Promoting Cooperative Behavior Between Cost-Based Planners

Carrie Rebhuhn, Ryan Skeele, Jen Jen Chung, Geoffrey A. Hollinger and Kagan Tumer

I. Introduction

Robots performing delivery tasks use cost-based planning in order to move from a starting point to a destination point. This can ensure obstacle-free and distance-optimal trajectories for individual robots. However, when many robots share the same space, their shortest routes may overlap, causing traffic. This is particularly relevant when managing unmanned aerial vehicles (UAV) in the airspace, where low-level conflict-avoidance maneuvers can cause potentially hazardous situations when the density of traffic becomes too high.

We propose a multiagent approach to traffic management, such that agents manipulate the sector-level cost space of UAVs traveling through an obstacle-filled environment. Results from testing on a simulated urban environment show that given only local UAV information, the team of sector agents was able to learn appropriate costing strategies to reduce the number of conflicts experienced in the entire airspace. Comparison to a uniform costing strategy showed on average a 16.4% reduction of conflicts in the airspace after 100 learning epochs.

II. Related Work

There are two primary ways of handling congestion in an airspace: conflict-avoidance and air traffic management. The task of conflict-avoidance includes how to avoid the trajectory of another aircraft in the system. This is essential to the safe operation of UAVs in the airspace. However, there are limits to the efficacy of conflict-avoidance techniques, and situations may arise where there are too many UAVs traveling through a particular corridor to effectively deconflict. Because of this risk, effective *traffic control* must exist at a higher level to manage the densities of UAVs across the airspace.

Previous research has explored congestion as a centralized controller scheduling and routing problem. A comprehensive survey of centralized scheduling methods for automated vehicles is given in [1]. Using centralized methods, every robot in the system is told when it can travel and where it can travel. In order to compute a solution, these methods require full state information and are often slow and computationally complex. Dynamic routing methods [2], [3] manage the time-window of each robot through time-expanded graphs without needing full state information [4]. However these methods become computationally expensive when applied to a fast-changing system.

This work was supported by NASA grant NNX14AI10G.

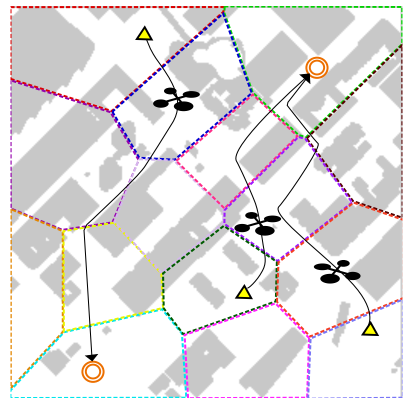

Fig. 1: The UAV traffic management problem. The airspace is divided into discrete sectors with a single UAV traffic management (UTM) agent controlling the cost of travel through each sector. UAVs traveling from start (triangle) to goal (circle) locations plan and execute paths to reduce their individual travel cost without explicitly coordinating with other UAVs in the airspace.

Congestion in the national airspace has been approached as a multiagent system. Using reinforcement learning agents to manage air traffic through geographical fixes, Agogino and Tumer [5] were able to reduce airspace congestion by over 90% when compared to current traffic control methods. In the current paper, we apply a similar network of routing agents to control the flow of UAV traffic. However, we consider the UAV Traffic Management (UTM) domain where platforms are not restricted to fly through particular fixes in the environment. Furthermore, our routing algorithm is based on the assumption that we have no direct control over the path planning aspect of the UAVs.

III. Problem Setup

We consider traffic management for UAVs performing delivery tasks in an urban airspace. The airspace is divided into 15 Voronoi partitions, as shown on Fig. 1, centered at hand-selected areas where UAV path intersections are expected. 20 locations, which we refer to as *fixes*, are randomly generated across this space, and serve as sources and destinations of UAV traffic. UAV traffic is generated randomly each timestep with a probability $p_{gen} = 5\%$. A *conflict* arises in the airspace when two UAVs are within a distance $d_{conflict} = 2$ pixels of one another.

There are 3 levels of control across the airspace:

Sector cost agents: These agents control the cost map at the sector level. These are neural networks that are coevolved using a cooperative coevolutionary algorithm, using a summation of conflicts in the whole airspace for the fitness function. They take as an input the tuple $\{n_N, n_S, n_E, n_W\}$,

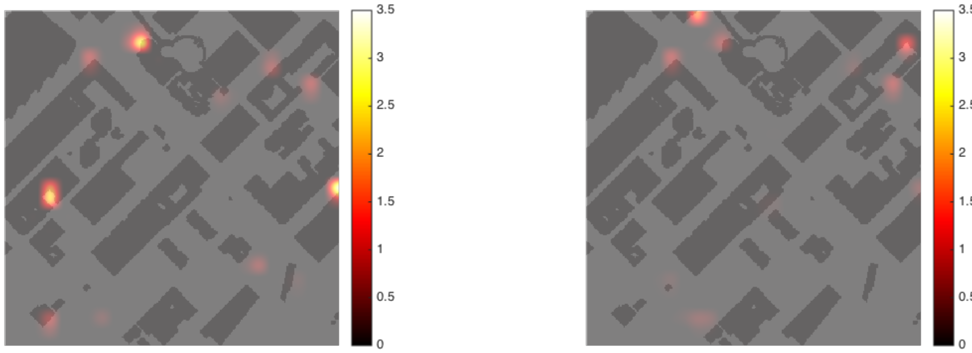

(a) (b)

Fig. 2: Change in congestion observed over one run of evolution. The overlaid heat map shows the number of conflicts experienced in an area for (a) the best evaluation trial for agents with random initialized sector travel costs and (b) the best evaluation trial for agents with evolved sector travel costs after 100 epochs. The overall number of conflict instances has reduced with some high congestion regions removed completely.

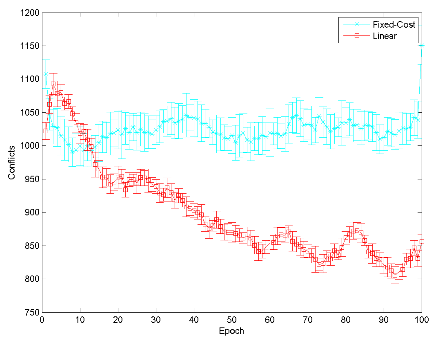

Fig. 3: Comparison of conflicft occurrence over epochs using fixed costs versus evolved costs.

which represents the number of UAVs traveling in each of the cardinal directions. They then output a cost for travel in each of the cardinal directions, $\{c_N, c_S, c_E, c_W\}$, which is used by the sector-level planner.

Sector-level planner: Each UAV has a sector-level planner that takes in sector cost information and plans across the sector map. This defines the sector visitation order for the UAV, and is used by the obstacle-level planner.

Obstacle-level planners: Each UAV also has an obstacle-level planner that plans a trajectory across the obstacle map. This is given as an 8-connected pixel image as shown in Fig. 1, where each pixel movement has unit cost. It is restricted to follow the sector visitation order dictated by the high-level planner, and plans obstacle-free trajectories across the airspace.

The path planners tested in this work use A^* to plan across their respective cost maps. However, this multiagent approach can accommodate any cost-based planner.

IV. Results

We initially performed a set of simulations using uniform fixed costs for travel between the sectors in order to compare to a distance-optimal path solution for each UAV. Each set of experiments contained 20 runs of 100 learning epochs.

Figure 2a shows the congestion observed in the first learning epoch. Sector agents initially assign travel costs randomly, so there are areas of high congestion that occur. Figure 2b shows the performance at the end of evolution. Comparing Fig. 2a and Fig. 2b we can see a definite reduction in the number of conflict instances. Figure 3 also shows this congestion reduction over the number of epochs in the simulation.

The number of conflicts in the baseline comparison algorithm remains high, with an overall average of 1023 conflicts. Comparing to the fixed-cost method, we see an average of 856 conflicts at the end of learning, which represents a 16.4% reduction in total conflicts after 100 epochs using evolution.

V. Conclusions and Future Work

The results presented in this paper show that a distributed UTM system can learn appropriate costs to apply to UAVs traveling in the airspace to incentivize implicit cooperation between the UAV planners. Using our method, we achieved a 16.4% reduction in conflicts compared to the baseline method with uniform sector costs. It is also worth noting here that agents do not require a model of the available airspace or knowledge of the obstacles in the sector. The number of conflicts experienced in the sector was sufficient information by which to evaluate the performance of the current policy.

The ability to manipulate high-level planners allows us to reduce occurrences of potentially dangerous congestion in the system. UAVs in the real world can handle encounters with other UAVs using low-level collision avoidance procedures, but by reducing the congestion in the airspace at a high level, we can permit safer travel by avoiding many of these conflicts before they occur.

As more companies become interested in sharing the urban airspace, we must accommodate heterogeneity in the airspace. Because of these differences in capabilities, algorithms developed for air traffic management must accommodate heterogeneity when routing. The heterogeneity between aircraft can affect the maximum safe congestion level in a sector. We are currently exploring ways to incorporate heterogeneous traffic flow into the evolution process outlined in this paper.

References

[1] L. Qiu, W.-J. Hsu, S.-Y. Huang, and H. Wang, "Scheduling and routing algorithms for AGVs: a survey," *International Journal of Production Research*, vol. 40, no. 3, pp. 745–760, 2002.

[2] F. Taghaboni-Dutta and J. M. A. Tanchoco, "Comparison of dynamic routeing techniques for automated guided vehicle system," *International Journal of Production Research*, vol. 33, no. 10, pp. 2653–2669, 1995.

[3] F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith, "Dynamic vehicle routing for robotic systems," *Proceedings of the IEEE*, vol. 99, no. 9, pp. 1482–1504, 2011.

[4] E. Gawrilow, E. Köhler, R. H. Möhring, and B. Stenzel, "Dynamic routing of automated guided vehicles in real-time," in *Mathematics - Key Technology for the Future*, H.-J. Krebs and W. Jäger, Eds. Springer, 2008, ch. 5, pp. 165–177.

[5] A. K. Agogino and K. Tumer, "A multiagent approach to managing air traffic flow," *Autonomous Agents and Multi-Agent Systems*, vol. 24, no. 1, pp. 1–25, 2012.