TRR
JOURNAL OF THE TRANSPORTATION RESEARCH BOARD

Transportation Research Record
2022, Vol. 2676(12) 85–95

DOI: 10.1177/03611981221095744
journals.sagepub.com/home/trr
SAGE

ASPIRES: Airport Shuttle Planning and Improved Routing Event-driven Simulation

Qichao Wang[1], Devon Sigler[1], Zhaocai Liu[1], Andrew Kotz[1], Kenneth Kelly[1], and Caleb Phillips[1]

Abstract
Most of the existing traffic simulation packages require significant calibration work to be able to reflect reality. To evaluate special operations including emerging technologies, a microscopic simulation that tracks detailed interactions of all the elements of the traffic systems is usually needed. This type of simulation is usually computationally demanding. This work developed an Airport Shuttle Planning and Improved Routing Event-driven Simulation (ASPIRES) package to simulate and evaluate current, potential, and future airport shuttle operations. The simulation was driven by data and thus did not require much calibration effort. The discrete-event simulation nature of ASPIRES makes the simulation computationally efficient. Simulating 1 day of shuttle operations takes less than 2 s. The study site of this work is the Dallas/Fort Worth International Airport in the U.S. The shuttle service that connects the five terminals of the airport and the rental car center was studied. Travel times, dwell times, and passenger arrivals were simulated using empirical distributions derived mainly from real data to capture the stochastic nature of the rental car center shuttle bus operations. Data on bus miles traveled, bus energy consumption, passenger wait times, and passengers left behind at stops were collected to study the trade-off between energy use and passenger experience. Electric bus and on-demand bus operations were also included. The simulation outputs can show passengers statistics at terminals, shuttles statistics, and charging station statistics. ASPIRES cannot be used to model a generic traffic system but is well-suited for fleet systems.

Keywords
operations, traffic simulation, mesoscopic traffic simulation, sustainability and resilience, transportation and sustainability, transportation energy, transportation network modeling and simulation

One challenge in simulating transportation systems is to calibrate simulation models to reflect reality. The stochastic nature of both microscopic-level drivers' behaviors and macroscopic-level road link properties requires the output statistics from a large amount of simulations align with the observations from the field. The dynamics of the simulations for such systems are usually driven by individual vehicles which are modeled with different parameters. The calibration of a large transportation system includes origin-destination (OD) flows, car-following model parameters, and lane-changing model parameters (*1*). To evaluate energy consumption in such simulations, the energy consumption model also needs to be calibrated. The calibration process is usually computationally costly and the calibrated simulation still may not reflect the reality.

This work focused on simulating airport shuttle bus operations, and was motivated by work from an ongoing collaboration with Dallas/Fort Worth (DFW) International Airport in the U.S. The shuttle operations at DFW International Airport were used as a case study for this work. In a previous work, the authors developed a mixed integer optimization model to optimize the routes and bus schedules for the DFW airport shuttle buses (*2*). The routes and bus schedule were optimized with different constraint parameters for decision-makers to choose from. These different constraint parameters

[1]National Renewable Energy Laboratory, Golden, CO

Corresponding Author:
Qichao Wang, Qichao.Wang@nrel.gov

indicated different levels of expectations for the shuttle bus passenger experience. This optimization work resulted in 756 optimized day-long bus route schedules. To gain insight into how these optimized routes would perform in a stochastic operating environment, it was determined they should be tested using a detailed simulation. In addition, DFW airport was evaluating the possibility of adding electric buses to their shuttle fleet which raised many operational and planning questions simulation could shed light on. Thus, a need was identified to have a simulation model with the following features:

- can evaluate the optimized shuttle bus routes and schedules
- can capture the complex shuttle dispatching operations in DFW airport
- can support electric shuttle planning
- can easily be calibrated
- can reflect the stochasticity of the traffic systems

Transportation systems have been widely simulated with discrete time simulations. The systems are usually simulated with certain time steps (usually 1 s or 0.1 s). Before each step, the status of each element in the simulation is updated by various models (e.g., car-following models, lane-changing models, traffic signal control models). Most available traffic simulation software packages can be customized to simulate airport shuttle operations and to evaluate electric shuttle impacts through application programming interfaces (APIs). However, the calibration of drivers' behaviors and vehicles' energy consumption takes significant research and computational efforts. Also, the detailed interaction models among all the elements in the transportation systems slow down the simulation computational time. In another work under the same project, simulating 1 day of DFW airport land-side traffic in the SUMO simulation package took up to 6 h using high performance computing (HPC) systems (*3*, *4*). When traffic became more congested, the simulations with SUMO took a longer time.

In recent years, discrete-event simulation has started to be implemented for simulating traffic systems. Two simulation models—POLARIS and Mobiliti—are actively being developed through U.S. Department of Energy (DOE)-funded projects (*5*, *6*). These two models implement parallel discrete-event simulations to simulate city-scale vehicle movements in a reasonable amount of time (*7*). While they can quickly simulate a city-scale traffic system, significant calibration efforts are still required before using these simulation packages, which can be a time-intensive process.

An early effort in bringing discrete-event simulation to transportation systems modeling is Burghout et al.'s discrete-event mesoscopic traffic simulation model (*8*). This model combined queue-server and speed-density modeling for the mesoscopic traffic modeling and has the ability to integrate with microscopic traffic models to address intelligent transportation systems (ITS) applications. Soh et al. developed a discrete-event traffic simulation model for multilane-multiple intersections based on queuing theory (*9*). Zhang et al. proposed a discrete-event and hybrid simulation framework based on SimEvents which facilitates testing for safety and performance evaluation of an ITS (*10*). Notably it has been used to build a traffic simulation model of a connected and automated vehicles test facility. Aimsun also introduced discrete-event simulations to their capabilities (*11*).

All the aforementioned discrete-event traffic simulation models were designed to be general purpose. None of them can be used to address the requirements for evaluation of shuttle bus operations at DFW airport without a significant amount of customization work. Additionally, the calibration effort is still significant. Therefore, an Airport Shuttle Planning and Improved Routing Event-driven Simulation (ASPIRES) package was developed to simulate the shuttle operations at DFW airport. It can simulate passenger arrivals, passenger pick-up/drop-off, on-demand operations, dispatching operations, and battery charging and discharging behaviors for electric shuttle buses. Compared with existing traffic simulation packages, ASPIRES is driven by empirical distributions of field data, therefore little calibration effort is needed. ASPIRES is cross-platform and can be run on local machines (e.g., personal laptops) or in parallel on HPC. The discrete-event nature of ASPIRES provides very high-performance simulation. Simulating one-day's operation took less than 2 s. Beyond its original purposes, ASPIRES can also be used to simulate any type of fleet operation with fixed stop locations and people/cargo to be carried between the stops. This work brings a new way to model fleet systems that is data-driven and high-performance.

Airport Shuttle Operation at DFW Airport

The airport shuttle services at DFW airport connect airport terminals, employee parking lots, remote parking lots, and the rental car center. This study focused on the shuttle service connecting the airport terminals and the rental car center. The airport has around 50 shuttle buses that can be used to move passengers to and from the rental car center. The shuttle buses serving the rental car center were mainly 43-seat buses at the time of the study. The shuttle buses loop between the rental car center and the airport terminals. A shuttle bus could go to multiple terminals within one trip. There are multiple stops within a terminal for different segments of the curbside. When a shuttle bus can not pick up all the passengers at a stop, the driver will inform the dispatcher to send a bus (called a hotshot bus) to pick up the rest of the passengers at the bus stop.

The airport is interested in acquiring electrified shuttle buses and wanted to evaluate the impact of shuttle electrification. This includes understanding the required battery size/range of each bus, the number of chargers required to keep the fleet sufficiently charged, and the speed at which those chargers must recharge buses. Additionally, in the case that the bus fleet is only part-electric, trade-offs between fleet emission reductions from electrification and battery/charging infrastructure costs emerge.

Airport Shuttle Operation Discrete-Event Modeling

In light of recent success in the literature and performance considerations, it was chosen to model and simulate the airport shuttle operations with discrete-event simulation. Instead of moving the simulation by a time step, discrete-event simulation moves the simulation by events. A discrete-event simulation can be modeled by a time clock $t \in T$, system states $\boldsymbol{S}(t)$, and a dynamic event list $\{\boldsymbol{e}_1, \boldsymbol{e}_2, ...\boldsymbol{e}_i, \boldsymbol{e}_{i+1}, ..., \boldsymbol{e}_n\}$. T is the simulated time duration. Each event $\boldsymbol{e}_i$ is associated with an event time t_i and an event action a_i, that is, $\boldsymbol{e}_i = [t_i, a_i]$. An event $\boldsymbol{e}_i$ can be triggered (created) by an earlier event $\boldsymbol{e}_j$ at Δt_j ago, that is, $t_i = t_j + \Delta t_j$. After an event $\boldsymbol{e}_i$ is processed, the simulator will find the very next event, $\boldsymbol{e}_{i^{Next}}$, in the event list. The system's state does not change between two events, that is, $\boldsymbol{S}(t) = \boldsymbol{S}(t_i)$, for $t \in [t_i, t_{i^{Next}})$. The quantity $t_{i^{Next}}$ can be found using Equation 1:

$$t_{i^{Next}} = \min\{t | t \geqslant t_i, t \in T\} \tag{1}$$

To maintain the event list effectively, a discrete-event simulation usually removes the events from the list after the events have been processed. For the same reason, the event list will be sorted by the event time when a new event is added. This way, the first event in the event list will always be the next event to simulate.

Figure 1 illustrates a single-process discrete-event simulation of a shuttle bus movements. Event e_i happened at time t_i with the action of the bus finishing loading passengers at the rental car center. The next event, e_{i+1}, was created at time $t_{i+1} = t_i + \Delta t_i$ to arrive at a terminal bus stop. The state $\boldsymbol{S}(t_i)$ could be the bus's location is at the rental car center with $x(t_i)$ number of passengers on the bus. Before t_i, $\boldsymbol{S}(t_i - \epsilon)$, with ε being a small number, could be the bus's location is at the rental car center with $x(t_i - \epsilon)$ number of passengers on the bus. Event e_i changed the number of passengers on the bus. $\boldsymbol{S}(t_{i+1})$ could be the bus's location is at the terminal bus stop still with x number of passengers on the bus. The event e_{i+1} was created when event e_i was processed. It is known that, after finishing loading the passengers at the rental car center, the next event is for the shuttle bus to arrive at a terminal bus stop, since it is part of the shuttle bus's route. Parameter Δt_i can be a constant number, for example, expected travel time from the rental car center to the terminal bus stop, that was set. To address the stochasticity of the simulation, Δt_i is usually drawn as a random number from a known distribution.

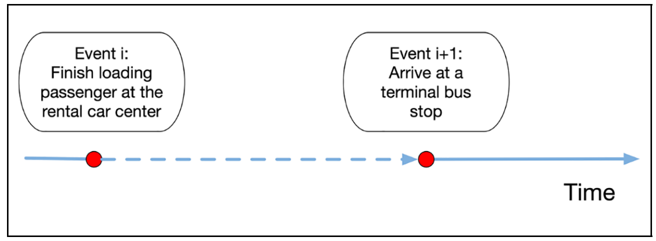

Figure 1. Illustration of one discrete-event simulation.

Figure 2 illustrates the overall ASPIRES model logic. Passengers arrives at different bus stops with their destinations in mind. The shuttle buses carry passengers between the airport terminals and the rental car center following the routes specified by the dispatching center.

The shuttle bus fleet includes electrified buses (noted as EV bus in the figure) and diesel engine buses (noted as Bus in the figure) with two different sizes: 43-seat buses and 14-seat buses. An electrified bus will go to the charging station to charge when the battery level is low. Under opportunity charging scenarios, the electrified shuttle buses will charge the batteries at each bus stop for the duration of the stopped time.

In most scenarios, on-demand operations were inspected. The on-demand operation comes in two forms: on-demand buses and hotshot buses. The on-demand buses are a fleet of buses serving during night-time when the passenger demand is low and no regular bus route is preset. There could be different on-demand policies to send on-demand buses based on the passenger arrivals. The hotshot buses are regular buses following their routes during regular times. They are called hotshot buses when they are sent by the dispatching center to a specific terminal to pick up excess passengers who were not picked up by a previous shuttle bus. This could happen when a shuttle bus is full and there are still passengers at the bus stop. The shuttle buses only get instructed by the dispatching center at the rental car center. Therefore, the hotshot buses only start from the rental car center to pick up passengers at the airport. The hotshot bus is a standard operation at DFW airport, while the on-demand bus is an exploration from the research team.

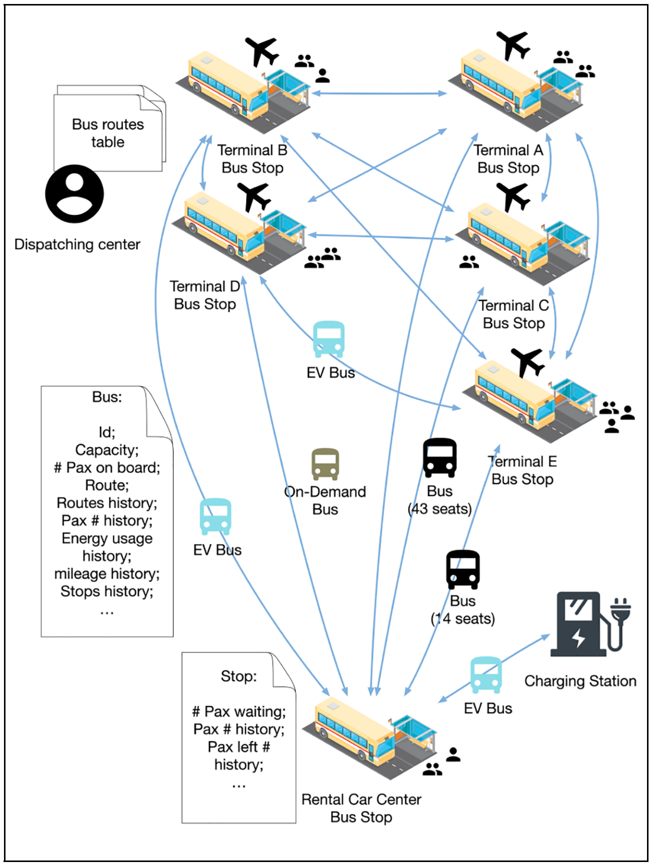

Figure 2. Airport Shuttle Planning and Improved Routing Event-Driven Simulation (ASPIRES) model.
Note: EV = electrified bus.

The ASPIRES package is developed with SimPy as the discrete-event simulation engine (*12*).

Passengers and Bus Stops Model

Passengers arrive at the bus stops and are carried by the shuttle buses to their destination. The passenger arrivals were modeled as Poisson processes. Aggregation was done by day of week and hour of day. The passenger arrivals for each stop, s_m, is an independent Poisson process with arrival rate of $\lambda_{s_m}(t)$. $\lambda_{s_m}(t)$ is approximated by the average number of boarding passengers at stop s_m in the hour of t. In the simulation, it was assumed that a non-existing passenger arrived at the beginning of the simulation as an initial event. The arrival time of the $\{j+1\}^{th}$ passenger at stop s_m, t_{j+1}^{p,s_m} can be calculated by Equation 2:

$$t_{j+1}^{p,s_m} = t_j^{p,s_m} + \Delta t_j^{p,s_m} \qquad \forall S_m, j \qquad (2)$$

The passenger arrival headway, $\Delta t_j^{p,s_m}$, at stop s_m follows exponential distribution and can be calculated by Equation 3:

$$\Delta t_j^{p,s_m} = \frac{-3600 \times \ln u}{\lambda_{s_m}(t)}, u \sim U(0,1) \qquad \forall S_m, j \qquad (3)$$

The passengers were taken by the shuttle buses when the shuttle buses arrived. The passengers boarded the buses on a first-come-first-served basis. The time when each passenger j arrived, t_j^{p,s_m}, at a certain bus stop, s_m, and the time when the passenger boarded the shuttle bus, $t_{j,d}^{p,s_m}$, were recorded for evaluating the waiting time statistics and queue length statistics at the bus stop. The method to derive waiting time and queue length will be presented in the Evaluation in Simulation section.

Also, a lock flag was set for each bus stop. When a shuttle bus started loading passengers, the bus stop will be locked to prevent another shuttle bus from loading passengers at the same time. This was designed because the system state for the bus stop does not change until the loading event is finished. Two shuttle buses loading passengers will cause one passenger boarding two different buses.

A bus stop tracks the current number of passengers, $N_{s_m}^p(t)$, waiting at the bus stop s_m. This information can be used to trigger hotshot buses or on-demand buses. The true headway during the simulation was usually different from the headway that was set for the simulation because of the bus bunching effect caused by randomness. The bus arrival records at the bus stops were tracked to evaluate the true headway.

An airport terminal has two or three bus stops at different locations at the terminal. Five different bus stops were set at the rental car center—each corresponded to one terminal at the airport. In the simulations, a passenger arrived at a bus stop at an airport terminal who had the destination of the rental car center, and a passenger arrived at a bus stop at the rental car center who had the destination of a certain bus stop within the corresponding airport terminal.

Basic Shuttle Model

The basic shuttle model captures the existing diesel engine shuttle bus at DFW airport and also serves as the base for electric shuttle buses and on-demand shuttle. The basic shuttle model has hotshot operation when the hotshot operation is enabled. Each basic shuttle bus is defined by certain parameters (e.g., vehicle ID [vid], capacity [C_{vid}]). The status of a basic shuttle bus $\boldsymbol{S}_{vid}(t)$ includes current route $\boldsymbol{S}_{vid}^r(t)$, accumulative boarding passengers number $S_{vid}^b(t)$, accumulative alighting passengers number $S_{vid}^a(t)$, number of passengers on board $S_{vid}^o(t)$, distance traveled $S_{vid}^d(t)$, and energy consumed $S_{vid}^e(t)$.

All buses are created in the beginning of a simulation based on a route schedule table. The route schedule table includes—at any hour of the day, for each day of the week—what are the bus routes and how many vehicles will be serving each of the routes. The simulation examines the route schedule table to find the maximum number of vehicles needed at any time of the simulated period. The simulation usually starts from midnight when only a few shuttle buses are used. The shuttle buses that are not serving passengers are set to be on a virtual garage parking route.

The current route is assigned by the dispatching center before the start of a trip. A trip is defined to start and end at the rental car center. A trip could be a service trip from the rental car center to certain airport terminals and then back to the rental car center, a charging trip from the rental car center to the charging station and then back to the rental car center, or a parking trip from the rental car center to the parking garage when the bus is not used and then back to the rental car center. The service trip could follow a predefined route or follow an ad hoc route generated by the dispatching center as an on-demand trip or hotshot trip.

Accumulative boarding passengers number, accumulative alighting passengers number, number of passengers on board, distance traveled, and energy consumed are updated by Equation 4:

$$\boldsymbol{S}(t_i) = \boldsymbol{S}(t_i - \epsilon) + \Delta\boldsymbol{S}(t_i) \tag{4}$$

The state accumulative alighting passengers number is updated when the bus unloads passengers at time t_i. All the passengers will leave the bus when the bus arrives at the rental car center. Only the passengers with the destination of the current stop will alight the bus if the bus stop is at an airport terminal.

The state accumulative boarding passengers number is updated when the bus loads passengers at time t_i. $\Delta S^b_{vid}(t)$ can be calculated with Equation 5:

$$\Delta S^b_{vid}(t) = \begin{cases} \min\left\{N^p_{s_m}(t), C_{vid} - S^b_{vid}(t-\varepsilon)\right\}, & t = t_i \\ 0, & o.w. \end{cases} \tag{5}$$

In the rental car center, the bus records the boarding passenger number at each stop to inform the alighting passenger number when the bus arrives the terminals.

The change in number of passengers on board $\Delta S^o_{vid}(t_i)$ is calculated by subtracting newly boarded passengers number and newly alighted passengers number, that is, $\Delta S^o_{vid}(t_i) = \Delta S^b_{vid}(t_i) - \Delta S^a_{vid}(t_i)$. $S^o_{vid}(t)$ is updated after both alighting and boarding. When alighting, $\Delta S^b_{vid}(t) = 0$; when boarding, $\Delta S^a_{vid}(t_i) = 0$. Alighting always happens before boarding to make room for the boarding passengers. In the early stage, the alighting/boarding time was calculated based on the number of alighting/boarding passengers and average alighting/boarding times. When there was access to more granular data, it was found that the dwell time at each stop was not correlated with the alighting/boarding passenger number. One possible explanation is that the bus drivers have various reasons to wait when they stop. For example, they might get out to stretch or may decide to always wait a few minutes to make sure no one else shows up regardless of being instructed to do so or not. Therefore, the dwell time was simply drawn from the empirical distribution from the field data.

Between stops in the service routes, the empirical distribution generated from combined field data and simulated data was used to inform the simulations. The distance between two stops does not change. The energy consumed is highly correlated with the travel time in each trip segment. During congested times, the travel time between two stops is usually higher than other times and the consumed energy is higher than other times. When randomness is introduced in the simulation, the empirical distributions from energy consumption and travel time should be associated with each other.

The refueling process for the diesel engine was not modeled, since the refueling process is much faster than electric shuttle charging, and the service time is longer between two refuels.

Electric Shuttle Model

The main difference between the electric shuttle model and the basic shuttle model is that electric shuttle buses have battery charging and discharging process. The major operational difference between electric fleet and diesel fleet is that the charging time for electric buses is usually much longer than the refueling time of diesel buses because of the limitation of charging and battery technology. The long charging time of electric buses will make it more challenge to provide reliable service with limited fleet size. To ensure that the passenger waiting time is not significantly increased when switching from diesel buses to electric buses, shuttle operators need to ensure enough buses, sufficient battery capacity, and enough chargers in the charging station.

In the simulations, it was assumed that the charging station is located within the rental car center. For other applications, the charging station can be modeled to be at any location. The impact of the charging station location in ASPIRES is modeled as the required distance (a constant number), travel time (from a distribution), and energy consumption (from a distribution) from a stop or the stops which are connected to the charging station in the bus routes.

Electric shuttle has battery capacity B^c_{vid} and charging rate B^r_{vid} parameters to describe the battery characteristics. An electric shuttle, *vid*, tracks its state of charge (SoC), $S^b_{vid}(t)$. In reality, B^r_{vid} is a function of $S^b_{vid}(t)$. Studies showed that when SoC is within a certain range,

$[B_{vid}^{-}, B_{vid}^{+}]$ (e.g., $[20\%, 80\%]$), the charging rate can be viewed as constant (*13*, *14*). In the simulated operations, the batteries are operated within $[B_{vid}^{-}, B_{vid}^{+}]$ so that B_{vid}^{r} is constant during the battery charging time. In the simulation, SoC is also updated by Equation 4. When energy consumption historic data is available, $\Delta S_{vid}^{b}(t_i)$ can be calculated by $S_{vid}^{e}(t_i)$ by Equation 6:

$$\Delta S_{vid}^{b}(t_i) = \begin{cases} \frac{-S_{vid}^{e}(t_i)}{B_{vid}^{c}}, & \boldsymbol{e}_i \text{ is a discharging event} \\ \Delta t_j^{vid} \times \frac{B_{vid}^{r}}{B_{vid}^{c}}, & \boldsymbol{e}_i \text{ is a charging event} \end{cases} \tag{6}$$

The charging time Δt_j^{vid} can be the dwell time for opportunity charging. If the shuttle uses the charging station, Δt_j can be calculated based on the SoC of the battery and charging rate, that is:

$$\Delta t_j = \frac{B_{vid}^{c} \times (B_{vid}^{+} - S_{vid}^{b}(t_j))}{B_{vid}^{r}} \tag{7}$$

Electric shuttle buses serve passengers the same as diesel engine shuttle buses. After an electric shuttle bus has dropped passengers at the rental car center, it will check the battery SoC and the charging station status to decide whether to go to the charging station or keep serving the next trip. A logic for electric shuttle buses was designed to reduce the chance of wasting time waiting in line at the charging station.

Two thresholds were defined: B_{charge}^{-} and B_{charge}^{+}, $B_{charge}^{-} <= B_{charge}^{+}$. When $S_{vid}^{b}(t_j) > B_{charge}^{+}$, the shuttle bus *vid* will keep serving passengers if needed. When $S_{vid}^{b}(t_j) \in [B_{charge}^{-}, B_{charge}^{+}]$, the shuttle bus *vid* will check the charging station to see if there is a charger available. If all the chargers in the charging station are occupied, the shuttle bus will keep serving passengers; otherwise, the shuttle bus will take a charger and start a charging event $\boldsymbol{e}_j$ for a duration Δt_j which can be calculated by Equation 7. This operation will not happen if the two thresholds are set to be the same, that is, $B_{charge}^{-} = B_{charge}^{+}$. When SoC is lower than the lower threshold, that is, $S_{vid}^{b}(t_j) < B_{charge}^{-}$, shuttle bus *vid* will go to the charging station and wait in the queue to get charged. The time in the queue is recorded separately from the charging time. This is designed intentionally so that ASPIRES can be used to also model the power grid impact during the electric vehicles' charging processes.

When an electric shuttle goes to the charging station, the route that the shuttle was serving needs another vehicle to replace the electric shuttle. This is covered by the bus dispatching function of the dispatching center.

When an electric shuttle is not needed (i.e., less buses are needed in the service), before going to the garage, the shuttle bus will take the opportunity to charge its battery to B_{vid}^{+}.

On-Demand Shuttle Model

The on-demand shuttle buses were modeled with the intention to serve night-time passengers with smaller shuttle buses (i.e., the 14-seat buses). The dispatching center set the route for an on-demand shuttle bus if the bus is available to answer the on-demand call (i.e., the bus is not serving passengers). There could be many implementations of the on-demand policy to decide how to generate an on-demand route based on the number of passengers waiting at each bus stop. A naive policy is to check the bus stops periodically and send an on-demand bus to visit all the bus stops with passengers waiting to serve all the passengers. The special setting of an airport shuttle bus makes it clear to see where the passenger's destination is, if it is known at which bus stop that passenger is waiting. For example, a passenger waiting at any bus stop in the terminal has the destination of the rental car center.

Charging Station Model

One charging station for the DFW airport shuttle fleet was defined in the simulation. The charging station is modeled as an N-server queuing system with each charger in the charging station as a server. The chargers' usage history, $N^c(t)$, $t \in \boldsymbol{T}$, through the simulation is recorded.

Data Collection

This study used data collected from the field to feed the simulation. For the routes that did not exist or with missing data, data evaluated from a high-fidelity simulation model of DFW airport was used (*3*).

Field Data

Passengers' demand data, existing routes, travel time, and energy consumption data were collected for each trip segment.

The demand data gives time-dependent passenger OD information. For the studied scenarios, the passengers either leave from the rental car center to go to airport terminals or leave from the terminals for the rental car center, making OD information easy to calculate. The number of passengers leaving a certain terminal to go to the rental car center can be determined by the number of passengers boarding the shuttle buses at the terminal. The number of passengers moving from the rental car center to a certain terminal can be reflected by the alighting passenger from the shuttle buses' number at that terminal. During the study, the passenger demand was first collected from the drivers' notes. The shuttle bus drivers manually wrote down the number of passengers loaded

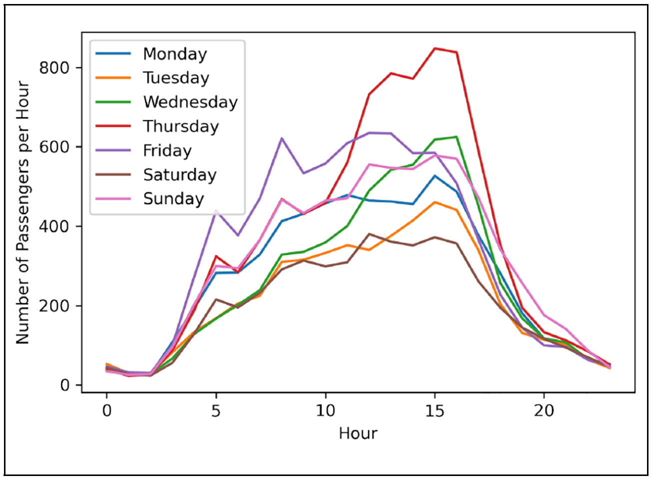

Figure 3. Passenger arrival rates at rental car center.

to the buses at each stop and aggregated every half an hour. A total of 16 months of these records were collected. In the later phase of the study, DFW airport implemented a global positioning system (GPS)-enabled automated passenger counter (APC) system that reports the time, location, route, and boarding and alighting passenger counts. Two months of data was collected from the bus APC system. In the early stage of the ASPIRES development, the manually collected data from drivers was used. Once access to APC data was available, a switch was made to APC data. For readers who want to implement ASPIRES-like simulations for other sites, if there is no APC system implemented, the drivers' notes can be used to provide passenger demand data to drive the simulations. Figure 3 shows the passenger arrival rates through different time of day and day of week at the rental car center stop.

ISAAC Instruments DRU-900 controller area network (CAN)-based data loggers with GPS were also used. Data collected from CAN conforms to the Society of Automotive Engineers (SAE) J1939/71 standard which provides instantaneous vehicle speed, fuel rate, and engine power that is recorded with GPS location at 1 s intervals (1Hz rate). Approximately 1 month of bus operation data were collected to capture representative vehicle operations and ensure all the various currently used routes were covered.

Vehicle location from the GPS was used to identify and segment the data by route. Once segmented, energy consumption and operational analysis metrics were calculated from the data, including distance, average speed, fuel consumption, engine output energy in kWh, and fuel economy. Vehicle speed and distance were captured using the GPS system, and fuel metrics were calculated using the fuel rate metric. Finally, the engine power output in kWh was calculated using the estimated engine torque and engine speed data channels. The instantaneous engine power is then multiplied by the data capture rate of 1 s and summed over the segmented data to get a trip positive tractive energy requirement in kWh. While this data was collected from conventional vehicles, the tractive energy requirement in kWh can be approximated to the energy requirement of an electric bus operating over the same route. In actuality, the electric buses will be able to recapture energy through regenerative braking, meaning the total trip energy of an electric bus will be slightly less than the positive tractive energy requirement of a conventional vehicle and have an overall higher efficiency.

Simulated Data

The ASPIRES framework evaluated different routing strategies, including the routes with non-existing trip segments. For example, the existing routes in the field do not have shuttle buses moving from terminal C to terminal E. Therefore, the data collected from the field cannot be used alone to execute the simulations. Advantage was taken of a calibrated high-fidelity simulation of the area done by NREL researchers to augment the field data. The high-fidelity simulation was modeled in the SUMO simulation package (*4*). The travel time and energy consumption data for each possible trip segment was extracted to augment the field data.

It was found that the overall mean and standard deviation for energy consumption are 2.29 kWh/mi and 1.04 kWh/mi, respectively.

Evaluation in Simulation

The purpose of developing ASPIRES was to evaluate different planning and routing scenarios. The focus was on two types of evaluation: passenger service level and utility usage.

Passenger Service Level

In this study, the passenger service level is indicated by the queue length and waiting time. From the passenger arrival and boarding time records, it is possible to generate a cumulative arrival function, $A_{s_m} : T \rightarrow N$, and a cumulative loading (departure) function, $D_{s_m} : T \rightarrow N$ for a bus stop s_m, as shown in Figure 4.

The queue length at bus stop s_m at time t is the vertical distance between the two curves at time t, that is, $q_{s_m}(t) = A_{s_m}(t) - D_{s_m}(t)$. For instance, the queue length at 10:20 in Figure 4 was 181 − 181 = 0. The waiting time for passenger j at stop s_m is the horizontal distance between the two curves defined by

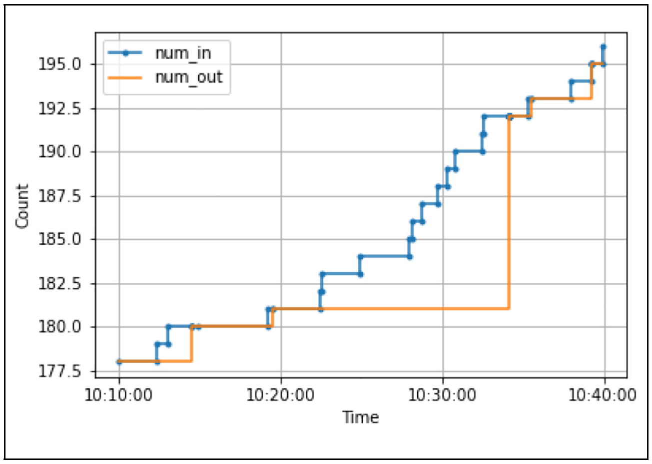

Figure 4. A cumulative arrival departure curve plot.

$w_{s_m}(j) = \min\{\mathbf{t_d}\} - \max\{\mathbf{t_a}\}$ where $\mathbf{t_d} = \{D_{s_m}(t_d) = j\}$ and $\mathbf{t_a} = \{A_{s_m}(t_a) = j\}$.

Utility Usage

The bus occupancy, total vehicle miles traveled, and total energy used during the simulation are evaluated. The bus occupancy comes from the number of passengers on board record, $S^o_{vid}(t)$. The vehicles miles traveled and total energy used come from the corresponding states records of the basic shuttle model.

For scenarios with electric shuttle buses in the fleet, the number of chargers being used through the simulation, $N^c(t)$, $t \in \boldsymbol{T}$, is also evaluated. If $\max_{t\in\mathbf{T}} N^c(t)$ is less than the total number of chargers set in the charging station, only $\max_{t\in\mathbf{T}} N^c(t)$ chargers were needed in the simulated setting.

Routing Strategies

The routing strategy came from a shuttle bus route optimization work for DFW airport. The shuttle bus route optimization used mix-integer optimization to generate the optimal route combinations, number of buses on each route, and the bus type for each given time window of each day of the week. The optimization goal was to minimize energy usage of the entire shuttle fleet during each time window. More details of the model was reported in Sigler et al. (*2*). There is a clear trade-off between passengers' service level provided by the shuttle and the energy consumption from serving the passengers. To allow the decision-makers to evaluate the optimal routes under different passenger service-level constraints, the routes and frequencies were optimized under different parameter settings (e.g., maximum headway, maximum in-vehicle travel time, and predicted arrival rates).

Also, a baseline simulation was conducted to simulate the current route schedule. The bus tracking system's report from DFW airport was processed, and the routes and schedule were re-created for the baseline simulation.

Results

ASPIRES can be run with *UNIX-like commands which enables single simulation run at a time on a local machine or massive parallel simulation runs on an HPC cluster through Slurm workload manager. Besides being able to run thousands of simulation scenarios at the same time, running on an HPC cluster also makes it possible to take advantage of the parallel file system for faster reading and writing. For users who do not have access to HPC, running ASPIRES on a local machine is also fast enough for what-if scenario evaluations. Thousands of simulation evaluations with ASPIRES were run on NREL's HPC system, Eagle. On average, simulating 1 day of operation took around 1 s of computational time. Some simulated results are presented here. First, the results for evaluating improved routes against a baseline simulation are showcased. Later, the results for evaluating electric shuttle operations are presented.

Routes Comparison

The existing routes configuration was simulated as the baseline simulation. To measure the accuracy of the ASPIRES simulation of bus fleet energy consumption, the baseline simulation was compared against DFW fuel and mileage logs from August, 2018, to August, 2019, for the rental car shuttle fleet. The data from these logs was used as the ground truth for estimating the average daily energy consumption of a DFW rental car center bus. ASPIRES was used to simulate the baseline scenario, which used the same routes the buses used during the time period the fuel and mileage logs were collected. It was found that ASPIRES overestimated the energy consumption of a DFW rental car center bus by less than 2% on average, when the data sources mentioned previously were used in the simulation. This result shows that the ASPIRES model can very accurately simulate energy consumption from a bus fleet provided there is sufficient data to drive the simulation.

The existing operation, most of the time, has dedicated routes between the rental car center and each of the terminals (as shown in Figure 5). The circled "R" represents the rental car center and the other circled letters represent different terminals. There are five existing routes, each connecting the rental car center and one of the terminals. The optimized operations could have one route visiting multiple terminals before getting back to the terminal (see Figure 6 for example).

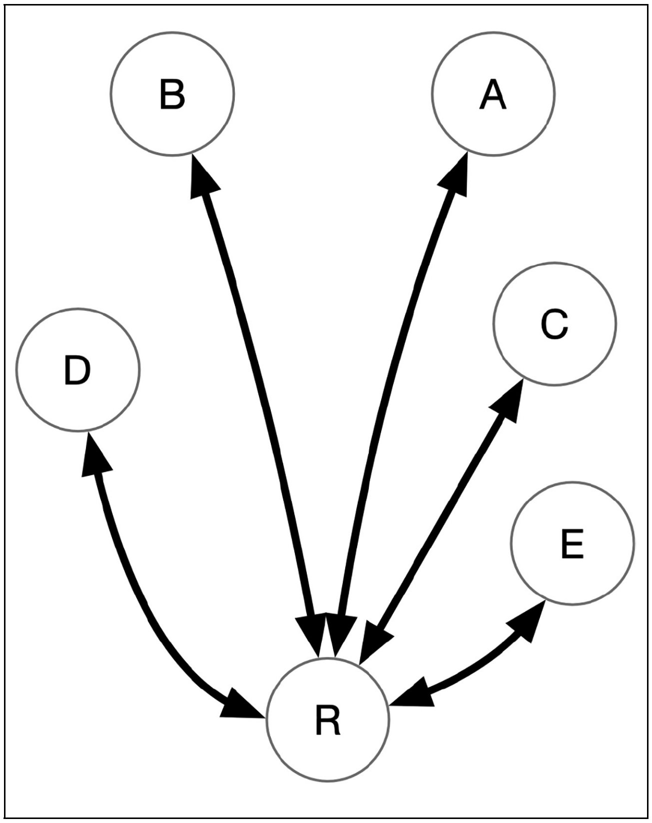

Figure 5. Existing routes.

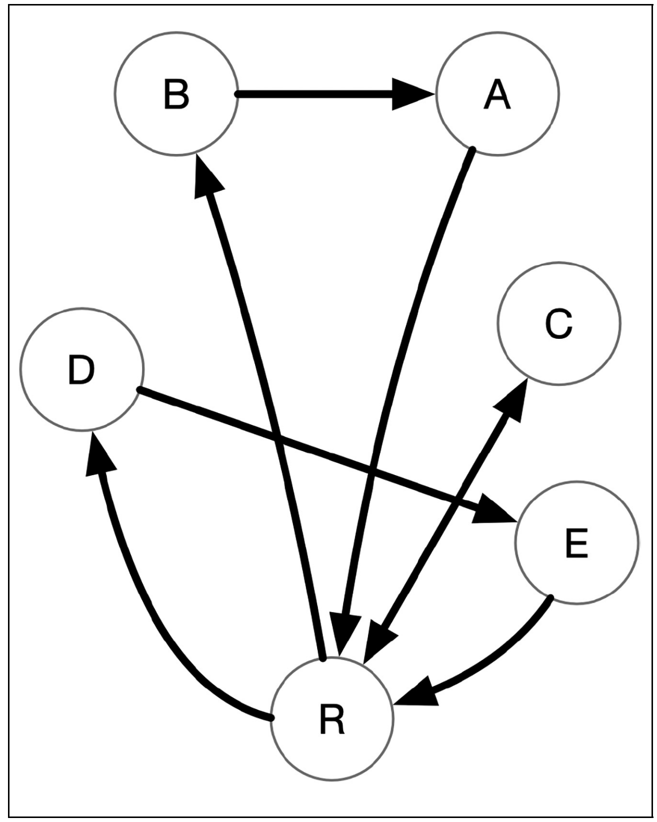

Figure 6. One set of optimized routes for 4 to 8 a.m. time window on a Monday.

The high computational performance of ASPIRES makes it possible to evaluate many operation settings. Figure 7 shows the trade-off between energy consumption of the fleet operation and the fleet service level (represented by passenger waiting time). The *x*-axis is the average system-wide passenger waiting time in minutes. The *y*-axis is the total energy consumption of the fleet operation in a typical week in gasoline gallon equivalent (GGE). Each blue point represent a certain route setting which has different bus routes and frequency across different time of day and different day of week. The route settings were generated from a mixed integer optimization model (see more details in Sigler et al.) under different optimization parameters (*2*). Under one set of optimization parameters, the optimization model gave one route setting. The red line in the figure gives the pareto front of minimizing energy consumption and mean passenger waiting time. It can be seen that, from setting A to setting B in the figure, with less than 2 min increase in mean passenger waiting time, over 25% savings in energy consumption were achieved. Note that all the points in the figure came from optimized settings. It is intuitive that reducing energy consumption (saving operational cost) and reducing passenger waiting time (improving customer service) are two competing objectives. Evaluating different (optimized) routes with ASPIRES enables decision-makers to see the trade-off between the two objectives quantitatively.

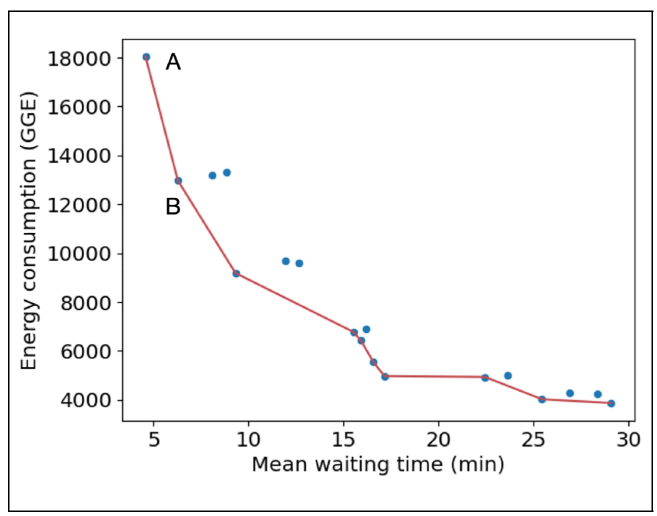

Figure 7. Trade-off between energy consumption and passenger waiting time under different route settings.
Note: GGE = gasoline gallon equivalent.

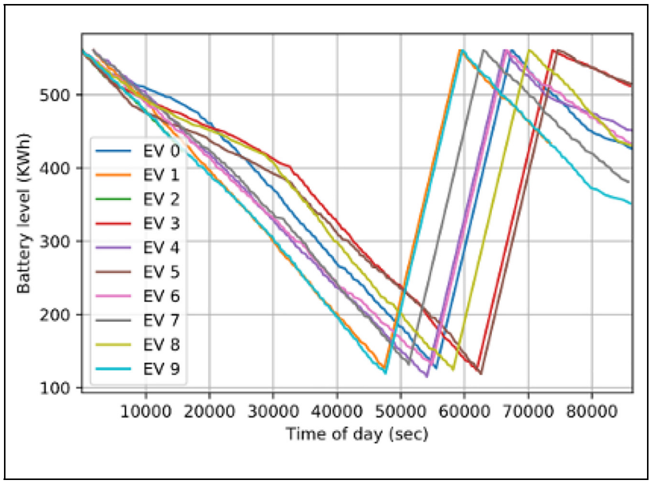

Figure 8. The battery level profile for charging the electrified bus (EV) shuttles when the state of charge (SoC) is lower than 20%.

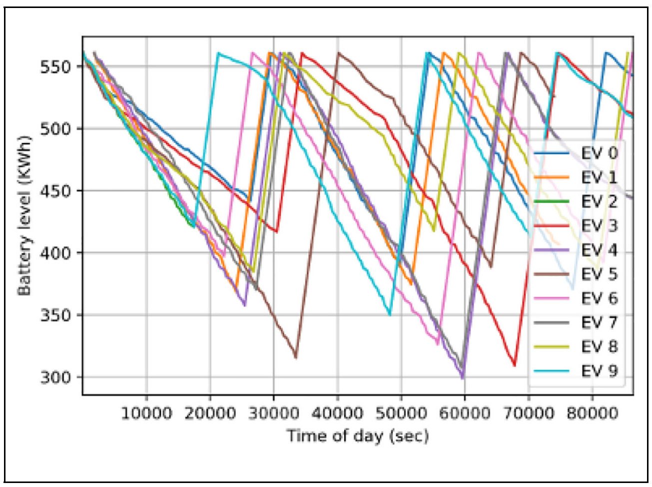

Figure 9. The battery level profile for charging the electrified bus (EV) shuttles randomly when the state of charge (SoC) is between 20% and 70%.

Electric Shuttle Operation

Two different electric shuttle charging strategies were compared: 1) charge the battery only when SoC is lower than 20%, and 2) charge the battery randomly when SoC is between 20% and 70%. Figures 8 and 9 show the battery level profiles of the two charging strategies. In both strategies, 10 electric shuttle buses were introduced to the airport shuttle fleet. The horizontal axis is the time of day in seconds during the simulation and the vertical axis is the battery level in KWh. The battery level decreased when the electric shuttle bus was serving trips. The different discharging rates of different vehicles came from different routes and the randomness of the simulations. The battery level increased when the electric shuttle bus was charging its battery. In both scenarios, the shuttle buses can only charge the batteries at the charging station.

It can be observed that vehicles that charge only at a low SoC tend to charge together for a fairly long time (as shown in Figure 8). This could cause queues building up at the charging station if the charger is limited. If the vehicles randomly choose to charge when SoC is between 20% and 70%, the charging requests were spread across the time (as shown in Figure 9).

Conclusions and Future Opportunities

This paper presented the ASPIRES model. Unlike other traffic discrete-event simulation models, ASPIRES cannot model a generic traffic system. However, ASPIRES is well-suited for a specific type of traffic system (i.e., fleet operations). The constrained targeting modeled systems enable ASPIRES to fully take advantage of discrete-event simulation. ASPIRES skips the modeling of detailed interaction among all vehicles on the road but focuses on the operations of the fleet vehicles. The impacts of the surrounding environment (including other vehicles on the road, road infrastructure conditions, traffic signal operations) are captured by historical data (and simulated data) which drives the simulations. Therefore, ASPIRES is computationally efficient. The data-driven nature of ASPIRES saves calibration effort while providing highly realistic evaluation.

The electric shuttle model enabled planning for transportation electrification. The on-demand shuttle model could be used for studying on-demand policies in preparation for connected and automated vehicles in the fleet operations. The charging station in ASPIRES can also serve as the interface to connect transportation simulation and power systems simulation. The fast and parallel simulation capability enables simulation-based optimization (e.g., optimal charger number at the charging station) and optimal control (shuttle routes control through reinforcement learning) for such systems.

ASPIRES can be generalized for other types of fleet systems. It can be directly adopted for other types of shuttle bus systems. Given passenger route choices, it is possible to extend the simulated scenarios from airport shuttle bus systems to urban transit bus systems. If passengers are changed to goods, ASPIRES can be used to simulate freight fleet systems including delivery/pick-up fleets. With different cargos, ASPIRES can be adapted for simulating special fleets such as simulating mining fleets. All the aforementioned fleets are, or will soon be, undergoing an electrification process. The generalized ASPIRES could be used to guide the electrification process (e.g., evaluating the impacts of different numbers of electrified fleet vehicles, different charging station configurations for the fleet systems, different vehicle and battery sizes, and different routing strategies under electrified vehicles).

Acknowledgments

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 and American Airlines. This research team acknowledges and appreciates particular guidance and technical support from Robert Horton, Esther Chitsinde, Kris Russell, Sarah Ziomek, Zoe Bolack, Jannette Benefee, Greg Royster, Richard Gurley, Smitha Radhakrishnan, and Stefan Hildebrand.

Author Contributions

The authors confirm contribution to the paper as follows: study conception and design: Q. Wang, D. Sigler, K. Kelly, C. Phillips; data collection: A. Kotz, K. Kelly, D. Sigler, Z. Liu; analysis and interpretation of results: Q. Wang, Z. Liu; draft manuscript preparation: Q. Wang, Z. Liu, A. Kotz. All authors reviewed the results and approved the final version of the manuscript.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by the U.S. Department of Energy (U.S. DOE) and Dallas Fort/Worth (DFW) International Airport.

ORCID iDs

Qichao Wang https://orcid.org/0000-0002-0863-4564
Devon Sigler https://orcid.org/0000-0001-7672-6658
Zhaocai Liu https://orcid.org/0000-0002-3016-6832
Andrew Kotz https://orcid.org/0000-0002-0865-7416
Caleb Phillips https://orcid.org/0000-0002-3665-4239

References

1. Balakrishna, R., C. Antoniou, M. Ben-Akiva, H. N. Koutsopoulos, and Y. Wen. Calibration of Microscopic Traffic Simulation Models: Methods and Application. *Transportation Research Record: Journal of the Transportation Research Board*, 2007. 1999: 198–207.
2. Sigler, D., Q. Wang, Z. Liu, V. Garikapati, A. Kotz, K. J. Kelly, M. Lunacek, and C. Phillips. Route Optimization for Energy Efficient Airport Shuttle Operations–A Case Study from Dallas Fort Worth International Airport. *Journal of Air Transport Management*, Vol. 94, 2021, p. 102077.
3. Ugirumurera, J., J. Severino, K. Ficenec, Y. Ge, Q. Wang, L. Williams, J. Chae, M. Lunacek, and C. Phillips. A Modeling Framework for Designing and Evaluating Curbside Traffic Management Policies at Dallas-Fort Worth International Airport. *Transportation Research Part A: Policy and Practice*, Vol. 153, 2021, pp. 130–150.
4. Lopez, P. A., M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner. Microscopic Traffic Simulation Using SUMO. *Proc., 21st International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, IEEE, New York, 2018, pp. 2575–2582.
5. Auld, J., M. Hope, H. Ley, V. Sokolov, B. Xu, and K. Zhang. POLARIS: Agent-Based Modeling Framework Development and Implementation for Integrated Travel Demand and Network and Operations Simulations. *Transportation Research Part C: Emerging Technologies*, Vol. 64, 2016, pp. 101–116. https://doi.org/10.1016/j.trc.2015.07.017.
6. Chan, C., B. Wang, J. Bachan, and J. Macfarlane. Mobiliti: Scalable Transportation Simulation Using High-Performance Parallel Computing. *Proc., 21st International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, IEEE, New York, 2018, pp. 634–641. https://doi.org/10.1109/ITSC.2018.8569397. 7.
7. Liu, J. Parallel Discrete-Event Simulation. In *Wiley Encyclopedia of Operations Research and Management Science* (Cochran, J. J., L. A. Cox, Jr., P. Keskinocak, J. P. Kharoufeh, and J. C. Smith, eds.), Vol. 8, 2011.
8. Burghout, W., H. N. Koutsopoulos, and I. Andreasson. A Discrete-Event Mesoscopic Traffic Simulation Model for Hybrid Traffic Simulation. *Proc., IEEE Intelligent Transportation Systems Conference*, Toronto, ON, Canada, IEEE, New York, 2006, pp. 1102–1107.
9. Soh, A. C., M. H. Marhaban, M. Khalid, and R. Yusof. A Discrete-Event Traffic Simulation Model for Multilane-Multiple Intersection. *Proc., 9th Asian Control Conference (ASCC)*, Istanbul, Turkey, IEEE, New York, 2013, pp. 1–7.
10. Zhang, Y., C. G. Cassandras, W. Li, and P. J. Mosterman. A Discrete-Event and Hybrid Traffic Simulation Model Based on SimEvents for Intelligent Transportation System Analysis in Mcity. *Discrete Event Dynamic Systems*, Vol. 29, No. 3, 2019, pp. 265–295.
11. Casas, J., J. L. Ferrer, D. Garcia, J. Perarnau, and A. Torday. *Traffic Simulation With Aimsun*. Springer, New York, NY, 2010, pp. 173–232. https://doi.org/10.1007/978-1-4419-6142-6_5.
12. Matloff, N. *Introduction to Discrete-Event Simulation and the Simpy Language*, Vol. 2. Department of Computer Science, University of California, Davis. August 2008, pp. 1–33.
13. Ko, Y. D., and Y. J. Jang. The Optimal System Design of the Online Electric Vehicle Utilizing Wireless Power Transmission Technology. *IEEE Transactions on Intelligent Transportation Systems*, Vol. 14, No. 3, 2013, pp. 1255–1265. https://doi.org/10.1109/TITS.2013.2259159.
14. Liu, K., X. Hu, Z. Yang, Y. Xie, and S. Feng. Lithium-Ion Battery Charging Management Considering Economic Costs of Electrical Energy Loss and Battery Degradation. *Energy Conversion and Management*, Vol. 195, 2019, pp. 167–179. https://doi.org/10.1016/j.enconman.2019.04.065.

The contents of this paper reflect the view of the authors who are responsible for the facts and accuracy of the data presented here. The contents do not necessarily reflect the official view or policies of U.S. DOE.