

Dynamic Routing of Aircraft in the Presence of Adverse Weather Using a POMDP Framework

Edward Balaban*

NASA Ames Research Center, Moffett Field, CA, 94035, USA

Indranil Roychoudhury†

Stinger Ghaffarian Technologies Inc., NASA Ames Research Center, Moffett Field, CA, 94035, USA

Lilly Spirkovska‡

NASA Ames Research Center, Moffett Field, CA, 94035, USA

Shankar Sankararaman§, Chetan Kulkarni¶ and Tomer Arnon¶

Stinger Ghaffarian Technologies Inc., NASA Ames Research Center, Moffett Field, CA, 94035, USA

Each year weather-related airline delays result in hundreds of millions of dollars in additional fuel burn, maintenance, and lost revenue, not to mention passenger inconvenience. The current approaches for aircraft route planning in the presence of adverse weather still mainly rely on deterministic methods. In contrast, this work aims to deal with the problem using a Partially Observable Markov Decision Processes (POMDPs) framework, which allows for reasoning over uncertainty (including uncertainty in weather evolution over time) and results in solutions that are more robust to disruptions. The POMDP-based decision support system is demonstrated on several scenarios involving convective weather cells and is benchmarked against a deterministic planning system with functionality similar to those currently in use or under development.

I. Introduction

Weather is a frequent cause of airline flight delays, with approximately 12% of airline flight delays attributable to non-extreme weather.¹ This number does not include delays or cancellations caused by extreme weather such as hurricanes, blizzards, or tornadoes that the airlines deem unflyable. Rather, it includes only delays attributed to airport or national airspace system (NAS) operations and represents delays that could be alleviated by improvements in the system. The Federal Aviation Administration (FAA) has significantly improved NAS operations, as evidenced by the contribution of all sources of NAS delays falling from causing over 36% of total delay in 2003 to less than 23% in 2015.¹ Still, that 23% results in a large loss of airline revenue. Delays and the consequent lack of predictability in NAS operations also has non-monetary costs, such as schedule padding (i.e., adding extra time to a schedule to protect against delays), decreasing the maximum number of daily flights assigned to an aircraft, as well as an inconvenience cost to passengers.

A major challenge in NAS operations is dealing with the various sources of uncertainty inherent in the system. It is difficult for airline dispatchers and pilots to optimize flight routes given the uncertainties in weather forecasts and traffic congestion. Similarly, it is difficult for air traffic control (ATC) traffic flow managers to optimize delay management initiatives and for flight controllers to optimize the sequence

*Intelligent Systems Division, NASA Ames Research Center, AIAA Senior Member

†SGT Inc., Intelligent Systems Division, NASA Ames Research Center, AIAA Member

‡Intelligent Systems Division, NASA Ames Research Center

§SGT Inc., Intelligent Systems Division, NASA Ames Research Center, AIAA Senior Member

¶SGT Inc., Intelligent Systems Division, NASA Ames Research Center

of aircraft given the uncertainties in weather forecasts, aircraft pushback readiness, and NAS equipment status.²

There are various methods under development to deal with these types of uncertainty, including improved weather forecasting, airline and ATC shared situational awareness, and collaborative decision making.³ Our work focuses on improving flight routing while dealing with uncertainties in convective weather (i.e., thunderstorm) forecasts.

Convective weather forecasts are improving. Weather models and meteorologists are fairly accurate in short range (half day to two-and-a-half day) forecasting of large-scale trends.^{4,5} Still, there remains uncertainty in the timing of a weather system and its precise local extent. If the severity and extent of a convective weather system is forecast to interfere with major traffic flows, the Air Traffic Control System Command Center (ATCSCC) will reroute flights using alternative routes detailed in the National Severe Weather Playbook.⁶ These strategic reroutes give the anticipated adverse weather areas a wide berth, sacrificing efficiency for perceived increase in safety. If the weather system does not develop as forecast, decision support tools aim to improve efficiency by dynamically shortening routes by getting closer to the current location of the weather system.⁷ If a convective weather system is forecast to be less severe and more localized, playbook routes are not activated. Instead, pilots are responsible for tactical weather avoidance using information from ground-based NEXRAD images⁸ and on board weather radar. Because processing and transmission of NEXRAD images is delayed by five to fifteen minutes, pilots must rely on onboard radar for visualizing current conditions. As we discuss later, even onboard radar is not without issues in providing an accurate depiction of the actual conditions. Indeed, actual weather conditions are only partially observable, impairing the pilot's efforts to optimize the route.

In this work, we explore how to provide pilots and ground personnel with decision support aids that can help optimize routing in the presence of convective weather by fusing information from both ground and onboard radars and taking the uncertainty of their indications into consideration. To accomplish this we propose using Partially Observable Markov Decision Processes (POMDPs)⁹ - a general mathematical framework for reasoning under uncertainty. For each potential system state (e.g., configuration of the airspace), a POMDP describes the actions that can be taken, the likely follow-on states, and the costs required to achieve these follow-on states. The current state is not assumed to be known exactly; it can, however, be estimated using observations (e.g., radar data). Typically, the objective when solving a POMDP problem is to determine how to achieve a goal state while minimizing costs (or, conversely, maximizing rewards). For a complex real-world system with a large (possibly infinite) number of states, numerous action options, and a variety of (not always consistent) observations, determining the best strategy becomes a challenging computational task.

The high-level goals of our work are, therefore, threefold: (1) explore the applicability and the benefits of probabilistic reasoning under uncertainty in aeronautical decision making; (2) create a prototype decision support system for dynamic routing of aircraft in presence of adverse weather; and (3) in the process, develop general-purpose decision-making algorithms capable of dealing with other complex real-life systems.

The rest of the paper is organized as follows. Section II contains the problem formulation and the set of assumptions used in this work. Section III describes the POMDP-based decision support system, with Section IV devoted to the model the system uses for this particular problem. Section V describes the experiments conducted to validate the performance of the decision support system against a deterministic flight planning system designed to approximate the functionality of a modern weather-aware planner, such as the Dynamic Weather Routes (DWR) system.⁷ Section VI concludes the paper and outlines future work.

II. Problem Formulation

The objective of this problem is to route an aircraft to its destination airport in presence of convective weather in its direct path, while observing safety rules and optimizing for flight time and fuel consumption. Another (simpler) variant of this problem is to guide an aircraft through a region of convective weather but not land, thus not requiring routing to a specific point. This latter formulation is not covered in this paper.

The following problem assumptions are currently made:

- the airspeed and the altitude of the aircraft are constant (until the landing phase); therefore, optimizing for route length currently also results in optimization for flight time and fuel consumption;

- weather is dynamic, governed by stochastic processes;
- aircraft position is known exactly;
- aircraft transitions between positions are deterministic;
- onboard radar performance depends on the distance from the target cell;
- information provided by ground-based weather radar is precise; and
- there is a time delay (on the order of minutes) associated with obtaining the most current ground-based weather radar information.

III. Decision Support System

The prototype decision support system developed in the course of this work, AvDMU (for Aviation Decision Making under Uncertainty), consists of two main parts: the POMDP solver and the model. The solver is general and can be used for reasoning under uncertainty in other problems and other problem domains. The model describes a particular problem for the solver (in this case routing in presence of adverse weather) and is covered in detail in Section IV.

The POMDP solver, Large Problem Decision Making (LPDM), is based on Monte-Carlo Tree Search (MCTS) principles.¹⁰ Similarly to DESPOT¹¹ (a recent MCTS-style POMDP solver), LPDM uses sparsely-sampled scenarios and generates online policies that interleave planning and execution (Figure 1). The process begins at some initial belief state b_0 . A tree branching on actions and observations is constructed using sparsely-sampled scenarios that allows to estimate the values of actions at the root node. The action with the highest expected value (a_i^*) is chosen and executed. The observation vector (e.g., sensor readings) obtained during the execution is used to update the belief state via a particle-filtering algorithm and the process begins anew.

LPDM incorporates double-progressive widening techniques¹² to handle large action spaces. It also features the *anytime* property, where CPU time limits on how long LPDM can take before producing a recommendation can be set (the limits can also be defined by the number of simulations executed or through uncertainty metrics). The algorithm can even be interrupted before the allowed time expires and still produce a valid, safe solution, just one that may not have been fully optimized yet. Both LPDM and AvDMU are implemented in Julia programming language (v0.5.0). LPDM is compatible with POMDPs.jl, a generalized POMDP modeling and solving framework.¹³

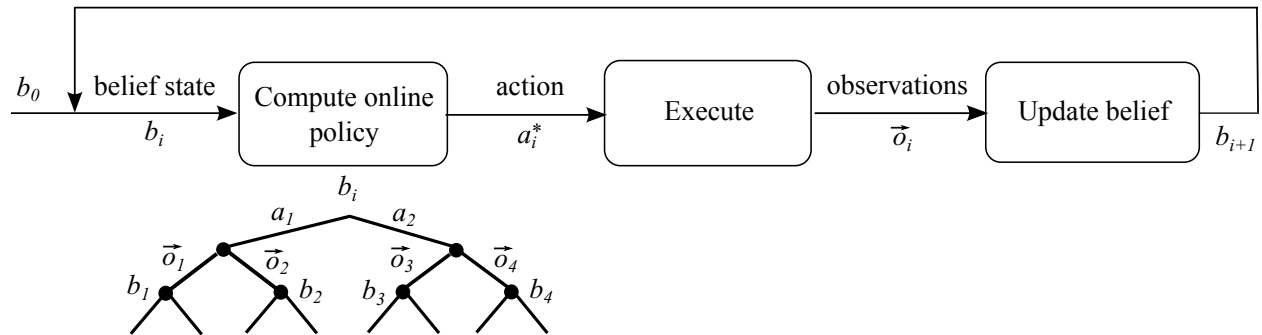


Figure 1: LPDM operation loop

IV. Modeling

A typical POMDP model consists of the following elements:

- states: S ;
- actions: A ;

- observations: O ;
- belief states: $b \in B$ (B is the set of all possible belief states);
- transition model: $T(s, a, s') \rightarrow [0, 1]$, where s is the current state, s' is the follow-on state, and a is the action taken;
- observation model: $O(s, a, s', o) \rightarrow [0, 1]$, where o is the observation received when taking action a in state s ;
- rewards/costs model: $R(s, a, s') \rightarrow \mathbb{R}$; and
- discount factor: $\gamma \in [0, 1]$.

The implementation of these elements for the AvDMU system is described in the subsections below. The objective when solving a POMDP is to find a policy $\pi^* : B \rightarrow A$ that maximizes reward (or minimizes cost). Formally this means mapping every possible belief state to an action. In practice this often means deriving this mapping for one of the following: a set of “binned” (discretized) belief states, a belief subspace of interest, or the current belief state. LPDM implements the last option.

A. State Space

A state is a complete configuration of weather and aircraft positions. In AvDMU a state is currently modeled using a two-dimensional grid, with grid squares corresponding to a region of airspace at a particular altitude (see Figure 2 for an example). States where the aircraft has reached the airport are denoted as terminal states. Gray-shaded squares indicate adverse weather and the numbers inside them indicate relative moisture content $\Theta_i \in [0.0, 10.0]$, where i is the square index. Weather in a square is considered adverse if $\Theta_i > \Theta_a$, where Θ_a is the adverse weather threshold ($\Theta_a = 1.0$ in the current work). $\Theta_i \leq \Theta_a$ for unfilled squares.




Figure 2: Example AvDMU state

B. Actions

The action space A currently includes headings in 90° increments: N, E, S, W , i.e., North, East, South, and West, respectively. This will be expanded to finer-grain headings (e.g., with 1° precision) and altitude changes in future work.

C. Observations

An observation estimates the relative moisture content (Θ) in a set of grid squares. The content is represented in a manner similar to the state model (with a number $\in [0.0, 10.0]$). A vector of four numbers is used in the current work (to be expanded in the future). These numbers model readings within the radar cone extending

from the aircraft (Figure 3). The areas shaded with darker gray on Figure 3 show the radar coverage area relative to the aircraft's position and orientation. Compared to the corresponding area in Figure 2 (the second convective area, in the lower half of the grid), the radar reading numbers for the same grid squares differ. They illustrate the possibility of inaccurate and erroneous readings. While in most cases inaccuracies will not impact the decision-making process significantly, in some cases (e.g. the square with the value of 1.2) they can take the value over the adverse weather threshold (Θ_a) and the square will be treated accordingly.




Figure 3: Example of an observation

D. Belief States

Belief states are represented via sets of particles, each containing a state instantiation and a weight value (used in computing the probability associated with the particle).

E. Transition Model

The transition model consists of two parts: aircraft transitions and weather transitions. Aircraft transitions are currently modeled in a deterministic fashion, where the aircraft in a state s always transitions to the same follow-on state s' , given a particular action a . The weather, on the other hand, is modeled stochastically. Each grid square has a real number between 0 and 10 associated with it, indicating the relative amount of moisture. At each time step, a randomly-distributed amount of moisture (or dry air) from each square is transferred to the adjacent squares (depending on the wind velocity vector); then the moisture indices for the squares affected are updated. The moisture content can change through maturation/dissipation processes as well, also modeled stochastically.

F. Observation Model

The observation model for the onboard weather radar converts the real numbers indicating the relative amount of moisture in the radar coverage area into a set of $16(2^4)$ discretized observations (indicating which of the four observation regions contain adverse weather) and returns a probability distribution of them being true, conditioned on a particular system state.

Confidence in onboard radar readings is currently approximated with a logistic function:

$$c = \frac{1}{1 + e^{k(d-d_0)}}, \quad (1)$$

where d_0 is the 50% confidence distance, currently set to 80 nautical miles (as the typical maximum effective onboard radar range¹⁴) and k is the steepness coefficient set to $k = 0.07$ (Figure 4). The latter parameter was picked simply to represent a gradual drop-off in confidence with distance. A more accurate model will be developed in future work.

The ground radar observation model is deterministic ($c \equiv 1$), thus providing “ground truth” information for every square of the grid.

G. Cost/Reward Model

The cost/reward model currently has the following components:




Figure 4: Radar indications confidence values vs. distance

- Transition cost: $c_t = -10$. The cost associated with flying (transitioning) to the next square. Since the airspeed and fuel burn are considered to be constant in the current problem formulation, this cost is the same for all transitions (wind velocity also currently does not play a role). This will be changed in future work, where we intend to incorporate airspeed, wind, and other factors into the state transition model.
- Safety violations cost: currently this only includes the cost of entering a convective weather cell i , $c_{cc,i} = -100 \times \Theta_i$, with Θ_i being the aforementioned cell moisture content. The intention of this formulation is to assign a high cost to flying the aircraft into a weather cell, while also make the cost proportional to the “severity” of the cell so that if the aircraft does find itself trapped by weather, it will find the least hazardous trajectory out of it. In future work other safety metrics will be added as well, such as aircraft-to-aircraft conflicts.
- Terminal state reward: $c_{ts} = 1000$. Obtained when the aircraft reaches its destination airport (any state with the aircraft coordinates matching destination airport coordinates is a terminal state).

H. Discount Factor

A discount factor less than 1 is typically used to discourage overly-complex solutions. In our model formulation the discount factor is set to 1, as a positive reward is only given upon reaching a terminal state and negative rewards (costs) provide a sufficient deterrent against undesirable policies.

I. Value Bounds

The LPDM algorithm will converge to a solution faster if good estimates of upper and lower bounds on the values of belief states are supplied (a value of a belief state is the expected reward obtainable from that state if a particular policy is followed). Domain-specific heuristics are often used to approximate the upper and lower bounds. For this model the lower bound is computed by summing up the rewards along a shortest (Manhattan) path to the airport (including penalties for going through the regions of adverse weather). The upper bound is computed similarly, only now with no adverse weather.

V. Experiments

The purpose of the experiments is to both verify the correctness of the advisory system’s output and to quantitatively compare its performance with a deterministic system, which operates similarly to contemporary routing systems.

A. Deterministic planner

The Deterministic Route Planner (DRP) uses the A* algorithm¹⁵ for optimizing flight path from the current location. DRP receives the same radar information (both ground and onboard) as the POMDP solver. It also utilizes the same weather model, however treats its output in a deterministic manner. At each time step, DRP computes the shortest distance from the current location of the aircraft to the destination airport using the knowledge of how the weather cells will be evolving based on the weather forecasts. Although this implementation works for a small grid, such as the 15×10 grid used in these experiments, work is in progress to replace the A* algorithm with D*-Lite in future work to scale up to real-world scenarios with much larger grids.

B. Test scenarios

In the current series of experiments a 15×10 grid is used, with each square representing 25×25 statute miles. Three different initial weather configurations (Figure 5) are used to construct ten test scenarios (Table 1). Scenarios based on the same initial weather configuration (the first number in the scenario ID denotes the weather configuration) differ in wind direction, wind velocity, initial aircraft position, or airport location.




Figure 5: Initial weather configurations

Table 1: Scenarios

| Scenario ID | Aircraft Coord. | Wind Dir. (deg) | Wind Velocity (mph) | Airport Coord. |
|-------------|-----------------|-----------------|---------------------|----------------|
| 1.1 | (1,5) | 315 | 15 | (15,3) |
| 1.2 | (9,6) | 315 | 15 | (15,3) |
| 1.3 | (9,10) | 315 | 15 | (15,3) |
| 1.4 | (1,10) | 45 | 15 | (15,8) |
| 2.1 | (1,5) | 315 | 20 | (15,3) |
| 2.2 | (1,5) | 90 | 50 | (15,3) |
| 2.3 | (1,10) | 180 | 20 | (15,3) |
| 3.1 | (1,10) | 270 | 20 | (15,1) |
| 3.2 | (1,5) | 235 | 20 | (15,10) |
| 3.3 | (1,1) | 270 | 30 | (15,10) |

Δt , the decision-making cycle, is set to 3 min, which is how long it would take, approximately, for an airliner flying at 500 mph to traverse a 25 mile square. Onboard weather readings are taken every 3 minutes as well. The field of view of the onboard radar is 135° , centered on the direction of travel. AvDMU receives confidence values along with onboard radar observations, generated according to Equation 1. DRP receives only the onboard radar observations above a certain confidence level (currently 0.5).

Ground radar readings are treated as the ground truth by both solvers and are considered to be global (providing information about the entire 15×10 grid). In the current experiments, they are updated every 9 minutes, i.e. every three decision cycles. This rate was chosen to be both reasonably close to real-life ground radar update rates (once every 4.5-10 minutes, not accounting for transmission delays¹⁶) and test AvDMU's

ability to reason for a few cycles without having the ground truth available. In future experiments the size of the grid squares will be reduced, the grid size increased, the decision-making interval and the onboard weather radar update rate set to once a minute, and the ground weather radar update rate set to once every 5 minutes.

Two metrics are used for performance evaluation: (i) the distance traveled by aircraft d , i.e., the number of grid squares traversed (which, with constant airspeed, is proportional to the travel time) and (ii) the cumulative reward r , which includes penalties for incursions into regions of adverse weather. Each scenario is executed 30 times by both AvDMU and DRP, then the mean and standard deviation of d and r are computed. AvDMU is used with two limits on the reasoning time per step: 2 seconds and 4 seconds (and the results are tabulated separately). This is done in order to test the *anytime* property of LPDM and identify scenarios benefiting from increased computation time^a. A Lenovo Thinkpad T530 computer (Intel Core i7-3820QM CPU @ 2.70Ghz, 16 GB RAM), running Red Hat Enterprise Linux 6 in a VMware Workstation 12 Pro virtual machine, was used in the experiments.

Unique random seeds are used to generate streams of random values for each of the runs (the seeds are exactly the same for both AvDMU and DRP) in order to ensure variability among the runs. Mersenne Twister algorithm¹⁷ is used in the random number generator. Belief states were represented in LPDM as sets of 100 weighted particles and the decision horizon of 50 steps was used by the solver.

C. Results

Table 2 summarizes the results for all scenarios. AvDMU results are shown separately for 2-second and 4-second computing time limits (per step). Mean and standard deviation are shown for both the number of steps taken by each algorithm to reach the terminal state and the amount of reward obtained (a lower number of steps and a higher reward value are desired). Given the current problem formulation, the reward metric captures the penalties for entering adverse weather regions (in addition to per-step costs).

Table 2: Experiment results

| Scenario | AvDMU 2 sec | | AvDMU 4 sec | | DRP | |
|----------|---------------|---------------|---------------|---------------|---------------|---------------|
| | steps, d | reward, r | steps, d | reward, r | steps, d | reward, r |
| | $\mu(\sigma)$ | $\mu(\sigma)$ | $\mu(\sigma)$ | $\mu(\sigma)$ | $\mu(\sigma)$ | $\mu(\sigma)$ |
| 1.1 | 18.0(00.0) | 820.0(00.0) | 18.0(00.0) | 820.0(00.0) | 18.1(00.4) | 819.3(003.7) |
| 1.2 | 11.0(00.0) | 890.0(00.0) | 11.0(00.0) | 890.0(00.0) | 11.2(00.6) | 888.0(006.1) |
| 1.3 | 13.0(00.0) | 870.0(00.0) | 13.0(00.0) | 870.0(00.0) | 14.1(01.9) | 858.7(019.4) |
| 1.4 | 30.9(01.4) | 691.3(13.6) | 30.9(01.4) | 691.3(13.6) | 22.6(01.3) | 767.3(031.7) |
| 2.1 | 18.1(00.6) | 819.3(06.4) | 18.1(00.6) | 819.3(06.4) | 18.2(00.6) | 814.7(022.2) |
| 2.2 | 22.4(01.0) | 776.0(09.7) | 21.7(01.7) | 782.7(17.2) | 22.4(01.4) | 726.3(062.2) |
| 2.3 | 23.3(00.9) | 763.3(19.9) | 23.2(00.6) | 764.7(18.9) | 27.9(01.6) | 671.3(066.6) |
| 3.1 | 23.0(00.0) | 770.0(00.0) | 23.0(00.0) | 770.0(00.0) | 23.5(01.0) | 764.7(010.4) |
| 3.2 | 19.0(00.0) | 810.0(00.0) | 19.0(00.0) | 810.0(00.0) | 47.9(13.3) | 467.3(179.8) |
| 3.3 | 23.0(00.0) | 770.0(00.0) | 23.0(00.0) | 770.0(00.0) | 35.7(01.9) | 592.7(073.3) |

AvDMU performs better than DRP (both by the number of steps and the reward value metrics) in all of the scenarios, except for Scenario 1.4 (discussed below). Scenarios 2.2 and 2.3, where more optimal solutions are obtained by giving LPDM more time to reason, illustrate the *anytime* property of LPDM. The aforementioned Scenario 1.4 (Figure 6), on the other hand, illustrates an interesting feature of AvDMU, where a more complex procedure is composed out of basic directional actions. AvDMU (black airplane symbol) decides, at some point, that the weather system blocking the aircraft path is likely to move southward and enters the aircraft into a holding pattern. In reality this strategy proves to be suboptimal (the weather system in the “real world” simulation takes longer to move) and a quicker route would have been the one around the system (taken by DRP - white airplane symbol), but the fact that a complex procedure was formed by the virtue of LPDM’s look-ahead reasoning is notable in its own right. It illustrates that with the

^aTo put the numbers in perspective, it would take a commercial airliner approximately 3 minutes to travel 25 statute miles, so even a 4-second decision time still represents only a negligible fraction of that.




Figure 6: Scenario 1.4

appropriate modeling of rewards, constraints, and objectives, AvDMU could automatically generate complex routing procedures used in air traffic management.

Scenario 3.2 (Figure 7), as well as 3.1 and 3.3, illustrate the value of reasoning over uncertainty and the value of generating a long-horizon policy. Here DRP (white airplane symbol) takes the more obvious eastern route and gets blocked by the moving storm front. It changes course several times as the weather in the “real world” simulation ends up evolving differently enough from the weather predicted by the deterministic model. In some of the runs, DRP directs the aircraft into a region of adverse weather (Step 11) - also because the weather model predictions between ground radar updates were treated in a deterministic manner and ended up being differing significantly enough from the “real world” to misroute the aircraft. The penalties for entering the region of adverse weather on some of the runs are reflected in a low mean reward value




Figure 7: Scenario 3.2

(and a very high variance) for this scenario. AvDMU, on the other hand, reasons that the likelihood of the eastern path getting blocked is sufficiently high to choose the southern route - which proves to be the correct decision.

VI. Conclusions and Future Work

In this work we take the first steps towards creating an aircraft routing advisory system (AvDMU) that reasons, in a systematic manner, over convective weather uncertainty and illustrate the potential advantages of this approach over the prevalent deterministic systems. We believe that these advantages include improved flight safety and route optimality. AvDMU features include realtime performance, *anytime* solution generation, long-horizon sequential reasoning, and handling of large state spaces. AvDMU is built on the foundation of LPDM - a general-purpose POMDP solver that uses the Monte Carlo Tree Search approach. In our validation experiments, AvDMU was benchmarked against a deterministic routing system, DRP, on a set of common scenarios. AvDMU performed better than DRP (according to the metrics measuring the number of steps to a terminal state and a route's cumulative reward) in all scenarios except for one, where, nevertheless, AvDMU generated a complex maneuver (a holding pattern) out of elementary directional actions - a development that illustrates the potential for automatically generating complex routing procedures

through the appropriate modeling of rewards, constraints, and objectives.

Our future work will include making improvements to the LPDM solver, such as better handling of large observation spaces, as well as taking the next steps in increasing the realism of the case study. For the latter, we plan to extend the state space to three dimensions, demonstrate the ability to simultaneously route multiple aircraft, and extend the action space to include finer-granularity headings and variable airspeeds. Other dynamic elements will be added as well, such as traffic and temporary flight restrictions. Further out, we intend to add additional sources of uncertainty (e.g., winds aloft).

Acknowledgments

The authors gratefully acknowledge funding from the Real-Time Safety Monitoring project of the Aviation Operations and Safety Program, Aeronautics Research Mission Directorate, NASA. They also express their appreciation to their colleagues Dr. Wendy Okolo and Dr. Matthew Daigle (both of Intelligent Systems Division, NASA Ames Research Center) and Mr. Xiaoge Zhang of Vanderbilt University for their assistance in performing the work described and preparing the manuscript.

References

- ¹Bureau of Transportation Statistics, “Understanding the Reporting of Causes of Flight Delays and Cancellations,” 2016, <https://www.rita.dot.gov/bts/help/aviation/html/understanding.html> (accessed 2017-05-01).
- ²Donohue, G. L., *Air transportation systems engineering*, Vol. 193, AIAA, 2001.
- ³Wambsganss, M. C., “Collaborative Decision Making in Air Traffic Management,” *New Concepts and Methods in Air Traffic Management*, Springer, 2001, pp. 1–15.
- ⁴FAA, *Federal Aviation Regulations/Aeronautical Information Manual (FAR/AIM series)*, ASA, 2016th ed., 2016.
- ⁵Mueller, C. K., Fidalgo, C. B., McCann, D. W., Meganhardt, D., Rehak, N., and Carty, T., “National Convective Weather Forecast Product,” *8th Conference on Aviation, Range, and Aerospace Meteorology*, 1999.
- ⁶FAA, “National Severe Weather Playbook,” 2017, <https://www.fly.faa.gov/Operations/playbook/current/current.pdf> (accessed 2017-04-27).
- ⁷McNally, D., Sheth, K., Gong, C., Love, J., Lee, C. H., Sahlman, S., and Cheng, J.-H., “Dynamic Weather Routes: A Weather Avoidance System for Near-Term Trajectory-Based Operations,” *28th International Congress of the Aeronautical Sciences*, 2012.
- ⁸Heiss, W., McGrew, D., and Sirmans, D., “NEXRAD: next generation weather radar (WSR-88D),” *Microwave Journal*, Vol. 33, No. 1, 1990, pp. 79–89.
- ⁹Kaelbling, L., Littman, M., and Cassandra, A., “Planning and Acting in Partially Observable Stochastic Domains,” *Artificial Intelligence*, Vol. 101, No. 1-2, 1998, pp. 99–134.
- ¹⁰Browne, C. and Powley, E., “A Survey of Monte Carlo Tree Search Methods,” *IEEE Transactions on Computational Intelligence and AI in Games*, Vol. 4, No. 1, 2012, pp. 1–49.
- ¹¹Somani, A., Ye, N., Hsu, D., and Lee, W., “DESPOT: Online POMDP Planning with Regularization,” *Advances In Neural Information Processing Systems*, 2013, pp. 1772–1780.
- ¹²Couëtoux, A., Hoock, J.-B., Sokolovska, N., Teytaud, O., and Bonnard, N., “Continuous upper confidence trees,” *International Conference on Learning and Intelligent Optimization*, Springer, 2011, pp. 433–445.
- ¹³Egorov, M., Sunberg, Z. N., Balaban, E., Wheeler, T. A., Gupta, J. K., and Kochenderfer, M. J., “POMDPs.jl: A Framework for Sequential Decision Making under Uncertainty,” *Journal of Machine Learning Research*, Vol. 18, No. 26, 2017.
- ¹⁴National Oceanic and Atmospheric Administration, “Airborne Weather Radar Limitations,” 2014, <https://www.weather.gov/media/publications/front/14dec-front.pdf> (accessed 2017-05-10).
- ¹⁵Russell, S. and Norvig, P., *Artificial Intelligence: A Modern Approach (3rd Edition)*, Prentice Hall, 2009.
- ¹⁶National Oceanic and Atmospheric Administration, “Volume Coverage Pattern Comparison Table,” 2008, <http://www.wdtd.noaa.gov/tools/RPS/VCPCmpTable.pdf> (accessed 2017-05-10).
- ¹⁷Matsumoto, M. and Nishimura, T., “Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator,” *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, Vol. 8, No. 1, 1998, pp. 3–30.