Extending the Range of Delivery Drones by Exploratory Learning of Energy Models

(Extended Abstract)

Ty Nguyen
School of Electrical and Computer Engineering
Ulsan National Institute of Science and Technology
Ulsan, South Korea 44919
tynguyen@unist.ac.kr

Tsz-Chiu Au
School of Electrical and Computer Engineering
Ulsan National Institute of Science and Technology
Ulsan, South Korea 44919
chiu@unist.ac.kr

ABSTRACT

Delivery drones have a fairly short range due to their limited battery life. We propose new exploration strategies to generate paths for a drone to reach its destination while learning about the energy consumption on each edge on its path so as to optimize its range in future missions. As the energy consumption mostly depends on the payload, the wind direction, and the wind speed, we developed an energy model to estimate the energy consumption based on these factors. We evaluated our exploration strategies for learning the energy model in order to identify the set of all reachable destinations. We found that adding a small amount of perturbation to encourage exploration can increase the learning rate.

Keywords

Exploration vs. Exploitation; Delivery Drones; Path Planning

1. INTRODUCTION

Some companies such as Amazon Prime Air and DHL Parcelcopter have started to use delivery drones to deliver packages to their customers. However, delivery drones have a very limited range due to their short battery life. The question of how to extend the range of drones is crucial to the deployment of drone-based delivery in the future. A drone is a lightweight machine and its power usage is greatly affected by the airflow, which in turn is affected by buildings, trees, etc. A drone flying in an urban area can take advantages of the structures in the environment by flying between them in order to fly downwind as much as possible while avoid flying the headwind. Therefore, we propose to extend the range of drones by finding paths to take advantages of environmental structures such as buildings and trees while taking wind direction and wind speed into account. The energy saving can be translated into an extended range which allows the drone to serve a larger area.

The estimation of the actual energy consumption is difficult since the interaction between the drone and the environment can be quite complicated. We therefore propose to learn a model of energy consumption by exploring the environment while sending drones to deliver packages. Then we can use this model to optimize routes for a drone to deliver packages to distant destinations. Ultimately, the model will be used to check the *maximum coverage* of the delivery service—finding the set of all possible destinations that are reachable from a distribution center given a finite amount of energy. In this paper, we propose two exploration strategies to speed up the learning process and compare them experimentally.

Previous research often considered the problem of optimizing static soaring trajectories to minimize energy loss and extend the UAV's operating coverage [1, 8]. More recent research has addressed the dynamic soaring problem by introducing computational optimization techniques [7, 9]. The increasing popularity of UAVs has attracted more research on the soaring problem. For example, Chakrabarty et al. [2, 3] introduced an energy mapping that indicates the lower bound of necessary energy to reach the goal from any starting point to any ending point in the map with known wind information. Exploiting this map can provide a path from a starting point to a goal with the optimal speed and the heading to fly over each cell in the path. An alternative is to use soaring mechanisms to utilize observations to determine optimal flight [4, 6]. A further alternative is to predefine a wind model and utilize on-the-fly observations to fit this model. However, there has been few research focusing on building a mapping from wind condition and payload to energy consumption and simultaneously using that information to generate paths from a base point to target points. Unlike these work, [5] presented a Gaussian process regression approach to estimate the wind map but they devised a reward function to automatically balance the tradeoff between exploring the wind field and exploiting the current map to gain energy. Moreover, their problem differs from ours in which it aims to drive gliders to predefined locations. In contrast, our work aims to find all locations at which a delivery drone can arrive (i.e., the maximum range) and return to the distribution center with different payload and wind conditions.

2. ENERGY-BOUNDED DELIVERY DRONE PROBLEMS

Suppose a company sets up a distribution center with one drone to serve a community of households. The drone, which has a maximum energy $\mathbb{E}$, must fly on some designated trajectories represented as a graph $G = (V, E)$ connecting the distribution center to the households. Let $v_0 \in V$ be the location of the distribution center and $D \subset V \setminus \{v_0\}$ be the locations of the households. From time to time, the distribution center receives requests from the households. A request is a pair (v_{dest}, l), where $v_{\mathsf{dest}} \in D$ is the destination and l is the payload. Upon receiving a request, the distribution center will first decide whether it is feasible to deliver a package using a drone, and then send out a drone if it is feasible.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

When a drone traverses an edge e in G, it consumes a certain amount of energy, which depends on 1) the payload l, 2) the wind speed s, and 3) the wind direction d. We use $\epsilon(e; l, s, d)$ to denote the energy consumption for traversing an edge e under a *configuration* (l, s, d). To simplify our analysis, we assume the wind speed and the wind direction do not change in a trip to the destination, but they can be different in other trips. Before sending a drone, the distribution center obtains the current wind speed and the current wind direction from an information source. Given a path ρ, let $\epsilon(\rho; l, s, d) = \sum_{e \in \rho} \epsilon(e; l, s, d)$ be the total energy a drone consumes when traversing ρ under (l, s, d). $\tau = \rho_1 \oplus \rho_2$ is a *trip* to a destination v_{dest} if ρ_1 is a path connecting v_0 to v_{dest} and ρ_2 is a path connecting v_{dest} to v_0. The total energy consumption of τ under (l, s, d) is $\epsilon(\tau; l, s, d) = \epsilon(\rho_1; l, s, d) + \epsilon(\rho_2; 0, s, d)$. We say a trip τ to v_{dest} is *successful* under (l, s, d) if $\epsilon(\tau; l, s, d) \leq \mathbb{E}$. A destination $v_{\mathsf{dest}} \in D$ is *reachable* under (l, s, d) (or, in short, (l, s, d)*-reachable*) if there *exists* a successful trip τ under (l, s, d). However, we do not know which nodes are reachable until we send a drone to explore the graph and measure the unknown energy consumption. Our goal is to find the set $D_{l,s,d} \subseteq D$ for every configuration (l, s, d) such that all $v \in D_{l,s,d}$ are (l, s, d)-reachable. We call $D_{l,s,d}$ the *reachable set* under (l, s, d), which represents the maximum coverage of the distribution center under (l, s, d).

3. TRIP GENERATION

We considered two algorithms for trip generation. The first algorithm is a randomized algorithm that randomly generates a trip to reach v_{dest}. It first generates a random path from v_0 to v_{dest} and then another random path from v_{dest} to v_0. The random path generation is biased towards v_{dest} or v_0 by a heuristic function that returns the straight-length distance between two nodes. This algorithm, however, has difficulty in generating trips for distant destinations. The second algorithm biases the trip generation along with the shortest paths between the distribution center and the destination using the energy model learnt until the previous trip. The algorithm also deliberately adds a small amount of perturbation to the path generation process by randomly modifying the shortest paths to encourage exploration. Both algorithms control the amount of exploration by an integer K—the algorithms return the K'th trips sorted by an ascending order of the number of unknowns on the trips.

4. EMPIRICAL EVALUATION

We developed a simulator that performs the following steps at the beginning of a simulation: 1) generate a 2-D map of size 1000m $\times$ 1000m and randomly choose 50 locations as nodes; 2) designate the center of the map as the node of the distribution center; 3) randomly connect the nodes by adding directed edges between adjacent locations based on a distance threshold; 4) set the energy consumption on all edges under all possible configurations (l, s, d); and 5) identify the set of reachable nodes given the maximum energy $\mathbb{E}$ of a drone. In a simulation, a series of delivery requests at random nodes were generated with a random configuration (l, s, d), and the simulator used the two algorithms to generate trips to handle the requests. After each request, the simulator computed the current reachable set based on the current energy model, and compared it with the truly reachable set. The main criteria of the comparison is recall—how many nodes in the truly reachable set are present in the current reachable set. Since the precision of both algorithms are quite high, our evaluation will mainly focus on recall.

We compared several combinations of the trip generation algorithms and the levels of perturbation. The combinations are: 1) RANDOM (the first algorithm) with a fixed $K = 1$; 2) RANDOM

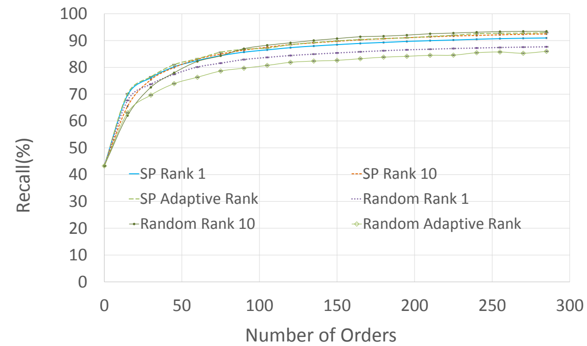

Figure 1: The percentage of reachable destinations of some combinations of the algorithms.

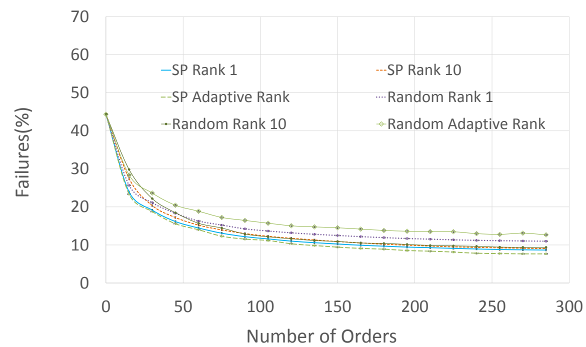

Figure 2: The percentage of delivery failures.

with $K = 10$; 3) RANDOM with K that is proportional to the length of the shortest path to the destination (adaptive rank); 4) SP (the second algorithm) with $K = 1$; 5) SP with $K = 10$; and 3) SP with K that is proportional to the length of the shortest path to the destination. We repeated the experiment 1000 times with 1000 random maps. The number of requests in each simulation is 300. The results are summarized in Figure 1 and Figure 2. Figure 1 shows the average recall at various numbers of requests, while Figure 2 shows the average numbers of failure in which the drone ran out of energy during a delivery. The 95% confidence intervals are shown as the tiny error bars in these figures. While all combinations have a similar performance, the second algorithm performed slightly better. Although the adaptive rank strategies were not superior to the fixed K strategies in terms of recalls, they have lower failure rates.

5. CONCLUSIONS AND FUTURE WORK

To extend the range of delivery drones, we propose to utilize environmental structures such as buildings and trees to avoid flying into a headwind and take advantages of flying downwind. In this paper, we presented the problem of how a delivery drone should learn about the energy consumption under different payloads, wind speeds and wind directions, in order to maximize the coverage of the service area of a distribution center. We evaluated two trip generation algorithms and proposed a perturbation strategy to speed up the learning process. In the future, we intend to devise new cooperative learning strategies to extend the ranges of multiple drones.

Acknowledgments

This work has taken place in the ART Lab at Ulsan National Institute of Science & Technology (UNIST). ART research is supported by NRF (2.170415.01 and 2.170511.01).

REFERENCES

[1] M. B. Boslough. Autonomous dynamic soaring platform for distributed mobile sensor arrays. *Sandia National Laboratories, Sandia National Laboratories, Tech. Rep. SAND2002-1896*, 2002.

[2] A. Chakrabarty and J. W. Langelaan. Energy maps for long-range path planning for small-and micro-uavs. In *Guidance, Navigation and Control Conference*, volume 2009, page 6113, 2009.

[3] A. Chakrabarty and J. W. Langelaan. Energy-Based Long-Range Path Planning for Soaring-Capable Unmanned Aerial Vehicles. *Journal of Guidance, Control, and Dynamics*, 34(4):1002–1015, 2011.

[4] J. W. Langelaan. Gust energy extraction for mini and micro uninhabited aerial vehicles. *Journal of guidance, control, and dynamics*, 32(2):464–473, 2009.

[5] N. R. Lawrance and S. Sukkarieh. Autonomous exploration of a wind field with a gliding aircraft. *Journal of Guidance, Control, and Dynamics*, 34(3):719–733, 2011.

[6] P. Lissaman. Wind energy extraction by birds and flight vehicles. *43rd AIAA Aerospace Sciences Meeting and Exhibit*, 241, 2005.

[7] G. Sachs. Minimum shear wind strength required for dynamic soaring of albatrosses. *International Journal of Avian Science*, 147(1):1–10, 2005.

[8] Y. J. Zhao. Optimal patterns of glider dynamic soaring. *Optimal control applications and methods*, 25(2):67–89, 2004.

[9] Y. J. Zhao and Y. C. Qi. Minimum fuel powered dynamic soaring of unmanned aerial vehicles utilizing wind gradients. *Optimal control applications and methods*, 25(5):211–233, 2004.