

Scheduling and Airport Taxiway Path Planning under Uncertainty

Jiaoyang Li* and Han Zhang*

University of Southern California, Los Angeles, CA, 90089

Mimi Gong[†], Zi Liang[†], Weizi Liu[†], Zhongyi Tong[†], and Liangchen Yi[†]

Carnegie Mellon University Silicon Valley, Moffett Field, CA, 94035

Robert Morris[‡] and Corina Pasareanu[§]

NASA Ames Research Center MS-269-1, Moffett Field, CA, 94035

Sven Koenig[¶]

University of Southern California, Los Angeles, CA, 90089

Congestion and uncertainty on the airport surface are major constraints to the available capacity of the air transport system. This project seeks to study the problem of planning and scheduling airport surface movement at large airports. Specifically, we focus on scheduling and taxiway path planning of multiple aircraft. In this paper we describe a simulation tool that is capable of simulating aircraft movement along the taxiway, including models of uncertainty during aircraft movement. We introduce a new approach to scheduling that includes models for predicting surface movement uncertainty.

I. Introduction

Airport surface operations present a difficult, large-scale logistics problem with a wide range of sub-problems, including: runway sequencing and scheduling; spot or gate release scheduling; gate allocation and taxi route planning and scheduling [1–4]. Surface movement planning and scheduling is dynamic, with aircraft continuously entering and leaving the operating space. Furthermore, surface movement is unpredictable and prone to unexpected changes in operating conditions due to external factors such as weather. In general, efficiency and safety are difficult objectives to achieve in practice, due to the challenges posed by the presence of uncertainties, human factors, and competing stakeholder interests.

Fast-time simulation of uncertainty in airport surface operations allows for the testing and analysis of modeling concepts and algorithms for planning and scheduling aircraft movement. This paper extends previous and current work in this area [5, 6] by evaluating the result of incorporating probabilistic models into the planning and scheduling process. Our goal is to compare probabilistic approaches to planning and scheduling with standard heuristic approaches that assume a deterministic world, and as a separate stage address uncertainty through continuous re-planning [1]. The goal is to determine whether incorporating recent advances in probabilistic reasoning in planning and scheduling yields more robust schedules through the anticipation of surface delays.

This work intersects previous efforts in at least three areas. First, the problem to be solved requires the coordinated planning and scheduling for multiple agents [7]. Second, this work expands upon work on so-called ‘rolling horizon’ approaches to solve complex scheduling problems under uncertainty [8]. Finally, this work contributes to recent work at building models of uncertainty to improve the robustness of solutions to planning and scheduling problems (for example [9]). This work is motivated in part by recent assessments [6] of the potential benefits of more predictable schedules in improving the efficiency of airport operations, as well as research in Integrated Arrival, Departure and Surface Operations (IADS), including the ATD-2 System [10].

In the remainder of this paper we will provide an overview of a fast-time simulation and scheduling tool for evaluating approaches to planning surface operations, including the technical approach to modeling the geometry and dynamics of

*PhD Student, Computer Science Department.

[†]Master Student, Information Networking Institute.

[‡]Senior Researcher, Intelligent Systems Division.

[§]Researcher, Intelligent Systems Division/Associate Research Professor, Carnegie Mellon University Silicon Valley.

[¶]Professor, Computer Science Department.

airport movement, and the approach to scheduling. We conclude with a summary of the current and expected results of this work for the final draft of the paper.

II. System Overview

The Airport Surface Planner and Simulation System is a fast time simulator for generating statistics about the performance of automated planning and scheduling tools for complex airports. Its overall architecture is summarized in Figure 1.

The inputs to the main system component, the Simulator, consists of the following:

- An airport *node-link model* automatically generated from a Google Map page;
- An *aircraft dynamics model*, for specifying speeds along the taxiway for different types of aircraft;
- A set of *designated path routes* to and from gates and runways, used by the scheduler to generate itineraries for each aircraft on the surface;
- A *scenario*, which describes a problem instance to the system, consisting of a set of arrivals and departure information, including scheduled arrival or departure times, as well as gate and runway information;
- A set of *clock parameters* to guide the progression of the simulation;
- A set of *input parameters* that describe the level of uncertainty in the problem. These parameters allow for the testing of different levels of uncertainty during simulation;
- A *scheduler* to be used in the simulation; and
- A set of *test parameters* that allow for the collection of statistics when the simulator is in batch mode (i.e. when multiple simulations are run in batch mode).

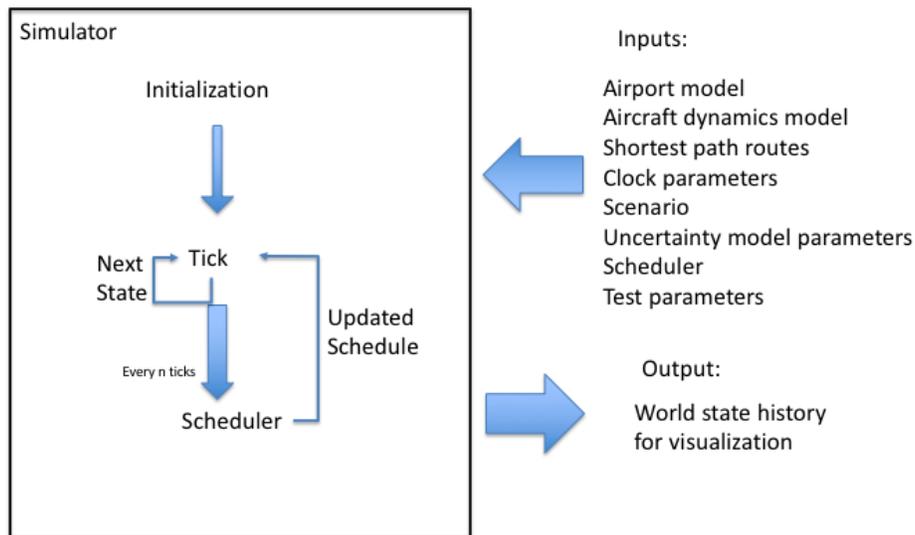


Fig. 1 Airport Surface Planner and Simulation Architecture

The Simulator is comprised of three main components: the initialization, the *Tick* module, and the Scheduler. The Initialization component loads the inputs and models, and invokes the main loop, called the *Tick* module. The *Tick* module advances a world state in time, where a world state is comprised of the locations and speeds of all the aircraft. A world state also consists of a set of itineraries (schedules) for each aircraft, generated from calls to the Scheduler. *Tick* uses the current state as well as the input data, models and parameters to add, delete or advance aircraft along the airport surface. The output of each run of *Tick* is a new state. The new state is either fed back directly into the *Tick*, or,

as designated intervals, into the Scheduler. Finally, the Scheduler updates the itineraries associated with the airport state. Specifically, an itinerary is a set of aircraft and associated *targets*, where the targets are nodes on the node-link graph of the airport. The targets are either generated by the pre-compiled shortest path routes, or incrementally inside the scheduler itself. The output of the scheduler is an updated set of conflict free itineraries, which are fed back into the next iteration of the Tick. An itinerary is said to be conflict-free if no pair of aircraft are ever in an unsafe state (violating separation constraints) as the result of executing their associated itineraries simultaneously.

The entire system was implemented in Python and C++. The remainder of this section will explore the main system components in more detail.

A. Node-Link Model and Scenario Generation

The system allows for a node-link model to be created for any airport from Google Maps. The steps required to build an airport model consist of, first, editing a Google map by adding the desired nodes and edges for runways, taxiways, spots (nodes linking ramp areas with taxiways), and gates; second, downloading the KML file from Google Maps; and third, build a graphical node-link model by running a Python script using the downloaded KML file. It should be noted that the system automatically creates a node-link model with two degrees of resolution: a coarse model that can be used by the Scheduler to generate itineraries; and a fine-grained model that allows for a more realistic simulation and visualization of airport movement.

For faster processing of large airport models, we performed some optimizations. In particular, the routing and scheduling requires a model with nodes for every intersection on the surface, in order to test for conflicts, as discussed below. In addition, for high resolution modeling of movement along a curved link, we split a link into a set of segments that were approximately joined by a straight line. Adding gates and runway terminal nodes results in a graph with a potentially large number of nodes. On the other hand, airports tend to be graphically sparse: each node typically has only one or two nodes coming in or out of it, and relatively few nodes are intersecting nodes. Therefore as a pre-processing step, for the purpose of fast routing, a set of optimizations was performed to limit the graph used for routing to nodes that formed intersections between links.

A random scenario can be generated from a node-link model. The scenario generator is not at this point fully automatic; rather the user must specify a set of arrival and departure *flight templates* using the nodes of the model (gates, runways) and a list of aircraft. The generator will randomly assign arrival or departure times to the templates. The user can set parameters that indicate the 'tightness' of the scenario, i.e., the spacing of the times. Current work on this project includes developing a scenario specification that is based on real surface movement data.

For this paper, we selected a specific flow pattern or arrivals and departures from San Francisco Airport (SFO). Figure 2 shows a node-link model extracted from Google Maps, containing 100 gates, all taxiways and 4 runways that allow for modeling flows on specific pattern called the Southeast Plan. In this flow, aircraft typically arrive on Runway 19L with Runway 19R used as an alternate arrival runway (red lines running northeast/southwest). Departure runways are 10L and 10R (red lines running northwest/southeast). Also modeled are exits and entrances to taxiways from runways, and taxiways (blue lines). Finally, unrestricted areas contain gates and node/links to spot nodes. The result, although a simplification from the complete model of airport flow, is complex enough to test our concepts of planning under uncertainty.

B. Simulator and Visualizer Models

Airport operations is a distributed logistics problem. Consequently, to simulate operations we model airport surfaces as directed graphs and model aircraft as points moving along graph edges with safety distance enforced between each pair of aircraft. Specifically, a Controller module (invokes during a Tick operation) observes the current flow to control the traffic in the intersections. It decides the order of aircraft to pass through each intersection while keeping safe separation between aircraft. The eventual goal of the Controller is to simulate common ATC instructions to pilots, such as "Taxi via...", "Cross Runway..." , "Hold short", etc [11].

An *agent model* is used to simulate aircraft decision making on the surface (this could be the pilot, or an autonomous aircraft towing vehicle, as proposed in [12]). Specifically, this model is used to simulate pilot adjustments to ensure safe distances. In developing the agent model we incorporated recent advances in car-following models for self-driving cars [13].

The *Uncertainty model* allows for the simulation of exogenous events (specifically delays) that affect run-time operations and allows for the testing of more sophisticated scheduling strategies. Delays are injected randomly into the simulator during a Tick operation based on a set of parameters specifying the frequency with which delays occur.



Fig. 2 SFO airport model.

Currently, delays can be injected at gate nodes and runway nodes. Injecting a delay during a simulation means that with some probability an aircraft will not advance from its current node (for example, a departure from a gate node) to a target node specified on its itinerary.

The output of the simulator is a sequence of states that provides a complete run of the scenario. This output can be either stored and played back, or streamed in 'real time' by the Visualizer. A window in the display shows the set of active aircraft and their progress along the surface. Figure 3 shows a snapshot of the visualizer replaying a scenario. We plan to add a capability to allow uncertainties to be manually injected during Visualization to allow more fine-tuned testing of scheduling algorithms.

C. Scheduler

As noted above, the Scheduler is called by the Simulator periodically (a period specified by an input parameter; to model typical real operations, we chose 15 minutes in world time). The inputs are the current state of the world and the scenario. It outputs a set of paths and departure times (the time to leave the gate nodes for departures), called *itineraries*, one for each active aircraft on the surface.

The scheduler algorithm is motivated by the idea that a model of uncertainty can be employed during scheduling to predict delays and so generate schedules that are more robust to unexpected events in the world. Specifically, the scheduling problem becomes a multi-agent path finding problem under uncertainty [9] using a prioritized planning framework. The travel time of an aircraft along a link is a random variable which is decided by the previous movement of this aircraft, the travel time of the aircraft that is in front of this aircraft, and the uncertainties of the environment. This is formulated as a probability propagation Markov Chain model which propagates travel time distribution from the aircraft with the highest priority to the aircraft with the lowest priority. We propose two prioritized algorithms: First-Come-First-Serve algorithm (FCFS) where agents with earlier release times have higher priorities, and First-Leave-First-Serve algorithm (FLFS) where agents with earlier desired finish times have higher priorities. Once we have the priorities, we use a modified Cooperative A* search [14] to plan paths. In particular, the state in the path-planning search space is specified by a node and a time distribution, and the solution is evaluated by wait times at the gates and the expected travel times.

III. Experimental Results

The scheduler is implemented in C++. Our experiments are conducted on a 2.80 GHz Intel core i7-7700 laptop with 8 GB RAM. For this paper we test on random scenarios which include only departure aircraft; future work will include

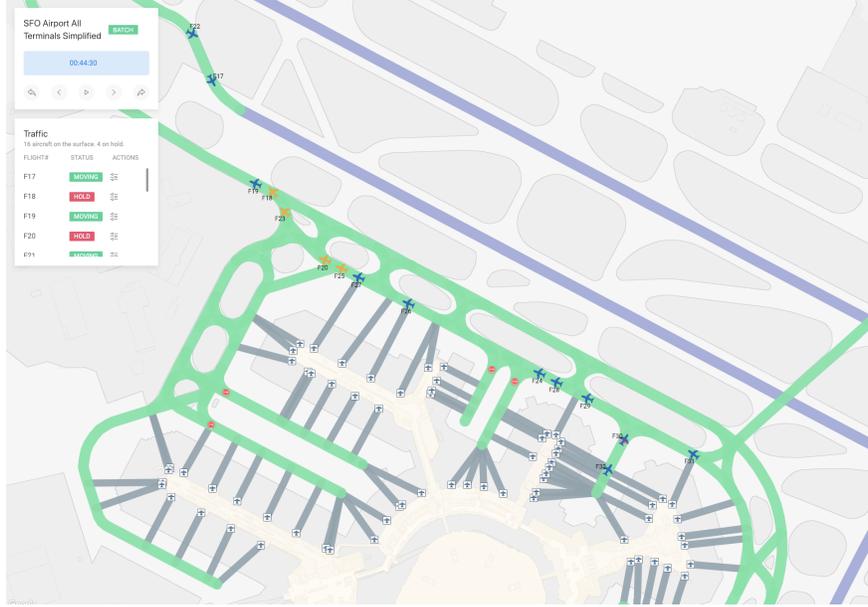


Fig. 3 The Visualizer playing back a scenario.

scenarios with arrivals.

A. Performance of the Scheduler

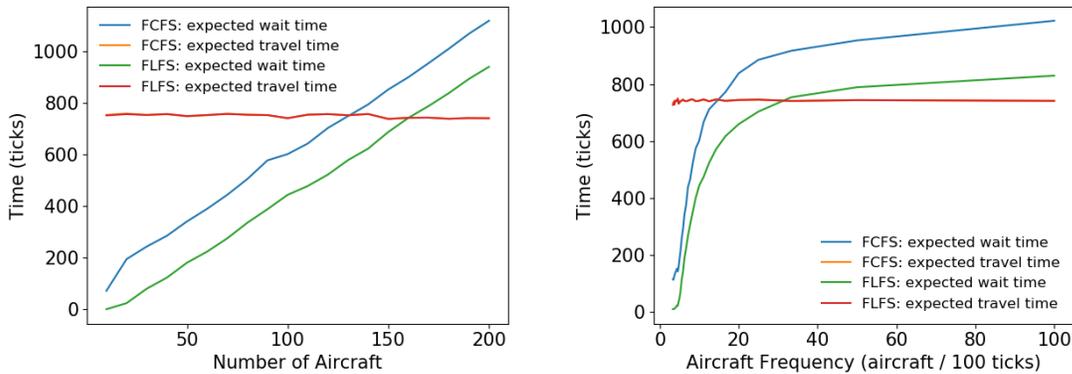


Fig. 4 Solution quality of the Scheduler. The left figure uses a frequency of 1 aircraft every 10 ticks, while the right figure uses 100 aircraft. (Most parts of the yellow line is hidden by the red line.)

Effectiveness: Figure 4 shows the solution qualities of FCFS and FLFS algorithms. We use a delay probability of 0.5 for delays at both gate nodes and runway nodes. “wait time” represents the expected wait time at gates before pushback, and “travel time” represents the expected travel time from the gate to the runway after pushback. The travel times of both algorithms are similar and stable, which indicates that both algorithms control the taxiway congestion very well. The wait times of both algorithms increase as the number of aircraft increases, and FLFS has much smaller wait times than FCFS. This is because FLFS first predicts desired finish times of aircraft and assign priorities accordingly. Therefore, FLFS is more effective than FCFS.

Efficiency: Figure 5 shows the efficiency of the Scheduler. For each number of aircraft and each aircraft frequency, we test 50 random Scenarios and reported the average results. Both algorithms are scalable and take less than 30 seconds to solve the problem. FLFS still runs faster than FCFS because FCFS usually generates longer paths than FLFS and

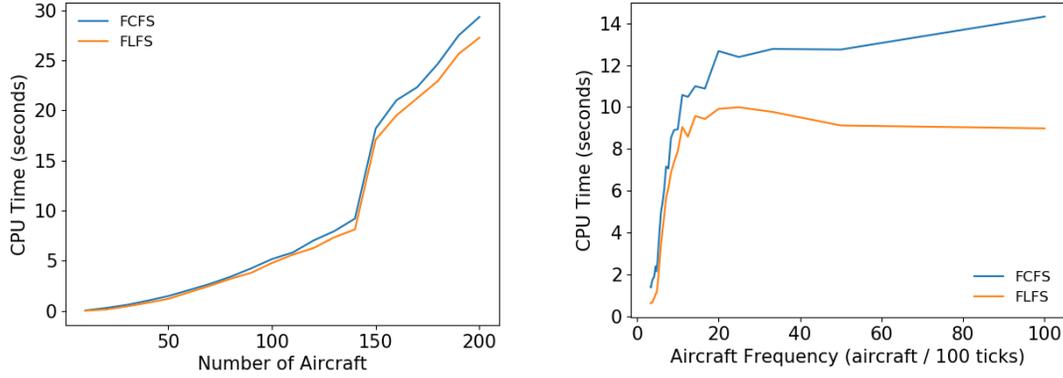


Fig. 5 The runtime of the Scheduler. The left figure uses a frequency of 1 aircraft every 10 ticks, while the right figure uses 100 aircraft.

longer paths needs much more search efforts than shorter paths. Therefore, FLFS is more efficient than FCFS.

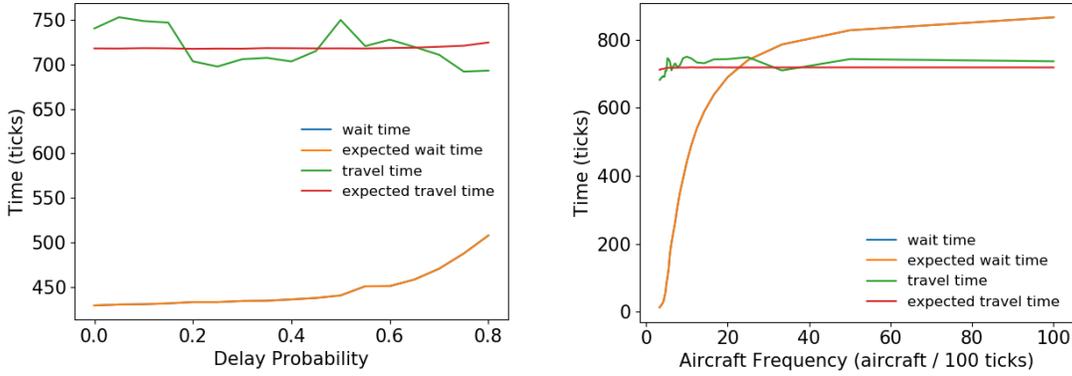


Fig. 6 Comparison between the results from the Scheduler with the results from the Simulator. The left figure uses a frequency of 1 aircraft every 10 ticks, while the right figure uses a delay probability of 0.5. (Most parts of the blue lines in both figures are hidden by the yellow lines.)

Accuracy: As discussed above, the Scheduler uses a probability propagation Markov Chain model to approximate travel time of each aircraft under uncertainty. Figure 6 compares the time ticks predicted by Scheduler with the real time ticks and simulated by the Simulator under various aircraft frequencies and uncertainties. The Scheduler uses the FLFS algorithm. For each Scenario, we run the simulator 50 times and report the average results. It turns out that the prediction for the wait time is always accurate, as we use a geometric distribution to approximate the gate delays. The prediction for the travel time is not as accurate as the prediction for the wait time, because the Scheduler uses an abstract velocity model in the Markov Chain model that ignores the limit of the acceleration of the aircraft and the behaviours of the aircraft along each link. But overall, both predictions are close to the real time ticks.

B. Performance of the system

Finally, we compare our Scheduler with a baseline scheduler which ignores uncertainties. In particular, the baseline scheduler plans a shortest path for each aircraft from its current location to its goal location while ignoring the presence of *conflicts* (two aircraft violating separation constraints). The conflicts on each link can be resolved by applying the car-following model [13], while the conflicts at each intersection can be resolved by the Controller, simulating tower-pilot communication. We use a time window (i.e., n in Figure 1) of 900 ticks, a delay probability of 0.5, and a scenario of 500 aircraft with a frequency of 1 aircraft every 10 ticks. As shown in Table 1, our Scheduler produces

solutions with smaller cost than the baseline scheduler, although it runs slower than the baseline scheduler. The solution of our Scheduler is also more robust since it leads to fewer conflicts (that need to be resolved by the Controller) than the baseline scheduler, thus potentially reducing human workload. Although preliminary, we feel these results will be part of a convincing case for developing probabilistic schedulers for real time deployment at busy airports.

Table 1 Performance of the system.

	Solution Cost		Scheduler	Controller
	Wait+Travel Time	Makespan	Runtime per call	#Conflicts resolved for each aircraft
BASE	9,601 ticks	28,049 ticks	15 ms	6.04
FLFS	8,916 ticks	27,079 ticks	3 s	4.05

IV. Summary

This paper has described a fast-time simulation system for aircraft surface movement expands on previous research by that allows for the quantification of the effects of delays on the robustness of solutions generated by surface movement planners and schedulers. The overall goal of this effort is to investigate and quantify the effects of using recent advances in probabilistic approaches to planning and scheduling to anticipate delays in order to produce more robust plans.

Acknowledgment

The first author was funded by a grant under The NASA Academic Mission Services (NAMS) contract, led by Universities Space Research Association (USRA). Thanks to Leighton Quon for helpful feedback on a earlier draft.

References

- [1] Malik, W., Gupta, G., and Jung, Y., "Spot release planner: Efficient solution for detailed airport surface traffic optimization," *AIAA Aviation Technology, Integration, and Operations (ATIO) Conference and AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, AIAA, Indianapolis, IN, 2012, p. 5498.
- [2] Rathinam, S., Wood, Z., Sridhar, B., and Jung, Y., "A generalized dynamic programming approach for a departure scheduling problem," *AIAA Guidance, Navigation, and Control Conference*, 2009, pp. 10–13.
- [3] Ravizza, S., Atkin, J., and Burke, E., "A more realistic approach for airport ground movement optimisation with stand holding," *Journal of Scheduling*, Vol. 17, No. 5, 2014, pp. 507–520. doi:10.1007/s10951-013-0323-3, URL <http://dx.doi.org/10.1007/s10951-013-0323-3>.
- [4] Roling, P. C., and Visser, H. G., "Optimal Airport Surface Traffic Planning Using Mixed-integer Linear Programming," *Int. J. Aero. Eng.*, Vol. 2008, No. 1, 2008, pp. 1:1–1:11. doi:10.1155/2008/732828, URL <http://dx.doi.org/10.1155/2008/732828>.
- [5] Montoya, J. V., Windhorst, R. D., Stroiney, S., Griffin, K., Saraf, A., Zhu, Z., and Gridnev, S., "Analysis of Airport Surface Schedulers Using Fast-time Simulation," *AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, AIAA, Los Angeles, CA, 2013, p. 4275.
- [6] Aponso1, B., Coppenbarger, R., Jung, Y., Quon, L., Lohr, G., OConnor, N., and Engelland, S., "Identifying Key Issues and Potential Solutions for Integrated Arrival, Departure, Surface Operations by Surveying Stakeholder Preferences," *15th AIAA Aviation Technology, Integration, and Operations Conference*, 2015.
- [7] Ma, H., and Koenig, S., "Optimal Target Assignment and Path Finding for Teams of Agents," *International Conference on Autonomous Agents & Multiagent Systems (AAMAS)*, ACM, Singapore, 2016, pp. 1144–1152.
- [8] Clare, G., and Richards, A. G., "Optimization of taxiway routing and runway scheduling," *IEEE Transactions on Intelligent Transportation Systems*, Vol. 12, No. 4, 2011, pp. 1000–1013.
- [9] Wagner, G., and Choset, H., "Path Planning for Multiple Agents under Uncertainty," *International Conference on Automated Planning and Scheduling (ICAPS)*, AAAI Press, Pittsburgh, PA, 2017, pp. 577–585.

- [10] Ging, A., Engelland, S., Capps, A., Eshow, M., Jung, Y., Sharma, S., Talebi, E., Downs, M., Freedman, C., Ngo, T., Sielski, H., Wang, E., Burke, J., Gorman, S., Phipps, B., and Ruszkowski, L., "Airspace Technology Demonstration 2 (ATD-2) Technology Description Document (TDD)," *NASA/TM-2018-219767*, 2018.
- [11] "Operations at Towered Airports," *AOPA Air Safety Foundation*, 2008. <https://www.aopa.org/-/media/files/aopa/home/pilot-resources/asi/safety-advisors/sa07.pdf>.
- [12] Morris, R., Chang, M. L., Archer, R., II, E. V. C., Thompson, S., Franke, J. L., Garrett, R. C., Malik, W., McGuire, K., and Hemann, G., "Self-Driving Aircraft Towing Vehicles: A Preliminary Report," *AAAI Workshop on Artificial Intelligence for Transportation: Advice, Interactivity, and Actor Modeling*, AAAI Press, Austin, TX, 2015.
- [13] Saifuzzaman, M., and Zheng, Z., "Incorporating human-factors in car-following models: A review of recent developments and research needs," *Transportation Research Part C: Emerging Technologies*, Vol. 48, 2014, pp. 379–403.
- [14] Silver, D., "Cooperative Pathfinding," *Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE)*, AAAI Press, Marina del Rey, CA, 2005, pp. 117–122.