et al., 2013). One unique characteristic of the rental car center shuttle bus service at an airport is that the origin and destination pair can only be terminals and rental car center or rental car center and terminals. This characteristic allows us to better model the optimization and simulation problem with the existing data sources that are available through the airports. This paper adds to the limited literature on airport shuttle route optimization with customer satisfaction, and energy reduction handled simultaneously. This paper stands out from most of the existing literature on VRP and airport shuttle route optimization in that it is not a purely academic exercise. The problem identified in this paper is part of an ongoing collaboration between the National Renewable Energy Laboratory (NREL) and the Dallas-Fort Worth International Airport (DFW). Utilizing data from DFW, this research develops generic, open-source tools that any airport can use for within airport shuttle operations planning.

For context, DFW is the fourth busiest airport in the world, with a record setting annual passenger traffic of 75 million passengers in 2019 (Dallas Fort Worth Airport, 2020a). DFW is also the world’s largest airport to achieve a carbon neutral status (Dallas Fort Worth Airport, 2020b). Air travel took a major hit due to the COVID-19 pandemic, and as many airports, DFW saw a 90% decrease in passenger traffic in March, and April 2020 (Brian New, 2020). As air travel resumes with enhanced safety measures, DFW has seen a faster increase in flight and passenger volumes compared to other airports and became the busiest airport in the world in May 2020 (Arnold, 2020). As the world heals from the pandemic, becoming more resilient and prepared for coming back to a ‘new normal’, DFW could see a continued rise in air travelers in the months and years to come. Though COVID-19 as a special use case is definitely of interest, much of the work presented in this paper was carried out before the onset of the pandemic, and uses data from typical passenger traffic observed at DFW from December 2019 to February 2020. We note however, that the methodology presented in this paper is valid for newer data and could be used on data collected after the onset of the COVID-19 pandemic.

This paper proposes an optimization model combined with a discrete event simulator, to solve the ‘travel within the airport premises’ shuttle route optimization problem and explore tradeoffs between passenger wait times, and energy efficiency of transporting passengers between the rental car center and the terminals at DFW. These efforts focus on the rental car center shuttle operations due to their significant measured energy expenditure and miles driven relative to other routes, however the work described here could be adapted to other multi-stop shuttle routes. Data used in this study comes directly from data loggers installed on the airport shuttles, as well as spatial positioning on transit data provided by DFW (Transit, 2020). The optimization model is run with various combinations of constraints on shuttle headways, maximum passenger ride times, and passenger arrival rates. The discrete-event simulator takes the optimal solution for each scenario generated by the optimization model and simulates DFW shuttle operations using that solution for a period of four weeks to generate passenger wait time and energy consumption outputs. Results from the optimization model and the discrete-event simulator provides airport authorities with shuttle routes that simulations indicate will lead to a net saving in energy, while still providing a quality passenger travel experience.

2. Literature review

Vehicle routing problems can be traced to the 1960s with some of the early research efforts focusing on the traveling salesman problem (Lin, 1965) and vehicle scheduling problems (Knight and Hofer, 1968; Christofides and Eilon, 1969). The necessity for vehicle routing research, and the resulting complexity of the solution approaches increased greatly over the past half century. Among recent efforts relevant to this research, Chien (I-Jy Chien, 2005) combined analytical and numerical techniques to optimize operational characteristics (headway, seat capacity, and route choice) of a feeder bus service, while satisfying constraints pertaining to vehicle schedules, bus availability, and budget. Yan et al. (2012) took route flexibility out of the equation to solve a schedule design problem for a fixed bus route. They developed a robust optimization model for the bus route schedule design problem addressing the bus travel time uncertainty and the bus drivers’ schedule recovery efforts. Their algorithm was tested using bus data from Suzhou city in China, and find that slack time, and the bus driver’s schedule recovery behavior are key factors for maintaining optimal bus schedules. Using data from Dalian City, China, Yu et al. (2011) used a parallel genetic algorithm based on Tabu search to optimize bus route headways. Using a hypothetical road network, Xiong et al. (2013) compare the performance of genetic algorithm (GA), and depth-first search algorithm (DFS) to find the optimal route and headway for metro stations community shuttles, and report that GA is more efficient and reliable in providing a solution. While some of the variables used in aforementioned studies are similar to the ones proposed in this paper, an important point of divergence here is that focus of this paper is on within airport shuttle operations (which deal with bursts of rider demand owing to flight schedules), whereas a majority of the studies presented above focus on traditional bus service operations (which see more uniform passenger demand).

While studies on bus schedule, and route optimization in the context of a city’s transit services go well beyond the body of literature presented above, optimization of shuttle services in the context of special generators such as airports are relatively scarce. Even within the context of research on airport shuttles, most of the efforts focus on solving shuttle routing problems for travel to/from airports, where the shuttle service caters to multiple origins but a common destination. For example, Bao et al. (2018) used a hybrid genetic algorithm to explore the factors that influence of the travel time reliability of airport bus routes to Nanjing Lukou International Airport in China. They used synthetically generated constraints on time, station, and service, and concluded that travel time reliability in the peak hour (7:00–8:00 a.m.) is greatly affected by the road conditions. While the reliability maximization framework proposed by Bao et al. (2018) is attractive, considering synthetic constraints is an important limitation of the study. Feng et al. (2014) formulated the airport shuttle pick up and drop off problem as a mixed integer program, and solved the problem using an exact approach as well as adaptations of two existing heuristic approaches. They tested the performance of their algorithms using real-world data from Washington, D.C., Dulles International Airport. Linqing et al. (2019) proposed a vehicle routing and scheduling method where network travel times can vary, which is a significant extension to general vehicle routing problems that consider travel time as a travel invariable parameter. This aspect is reflected in the current work by using network travel times pertaining to specific times of day in the optimization model, and also simulating the optimized routes in a stochastic environment.

Linqing et al. (2019) proposed using passenger ratings of pick up and arrival time to define passenger satisfaction constraints. Pei-Ying et al. (2013) adopted mixed integer programming to solve the vehicle routing and scheduling problem, with an objective to minimize the emission footprint of airport shuttle services. Specifically, they analyzed the impact of customer position distribution, passenger demand, vehicle capacity, and degree of customer satisfaction on the fuel consumption outcomes of the optimized shuttle routes. As seen from the body of literature on airport shuttle operations (i.e., to travel to/from airports), there is evidence of using customer satisfaction as one of the measures in evaluating the performance of the proposed algorithms. In general, passengers' expectation of travel time and waiting time are given lesser weightage (compared to reliability, or personal security) in the context of traveling to/from airport, as such shuttle trips are scheduled well in advance of the actual travel date and time. However, for within airport shuttles (i.e., transfers from rental car centers and parking lots to the terminal), passenger's expectation of in-vehicle travel time and waiting time becomes higher as time is of the essence when passengers are trying to reach the terminal with a few minutes to an hour to catch the plane. Therefore, the routing and frequency settings of the shuttle bus in the study have higher impact on passengers' travel experiences. Instead of using a categorical satisfaction parameter (i.e., 1-extremely dissatisfied – 5-extremely satisfied), we propose using the continuous quantities in-vehicle travel time and waiting time as the evaluation measures in this research effort.

Exploration of the literature on vehicle routing and scheduling reveals some interesting themes. *First*, vehicle routing problems have been widely used in the context of city-wide transit schedule optimization. As for airport shuttle applications, most studies focused on airport access from city centers as opposed to shuttle services within an airport. Even among the studies that explore airport shuttle service route optimization (or general bus schedule optimization for that matter), only one study focused on an energy/emission related outcome (Pei-Ying et al., 2013). *Second*, all the studies identify a problem, solve it using an exact solution or a heuristic approach, and evaluate the performance of the optimization algorithm. We found no studies that did in depth explorations of the tradeoffs between energy consumption and passenger waiting time via stochastic simulation of a massive numbers of optimization runs with different model parameters (e.g., headway, maximum in-vehicle travel time, and passenger arrival rates). In the real world, the stochasticity in passenger arrivals and loading/unloading shift the dynamics of the system away from the exact mathematical optimization model assumptions. Therefore, stochastic simulation can be a useful tool for measuring vehicle fleet energy consumption and passenger waiting time. *Finally*, most of the studies reviewed used data from a real-world context for optimization and validation purposes, but none of the studies report on their solution being implemented in the field to demonstrate the benefit realized from optimized routes and schedules. This research effort contributes to the state-of-practice on airport shuttle route optimization by:

1. Focusing on within airport shuttle operations as opposed to shuttle operations connecting city centers and airports. The in-airport shuttle service can be adopted for future "remote curb" operations to reduce the curbside congestion at the airport terminals.
2. Providing a novel mathematical model for selecting routes, shuttles per route, and shuttle size per route, while evaluating for travel as well as energy related parameters of airport shuttle operations. This work adds to the sparse body of literature (Pei-Ying et al., 2013) that considers energy/emission reductions as an objective in shuttle route optimization.
3. Integrating the optimization model with a discrete-event simulator to explore the tradeoff between passenger travel experience and energy outcomes of the optimization solution.

3. Methodology

As shown in Fig. 1, the model system being proposed in this paper consists of two modules: i) a route optimization model which solves the dispatching problem to provide a set of shuttle routes and determine the number of shuttles and shuttle type to serve each route, such that the energy consumption of the fleet is minimized, and ii) a discrete-event simulator that tests the performance of the solution provided by the optimization model in a stochastic environment (with varying dwell times, travel times, and arrival rates).

With the route optimization model, our intent is to go one step beyond traditional vehicle routing solutions by incorporating energy consumption of the fleet in the decision space. The optimization model assumes that passengers arrive at a constant rate at bus stops, buses moving along routes are able to keep themselves evenly spaced along those routes, and that travel times between origin and destination pairs are constant. In reality passenger arrival at bus stops is stochastic and irregular, buses bunch up along routes, and travel times vary depending on many factors. Hence, we felt it was important to simulate the routes computed by the optimization model in stochastic setting over a reasonably long period of time to capture how those routes would really perform in a non-idealized setting. Therefore, with the discrete-event simulator, we aim to rigorously test the robustness of the solutions provided by the optimization model in order to reinforce confidence for the airport ground transport team that is envisioned to implement the optimized shuttle routing in the real world. Both these modules have been developed with effortless and intuitive implementation goals in mind.

3.1. Optimization model formulation

The optimization model used in this study is a mixed integer linear program where the user can specify the number of routes, capacity of the buses, allowable headways, and maximum in-vehicle travel time parameters. The solution generated by the model consists of a set of routes, each with a specified number of buses, and the capacity of the bus servicing it. As such, the optimization model presented here uses some concepts from the classical VRP with pick-up and delivery time windows, but has been augmented and modified in many ways to accomplish the goals associated with the 'travel within the airport premises' problem. Table 1 presents the sets, parameters, and variables used in the model. The optimization model, and the adjoining constraints are presented in equations (1-24).; ; ; .

$$\text{minimize} \sum_{k \in \mathcal{T}} \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{R}} h_{k,\,s,\,r}$$

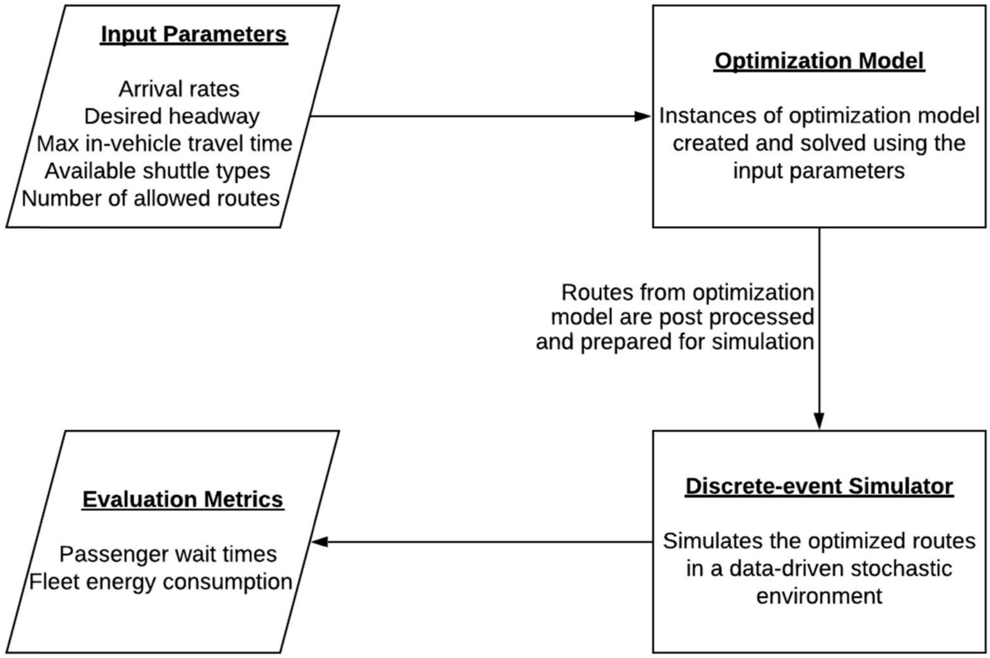

Fig. 1. Route optimization for energy efficient airport shuttle operations: process flow chart.

subject to

$$-M\left(2-y_{s,\,r}-w_{k,\,r}\right)+E_k n_s \le h_{k,\,s,\,r}\ \forall k\in\mathcal{T}\ s\in\mathcal{S}\ r\in\mathcal{R} \tag{1}$$

$$\sum_{s\in\mathcal{S}} y_{s,\,r}\le 1\ \forall r\in\mathcal{R} \tag{2}$$

$$\frac{T_{2n+1,\,r}}{\gamma}\le\sum_{s\in\mathcal{S}} n_s y_{s,\,r},\forall r\in\mathcal{R} \tag{3}$$

$$\sum_{k\in\mathcal{T}} w_{k,\,r}=1\ \forall r\in\mathcal{R} \tag{4}$$

$$\sum_{r\in\mathcal{R}}\sum_{j\in\mathcal{N}} x_{i,\,j,\,r}=1\ \forall i\in\mathcal{P} \tag{5}$$

$$\sum_{j\in\mathcal{N}} x_{i,\,j,\,r}-\sum_{j\in\mathcal{N}} x_{n+i,\,j,\,r}=0\ \forall i\in\mathcal{P}, r\in\mathcal{R} \tag{6}$$

$$\sum_{j\in\mathcal{N}} x_{0,\,j,\,r}=1\ \forall r\in\mathcal{R} \tag{7}$$

$$\sum_{j\in\mathcal{N}} x_{j,\,i,\,r}-\sum_{j\in\mathcal{N}} x_{i,\,j,\,r}=0\ \forall i\in\mathcal{P}\cup\mathcal{D}, r\in\mathcal{R} \tag{8}$$

$$\sum_{i\in\mathcal{N}} x_{i,\,2n+1,\,r}=1\ \forall r\in\mathcal{R} \tag{9}$$

$$\sum_{j\in\mathcal{N}} x_{j,\,i',\,r}=x_{i',\,j',\,r}\ \forall l\in\mathcal{L}, r\in\mathcal{R},\ i'\in\mathcal{D}_l,\ j'\in\mathcal{P}_l \tag{10}$$

$$T_{j,\,r}\ge T_{i,\,r}+d_i+t_{i,\,j}-M\left(1-x_{i,\,j,\,r}\right)\ \forall i\in\mathcal{N},\ j\in\mathcal{N},\ r\in\mathcal{R} \tag{11}$$

$$T_{n+i,\,r}-\left(T_{i,\,r}+d_i+t_{i,\,n+i}\right)\ge 0\ \forall i\in\mathcal{P}, r\in\mathcal{R} \tag{12}$$

$$T_{n+i,\,r}-\left(T_{i,\,r}+d_i\right)\le\delta-M\left(1-\sum_{j\in\mathcal{N}} x_{i,\,j,\,r}\right)\ \forall i\in\mathcal{P}, r\in\mathcal{R} \tag{13}$$

$$-M\left(1-y_{s,\,r}\right)+\frac{\overline{q}_i T_{2n+1,\,r}}{n_s}\le q_{i,\,r}\ \forall i\in\mathcal{N},\ r\in\mathcal{R},\ s\in\mathcal{S} \tag{14}$$

$$q_{i,\,r}+q_{n+i,\,r}=0\ \forall i\in\mathcal{P}, r\in\mathcal{R} \tag{15}$$

$$Q_{j,\,r}\ge Q_{i,\,r}+q_{j,\,r}-M\left(1-x_{i,\,j,\,r}\right)\ \forall i\in\mathcal{N},\ j\in\mathcal{N},\ r\in\mathcal{R} \tag{16}$$

$$\max\{0,\ q_{i,\,r}\}\le Q_{i,\,r}\le\min\left\{\sum_{k\in\mathcal{T}}\overline{Q}_k w_{k,\,r},\sum_{k\in\mathcal{T}}\overline{Q}_k w_{k,\,r}+q_{i,\,r}\right\}\ \forall i\in\mathcal{N},\ r\in\mathcal{R} \tag{17}$$

$$T_{i,\,r},\ Q_{i,\,r},\ h_{k,\,s,\,r}\ge 0 \tag{18}$$

$$x_{i,\,j,\,r}, y_{s,\,r},\ w_{k,\,r}\in\{0,\ 1\} \tag{19}$$

$$x_{i,\,j,\,r}\le z_{i,\,j,\,r}\ \forall i\in\mathcal{N},\ j\in\mathcal{N},\ r\in\mathcal{R} \tag{20}$$

$$z_{i,\,j,\,r}+z_{j,\,i,\,r}\le 1,\ \forall i\in\mathcal{N},\ j\in\mathcal{N},\ r\in\mathcal{R} \tag{21}$$

$$\sum_{k\in\mathcal{N}} x_{i,\,k,\,r}+\sum_{k\in\mathcal{N}} x_{j,\,k,\,r}-1\le z_{i,\,j,\,r}+z_{j,\,i,\,r}\ \forall i\in\mathcal{N},\ j\in\mathcal{N},\ r\in\mathcal{R} \tag{22}$$

$$z_{i,\,j,\,r}+z_{j,\,k,\,r}+z_{k,\,i,\,r}\le 2\ \forall i,\ j,\ k\in\mathcal{N},\ r\in\mathcal{R} \tag{23}$$

$$z_{i,\,j,\,r}\in\{0,\ 1\} \tag{24}$$

The objective function presented in equation (1) seeks to minimize the average hourly energy consumption of the routes and shuttle fleet chosen by the optimization model. Constraints (2) and (3) ensure that each route $r\in R$ available to the model is either not used and has no buses assigned to it, or it selects a value of 1 for exactly one of the $y_{s,r}$ variables which commits a number of buses to route r that satisfies the maximum allowable headway requirement set by the user. Constraint (4) ensures each route r is assigned exactly one bus type to service it (even if it is an unused route). The objective function (1) works in conjunction with constraints (2-4) to ensure the following condition: *if n_s buses of type k are assigned to route r, then the auxiliary variable $h_{k,s,r}$ is bounded by the resulting average energy consumption* per *hour of the buses assigned to route r, otherwise* (1) *is allowed to be zero*. In other words, since the objective is to minimize the sum of the $h_{k,s,r}$, equations (1-4) ensure that each $h_{k,s,r}=E_k n_s$ or $h_{k,s,r}=0$. Constraints (5) and (6) make sure that each request type is served exactly once by one route and that it is the

Table 1
Sets, parameters, and variables.

Sets	
$\mathscr{P}$	Set of nodes representing pickup requests, $\mathscr{P} = \{1, \cdots, n\}$
$\mathscr{D}$	Set of nodes representing drop-offs corresponding to a specific pickup, $\mathscr{D} = \{n+1, \cdots, 2n\}$
$\{o, d\}$	Origin and destination of all buses, $\{o, d\} = \{0, 2n+1\}$
$\mathscr{N}$	$\mathscr{P} \cup \mathscr{D} \cup \{o, d\}$
$\mathscr{R}$	Set of route ids
$\mathscr{S}$	Set indexing the possible allowed number of buses on a route
$\mathscr{T}$	Set of bus types
$\mathscr{L}$	Set of terminals
$\mathscr{D}_l$	The drop-off node $i \in \mathscr{D}$ at terminal $l \in \mathscr{L}$
$\mathscr{P}_l$	The pickup node $i \in \mathscr{P}$ at terminal $l \in \mathscr{L}$
Parameters	
$\overline{q}_i$	Arrivals per hour at i for each $i \in \mathscr{P}$. We note that $\overline{q}_{n+1} = -\overline{q}_i$ for each $i \in \mathscr{P}$, and $\overline{q}_{2n+1} = \overline{q}_0 = 0$.
d_i	Dwell time at node i for each $i \in \mathscr{P} \cup \mathscr{D}$
$t_{i,j}$	Travel time from node i to j where $i, j \in \mathscr{N}$
$\overline{Q}_k$	Capacity of bus type k for each $k \in \mathscr{T}$
E_k	Hourly energy consumption of bus type k for each $k \in \mathscr{T}$
n_s	Number of buses in case s for each $s \in \mathscr{S}$
M	Large positive constant
δ	Maximum allowable ride time for a passenger
γ	Maximum allowable headway for buses
Variables	
$x_{i,j,r}$	Binary decision variable determining if a bus on route r traverses from node i to j where $i \neq j$, $i, j \in \mathscr{N}$, and $r \in \mathscr{R}$
$z_{i,j,r}$	Binary decision variable determining if a bus on route r visits node i before j where $i \neq j$, $i, j \in \mathscr{N}$, and $r \in \mathscr{R}$
$y_{s,r}$	Binary decision variable for choosing the number of buses on route r where $s \in \mathscr{S}$, and $r \in \mathscr{R}$
$w_{k,r}$	Binary decision variable for choosing the type of bus on route r where $k \in \mathscr{T}$, and $r \in \mathscr{R}$
$Q_{i,r}$	Continuous non-negative decision variable indicating the capacity of a bus on route r after visiting node i where $i \in \mathscr{N}$, and $r \in \mathscr{R}$
$T_{i,r}$	Continuous non-negative decision variable indicating the time in hours after the route starts at which service begins at node i on route r where $i \in \mathscr{N}$, and $r \in \mathscr{R}$
$q_{i,r}$	Demand at node i on route r where $i \in \mathscr{N}$, and $r \in \mathscr{R}$
$h_{k,s,r}$	Continuous non-negative auxiliary decision variable used to construct the objective function where $k \in \mathscr{T}$, $s \in \mathscr{S}$, and $r \in \mathscr{R}$

same route that visits the pickup and drop-off nodes associated with the request type. Constraints (7-9) make sure that each route starts at the origin node and terminates at the destination node.

Constraint (10) ensures that if a route does a drop-off at a terminal's drop-off node, it must immediately visit the pickup node for that terminal. Consistency and precedence of time variables on routes are enforced through constraints (11) and (12) . The maximum in-vehicle travel time for a passenger on each route is enforced via constraint (13). Constraint (14) uses the time to complete a route, the number of buses servicing that route, and the hourly arrival rate at a node to determine a lower bound to enforce on the variable $q_{i,r}$. Constraint (15) ensures that if $q_{i,r}$ people are picked up at node i,the same number are dropped off at node $i+n$. Constraints (16) and (17) ensure that the capacity of passengers on the bus is tracked from node to node, and that the occupancy of the bus never exceeds its capacity. Finally (18), and (19) provide basic constraints on the domains of certain variables. Constraints in (11) function as sub-tour elimination constraints for the model in addition to tracking the model's time variables. Through computational experiments, we found that the solver found better solutions when additional sub-tour elimination constraints (20-24) were included in the model.

3.2. Discrete-event simulator

To test the robustness of the solutions generated by the optimization model, we developed an open-source event-driven simulator named Airport Shuttle Planning and Improved Routing Event-Driven Simulation (ASPIRES). ASPIRES was developed as a Python module to simulate and evaluate the current as well as optimized airport shuttle operations. ASPIRES takes the output of the optimization model and simulates airport shuttle operations using empirical probability distributions of travel times, dwell times, and passenger arrivals. The ASPIRES module addresses calibration issues faced by most traffic simulation packages by carrying out event-driven simulations based on empirical distributions of real-world data. Calibration is a difficult task in transportation simulation. The calibration result may not always reach a satisfying performance. The intent here is to let the data drive the simulation rather than calibrating or ground truthing the simulation. The ASPIRES module has been developed with rapid decision-making goals in mind and is optimized to simulate a days' worth of airport shuttle operations in about a second. In contrast, microsimulation of the airport shuttle operations without considering the passengers statistics of the same magnitude using open-source tool.

Simulation of Urban MObility SUMO (Krajzewicz, 2010) took about 30 min.

ASPIRES can derive empirical distributions of any parameter using real-world data. These distributions inform the event-driven simulation which executes all the events in airport shuttle operations at a rapid pace, and outputs performance statistics at the passenger- or the shuttle-level. Passenger-level statistics include wait time, queue length, unserved demand (number of passengers left after a shuttle bus pickup), and total travel time (wait time + in-vehicle travel time). Shuttle-level statistics constitute occupancy, cumulative distance traveled, and energy consumed, location of each bus, and a record of on-demand routes used to pick up any passengers that could not be served by the regular shuttles. Passenger arrival is simulated using a Poisson process (with real world non-stationary arrival rates) in ASPIRES. Poisson process assumes that arrivals could happen at any point of time with a pre-defined probability during the simulation period. Poisson distribution is adopted here as it allows for randomness in passenger arrivals and is one of the commonly used functions to simulate arrival processes with given arrival rates (Ross, 2014). Poisson process considers arrivals to happen at any time with a given probability, and does not assume that passengers are arriving with a pre-defined transit timetable in mind. While air travelers do come to the airport with flight schedules in mind, it is very unlikely that they have information regarding airport shuttle timetable. Hence, the randomness assumption of the Poisson process fits our use case well. It should be noted that other distributions such as Weibull, Gamma, and Beta are also used to capture passenger arrivals, and could be explored in future extension of this research effort. Since passenger arrival rate varies across the day, aggregated arrivals by hour of the day were fed into the simulation.

The movement of the buses is driven by the empirical distribution derived from the bus logger data at DFW. Instead of evolving from one time step to the next time step, ASPIRES jumps from one critical state of the system to the next critical state (a new event). When a bus arrives at a stop, ASPIRES will schedule the next event (i.e., leaving the bus stop) after a certain dwell time which is drawn from the empirical dwell time distribution at that time window. ASPIRES will mark the state of the bus to "leaving the bus stop" at the next event time (i.e., current simulation time + dwell time). The number of remaining passengers at the bus stop will be updated based on the number of passengers boarding the bus. When the bus leaves from one bus stop to the next bus stop, the simulation progressed to the next event (i.e., arriving at a bus stop) based on a travel time number drawn from the empirical distribution of the travel times between the two stops from that time window of the day in the week. ASPIRES also captures the energy consumption associated with the travel time from the empirical distribution. Therefore, ASPIRES is capable of addressing the impact of congestion on travel time and energy consumption during different times of day and days of the week.

Shuttle bus routes are provided from the solution generated by the optimization model. As the fleet size and routes are expected to vary across the day (owing to variation in flight schedules and passenger demand), ASPIRES also constitutes a dispatcher that manages the number of buses on different routes. The dispatcher routine changes buses over to new routes as prescribed by the optimization model. Additionally, the dispatcher can be used to assign on-demand shuttles to cater to unserved demand (i.e., passengers left behind due to shuttle occupancy constraints). The aspects covered by ASPIRES make it an attractive tool for airports to test the solutions generated by optimization routines before they are implemented in the field.

4. Data description, simulation setup, and validation

DFW provided access to their vehicles and data systems in support of this research effort. Firstly, DFW allowed NREL researchers to collect Controller Area Network (CAN) bus data from the airport rental car shuttles using vehicle data loggers resulting in approximately 100,000 miles of 1 Hz data from 14 buses over a period of one month of shuttle operations (Kotz et al., 2020). Secondly, DFW provided Spatial Positioning on Transit (SPOT) data, which uses commercial hardware to capture information pertaining to shuttle operations. Some noteworthy functionalities of SPOT data are: i) providing automatic vehicle location updates; ii) calculating bus arrival predictions and assessing on-time performance of the shuttles; iii) managing vehicle headways and viewing shuttle schedules. SPOT data combined with CAN data provided the information required for the optimization model and the discrete-event simulator.

4.1. Analysis of the CAN-based bus data

The CAN data provides detailed information on the geolocation of the bus, which can be translated into various travel parameters such as trip distance, trip time, and dwell time at each shuttle stop. In addition, CAN data provides information regarding vehicle parameters such as speed, acceleration, engine power, and fuel consumption. Though the CAN data provides a wealth of information regarding shuttle operations, it is not fully immune to outliers. Data collected was processed using trip time as the criteria for identifying outliers. Trip times less than the 5th percentile, and greater than the 95th percentile were treated as outliers and excluded for use in the optimization and simulation modules. Since the CAN bus data did not have enough data points for each hour in the day, the 24-h time period was divided into three aggregations of 12am–8am, 8am-7pm, and 7pm-12am, which were motivated by passenger demand trends in the SPOT data. Missing data from the bus loggers was supplemented using a microscopic simulation model of the DFW road network (developed using SUMO).

4.2. Analysis of the SPOT data

SPOT data is obtained from a centralized service as a continuous data stream. Among other functionalities, SPOT tracks the time, location, boarding and alighting, stop, and route served by each bus. SPOT data from December 1, 2019 to February 17, 2020 was used for the purposes of this analysis. Though additional data is available, a conscious decision was made to test the algorithms for pre-COVID air travel and shuttle operations (which are representative of long-term operations at DFW). Like data collected from the CAN bus, SPOT data is prone to outliers owing primarily to errors introduced during the post-processing of sensor data. SPOT data provided by DFW was subjected to extensive cleaning and filtering before being used in the analysis. As one example of the cleaning and processing done, the total number of passengers on-board (as reported in the SPOT data) were tracked and the number of passengers boarding a shuttle were scaled down or curbed, when the number of passengers on-board reached the capacity of the bus. From the cleaned SPOT data, fleet size and frequency for each route as well as the passenger demand information at each stop was extracted. Fig. 2 shows the service frequency for all DFW terminals across different times of day. From the figure, it can be observed that the current shuttle services vary over the course of the day, owing to changes in flight traffic and resulting passenger demand.

4.3. Computational experiments

To explore the tradeoff between energy efficiency and passenger travel experience (i.e., passenger wait times), we carried out shuttle route optimizations for all permutations of the following parameters, leading to a total of 2268 model runs.

- Days of the Week: M, T, W, Th, F, Sa, Su
- Time Windows: 12am–8am, 8am-7pm, 7pm-12am
- Arrivals Standard Deviations: 0, 1, 2
- Maximum In-vehicle Travel Times: 15, 20, 25
- Headways: 5, 7, 10, 15, 20, 25
- Available Bus Type Cases: 43 seats, 14 and 43 seats

Each model run (or scenario) could specify as many as five routes if desired. The arrivals standard deviation parameter indicates the number of standard deviations added to each mean arrival rate based on empirical distributions of arrivals. The optimization model detailed in equations (1-24); was implemented using the open-source mathematical modeling software Pyomo (Hart et al., 2017). Instances of the optimization model where run on the NREL High Performance Computer (HPC), Eagle, and the optimization was performed with the commercial solver, Gurobi (Optimization and Gurobi. "In, 2014). Through the course of the implementation, it was noticed that Gurobi was not able to prove optimality even after an hour of run time when the number of possible routes was increased to two or more. We believe this behavior is due to the expansive solution set and multiple distinct optimal solutions in the branch and bound tree causing difficulty in providing a certificate of optimality quickly. Thus, Gurobi was used as a heuristic by allowing it to run for an hour on each problem instance, and the best solution Gurobi found within that hour was taken as the optimal solution. Future research efforts could focus on custom algorithms designed to solve this type of combinatorial optimization model when larger number of possible routes must be considered.

Using the optimized routes, we were able to construct a weeklong schedule of routes for each combination of maximum in-vehicle travel time, headway, arrivals standard deviation, and available bus types. This led to 108 distinct weeklong schedules of optimized bus routes, which were then simulated in a stochastic environment using ASPIRES for four consecutive weeks. We note that each four-week simulation began with a one-day warm-up period to allow the shuttle system to evolve to a typical operating state.

4.4. Validation of baseline results

In addition to simulating the 108 weeklong schedules of routes, a baseline case was run in ASPIRES to reproduce current shuttle operations at the airport. Parameters for the baseline case were agreed upon in close consultation with the DFW shuttle bus operation team. This simulation was also done for a four consecutive week period. The baseline simulation provided the basis for comparing optimized routing results and helped illustrate the tradeoffs between travel and energy efficiency for shuttle operations at the DFW airport. The baseline case in ASPIRES resulted in an average passenger wait time of 5.7 min to be picked up by a bus. We note that passenger wait time refers to the time a passenger waits at the curb to be picked up by a bus, unless specified otherwise. During typical operations DFW bus operations seeks to have a 7-min headway between buses, which suggests the average wait time

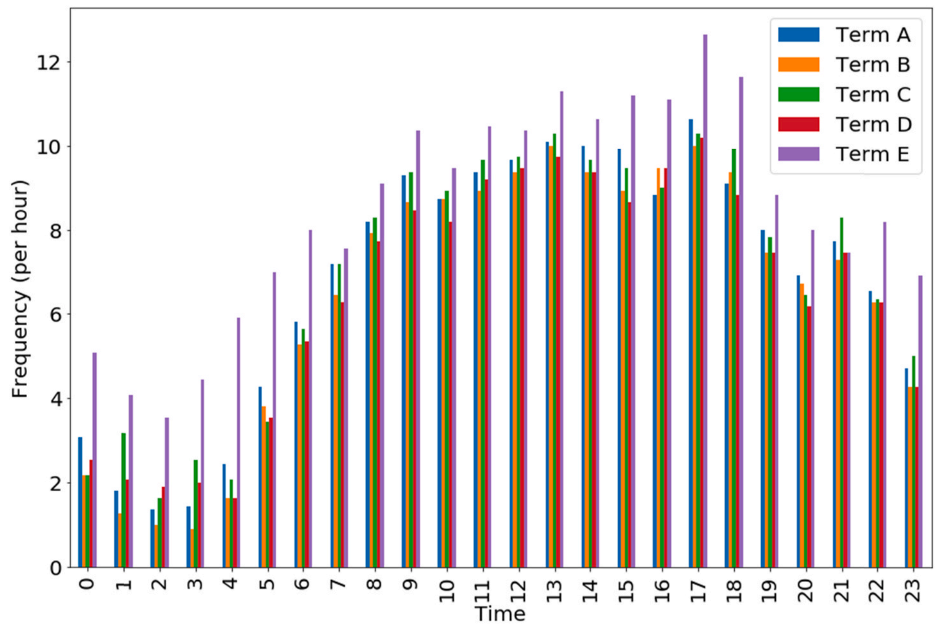

Fig. 2. Service frequency for each terminal over a day.

from our baseline simulation is reasonable. The total energy consumption of the rental car shuttle fleet for the baseline scenario came out to be 11,895 GGE per week, with about 20 shuttles in operation at any given time.

From DFW fuel and mileage logs, it was observed that the rental car shuttle fleet consumed on average a total of 13,342 GGEs each week from August 2018 to August 2019. This was taken as the ground truth (as this information was collected first-hand by DFW and provided to NREL researchers) to validate the output of the simulation module. For the baseline case described above, it was observed that the energy estimates based on a week-long simulation (using the ASPIRES module) of DFW's rental car shuttle fleet resulted in an energy consumption of 11,895 GGEs, which is about ~11% less than the ground truth. While the simulated value deviates slightly from the ground truth, an in-depth exploration revealed an interesting finding. The baseline simulation is based on collected CAN-based bus data, as well as boarding and alighting information from the SPOT data provided by DFW. From the baseline simulation, it was observed that there are on average ~20 shuttles in operation at any given time. However, a quick analysis using the DFW fuel and mileage logs from August 2018 to August 2019 and the collected CAN-based bus data revealed that during that period there where on average ~23 shuttles in operation at any given time of the day. Since the SPOT data-based simulation resulted in a lower number of shuttles being in operation than the DFW fuel and mileage logs indicate, it is intuitive that the energy consumption estimate from the simulation is slightly lower than that of the ground truth. If the energy consumption estimates are scaled up from 20 to 23 buses, the relative error between observed and simulated values reduces to ~1.18% which can simply be attributed to noise in the simulation. This exercise reinforces confidence in the efficacy of the discrete event simulator, and its ability to accurately simulate the bus system of interest in this paper. Also, this analysis reveals that there might be some discrepancies associated with the SPOT data that warrant further attention, or that bus operations at DFW from December 1, 2019 to February 17, 2020 had changed compared to operations from August 2018 to August 2019.

5. Model results

5.1. Optimization model

Fig. 3 presents the reference solutions for optimized routes on a typical Monday (using 43- seater shuttle buses). The maximum in-vehicle travel time (*ivtt*), headway (*hw*), and arrivals standard deviation (*sd*) used for each of the scenarios is shown towards the left-hand side of each figure. In this figure, the rental car center is denoted with the letter "R", and each of the terminals at DFW airport are denoted with the corresponding letters. The routes and the number of vehicles assigned to them resulting from each of the optimization solutions is noted at the bottom of the corresponding figure (for example, R-A-R; R-C-R etc.). Panels A and B of Fig. 3 show that when the maximum in-vehicle travel time is held at 15 min, shuttles take the shortest route possible between terminals and the rental car center (which are the standard shuttle routes currently used at DFW). Panels A and B also illustrate that changes to the *hw* and *sd* parameters can have a significant impact on the number of shuttles used. The combination of parameters presented in Panel A result in a maximum fleet size of 22, and a minimum fleet size of 12 for the most and least busy times of the day, respectively. However, in Panel B when the headway is reduced to 5 min (meaning a greater number of shuttles per hour), 30–31 buses are required throughout the day. This increase occurs in spite of the fact that the *sd* parameter (which governs the passenger arrival rates) has been reduced from 2 to 1 which in theory is expected to reduce the required number of shuttles.

From Panels C and D, we see that relaxing the maximum in-vehicle travel time constraint further to 20 min provides additional slack in the system resulting in optimal routes visiting multiple terminals before returning to the rental car center. Finally, when comparing Panels C and D, we see that a relaxation in headway from ten to 15 min and a reduction in the *sd* parameter from 2 to 1 result in a smaller fleet size across the board.

5.2. Discrete-event simulator

As was mentioned above, the base case simulation carried out on four weeks of data in ASPIRES resulted in an average passenger wait time of 5.7 min to be picked up by a bus, on average ~20 shuttles in operation at

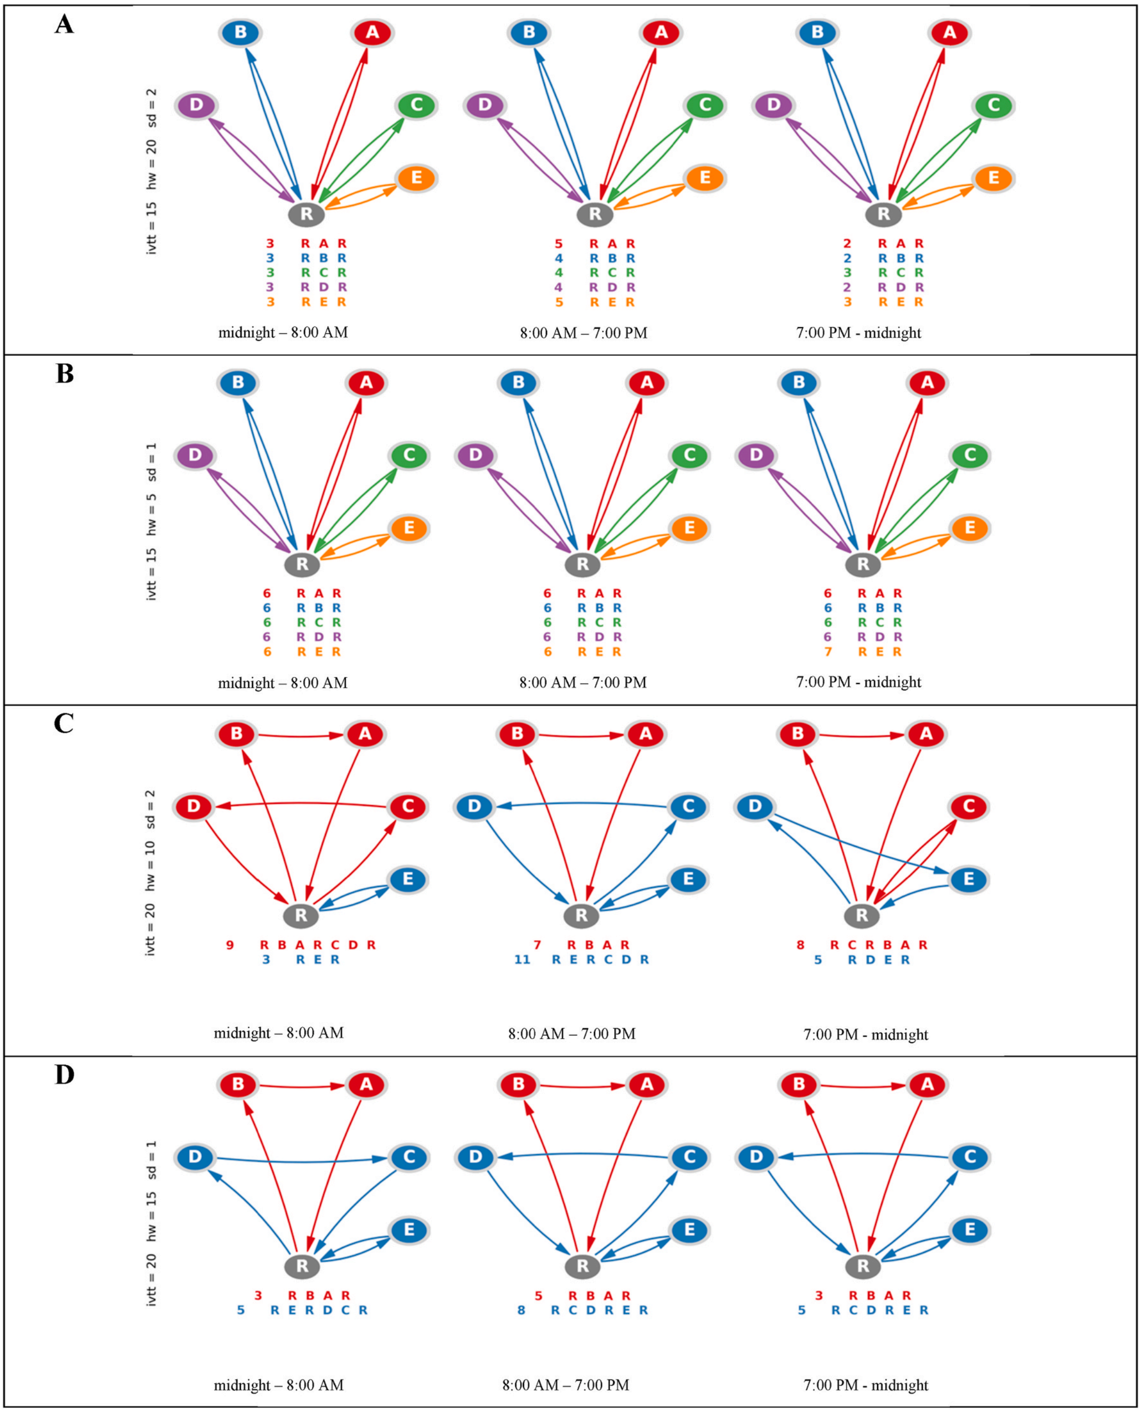

Fig. 3. Reference solution generated by the optimization model for a typical weekday using only 43-passenger buses.

any given time, and a total energy consumption of 11,895 GGE per week. Fig. 4 shows the average weekly energy consumption (GGE), and passenger wait times for the optimal routes when varying the headway, maximum in-vehicle travel time, and arrivals standard deviation parameters. Waiting times and energy consumption were similar between scenarios with only 43-seater buses compared to scenarios with a combination of 43- and 14-seater buses. Hence, we only present results for the 43-seater bus scenarios, which reduces the total number of scenarios presented in Fig. 4 to fifty-four. Panel A in the figure presents heatmaps for the average weekly energy consumption from each four-week simulation carried out in ASPIRES, while Panel B presents the mean passenger waiting time resulting from each simulation. Higher energy consumption and wait time values are presented with darker shading in both panels. A scenario with a lighter shade for energy consumption as well as mean waiting time is the ideal scenario where the solution produced by the optimization model achieves good performance with respect to passenger wait time as well as energy efficiency.

Generation of results shown in Fig. 4 involved a post processing step to remedy some unintuitive results produced by the optimization model. For example, it was observed that the optimization model would prescribe the route R-A-R-C-R-D-R-E-R for a given time window instead of the following four routes: R-A-R; R-C-R; R-D-R; R-E-R. As such, the solution identified by the optimization model meets all of the constraints defined in equations (1-24); ; ; , but it would not make for an attractive

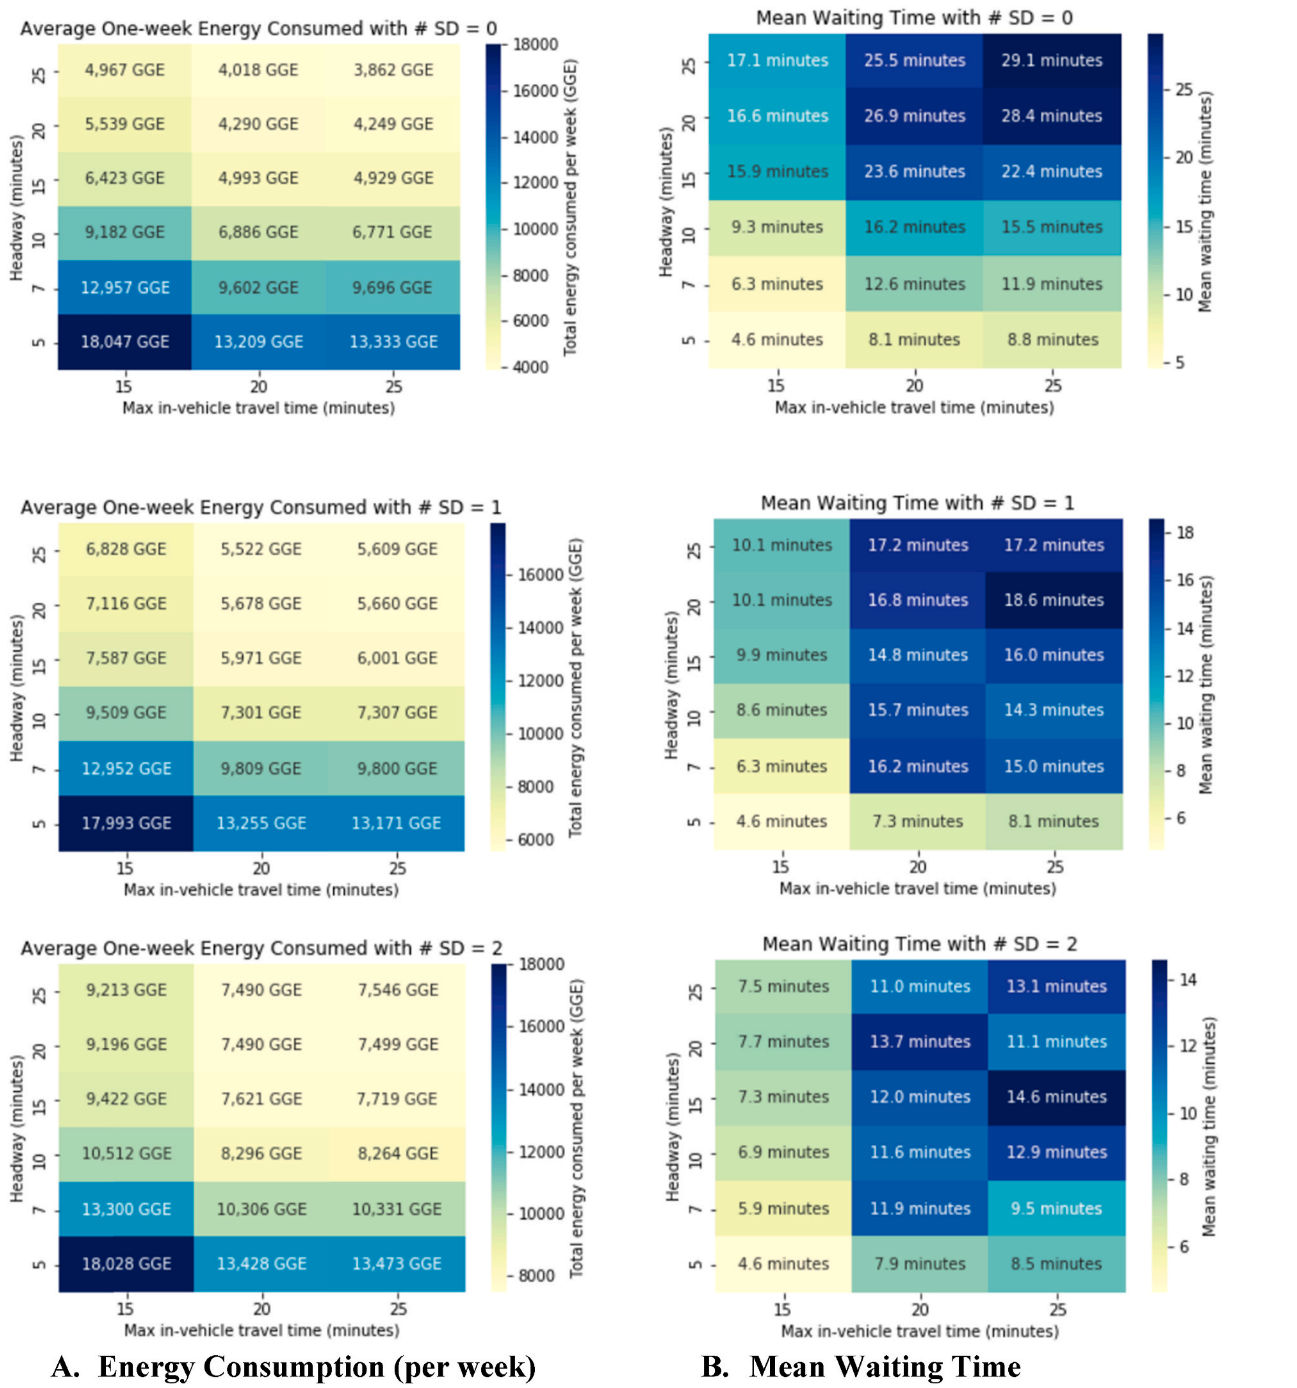

Fig. 4. Shuttle energy consumption and passenger wait time outputs from the ASPIRES module.

solution in the real world. To remedy this, we introduced a post processing step that would break up any routes such as R-A-R-C-R-D-R-E-R (that did not have any terminals bundled together) and split the number of buses prescribed to the single longer route across the new split terminal routes. If four buses are assigned to the route R-A-R-C-R-D-R-E-R, the post processing step would split this into four routes (R-A-R; R-C-R; R-D-R; R-E-R) and assign a single bus to each route. This post processing step is not applied when terminals are bundled together on a single route (such as R-**C-D**-R-**B-A**-R-E-R).

From Fig. 4, it can be observed that highest energy is consumed in cases where the constraints on maximum in-vehicle travel time, and headway are the most stringent (see lower left corner of the heatmaps in Panel A). These cases resulted in the highest number of buses across all the scenarios (for a visual of the optimal routes pertaining to these scenarios, see Panel B of Fig. 3). It is intuitive that the greater number of buses resulting from the optimization solution result in the lowest wait times across all scenarios (see lower left corner of the heatmaps in Panel B of Fig. 4). These scenarios lead to a ~20% reduction in passenger wait time at the expense of 50% increase in energy consumption. Conversely, routes where the headway and maximum in-vehicle travel time constraints are relaxed heavily resulted in lowest energy consumption and higher waiting times (see results towards the top right corner of Panels A and B), as lesser number of buses were prescribed by the optimization model for these scenarios (for example, Panel D of Fig. 3).

Another noteworthy observation from Fig. 4 is that the scenarios where the standard deviation parameter was higher (i.e., higher passenger arrival rates used in the optimization model) resulted in greater energy consumption, yet lower wait times. It can be conjectured that a greater passenger arrival rate causes the optimization model to assign a greater number of buses to a given route and possibly change the routes prescribed (as it might become difficult to bundle the demand and hence the different terminal stops together). With a greater number of buses, it is straightforward to expect a higher energy consumption coupled with lower passenger wait times. Based on the energy consumption and passenger wait time results, it is safe to say that energy efficiency and reduction in passenger wait time are conflicting goals that need to be balanced delicately by the airport authorities. The optimization model, and discrete-event simulation module presented here are developed to serve as decision support tools to the airport authorities to meet such competing objectives.

From Fig. 4, it can be observed that there are some scenarios which provide useful tradeoffs between energy consumption and passenger wait times. For the scenario with a 20-min headway constraint, 15-min maximum in-vehicle travel time constraint, and $sd = 2$, the optimized routes resulted in a mean passenger waiting time of 7.7 min and total energy consumption of 9196 GGE. This translates to a 22.7% reduction in energy consumption from the baseline (where the energy consumption was 11,895 GGE), with a modest 2-min increase in mean passenger wait times. As another corner case, the scenario with a 10-min headway constraint, 20-min max ride time constraint, and $sd = 2s$, sees a 30% energy reduction from the baseline (11,895 GGE → 8296 GGE), but this comes at a cost of 100% increase in mean passenger wait time. Based on these results, DFW might choose to go with the former case where sizeable energy reductions are made possible with a modest increase in passenger wait times.

In Fig. 5 plots of the total energy consumption against the mean passenger waiting time are provided to help illustrate the tradeoff between these quantities. The red polylines indicate the Pareto frontier of each standard deviation setting explored in the optimization process. Although each point came from a setting that has already been optimized, we can still see the trade-off between energy consumption and passenger service level. The results in Fig. 5 correspond to the data in Fig. 4. For reference, we circled the "knee" in the results. After the "knee," the saving in energy becomes less significant as we compromise the passenger service level. Pareto frontier plots like the ones in Fig. 5 can be used to support policy decision making by allowing the decision maker to clearly see the tradeoffs present between different decisions. It is intuitive that the energy consumption and the passenger service level compete against each other. We also observe that as the standard deviation in the optimization settings changes, the shape of the Pareto frontier changes as well.

A caveat worth noting here is that the average passenger wait time sometimes simplifies underlying distribution of wait times, leading to corner cases that might be unacceptable in real world implementations. One such case is shown in Panel A of Fig. 6, which shows the passenger wait time distribution for the optimized routes with a 15-min headway constraint, 20-min maximum in-vehicle travel time constraint, and $sd = 1$.

For this scenario, it can be observed from Fig. 4 that the optimized routes simulated in ASPIRES result in a 50% reduction in energy consumption (11,895 GGE → 5971 GGE) with a 9.1-min increase in passenger wait time from the baseline (5.7 → 14.8). While the tradeoff might seem reasonable from an energy reduction perspective, the average wait time increase only tells half the story. From the histogram presented in Panel A of Fig. 6, the distribution of the passenger wait times has a very long tail, which might be unacceptable for real world implementation. A converse case to this is with optimized routes where headway is set to 5 min, max ride time is set to 15 min, and standard deviation of passenger arrival rates is set to 1 (Panel B of Fig. 6). The passenger wait time distribution for the test case is better than the baseline scenario, but this comes at a 51% increase in energy consumption w.r.t baseline (11,895 GGE →17,993 GGE).

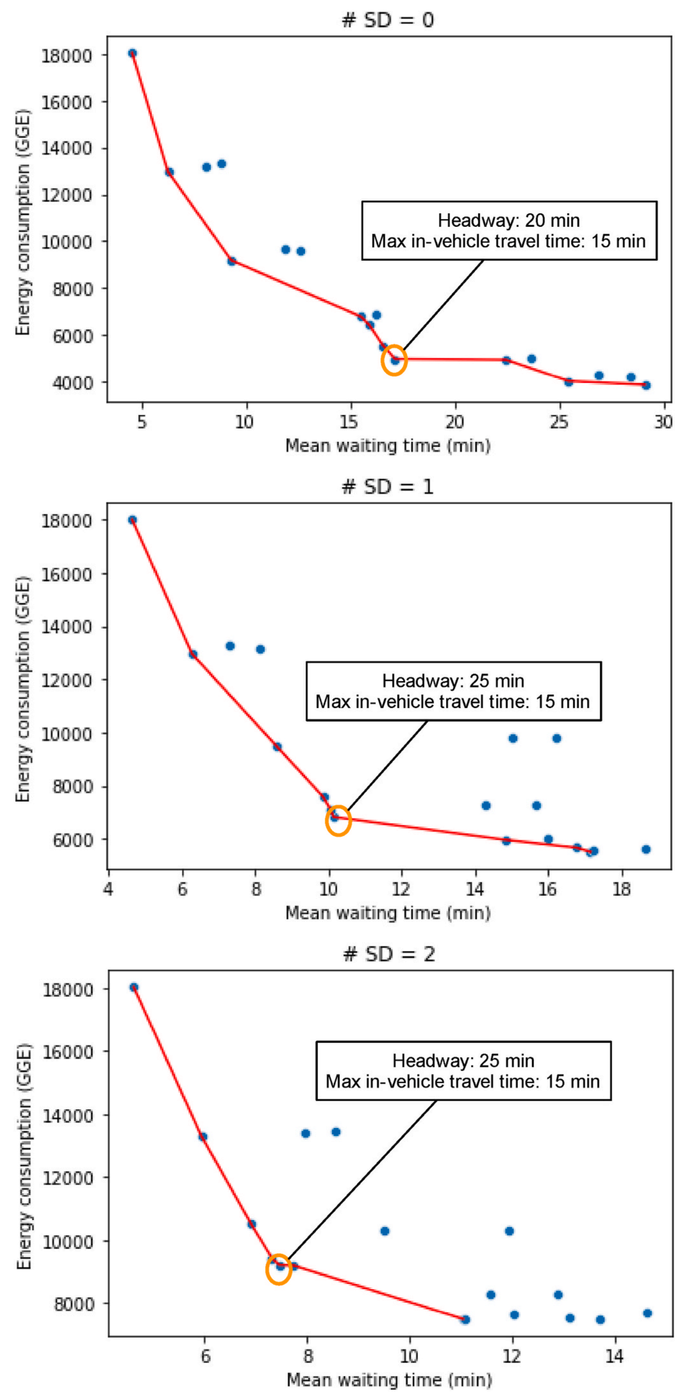

Fig. 5. Shuttle energy consumption versus passenger wait time.

It is likely that the long tail in Panel A of Fig. 6 occurs due to the stochastic nature of the environment the routes are simulated in. In particular, passengers do not arrive at a uniform rate. This leads to instances where passenger arrive in clusters and buses may need to leave people behind. Additionally, due to imperfect spacing, long gaps can occur between bus arrivals. These factors likely explain why, when simulated, routes can have mean passenger waiting times larger than the headway specified in the optimization model that produced the routes. Hence, these distributions provide strong evidence of the importance the ASPIRES module plays in the analysis of the routes produced by the optimization model.

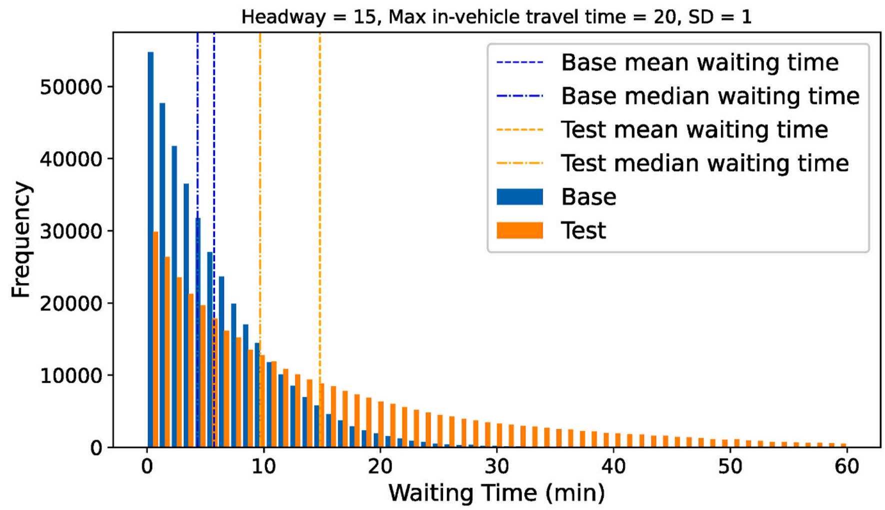

A

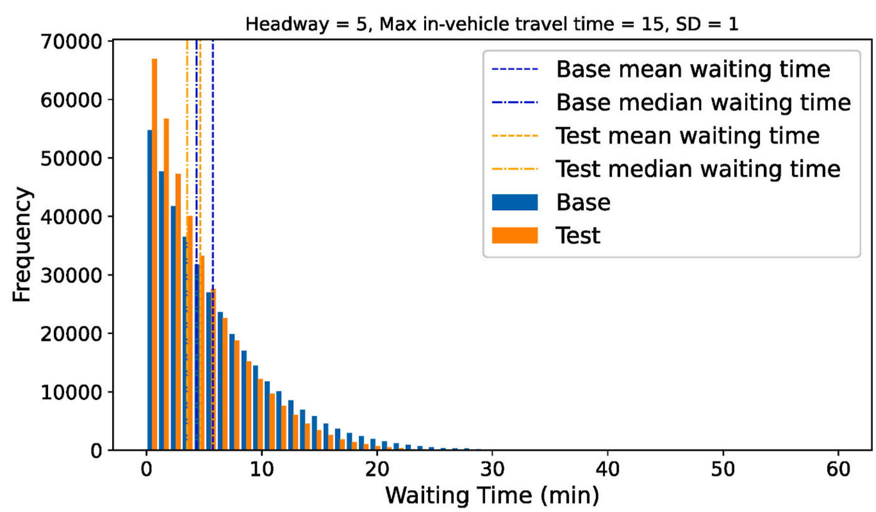

B

Fig. 6. Histograms of passenger waiting times using routes obtained with model parameters specified.

6. Conclusion and future work

Air travel has seen a steady growth for the past decade until the onset of the COVID-19 pandemic and is likely to continue to rise as restrictions and safety concerns ease. A majority of air travelers in the United States access airports by: i) getting dropped off by a family member or a taxi/ride hail/hotel shuttle service; ii) using transit where available; iii) driving their own vehicle and parking at the airport; iv) driving a rental car. While passengers get dropped off at the curb in the first two cases, the latter two modes involve a transfer from a parking lot or a rental car center to the terminal using an airport shuttle. Satisfying their customers while also balancing against operational costs and growing emissions is a significant challenge for modern airports. While airports exercise some discretion in varying fleet sizes to service varying passenger demand across the day, the routes used by airport shuttles are often pre-determined and do not change significantly with demand, leading to travel time as well as energy inefficiencies. While airports have a great deal of flexibility in altering the shuttle schedules to improve passengers experience, they often lack the tools that can help them in making such decisions while considering the frontier of potential travel-energy tradeoffs.

To approach this challenge, this paper presents a novel optimization model that builds on the traditional vehicle routing problem with pick-up and delivery time windows. To support this work, we leveraged data from DFW airport using both CAN-bus and SPOT real time vehicle tracking. The optimization model generates optimal routes for within-airport shuttle operations for a given set of headway, in-vehicle travel time, and passenger arrival rate constraints. To make the optimization problem solvable with reasonable computational resources, stochasticity in the passengers' arrivals and bus bunching effects were not captured in the optimization model. This necessitated evaluation of the optimized routes under a stochastic simulation environment, which led to the development of a high-performance stochastic simulation environment (ASPIRES). Together, these tools allow for efficient exploration of a vast set of potential operational scenarios while considering the resulting benefits and costs of each potential configuration.

The work presented here improves on classic solutions to routing problems by running the optimized routes through a discrete-event simulator. This gives airports the flexibility to test the robustness of the solutions generated even before they are implemented in the real world. From the results of the optimization model and the ASPIRES module, it was observed that stringent constrains on headway and maximum ride times lead to a greater number of shuttles required to serve the demand. This translates into a 20% reduction in passenger wait times, which comes at a cost of 50% increase in energy consumption from the shuttle operations. On the other hand, relaxing the constraints on headway and maximum ride times results in fewer shuttles operating on optimized routes. When run through the discrete event simulator, this solution leads to a 20% reduction in energy consumption with a modest 2-min increase in passenger wait times. These results can help DFW (or any airport) evaluate and implement solutions that balance the tradeoff between energy reduction and customer satisfaction.

While preliminary results from application of the proposed model system to data from the DFW airport are encouraging, there are some shortcomings that need to be addressed in future research. As noted in the results section, the optimization model generates some unintuitive (and long) routes (such as R-A-R-C-R-E-R) that satisfy all the constraints of the model. This was remedied using a stop gap measure (with a post processing step) for now, but efforts are currently underway to internalize the constraints that avoid such solutions which perform poorly when simulated. Another limitation worth noting here is that the objective function in the optimization model included fleet energy consumption as the only factor being minimized. Future research efforts should expand the minimization function to include generalized costs that incorporate capital and maintenance as well as fuel costs of shuttle operations. Also, the solutions generated from this work are of little consequence if they do not make it to field implementation. The results from this work were presented to the DFW authorities, and we are currently in discussions with them to implement the best of the solutions generated from this work in order to determine the actual extent of travel and energy implications of the optimized shuttle routes.

Author contributions

The authors confirm contribution to the paper as follows: study conception and design: All authors; data preparation: Andrew Kotz, Devon Sigler, Qichao Wang, Zhaocai Liu; analysis and interpretation of results: Devon Sigler, Qichao Wang, Venu Garikapati, Zhaocai Liu, Andrew Kotz, and Monte Lunacek; draft manuscript preparation: Venu Garikapati, Devon Sigler, Qichao Wang, Zhaocai Liu, Andrew Kotz, Kenneth Kelly, and Caleb Phillips. All authors reviewed the results and approved the final version of the manuscript.

Acknowledgements

This work was supported by the U.S. Department of Energy (U.S. DOE) and Dallas Fort-Worth International Airport (DFW). The contents of this paper reflect the view of the authors who are responsible for the facts and accuracy of the data presented herein. The contents do not necessarily reflect the official view or policies of the U.S. DOE. This work was only made possible through the close cooperation of Dallas Fort Worth International Airport and other partners including North Central Texas Council of Governments. This research team acknowledges and appreciates particular guidance and technical support from Mark Smith, Robert Horton, Esther Chitsinde, Kris Russell, Sarah Ziomek, Zoe Bolack, Jannette Benefee, Greg Royster, Richard Gurley, Smitha Radhakrishnan, Naveen Bandla, Armin Cruz, and Stefan Hildebrand. We thank Jeffrey Gonder, and Dave Biagioni from NREL for providing their valuable feedback on a draft version of this paper.

References

Arnold, Kyle, 2020. DFW Becomes the World's Busiest Airport during COVID-19 Downturn [Internet]. [place Unknown]. The Dallas Morning News; 2020 [updated 2020 June 4, cited 2020 July 23]. Available from: https://www.dallasnews.com/business/airlines/2020/06/04/dfw-becomes-the-worlds-busiest-airport-during-covid-19-downturn/.

Bao, D., Gu, J., Di, Z., Zhang, T., 2018. Optimization of airport shuttle bus routes based on travel time reliability. Math. Probl Eng. 2018 Jan 1.

Brian New, 2020. Summer Travel Season Expected to Be 'Very Strange' at DFW, Love Field Airports [Internet]. CBSDFW.COM, Dallas, TX [updated 2020 May 22, cited 2020 July 23]. Available from: https://dfw.cbslocal.com/2020/05/22/summer-travel-season-very-strange-dfw-love-field-airports/.

Budd, T., Ryley, T., Ison, S., 2014 Apr 1. Airport ground access and private car use: a segmentation analysis. J. Transport Geogr. 36, 106–115.

Budd, L., Ison, S., Budd, T., 2016 Nov 1. Improving the environmental performance of airport surface access in the UK: the role of public transport. Res. Transport. Econ. 59, 185–195.

Bureau of Transportation Statistics, 2017. Long-distance travel [internet]. [place unknown]: bureau of transportation statistics [updated 2017 May 20, cited 2020 July 23]. Available from: https://www.bts.gov/bts/archive/publications/highlights_of_the_2001_national_household_travel_survey/section_03.

Chokshi, Niraj, 2020. 'People Are Getting in Planes': the Travel Business Is Picking up [Internet]. [place Unknown]. The New York Times [updated 2020 June 8, cited 2020 July 23]. Available from: https://www.nytimes.com/2020/06/08/business/airlines-summer-travel-coronavirus.html.

Christofides, N., Eilon, S., 1969 Sep 1. An algorithm for the vehicle-dispatching problem. J. Oper. Res. Soc. 20 (3), 309–318.

Dallas Fort Worth Airport, 2020a. Traffic Statistics [Internet]. [Dallas, TX]. DFW Airport [updated 2020, cited 2020 July 23]. Available from: https://www.dfwairport.com/stats/.

Dallas Fort Worth Airport, 2020b. DFW Airport Achieves Carbon Neutral Accreditation for Third Year [Internet]. [place Unknown]. DFW Airport [updated 2020, cited 2020 July 23]. Available from: http://dfwairport.mediaroom.com/181108-DFW-Achieves-Carbon-Neutral-Accreditation-for-Third-Year.

Feng, L., Miller-Hooks, E., Schonfeld, P., Mohebbi, M., 2014 Jan. Optimizing ridesharing services for airport access. Transport. Res. Rec. 2467 (1), 157–167.

Hart, W.E., Laird, C.D., Watson, J.P., Woodruff, D.L., Hackebeil, G.A., Nicholson, B.L., Siirola, J.D., 2017 May 27. Pyomo-optimization Modeling in python. Springer, Berlin.

Henao, A., Sperling, J., Young, S.E., Garikapati, V., Hou, Y., 2018 Jun 11. Airport Analyses Informing New Mobility Shifts: Opportunities to Adapt Energy-Efficient Mobility Services and Infrastructure. National Renewable Energy Lab.(NREL), Golden, CO (United States).

I-Jy Chien, S., 2005 Oct 1. Optimization of headway, vehicle size and route choice for minimum cost feeder service. Transport. Plann. Technol. 28 (5), 359–380.

Jou, R.C., Hensher, D.A., Hsu, T.L., 2011 May 1. Airport ground access mode choice behavior after the introduction of a new mode: a case study of Taoyuan International Airport in Taiwan. Transport. Res. E Logist. Transport. Rev. 47 (3), 371–381.

Knight, K.W., Hofer, J.P., 1968 Sep 1. Vehicle scheduling with timed and connected calls: a case study. J. Oper. Res. Soc. 19 (3), 299–310.

Kotz A, Ficenec K, Kelly K, Berger S. Electrification of Airport Shuttle Operations. InProceedings of the 33rd Electric Vehicle Symposium 2020 Sep. 13.

Krajzewicz, D., 2010. Traffic simulation with SUMO–simulation of urban mobility. In: Fundamentals of Traffic Simulation. Springer, New York, NY, pp. 269–293.

Lin, S., 1965 Dec. Computer solutions of the traveling salesman problem. Bell System Technical Journal 44 (10), 2245–2269.

Linqing, W., Jun, Z., Wei, W., 2019 Dec 1. Hyperpath-based vehicle routing and scheduling method in time-varying networks for airport shuttle service. Nat. Comput. 18 (4), 769–784.

Malandri, C., Mantecchini, L., Postorino, M.N., 2017 Jan 1. Airport ground access reliability and resilience of transit networks: a case study. Transportation Research Procedia 27, 1129–1136.

Nunes, J., Matos, L., Trigo, A., 2011 Apr 14. Taxi Pick-Ups Route Optimization Using Genetic Algorithms. Springer, Berlin, Heidelberg, pp. 410–419. InInternational Conference on Adaptive and Natural Computing Algorithms.

Optimization, Gurobi, 2014. Inc.,"Gurobi Optimizer Reference Manual," 2015.

Pei-Ying, Y.A., Jia-Fu, T.A., Yang, Y.U., Jin-Xiang, P.E., 2013 Apr 1. Minimizing carbon emissions through vehicle routing and scheduling in the shuttle service of picking up and delivering customers to the airport. Acta Autom. Sin. 39 (4), 424–432.

Ross, S.M., 2014 Jan 8. Introduction to Probability Models. Academic press.

Transit, E.T.A., 2020. Spatial positioning on transit (SPOT). https://etatransit.com/meet-spot/. July.

Xiong, J., Guan, W., Song, L., Huang, A., Shao, C., 2013 Dec 1. Optimal routing design of a community shuttle for metro stations. J. Transport. Eng. 139 (12), 1211–1223.

Yan, Y., Meng, Q., Wang, S., Guo, X., 2012 Dec 1. Robust optimization model of schedule design for a fixed bus route. Transport. Res. C Emerg. Technol. 25, 113–121.

Yu, B., Yang, Z., Sun, X., Yao, B., Zeng, Q., Jeppesen, E., 2011 Dec 1. Parallel genetic algorithm in bus route headway optimization. Appl. Soft Comput. 11 (8), 5081–5091.

Yu, B., Yang, Z.Z., Jin, P.H., Wu, S.H., Yao, B.Z., 2012 Jun 1. Transit route network design-maximizing direct and transfer demand density. Transport. Res. C Emerg. Technol. 22, 58–75.

Zhao, F., Zeng, X., 2008 Apr 16. Optimization of transit route network, vehicle headways and timetables for large-scale transit networks. Eur. J. Oper. Res. 186 (2), 841–855.