Velocity Obstacles and Emergent Rules-of-the-Air for Autonomous Drone Traffic Management

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

14-10-2021 / 27-10-2021

CITATION

Bonnell, William; Wilson, R. Eddie (2021): Velocity Obstacles and Emergent Rules-of-the-Air for Autonomous Drone Traffic Management. TechRxiv. Preprint. <https://doi.org/10.36227/techrxiv.16810789.v1>

DOI

[10.36227/techrxiv.16810789.v1](https://doi.org/10.36227/techrxiv.16810789.v1)

Velocity Obstacles and Emergent Rules-of-the-Air for Autonomous Drone Traffic Management

William Bonnell and R. E Wilson
University of Bristol

October 14, 2021

Abstract

As the use of unmanned aerial vehicles (UAVs) becomes ever more widespread there is a growing need to develop traffic management and flight rules, in particular for autonomous UAVs or where the predicted traffic densities far exceed those of traditional manned aviation. Inspired by the current rules of the air and multi-agent systems (e.g., pedestrians and swarm robotics) we outline a set of flight rules for autonomous UAVs that consist of waypoint following and conflict avoidance schemes. These flight rules are then explored in small, pairwise simulations and thus refined to allow a UAV to choose from three potential avoidance behaviors based on it and its neighbors velocities and positions. Finally we compare the original and modified flight rules in larger scale simulations modelling two streams of UAV traffic crossing at a point. We show that the modified rules significantly reduce the mean transit time by reducing the impact of UAVs avoiding other UAVs from the same stream.

1 Introduction

Unmanned Aerial Vehicles (UAVs) are poised to become an integral part of the fourth industrial revolution with the global UAV service market predicted to be worth \$63.6 billion by 2025 [1]. Amongst a variety of civilian applications [2], goods delivery has attracted a great deal of media attention, for example a recent article [3] about drone delivery service Manna operating in Ireland. In future, these services will require unmanned traffic management (UTM) systems, and there are a range of proposals [4, 5, 6] for such systems based on UTM service providers that will schedule and share flight plans so as to avoid *conflicts* (i.e., maintain safe separation) between UAVs. Future demand modelling is highly uncertain, but some projections suggest very large numbers of autonomous craft in the air at once (e.g., 87,000 delivery drones on average operating above Paris by 2035 [7] or 32,887 packages a day deliverable by drone for the city of Sendai [8]). In these circumstances, a centralised UTM approach would seem infeasible.

Therefore, we propose a distributed approach to UTM, where UAVs are dispatched without centralised trajectory planning, but rather use autonomous sense-and-avoid capabilities to avoid conflicts [9, 10, 11]. The problem is thus similar to a range of other multi-agent applications, for example Reynolds flocks [12] and pedestrian flow [13] where many agents must navigate through an environment and each other. A recurring problem in swarm robotics, there are a variety of methods to maintain separation, including potential fields [14, 15, 16] and model predictive control [17, 18, 19].

However, for simplicity we adopt a pairwise velocity-obstacle (VO) approach [20, 21, 22, 23, 24], where inspired by current rules-of-the-air, UAVs on a ‘conflict course’ each turn to the right to avoid each other.

For simplicity, we constrain dynamics to a 2D layer that models the flight ceiling (presently 120m in the EU [25] and UK [26]) that we assume most of these operations will take place at, to maximise separation with buildings, people on the ground, etc. Furthermore, we will adopt a parsimonious approach to flight mechanics, which best represents quadcopter drones, henceforth referred to as *agents* — the acceleration

capability is assumed independent of the velocity, and there is no minimum stall speed, i.e., agents may hover. The rules of motion are set out in detail in Section 2.

In Section 3 then we outline an experimental setup consisting of two agents, which implement the flight rules in order to avoid each other, and detail a method to measure the performance of the flight rules in terms of the delay an agent incurs while performing the avoidance maneuvers. The results of simulations, swept over various parameter settings, are then presented and show that the ‘turn to the right’ rule can produce large delays for one of the agents when the angle of approach between them is small.

As a result, Section 4 explores similar simulation experiments but with one of the agents deviating from the ‘turn to the right’ rule by either turning to the left or not undergoing avoidance maneuvers. We show then that the overall delay can be improved by allowing one of the agents to deviate and use this to develop a hybrid avoidance rule which uses a simple heuristic to allow agents to choose the most appropriate avoidance behavior.

In order to compare the original and hybrid rules Section 5 will outline a different experimental setup where many simultaneous agents form two traffic streams that cross at a point. The results from several simulations at varying traffic demands are presented and show that the hybrid rule can significantly improve performance, in terms of transit time, for larger traffic demand levels.

The results and use case implications are then discussed in Section 6 before a brief conclusion is provided in Section 7.

2 Agent based model for quadcopter UAVs

We consider a system of (for simplicity) identical agents modelled as point masses moving in continuous time and 2D space, with displacements, velocities and accelerations denoted $\mathbf{r}_i$, $\mathbf{v}_i$, and $\mathbf{a}_i$ respectively. Inspired by the Social Force Model [13] for pedestrian dynamics, we suppose that each acceleration $\mathbf{a}_i$ is composed additively of desired contributions $\mathbf{a}_i^k$, which model e.g., waypoint-following and various pairwise interaction terms respectively, as described below.

However, to model each agent’s limited capabilities, we suppose there is a maximum amplitude acceleration $a_{\max}$ (assumed common to all agents). Furthermore, we assume $a_{\max}$ is independent of velocity, as a simple model of a quadcopter drone. If then the magnitude of the sum of all acceleration contributions $\mathbf{a}_i'$ is less than $a_{\max}$ it is unchanged. However, if the magnitude of $\mathbf{a}_i'$ exceeds $a_{\max}$ then the acceleration experienced by the agent is $\hat{\mathbf{a}}_i' a_{\max}$.

The first acceleration contribution for each agent applies at all times and takes the form

$$\mathbf{a}_i^0 = \frac{1}{\tau} (v_{\text{CS}} \hat{\mathbf{t}}_i - \mathbf{v}_i), \quad (1)$$

which models first order relaxation towards a (common) desired cruising speed v_{CS} in direction $\hat{\mathbf{t}}_i$, i.e., an *optimal velocity model* [27]. In Section 3 and 4, $\hat{\mathbf{t}}_i$ is a constant vector, different for each agent, which prescribes their overall desired direction of motion, whereas in Section 5, we prescribe

$$\hat{\mathbf{t}}_i = \frac{\mathbf{r}_i^{\text{WP}} - \mathbf{r}_i}{|\mathbf{r}_i^{\text{WP}} - \mathbf{r}_i|}, \quad (2)$$

so that agent i flies towards a waypoint at $\mathbf{r}_i^{\text{WP}}$, which in this paper, corresponds to their final destination. Note that the two models for $\hat{\mathbf{t}}_i$ approach each other when the waypoint is very distant from the agent.

Here, for illustration, the rate parameter τ is given by the natural timescale

$$\tau = \frac{2v_{\text{CS}}}{a_{\max}}, \quad (3)$$

so that an agent flying at the cruising speed v_{CS} in exactly the opposite direction to that required will initially retard itself with the maximum allowed acceleration $a_{\max}$.

Figure 1: The velocity obstacle of two converging agents. The dashed circle centred on j is used to construct the velocity obstacle. The relative velocity lies within the velocity obstacle so these two agents are on a ‘conflict course’. The red arrows show the unit vectors of the velocity obstacle, $\hat{\mathbf{r}}_A^{\text{vo}}$ and $\hat{\mathbf{r}}_B^{\text{vo}}$ respectively.

Next, we consider acceleration contributions that are introduced to help avoid conflicts. The goal is to ensure, that at all times,

$$|\mathbf{r}_{ij}| = |\mathbf{r}_j - \mathbf{r}_i| \geq S, \quad (4)$$

where S models a safety distance inside which no pair of agents should usually encroach. Thus if at some future time inequality (4) will be violated for a given pair i, j , they are on a ‘conflict course’ and should maneuver to avoid each other. Ideally, to minimise flight times, the maneuvers should involve the minimum possible detours that respect Eq. (4) unless the two agents have already violated it, then an emergency avoidance acceleration rule is introduced where

$$\mathbf{a}_i^j = -\hat{\mathbf{r}}_{ij} a_{\max}, \quad (5)$$

such that the agent i will accelerate directly away from its neighbor.

To achieve nominal conflict avoidance we adopt a scheme based on the velocity obstacle method [20], see Fig. 1. At each instant, the ego agent supposes that the alter will continue with constant velocity and thus computes a velocity $\mathbf{v}_i^j$ for itself, that if adopted instantly, would cause the two agents to pass each other at minimum separation S . Unfortunately, due to inertia, the velocity $\mathbf{v}_i^j$ cannot be adopted instantly, so instead we adopt an acceleration contribution

$$\mathbf{a}_i^j = \frac{1}{t_C} (\mathbf{v}_i^j - \mathbf{v}_i), \quad (6)$$

which prescribes first order relaxation to $\mathbf{v}_i^j$. Here the timescale t_C is given by the time to conflict, i.e., the time that remains until (4) would be broken, if the agents maintained their current velocities. In this work we also introduce the constraint that

$$|\mathbf{v}_i^j| = v_{\text{CS}}, \quad (7)$$

in an effort to ensure traffic continues to flow while stopping agents from continually speeding up during avoidance maneuvers.

In fact, the alter is applying the same logic so that the computed $\mathbf{v}_i^j$ and $\mathbf{v}_j^i$ would, if adopted, cause the agents to approach at nearest distance $2S$. This does not happen in practice, however because, firstly, v_i^j and v_j^i are not instantly adopted, secondly they change in time, and thirdly, because if at any point the agents came off ‘conflict course’, the manoeuvre would cease to apply. Finally, note that if despite earlier efforts, the agents continue on conflicting courses, the time to conflict t_C will reduce and increase the acceleration contribution in an effort to avoid conflict.

We now consider how to determine the avoidance velocity $\mathbf{v}_i^j$. At any instant the ego agent, i , will produce a velocity obstacle induced by the alter, j , which is defined by two lines that originate from i and

Figure 2: (a) Experimental setup (b) Comparison of linear flight path and representative flight path. The longitudinal displacement d is used to derive a delay metric t^d to compare different starting conditions.

are tangential to a circle of radius S centred on j , with unit vectors $\hat{\mathbf{r}}_A^{\text{vo}}$ and $\hat{\mathbf{r}}_B^{\text{vo}}$ in the frame of i . If the relative velocity of i and j lies within this cone then (4) will be broken at some point in the future. To prevent this, and minimise the detour i undertakes, we consider possible $\mathbf{v}_i^j$ that produce a new relative velocity

$$\mathbf{v}_{ij}^* = \mathbf{v}_i^j - \mathbf{v}_j, \quad (8)$$

that is parallel with either of the velocity obstacle's unit vectors.

Given Eq. (7) and the fact that $\mathbf{v}_{ij}^*$ is parallel to $\hat{\mathbf{r}}_A^{\text{vo}}$ or $\hat{\mathbf{r}}_B^{\text{vo}}$ we can find a value for $|\mathbf{v}_{ij}^*|$ by finding the roots of

$$|\mathbf{v}_{ij}^*|^2 + 2|\mathbf{v}_{ij}^*|\hat{\mathbf{r}}_A^{\text{vo}} \cdot \mathbf{v}_j + (|\mathbf{v}_j|^2 - (v_{\text{CS}})^2) = 0, \quad (9)$$

and a similar equation for $\hat{\mathbf{r}}_B^{\text{vo}}$, resulting in four possible roots. The roots can be used with Eq. (8) to find possible values for $\mathbf{v}_i^j$. We only consider positive real roots when determining $\mathbf{v}_i^j$ since these correspond to maneuvers where agents pass each other at some point in the future. To choose between possible maneuvers, and inspired by the current rules of the air, we prescribe that all agents must turn to their right, which ensures that two agents approaching each other head on do not pick mirrored velocities that lead to subsequent conflicts. If there is more than one $\mathbf{v}_i^j$ that results in a right hand turn then the agent will pick the one with the smallest angular deviation from its current velocity $\mathbf{v}_i$. If no viable $\mathbf{v}_i^j$ is found then the agent will employ the emergency acceleration from Eq. 5.

3 Experiments on pairwise interaction

Here we propose an experimental setup, implemented in simulation, to explore how the rules for agent motion in Section 2 perform for pairs of agents.

We suppose that two agents $i = 1, 2$ are initialised with positions $\mathbf{r}_i(0) = -R_i \hat{\mathbf{e}}_i$ and velocities $\mathbf{v}_i(0) = v_{\text{CS}} \hat{\mathbf{e}}_i$, with $\hat{\mathbf{e}}_1 = (1, 0)^T$ and $\hat{\mathbf{e}}_2 = (\cos \theta, \sin \theta)^T$, see Fig. 2(a). Furthermore, we set $\hat{\mathbf{t}}_i = \hat{\mathbf{e}}_i$ so that, in the absence of collision avoidance maneuvers, Eq. (1) implies the agents will continue at constant velocity, and their paths will cross at angle θ at the origin.

Note that the difference $\Delta R := R_2 - R_1$ constitutes a sort-of ‘distance-phase’ parameter. Clearly, if $|\Delta R| < S$, the agents will come into conflict (i.e., inequality (4) is violated at some point along their trajectories) in the neighborhood of the origin, whatever the approach angle θ . In contrast, in the perfect ‘head-on’ case $\theta = \pi$, the agents will come in to conflict irrespective of $|\Delta R|$. In fact, it may be shown that $|\Delta R| \leq S \sec(\theta/2)$ is a sufficient condition for a ‘conflict course’, in which case, the conflict avoidance scheme described in Section 2 is activated, and the agents will thus deviate from straight line trajectories for $t > 0$.

Here, we set R_1 and R_2 very large, to model agents approaching each other ‘from infinity’, so initially the time to conflict t_C is large, and the initial corrective motion is small. Subsequently, agents on a ‘conflict

Parameter	Value	Units
Cruising velocity (v_{CS})	20	ms^{-1}
Max possible acceleration ($a_{\max}$)	5	ms^{-2}
Desired avoidance separation (S)	30	m
Nominal distance to origin (R_1)	10000	m

Figure 3: Delay experienced by agent 1 when θ is (a) positive and (b) negative. The red lines show the maximum and minimum values of ΔR for which the agents begin on a ‘conflict course’.

course’ will avoid each other to the right, approaching at the minimum distance S , and then via Eq. (1), equilibrate to a path that is parallel to their original course, see Fig. 2(b). The net effect of the interaction is to displace each agent laterally and longitudinally with respect to its original trajectory. In effect, the interaction costs each agent a delay $T_i := d_i/v_{CS}$, where d_i is the respective longitudinal deficit

We conduct a set of experiments that sweeps through values of θ in the range $(-\pi, +\pi)$, with $|\theta| \geq 2S/R_1$ so that the agents are not in conflict at $t = 0$. Correspondingly, ΔR is swept through the range $(-3S/2, +3S/2)$ which captures a variety of settings in which conflicts between the agents occur. In these experiments the other problem parameters, see Table 1, are held at constant dimensional values inspired by previous work such as [8, 10, 28].

Although at this point agents apply the same rules, for the sake of discussion we shall view $i = 1$ as the ego agent and $i = 2$ as the alter agent. Fig. 3 shows the delay $T_1(\Delta R, \theta)$ experienced by agent $i = 1$. The geometric symmetry in the setup implies that $T_2(\Delta R, \theta) = T_1(-\Delta R, -\theta)$ and thus T_2 may be recovered from the T_1 plot. This means we can view agent $i = 2$ as the ego and $i = 1$ as the alter by flipping the signs of ΔR and θ .

The key observation from Fig. 3 is that the problematic situations correspond to small values of $|\theta|$, where the vast majority of the delay is experienced by the agent who is ego with positive θ . Thus a collaborative maneuver is giving rise to an unfair and thus undesirable outcome in the delays experienced. This result is explained in the trajectory plot of Fig. 4. Here, although agent 1 has a ‘head start’ in the initial setup and therefore might be expected to pass in front of agent 2, it turns to the right in the avoidance maneuver, and thus crosses behind agent 2, incurring delay. In contrast, agent 2 follows an almost perfectly linear trajectory. This effect is most severe as the angle θ tends to zero and ΔR approaches its maximum positive value i.e., when agent 1 is far enough ahead to almost be clear of a conflict course.

This finding suggests that other avoidance rules, for example in which agents can choose an alternative direction of avoidance, might reduce delays for small $|\theta|$ situations. Clearly, large angles of approach do not

Figure 4: Agent trajectories for a small angle setup. Agent 1 (red) is initially closer to the origin but deviates to pass behind agent 2 (blue).

Figure 5: Three avoidance behaviors. In (a) both agents turn to their right. In (b) agent 1 turns to its left. In (c) agent 1 does not employ a VO avoidance scheme.

present such problems.

4 Experiments with modified rules for pairwise interaction

We now consider adaptations to the conflict avoidance scheme analysed in Section 3, which employed a ‘turn to the right’ rule that we now denote R_A , where agent 1 was shown to experience large delay for small positive approach angles θ . Alternatively, Fig. 5(b) describes a rule R_B where agent 1 turns to the left to avoid conflict and Fig. 5(c) describes a rule R_C where agent 1 does not employ a conflict avoidance scheme at all, but rather continues in a straight line.

In both rules R_B and R_C , agent 2 continues to implement conflict avoidance by turning to the right, just as in rule R_A .

See Fig. 6(a), which shows the delay $T_1^B(\Delta R, \theta)$ experienced by agent 1 for rule R_B . In this rule, the agents turn away from each other, so that there is no longer left-right handedness in their interaction. Consequently, we may consider $\theta > 0$ without loss of generality, and the delay T_2^B experienced by agent 2 is given by $T_2^B(\Delta R, \theta) = T_1^B(-\Delta R, \theta)$. Compare with Fig. 3(a) and observe that rule R_B dramatically reduces the delay for agent 1 for combinations $\Delta R \lesssim S$ and $\theta \simeq 0$ for which it was worst for rule R_A , at the expense, see Fig. 3(b) of a modest increase in delay for agent 2. Note, however, that rule R_B performs badly at large θ values compared to rule R_A .

For rule R_C agent 1 experiences zero delay while agent 2 experiences similar delays to those of rule R_B , compare Fig. 6(b) with Fig. 6(a) reflected in the $\Delta R = 0$ line. Therefore, both R_B and R_C can be used to reduce the delay incurred by agent 1, but at the cost of additional delay to agent 2.

We therefore propose a hybrid rule where agent 1 chooses from rules A, B, or C, broadly according to the following criteria:

- Adopt rule B or C if

$$T_1^{B,C} \leq T_1^A, \quad (10)$$

that is, agent 1 improves its own experience by deviating from the default rule; and

Figure 6: Delay experienced by (a) agent 1 using R_B and (b) agent 2 using R_C . The red lines show the maximum and minimum values of ΔR for which the agents begin on a ‘conflict course’. The area under the dashed lines are where we propose agent 1 should adopt R_B and R_C respectively.

- We require

$$\left[T_1^A - T_1^{B,C} \right] \geq \left[T_2^{B,C} - T_2^A \right], \quad (11)$$

so that the reduction in agent 1’s delay is greater than any increase in agent 2’s delay and so that there is overall net system benefit.

If R_B and R_C satisfy these conditions at the same time then we suggest that agent 1 should adopt the rule which results in the smallest total delay for the pair of agents. According to this, the areas under the dashed lines in Fig. 6(a) and 6(b) approximate where agent 1 should deviate from R_A by adopting R_B and R_C respectively. In practice, this provides a heuristic which allows agents to choose their avoidance behavior, assuming R_A as the default, based on the parameters θ and ΔR . It is possible to apply this heuristic in more general scenarios by taking a pair of agents’ current states and recovering the setup from Fig. 2(a) in the ego agent’s frame of reference, for example by assuming the point at which their paths cross is based on their instantaneous velocities and finding ΔR with respect to that point.

5 Large scale simulation setup and results

We now perform simulations closer to the real UTM use case, in which large numbers of agents are ‘in the air’ simultaneously. The idea is to test and compare the original conflict avoidance rule R_A with the hybrid rule proposed in Section 4.

In this multi-agent setting, the number of potential conflicts scales quadratically with the number of agents, which is not practically feasible in the application setting. Therefore, we suppose that collision avoidance maneuvers only begin when $t_C < \tau$, i.e., when the time-to-conflict is less than the natural time scale introduced in Eq. (3). Thus typically each agent will only be involved in at most a handful of collision avoidance maneuvers at any one time, which are combined by the additive method described in Section 2.

In our simulations, there are two unidirectional streams of agents that travel between ‘ports’ at fixed locations $(\pm L/2, 0)$ and $(0, \pm L/2)$ respectively, see Fig. 7. This setup generates a sort-of crossroads at $(0, 0)$ where we expect agents to encounter each other, in a similar manner as described in Section 3, with $|\theta| = \pi/2$. Agents are generated at the origin ports $(-L/2, 0)$ and $(0, -L/2)$ with initial velocities $v_{CS}(1, 0)^T$ and $v_{CS}(0, 1)^T$, and waypoints $(+L/2, 0)$ and $(0, +L/2)$ respectively, according to independent

Figure 7: A snapshot of simulations using the experimental setup described in this section when the demand at both origin ports is (a) $0.4\lambda_{\max}$ and (b) $\lambda_{\max}$

Poisson processes of the same rate λ . When an agent is generated, it is added to a queue for take-off, which is served deterministically to maintain a minimum spatial separation $S_{\text{takeoff}} := 3S/2$ that ensures agents do not come into conflict in the early stages of their flight. It follows that when $\lambda \ll S_{\text{takeoff}}/v_{\text{CS}}$, most agents take-off at the same instant they are generated, but when $\lambda > S_{\text{takeoff}}/v_{\text{CS}}$, the agents are deterministically spaced.

During their subsequent journey, we expect agents to encounter agents from the crossing traffic stream somewhere in the neighborhood of the origin. In contrast to the controlled experiments described in Section 3, a single agent might become involved in a sequence of such encounters; there are also secondary encounters where an avoidance maneuver places an agent into a ‘conflict course’ that would not otherwise have occurred. Finally, as agents progress past the crossroads, they realign towards their destination and agents in the same stream may come into conflict with each other as they sort themselves out.

As agents converge upon their respective destination ‘ports’, there is the potential for conflict to arise that could only be resolved by centralised traffic management, that is outside the scope of this paper. Therefore, to simplify matters, we suppose that agents are removed from the simulation when they first enter a ‘landing zone’ of radius R_{LZ} centred upon their destination. The key metric that we will then analyse is each agent’s flight time, from takeoff to entering the landing zone, compared to the time $(L - R_{\text{LZ}})/v_{\text{CS}}$ they would have taken if no other agents were present.

Our experiments sweep over a range of traffic demands λ , from low values in which interactions are rare and delay is small, up to

$$\lambda_{\max} := \frac{v_{\text{CS}}}{2\sqrt{2}S}, \quad (12)$$

which may be shown to be the maximum capacity that the intersection could sustain, without interaction between the agents, if the traffic streams were evenly spaced and perfectly phased. Inspired by the real-world application, we set $L = 2,000$ m and $R_{\text{LZ}} = 60$ m $\ll L$.

Each simulation continues until 1,000 agents have taken off from each of the origin ports. Because the first agents to enter the simulation encounter generally smoother transit, we employ standard statistical procedures to discard the simulation run-up, after which the transit times for more than 900 agents from each stream remain for analysis. Since the overall numbers are small, a bootstrap procedure is used to estimate a distribution for the mean delay $\bar{T}$ for the agents from each stream.

Unsurprisingly, for both avoidance rules used, the mean delay increases along with the total demand, Fig. 8. However, once the demand reaches 40% of $\lambda_{\max}$ the hybrid avoidance rule begins to significantly shorten the mean time delay. While the two traffic streams approach at right angles to each other, and therefore we would not expect the hybrid rule to cause agents to deviate from R_A , agents that are displaced

Figure 8: The mean delay incurred during a simulation for a variety of demand levels. Red circle corresponds to the original flight rule. Blue asterisk corresponds to the heuristic avoidance rule. The boxes show the 95% confidence interval based on 1000 bootstrapped means for each demand level.

from their straight line path may subsequently come in to a conflict course with another agent with which it shares a destination. If this situation occurs the two agents are more likely to have a small angle θ between them and therefore one of the agents is more likely to alter its avoidance behaviour. This effect has a greater impact on the mean delay in scenarios with larger demands since more agents are deviated from their straight line paths and these secondary conflict avoidance maneuvers become more common.

6 Discussion

The results in Section 5 show that even when both streams experience maximum demand the mean delay, when all agents obey R_A , is only an extension of about 15% compared with the ideal, linear flight path. The hybrid rule provides a small but significant performance boost with the mean extension being closer to 12%. However, while the single crossroad setup is far more complex than the initial two-body setup described in Section 3, we expect the UAV use case of goods delivery to result in even more complex traffic scenarios in which agents may have to pass through multiple such intersections and it is unclear how these may interact with each other. Furthermore the crossroads themselves may be more complex with different angles between streams, counterflows travelling in the opposite direction or an intersection of more than two streams. It should also be noted that both rules perform extremely well at lower demands and we will explore in future work how higher level traffic management techniques can exploit this by affecting system level features such as the traffic layout.

Another aspect of this work that is under explored here is the way in which the various acceleration components described in Section 2 are combined additively and how this affects agents' behavior. The pairwise setup is thoroughly explored, well understood and through the use of the hybrid-rule scheme can resolve potential conflicts with small delays for agents approaching from far away. In contrast, the multi-agent scenario is not explored in the same detail and we can provide no guarantees that undesirable edge cases do not exist, e.g., interactions between many agents may lead to deadlocks or driving agents in to obstacles. Future work could explore other ways to combine the acceleration components, for example via a weighted sum, or turning off certain interactions.

An obvious point of comparison for any possible flight rules for autonomous agents are the current right-of-way rules for manned aircraft, defined by the International Civil Aviation Organization (ICAO) in Section 3.2.1 of [29]. These rules define three kinds of conflict course for a pair of aircraft; head on, converging, and overtaking, as well as prescribing which aircraft has the right-of-way and can therefore maintain its velocity.

In a head on scenario neither aircraft is given right-of-way and both must turn to their right similar to R_A . In a converging scenario the aircraft with its neighbor on its left is given the right-of-way, unlike in R_A which again causes both agents to avoid each other in an attempt to more evenly distribute the negative impact of avoidance. Finally, in an overtaking scenario, which [29] defines as one aircraft approaching another from behind within a certain range of angles, the aircraft being overtaken is given right-of-way while the other must turn to its right. Since in our scenario agents attempt to maintain some constant cruising speed overtaking does not occur in the same way as envisioned in the ICAO rules however the effect of both R_B and R_C is to stop craft from overtaking one another when the angle of approach is small, in other words to conserve the order in which agents pass through the point where their paths intersect.

We tentatively propose the flight rules outlined in Section 2 be combined with the heuristic method for determining avoidance behavior from Section 4, in order to form the basis of rules-of-the-air for future autonomous UAVs. However, before such a set of rules could be implemented there are a number of potential real-world issues that have not been addressed in the simulation work presented here. This includes the effect that noisy localisation or communication might have on the rules, particularly in high density traffic situations where many agents broadcasting simultaneously can degrade communication performance further. Similarly, this paper only considers a "sunny day" scenario where agents do not experience failures and all of them conform to the rules. It should be pointed out that the rules presented here can not prevent conflicts with malicious agents that actively pursues a conflict course, especially if it is travels faster than v_{CS} .

7 Conclusion

In this paper we have presented a set of rules that enable agents to navigate towards a destination while maintaining some safe separation with each other. The safe separation is maintained through the implementation of a velocity obstacle based method which enables UAVs to modify their velocity. We have shown that a simple 'avoid to the right' rule is effective for a wide range of setups for a pair of agents but that it can produce undesirable behaviour for small angles of approach. As such we presented two other avoidance rules and developed a heuristic based on the relative positions and velocities of a pair of agents to determine which of the three rules should be applied for any given pairwise interaction. We finally show that the resulting hybrid avoidance rule reduces the mean delay experienced by agents that travel through a crossroads setup.

8 Acknowledgements

This work is funded and delivered in partnership between the Thales Group and the University of Bristol, and with the support of the UK Engineering and Physical Sciences Research Council Grant Award EP/R004757/1 entitled 'Thales-Bristol Partnership in Hybrid Autonomous Systems Engineering (T-B PHASE)'.

References

- [1] Drone Service Market - Global Forecast to 2025, Markets and Markets (2019).
- [2] H. Shakhathreh, A. H. Sawalmeh, A. Al-Fuqaha, Z. Dou, E. Almaita, I. Khalil, N. S. Othman, A. Khreishah, M. Guizani, Unmanned Aerial Vehicles (UAVs): A Survey on Civil Applications and Key Research Challenges, IEEE Access 7 (2019) 48572–48634. [arXiv:1805.00881](#), [doi:10.1109/ACCESS.2019.2909530](#).
- [3] J. Koetsier, Drone Delivery Is Live Today, And It's 90% Cheaper Than Car-Based Services, Forbes (2021).
- [4] Connected Places Catapult, Enabling UTM in the UK, Tech. Rep. May (2020).
- [5] Federal Aviation Administration, Unmanned Aircraft System (UAS) Traffic Management (UTM): Concept of Operations v2.0, Tech. rep. (2020).

- [6] European Union Aviation Safety Agency, Opinion No 01 / 2020: High-level regulatory framework for the U-space, Tech. Rep. 01 (2020).
- [7] M. Doole, J. Ellerbroek, J. Hoekstra, Drone Delivery: Urban airspace traffic density estimation, in: 8th SESAR Innovation Days, 2018.
- [8] A. Oosedo, H. Hattori, I. Yasui, K. Harada, Unmanned Aircraft System Traffic Management (UTM) Simulation of Drone Delivery Models in 2030 Japan, *Journal of Robotics and Mechatronics* 33 (2) (2021) 348–362. doi:10.20965/jrm.2021.p0348.
- [9] L. Sedov, V. Polishchuk, Centralized and Distributed UTM in Layered Airspace, in: 8th International Conference on Research in Air Transportation, 2018, pp. 1–8.
- [10] A. Battista, D. Ni, A Comparison of Traffic Organization Methods for Small Unmanned Aircraft Systems, *Transportation Research Record* 2672 (23) (2018) 20–30. doi:10.1177/0361198118757995.
- [11] M. Liu, Y. Wan, F. L. Lewis, E. Atkins, D. O. Wu, Statistical Properties and Airspace Capacity for Unmanned Aerial Vehicle Networks Subject to Sense-and-Avoid Safety Protocols, *IEEE Transactions on Intelligent Transportation Systems* (2020) 1–14doi:10.1109/TITS.2020.3040520.
- [12] C. W. Reynolds, Flocks, herds, and schools: A distributed behavioral model, in: *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1987*, Vol. 21, 1987, pp. 25–34. doi:10.1145/37401.37406.
- [13] D. Helbing, P. Molnár, Social force model for pedestrian dynamics, *Phys. Rev. E* 51 (5) (1995) 4282–4286. doi:10.1103/PhysRevE.51.4282.
URL <https://link.aps.org/doi/10.1103/PhysRevE.51.4282>
- [14] O. Khatib, Real-time obstacle avoidance for manipulators and mobile robots, in: *IEEE International Conference on Robotics and Automation*, 1985, pp. 500–505.
- [15] J. Barraquand, B. Langlois, J.-C. Latombe, Numerical potential field techniques for robot path planning, *IEEE Transactions on Systems, Man, and Cybernetics* 22 (2) (1992) 224–241.
- [16] K. Sigurd, J. How, UAV trajectory design using total field collision avoidance, *AIAA Guidance, Navigation, and Control Conference and Exhibit* (2003) 1–11doi:10.2514/6.2003-5728.
- [17] T. Schouwenaars, B. De Moor, E. Feron, J. How, Mixed integer programming for multi-vehicle path planning, 2001 European Control Conference, ECC 2001 (2001) 2603–2608doi:10.23919/ecc.2001.7076321.
- [18] A. Richards, T. Schouwenaars, J. P. How, E. Feron, Spacecraft trajectory planning with avoidance constraints using mixed-integer linear programming, *Journal of Guidance, Control, and Dynamics* 25 (4) (2002) 755–764. doi:10.2514/2.4943.
- [19] D. Morgan, S. J. Chung, F. Y. Hadaegh, Model predictive control of swarms of spacecraft using sequential convex programming, *Journal of Guidance, Control, and Dynamics* 37 (6) (2014) 1725–1740. doi:10.2514/1.G000218.
- [20] P. Fiorini, Z. Shiller, Motion Planning in Dynamic Environments using Velocity Obstacles, *The International Journal of Robotics Research* 17 (7) (1998) 760–772.
- [21] J. D. Van Berg, M. Lin, D. Manocha, Reciprocal velocity obstacles for real-time multi-agent navigation, *Proceedings - IEEE International Conference on Robotics and Automation* (2008) 1928–1935doi:10.1109/ROBOT.2008.4543489.
- [22] J. Snape, J. V. D. Berg, S. J. Guy, D. Manocha, The hybrid reciprocal velocity obstacle, *IEEE Transactions on Robotics* 27 (4) (2011) 696–706. doi:10.1109/TR0.2011.2120810.

- 335 [23] J. Van Den Berg, S. J. Guy, M. Lin, D. Manocha, Reciprocal n-body collision avoidance, Springer Tracts in Advanced Robotics 70 (STAR) (2011) 3–19. doi:10.1007/978-3-642-19457-3_1.
- [24] J. A. Douthwaite, S. Zhao, L. S. Mihaylova, Velocity Obstacle Approaches for Multi-Agent Collision Avoidance, Unmanned Systems 7 (1) (2019) 55–64. doi:10.1142/S2301385019400065.
- [25] European Union Aviation Safety Agency, Easy Access Rules for Unmanned Aircraft Systems, Tech. rep. (2021).
- 340 [26] Civil Aviation Authority (UK), The Drone and Model Aircraft Code, Tech. rep. (2019). URL <https://register-drones.caa.co.uk/drone-code>
- [27] M. Bando, K. Hasebe, A. Nakayama, A. Shibata, Y. Sugiyama, Dynamical model of traffic congestion and numerical simulation, Physical Review E 51 (2) (1995) 1035–1042. doi:10.1103/PhysRevE.51.1035.
- 345 [28] M. F. B. M. Salleh, D. Y. Tan, C. H. Koh, K. H. Low, Preliminary concept of operations (ConOps) for traffic management of unmanned aircraft systems (TM-UAS) in urban environment, AIAA Information Systems-AIAA Infotech at Aerospace, 2017 (January) (2017) 1–13. doi:10.2514/6.2017-0223.
- [29] ICAO, Annex 2 - Rules of the Air - Tenth Edition (November) (2005) 1–74. URL <http://www.icao.int>