ALTISCOPE

Metrics to characterize dense airspace traffic

Richard Golding

Report TR-4
7 June 2018

Table of Contents

Abstract 1

1. Introduction 1

2. Metrics 2

3. Evaluation approach 3

3.1. Traffic patterns 5

4. Results 6

4.1. Uniform random traffic 6

4.2. Traffic with more order 8

4.3. Effect of distance versus speed 10

4.4. Interpreting closing time 11

4.5. Fraction of time maneuvering 13

4.5.1. Estimating encounter rate 16

5. Conclusions 17

Acknowledgements 19

References 19

Abstract

Many studies and design discussions are concerned with "dense" or "high density" UAV operations—but there is no definition of what this term means. We propose two metrics that can be used to determine when the traffic in flight is dense. These metrics are based on an intuition that density matters when the vehicles in flight interact with each other. We find that the absolute number of vehicles in a volume is not, by itself, a good metric for determining dense operations, but that instead our proposed metrics reflect the effects of varying degrees of order in traffic flows. Analysis of the results suggests that traffic can become "dense" at low traffic volumes, including levels much lower than anticipated demand in urban areas.

1. Introduction

Many studies and designs for UAS traffic management (UTM) have discussed that when the traffic in flight is "dense" it will require more control than when the traffic is "not dense". This usually is taken to mean that when traffic is not dense, few air traffic management services are required, but that when the traffic does become dense, a system that provides services analogous to the current Air Traffic Services [1] will be required to maintain safe and efficient operation. For example, the European U-space project [2] expects first to deploy services that only handle low-density operation, and follow on with a system that can handle dense operations some years later.

The investment required to define, analyze, and deploy these services is high, so it is mandatory to have a clear understanding of the conditions that make a traffic flow "dense". This report discusses two proposed metrics for determining when traffic should be treated as dense.

We start with the intuition that it is interactions between vehicles in flight that makes a traffic flow dense. If no aircraft are having to change their behavior—such as changing their path to avoid each other—then intuitively there is no issue with density. When aircraft are constantly tracking or avoiding other aircraft, then our intuition is that density is high.

The two metrics we propose are based on this operational understanding of density. Informally, these metrics are (1) how close the closest aircraft is on average, and (2) how many aircraft are in the immediate vicinity on average. These metrics provide an indication of how often the typical aircraft is having to maneuver to avoid collision, and how many nearby aircraft must be tracked as potential collision threats.

This approach to defining dense traffic makes the result dependent on patterns or order within the traffic flow. For example, if aircraft are travelling randomly in all directions, there are more interactions than if the same number of aircraft are all travelling in the same direction at about the same speed. This reflects intuitive understanding people have from walking in crowds or driving cars.

We find that by these metrics, traffic flows expected in urban areas by 2030 will regularly exhibit dense areas that require traffic management solutions to maintain safe operations.

2. Metrics

The two metrics we propose for evaluating effects of density both derive from ways of looking at how much vehicles in flight interact: if they don't interact, then the traffic flow is not dense, and if they interact a lot, it is dense.

We have focused on vehicle interactions that relate to collision avoidance: how soon a vehicle has to maneuver to avoid collision, and how many other vehicles are potential collision threats.

The metrics are computed from the point of view of the *ownship,* that is, from the point of view of one particular vehicle that has other vehicles around it. This is a standard approach for the design of collision avoidance mechanisms.

The metrics are based on the *closing time*. This is the distance from the ownship to another aircraft, divided by the speed at which the other aircraft is moving toward the ownship, which gives the amount of time required for the other aircraft to reach the ownship. This a simplification of the tau measures used in TCAS; we compute a single measure rather than making separate measurements in the horizontal plane and along the vertical axis.[3]

In most cases, the other aircraft will not be moving directly toward the ownship, but will be approaching at a slant. The closing time metric uses the component of its velocity vector that is in the direction of the ownship for the closing speed. This is found by projecting the other aircraft's velocity vector onto the unit vector from the other aircraft to the ownship:

$$s_{\text{closing}} = \vec{v}_{\text{aircraft}} \cdot \vec{c}$$

where $\vec{c}$ is the unit vector from the aircraft to the ownship, and $\vec{v}_{\text{aircraft}}$ is the net velocity of the other aircraft relative to the ownship: $\vec{v}_{\text{aircraft}} = \vec{v}_{\text{other}} - \vec{v}_{\text{ownship}}$. If the other aircraft is moving away from the ownship, this speed will be negative.

The first metric is the *minimum closing time*. This is computed by looking at all the aircraft in the airspace, measuring their closing time, and selecting the smallest. This measure approximates the amount of time that the ownship will have to react and maneuver to the situation when it has to maneuver to avoid another aircraft.

This metric slightly overestimates the actual time available for reacting and maneuvering, so that the actual situation is on average slightly worse than the time computed; however, that is not important as a metric of density. The problem can show up when there is another aircraft that is close to the ownship but travelling in close to the same direction, so that the closing speed is small. For this case, TCAS adds a second condition for issuing resolution advisories (RAs) that is based solely on distance, regardless of closing speed.[3] Since the interest in this study is for a metric that estimates density, we do not attempt to add a correction to the minimum closing time metric and accept the possibility of an imperfect metric.

The second metric is the *number of close aircraft*. This is determined by computing the closing time for all aircraft in the airspace, and counting the number that have a closing time less than 15 seconds. This gives an estimate of the number of aircraft the ownship will be monitoring as potential collision threats, which measure drives the complexity of guidance computation on the ownship and the complexity of making avoidance decisions. In Section 4.4 show that the 15 second parameter is a reasonable bound for different types of aircraft, and discuss how that parameter might change based on aircraft capability.

The 15 second threshold is what TCAS uses for issuing a resolution advisory at 1000 feet AGL.[3, p. 23] TCAS uses larger thresholds at higher altitudes, but most UAV operations will be at low altitudes and UAVs are not expected to be spaced multiple miles apart. The TCAS thresholds were derived from estimates of pilot reaction time and the time for a large commercial aircraft to perform a vertical maneuver with the expected acceleration. The thresholds for UAVs are not yet determined, and likely vary widely: a remotely-piloted fixed-wing UAV being operated using VLOS rules at the extreme of visual line-of-sight distance seems likely to have reaction times much slower than those assumed by TCAS, while an autonomous UAV with high-performance on-board detection capabilities is likely to react much faster.

3. Evaluation approach

This study evaluates the minimum closing time and number of close aircraft metrics for randomly-generated scenarios, to determine how these metrics vary with the number of aircraft in flight and the pattern of the traffic flowing in the region.

Both metrics are computed from the point of view of a single ownship, which is located at the center of the airspace (Figure 2). The analysis code generates random, instantaneous snapshots of an airspace and computing the metrics at that instant.

The analysis code generates random scenarios in a 10 km by 10 km airspace, with a 300 m thick vertical region. The ownship is travelling at 20 m/s in the $+y$ (northward) direction. The

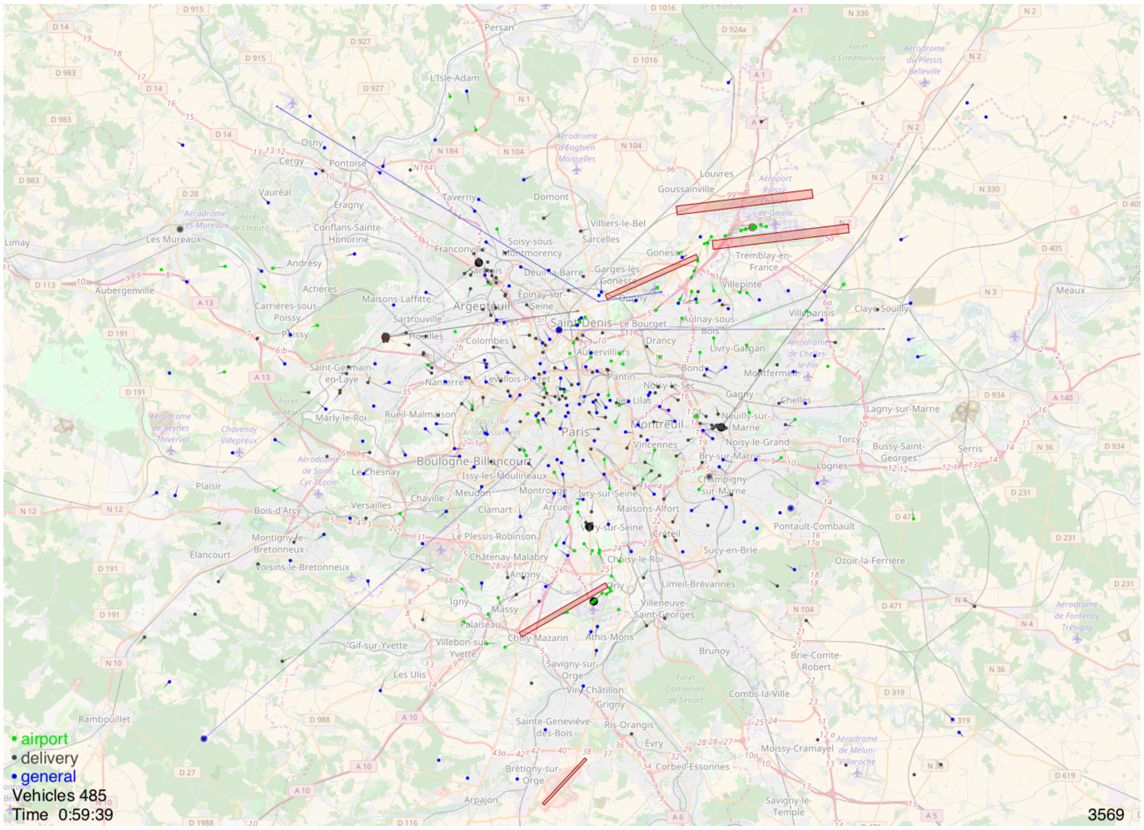

Figure 1: Snapshot of projected traffic over the Paris metropolitan region in 2030. Note how vehicles are clustered in the center of the region, and how no-fly zones create local concentrations of flights around airports.

other flights have start and end points picked randomly in the region, with speeds chosen uniformly randomly between 2 m/s and 40 m/s; the aircraft is at a randomly-chosen point on its path.

The analysis proceeds by varying the number of aircraft (besides the ownship) that are in flight between 25 and 10,000 aircraft.

For comparison, some of our traffic projections suggest an aggregate 2,000 flights per hour across all classes of traffic in a typical similarly-sized region; accounting for flight durations, we expect between 450 and 500 aircraft in flight at any given moment over the entire region. The density is not at all uniform, with greater density over the center of the region and a few points around no-fly areas that concentrate traffic even further.[1] We use 500 vehicles in flight as a comparison point in our evaluations below.

[1] Based on projections for UAV flights over the Paris metropolitan area in 2030, including Urban Air Mobility and package/cargo delivery traffic. Our projections for other metropolitan areas are similar.

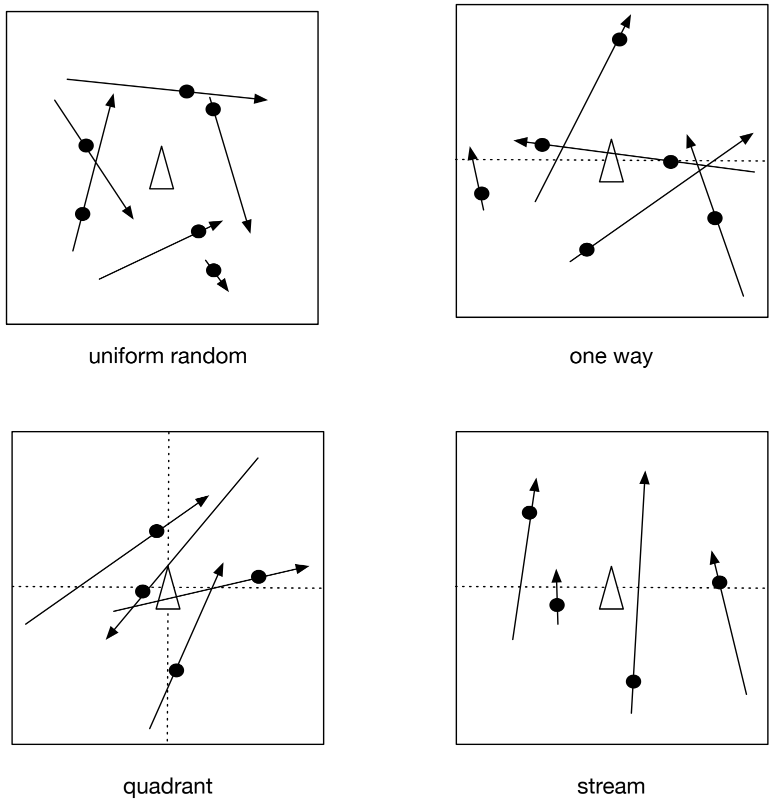

Figure 2: Traffic patterns evaluated. The triangle in the center depicts the ownship.

For each case of number of aircraft and traffic pattern, the analysis generates 10,000 random scenarios and computes the minimum closing time and number of close aircraft metrics for each. The analysis reports the mean and standard deviation of the 10,000 samples of both metrics.

3.1. Traffic patterns

The analysis evaluates the metrics using four different traffic patterns (Figure 2):

- *Uniform random* traffic, where the start and end points are selected using a uniform random distribution over the entire *x*, *y*, and *z* ranges. This pattern has the least order of the patterns evaluated.
- *One-way* traffic, where all traffic is going from the lesser *y* ("south") half of the airspace to the greater *y* ("north") half of the airspace. The *x* and *z* values are chosen uniformly at random over the full range of those dimensions. The *y* values are chosen uniformly at random within the appropriate half of the space. This pattern has a general direction to the flow, and the ownship is going in that same general direction, but some traffic can be headed at close to 90º off from the ownship's direction. This traffic pattern models what

would be observed if aircraft were segregated into two altitude bands by east-west heading.

- *Quadrant* traffic, where all traffic is going between the southwest and northeast quadrants of the airspace, in both directions. This traffic is generally at an angle to the direction of the ownship, with a bias toward a smaller range of encounter angles. This pattern is similar to what is observed when a bidirectional stream of aircraft is squeezed into a gap between obstacles or no-fly zones.
- *Stream* traffic, where all the traffic is going from the south half of the airspace to the north half; the paths are limited in how far east or west they will travel (to +/- 1 km). This produces the most ordered stream of traffic, and in the same direction as the ownship is travelling. This traffic pattern models what would be observed if aircraft are segregated into many altitude bands by narrow ranges of headings.

4. Results

Using the approach defined above, we ran the analysis to get the results that follow. We first present the results using uniform random traffic patterns, which presents the most difficult situation, followed by results for more ordered traffic patterns.

4.1. Uniform random traffic

This traffic model presents the most chaotic environment amongst the models evaluated. It models free flight with no mechanisms that segregate traffic to improve flow. Table 1 shows the results; Figures 3 and 4 illustrate them.

The metrics show smooth changes as the number of aircraft in flight increases. As these metrics are based on the geometric relationships between random flight segments, this is expected.

First consider the minimum closing time metric, shown in in Figure 3. In an average instant, when there are 250 flights in the air, the nearest aircraft will be close enough to trigger a TCAS resolution advisory. For UAV traffic, this means that at 250 aircraft in flight, a UAV using TCAS collision avoidance rules would be almost continuously maneuvering to avoid other traffic if flight plans are not coordinated. This can be resolved by either generating flight plans that reduce the number of encounters, or by using collision avoidance mechanisms that have a tighter tolerance.

This result indicates that a region where the traffic is uniformly random must be treated as dense at 500 aircraft in flight, the volume projected to occur over the Paris region. This can be handled by generating flight plans that reduce the number of encounters, especially by introducing greater order (as the evaluations of the next traffic patterns will show). Deploying collision avoidance mechanisms that can respond in a shorter time frame will also help, but will not solve the problem—if automation can reduce the reaction time from 15 seconds to 10 seconds, at 500 active flights the aircraft would still spend most of its time in avoidance maneuvers.

The number of close aircraft metric shows that at 500 active flights, a UAV will need to be tracking 1-2 other vehicles on average at all times for collision avoidance. The collision avoidance system will need to be designed for tracking significantly more vehicles than that, as the number in this analysis is only an average, and the peak numbers are surely much higher.

Table 1: Results using a uniform random traffic pattern

Number of flights	Average minimum closing time (seconds)	Average number close aircraft
25	46.8	0.0848
50	32.8	0.158
75	27.0	0.241
100	23.4	0.311
170	18.1	0.537
250	15.1	0.796
500	11.0	1.57
750	9.23	2.34
1000	8.21	3.12
1700	6.58	5.28
2500	5.63	7.74
5000	4.38	15.5
7500	3.79	23.3
10000	3.42	31.0

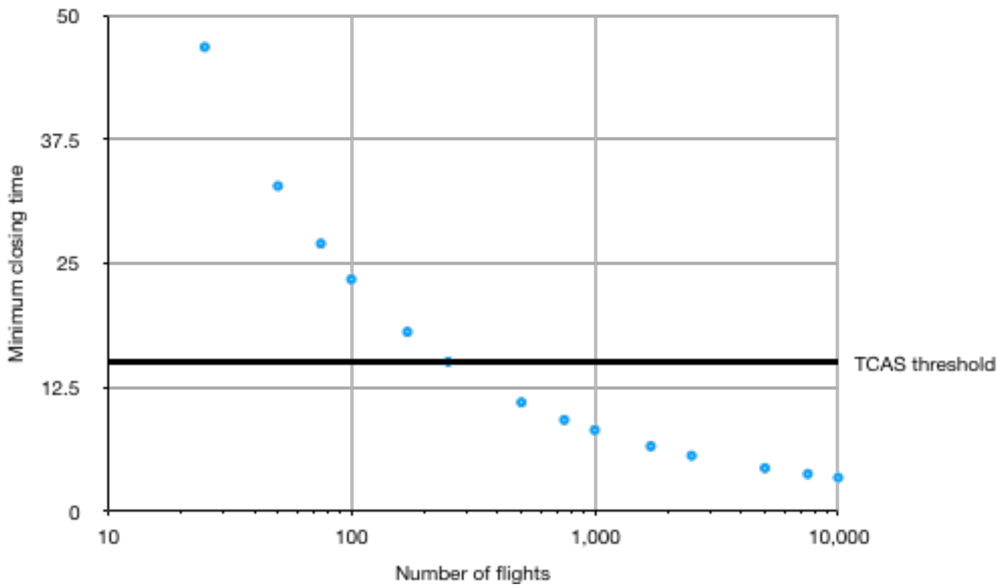

Figure 3: Minimum closing time for uniform random traffic

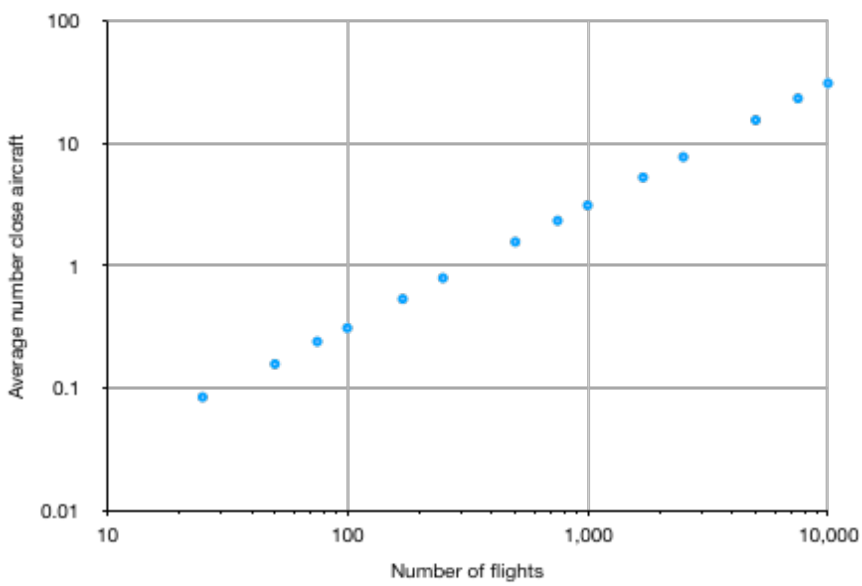

Figure 4: Number of close aircraft for uniform random traffic

4.2. Traffic with more order

The other three traffic patterns we modelled—quadrant, one way, and stream—have greater degrees of order in their flow. The quadrant pattern has traffic going both ways between the northeast and southwest quadrants; the one way pattern has traffic going arbitrarily from the southern half of the airspace to the northern; and the stream pattern going south to north with little east-west variation.

Figure 5 compares the results we observed for these patterns.

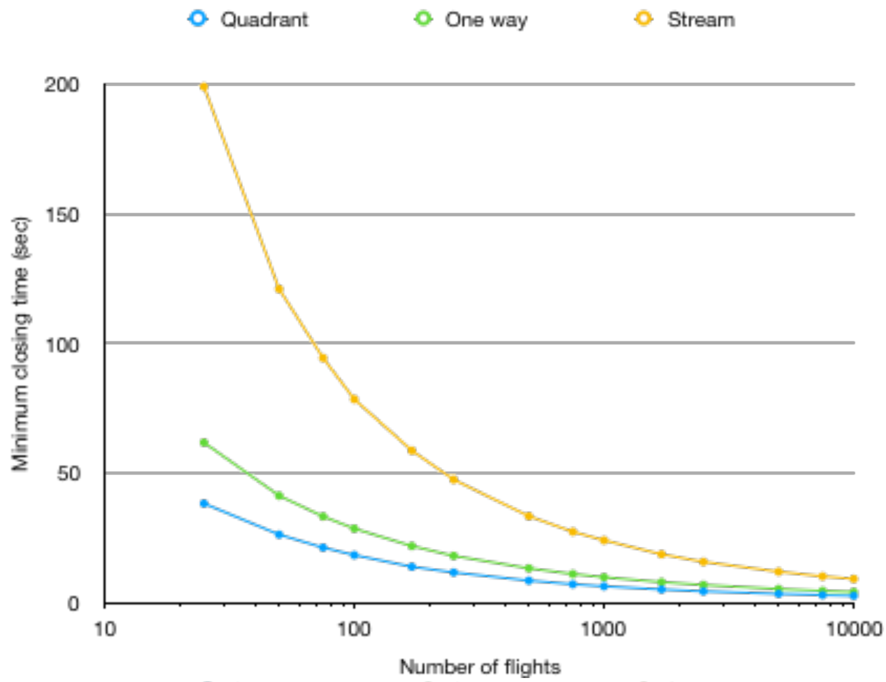

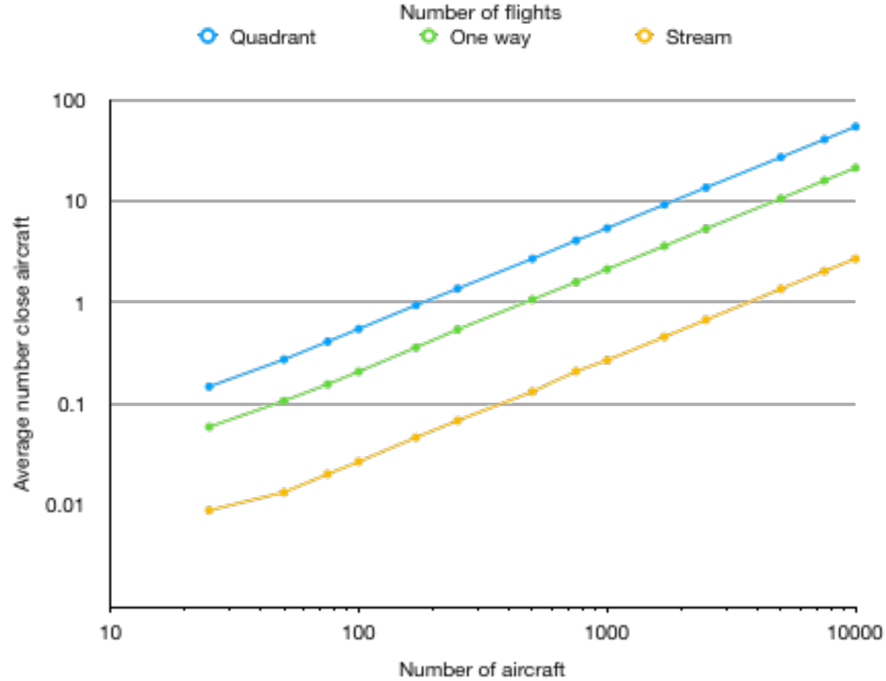

Figure 5: Metrics for traffic patterns with more order

The quadrant pattern has the least order: other flights can be traveling both in a similar direction as the ownship and in a nearly opposite direction, from nearly 180º horizontal range. It shows the most conflicts of the three, with minimum closing time a bit more than half of what the one way pattern yields.

The one way pattern also has flights from nearly 180º range as well, but all traveling south to north, the same as the ownship. This results in more time to react to nearby flights compared to the quadrant pattern.

The stream pattern has the greatest order: all flights are travelling south to north on nearly the same headings. This results in a dramatic increase in the time that the ownship has to react to other aircraft. Compared to random flight paths, having all traffic moving northward allows more than ten times the number of aircraft in flight before reaching the point where the ownship is nearly continuously avoiding other aircraft.

The first conclusion to draw from these results is that for a given number of aircraft in flight, the more that all these aircraft are on similar headings the better these measures of the effect of density are. We have observed this effect in simulation results before, and these metrics match those observations.

A second conclusion is that when a region has an area where traffic is pinched by obstacles, it is important to take management steps to reduce the entropy in the traffic flow. The quadrant traffic pattern is similar to a situation we observe in the Paris simulations, where traffic is forced in to a gap between the UAV no-fly zones around De Gaulle and Le Bourget airports. This situation compresses the traffic to a higher local number of aircraft per area while making more encounters closer to head-on with greater closing speeds and thus less time to react.

4.3. Effect of distance versus speed

The metrics by which we measure the effects of density are based on closing time. The closing time between the ownship and another aircraft depends on both the distance between the two aircraft and their relative velocities. We wanted to determine how much of the effects we observe are due to differences in distance and how much from differences in velocity.

Figure 6 compares the distributions of distance between the ownship and other aircraft for different patterns. This graph shows the fraction of all the other aircraft at different distance ranges: in buckets of 100 meters up to 500 meters, then in 500-meter buckets thereafter. (The bucket for distances over 5000 meters is not shown for clarity.) We note that all the traffic patterns except the stream pattern will tend to produce a concentration of flights near the center of the airspace, and indeed in the limit should tend toward a normal distribution of distances. The stream pattern should produce a near-uniform distribution of distances on the x axis and a normal distribution of aircraft locations on the y axis.

Our metrics only focus on the vehicles that are close to the ownship, so almost all the aircraft reported in these distributions do not affect the metrics. For the closest bins, 500 m and closer,

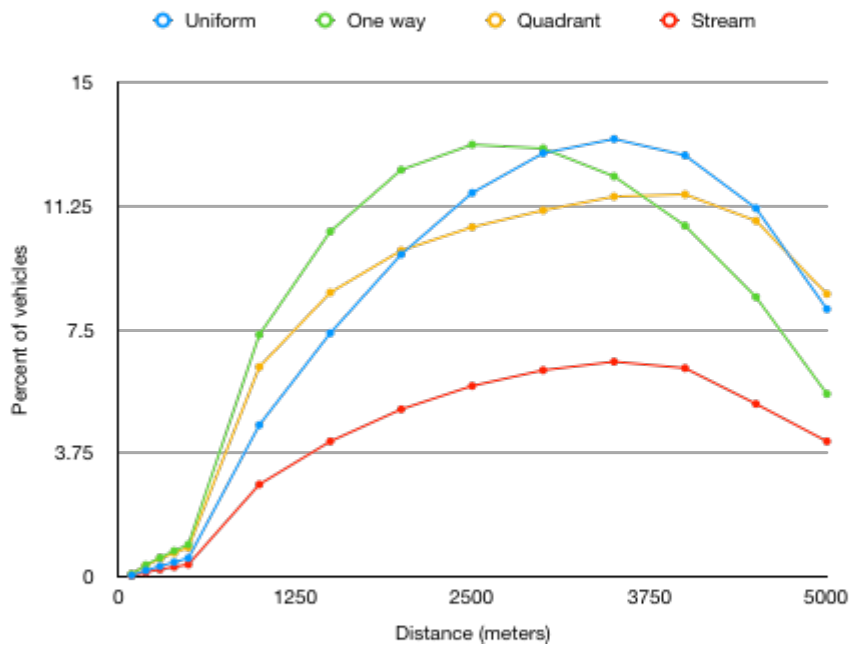

Figure 6: Distribution of other vehicles by distance from the ownship. Note that the graph shows values at every 100 m up to 500 m in order to detail the structure near the ownship, and in increments of 500 m thereafter.

the distance measure orders the traffic patterns in nearly the same way as the metrics: the stream pattern has the least aircraft less than 500 m away, just as it has the furthest average minimum closing time. However, the one way and quadrant traffic patterns generate similar numbers of aircraft less than 500 m distant, but show significantly different minimum closing time or number of close (measured by closing time) aircraft. This difference comes from the different relative velocities in the one way and quadrant traffic patterns: the quadrant pattern has half the other aircraft traveling north to south, in the opposite general direction to the ownship.

From this we conclude that neither distance nor relative velocity by themselves dominate the results we see in the metrics; instead, using closing time usefully combines these effects into a single measure.

4.4. Interpreting closing time

We have used 15 seconds, the TCAS resolution advisory threshold at 1,000 feet, as the threshold for interpreting closing time results so far. However, the actual time required to respond to potential collisions varies widely for different kinds of UAVs and missions. TCAS also uses much larger thresholds at higher altitudes, up to 35 seconds.

Most maneuvers involving civilian manned aircraft, where the people on board are not specially equipped or trained, are limited to 0.25 or 0.33 *g* acceleration. We have confirmed this through our own measurement of commercial transport aircraft during takeoff, turns, and maneuvering, and from the design rules for TCAS II.[3] We expect this standard to hold for passenger-carrying urban air mobility vehicles.

Maneuvers for unmanned UAVs can potentially experience higher accelerations. We have anecdotal reports of 2 *g* maneuvers for small quadcopters and fixed-wing UAVs flying non-aerobatic courses.

Based on these estimates, we can construct some best-cast estimates of the reaction time required for collision avoidance. We consider two cases, where the ownship must perform a 50 m deflection or a 100 m deflection in order to resolve a potential problem:

Type	50 m deflection	100 m deflection
Unmanned (2 g)	2.3 s	3.2 s
Manned (0.25 g)	6.4 s	9.0 s

These times assume that the vehicle performs a constant acceleration at 0.25 or 2.0 *g* until it has reached the deflection distance. The time to slow and return to the original course is not considered.

In addition, it will take some time for the vehicle to detect the need for performing the avoidance maneuver and for its control system to perform actuation to achieve the desired acceleration. For autonomous collision avoidance systems, we assume best case 3 sec lag total for these steps, based on receiving two sensing updates at one-second intervals plus a one-second lag in the controller, speed controllers, and motors. We ignore this overhead in the remainder of this analysis.

4.5. Fraction of time maneuvering

We combine the results so far to answer a final question: how much of the time will the ownship spend maneuvering around potential conflicts, if no deconfliction is provided by a traffic management system?

We model the ownship's status as a continuous-time Markov process with three states:

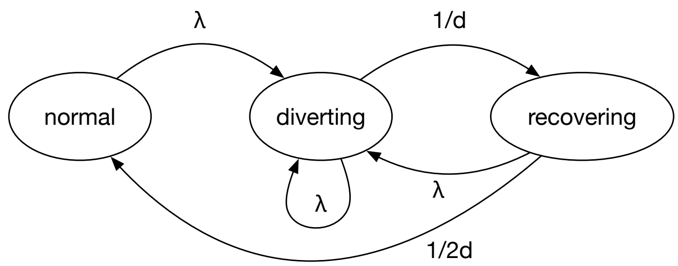

where the vehicle starts in normal flight, following its flight path. When another vehicle comes close enough (at rate **λ**), the vehicle begins diverting. While diverting, the vehicle may encounter another vehicle (at rate **λ**) or finish the evasive maneuver (at rate $1/d$). The vehicle then performs a second maneuver to return back to its normal flight path. It either has to evade another vehicle (at rate **λ**) while recovering, or it completes the return to its normal path (at rate $1/2d$). Using a continuous-time Markov process implies exponentially-distributed times for encountering other aircraft and for the ownship completing a diversion maneuver.

We computed the encounter rate **λ** from the observed durations of encounters and the distribution of times when there are zero, one, or n nearby aircraft, as detailed in Section 4.5.1 below.

The evasive maneuver takes on average d seconds, from Section 4.4. This translates to a rate of $1/d$.

We assume that the maneuver to return to the original flight path takes twice as long as the evasive maneuver: d seconds to decelerate from the evasion, and a similar amount of time to move to a new velocity that will move in an acceptable way to the next navigation waypoint on its flight path.[2]

---

[2] There is significant uncertainty in this assumption. A naïve recovery assumption would involve an average recover time of $3d$: deceleration at the same level to eliminate the Δv from the evasion, plus $2d$ at the same acceleration and deceleration to return to the previous path. The $2d$ value is based on assumption that the system will do something cleverer. The naïve approach will result in longer recover times and less time in normal flight; recovering at gentler accelerations will result in even worse results.

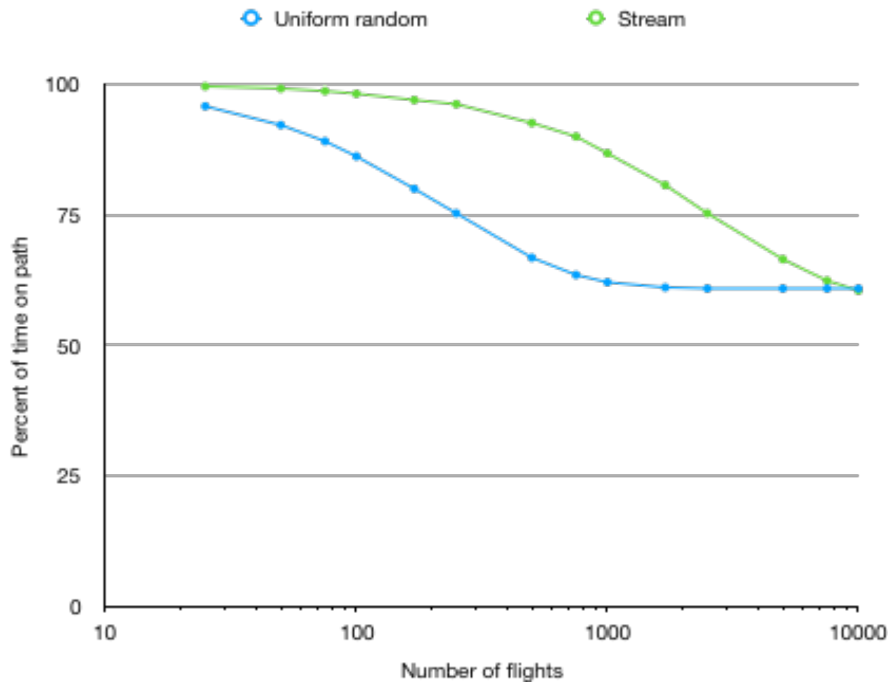

Figure 7: Fraction of time the ownship will be in normal state, for an unmanned UAV diverting 50 m

Using this model, we compute the fraction of a flight that the ownship will spend in the normal state for the uniform random and stream traffic patterns (Figures 7 and 8).

As we note in the analysis in Section 4.5.1, it is likely that the lower limits reached in Figures 7 and 8 are overoptimistic, and that the ownship spends more time in avoidance maneuvers than is indicated. These results should therefore be taken as an upper bound on how well aircraft can perform.

For uniform random traffic, at 500 active flights (the steady-state number in our expected urban scenario), an unmanned cargo UAV will be in normal flight about 2/3 of the time, and a passenger-carrying UAV will be flying normally only 29% of the time. That means that all vehicles will be spending a significant amount of time in maneuvering.

For the stream traffic pattern at 500 active flights, where all aircraft are moving in a similar direction though at different speeds, the situation is much better: an unmanned UAV will fly normally 93% of the time, spending less than 7% of its time in evasion, while a passenger-carrying UAV will be flying normally 75% of the time.

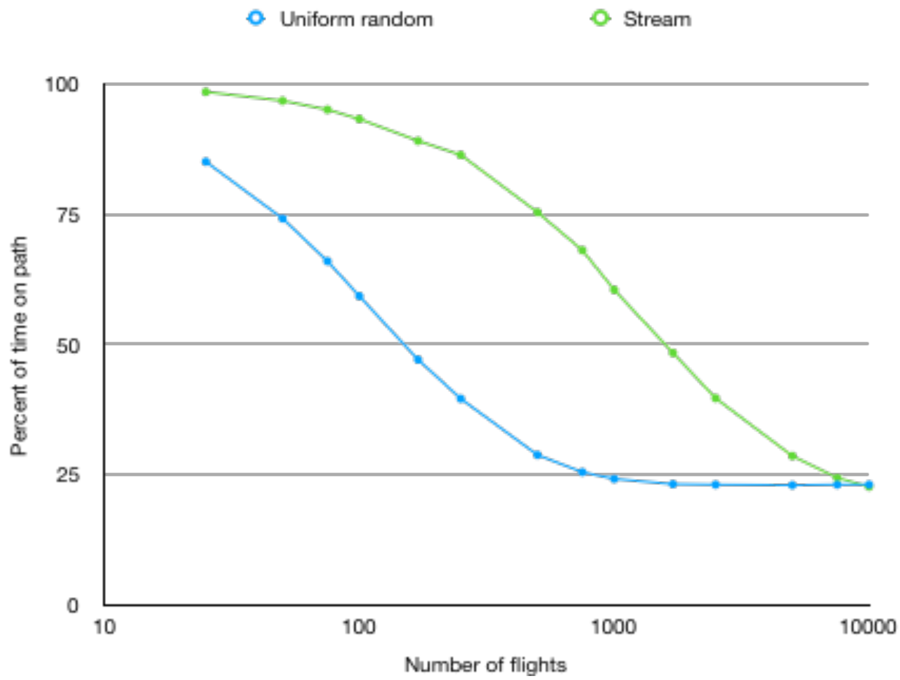

Figure 8: Fraction of time the ownship will be in normal state, for a manned UAV diverting 100 m

At 25 active flights, the lowest value we modeled, we find that vehicles will spend the following percentage of their time in normal flight:

	Uniform random traffic pattern		Stream traffic pattern	
Type	**50 m deflection**	**100 m deflection**	**50 m deflection**	**100 m deflection**
Unmanned (2 g)	95.8%	94.2%	99.6%	99.5%
Manned (0.25 g)	89.0%	85.1%	98.9%	98.5%

For 500 active flights:

	Uniform random traffic pattern		Stream traffic pattern	
Type	**50 m deflection**	**100 m deflection**	**50 m deflection**	**100 m deflection**
Unmanned (2 g)	66.8%	58.3%	92.6%	90.0%
Manned (0.25 g)	38.3%	28.8%	81.5%	75.4%

An analysis of TCAS system performance in the US reported one resolution advisory every 116 flight hours.[4] By comparison, the best result above is equivalent to roughly two resolution advisories per hour,[3] or about 230 times higher rate of collision avoidance events than what is currently observed in manned aviation systems.

These results suggest that "dense" traffic, meaning traffic where flights interact often enough that they spend more than perhaps 10% of their time maneuvering to avoid collision, occurs even at very low numbers of aircraft when the traffic is not organized. Conversely, if a UTM system is organizing traffic flows to reduce the number of interactions by grouping aircraft on similar headings together—the stream traffic model here—then the vehicles will spend far less time avoiding collision.

4.5.1. Estimating encounter rate

The analysis above uses the estimated rate λ at which the ownship will encounter other aircraft and begin an avoidance maneuver. We estimate this rate from the measured duration of encounters and the amount of time that no other aircraft is within the 15 second limit of the ownship. This estimation method is similar to that used by Maki *et al.*[5]

We model the number of other aircraft close to the ownship as a continuous-time birth-death Markov process:

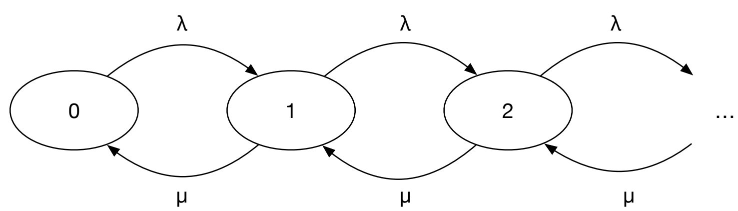

where states are labeled with the number of other aircraft within the 15 second limit; λ is the rate at which some other aircraft enters the limit zone; and μ is the rate at which other aircraft complete their encounter and move out of the zone.

We can then use the rule $\pi_i = (1 - \rho)\rho^i$ for the amount of time spent in state i, where $\rho = \frac{\lambda}{\mu}$. Solving for λ we get $\lambda = \mu(1 - \pi_0)$.

We measured the average duration d of an encounter for each traffic pattern and number of flights; $\mu = 1/d$ for exponentially-distributed durations. We also measured π_0, the fraction of time that no aircraft are close by the ownship.

[3] The average flight duration in our simulations lasts about 15 minutes. 99.6% in normal flight means about 3.6 seconds per flight in evasion on average; each evasive maneuver lasts on average 6.9 seconds, for ~0.52 maneuvers per 15-minute flight.

The resulting mean time to encounter values, in seconds, for different traffic patterns are:

Number of flights	Uniform random	One way	Quadrant	Stream
25	162	220	95.7	2090
50	81.3	117	51.5	738
75	58.3	78.2	36.1	536
100	45.1	63.2	29.0	409
170	29.5	39.1	20.1	249
250	22.5	28.5	16.4	169
500	15.3	18.1	13.1	88.9
750	13.4	14.8	12.5	61.3
1000	12.7	13.4	12.3	46.4
1700	12.2	12.2	12.3	29.9
2500	12.1	11.9	12.2	22.6
5000	12.1	11.9	12.2	15.0
7500	12.1	11.8	12.3	12.8
10000	12.1	11.9	12.3	11.9

At high numbers of flights, we did not reliably measure any time where there were no other aircraft close to the ownship. This suggests that we are underestimating π_0 for those cases and thus overestimating the time between encounters. Consequently the results for high numbers of aircraft in flight are likely significantly worse than our analysis shows.

5. Conclusions

We set out to find a metric that would allow us to determine when traffic in a region is "dense". We chose to follow the intuition that the effects of density would appear in how vehicles in flight interact with each other, and used collision avoidance as inspiration for defining "interaction".

Our first metric is the *minimum closing time:* how long does the ownship have before another vehicle might collide with it? We find that the traffic pattern—the general direction of flow of the other aircraft—affects this measure, with traffic on random headings resulting in much lower closing time than for traffic all headed in similar directions. This is largely due to the number of aircraft on nearly opposite headings from the ownship when traffic headings are random.

Our second metric is the *number of close aircraft*. This provides some insight into how often the ownship will avoid one aircraft only to have to maneuver again to avoid a different one, and how many other aircraft the ownship needs to be tracking. In other work we have found that collision avoidance systems reduce overall system safety if traffic goes above a threshold where one avoidance maneuver causes a cascade of other maneuvers in response; this metric helps show what regions and traffic patterns result in such situations. Once again, the degree to which all the traffic are flowing in the same direction affects this measure.

We found that both measures scale smoothly with the number of aircraft in flight. As a result, there is no obvious knee in a curve where we can declare a limit.

We also found that both heading and speed of the other aircraft affect the results, and that metrics based on closing time usefully combine these into one measure. This stands in contrast to the use of separation distance alone as the measure for determining when two aircraft are too close.

These results highlight the importance of certain kinds of order in traffic flows, such as organizing UAV traffic by heading. This suggests two conclusions: first, that it will be important to employ traffic management mechanisms in the airspace to reduce entropy in the traffic flow in volumes where traffic flows will be concentrated; and second, that the behavior of a dense air traffic flow is highly sensitive to perturbations in ordered flows—such as a wind gust, an intruding aircraft, or an aircraft that experiences a failure. Many ordered systems exhibit this kind of behavior, so this sensitivity is not a surprise: crystal structures, road traffic, and high-bandwidth network flows all exhibit sharp transition changes when disorder in a system crosses a threshold. This suggests that traffic management systems must take these sensitivities into account, and build resilience into traffic flows to absorb some amount of perturbation.

Each aircraft's collision avoidance system must also be sized to track several aircraft concurrently in urban settings. Except for when all traffic is going in the same direction, at the average moment an aircraft will be tracking at least one other nearby aircraft and often more. (We noted earlier that in the case of the Stream traffic model, the results likely underestimate the amount of interaction between aircraft and so may be overly optimistic.)

We also modeled how often aircraft are likely to be maneuvering, and how much of their flight time they will spend in such maneuvers. If we use a rule of thumb that no aircraft should spend more than 10% of its flight time in maneuvers, then only agile aircraft that can maneuver at high accelerations can meet this goal at traffic levels we expect to see over urban areas.

In the end, we conclude that "dense" operations occur at low absolute numbers of aircraft, if those aircraft are all on different headings. Applying a traffic management system that organizes traffic into more unidirectional flows helps, but even then traffic will be "dense" at levels below what we expect to see in urban areas. It will be necessary to deploy a traffic management system with capabilities for handling dense operations within a very few years' time.

Acknowledgements

Thanks to Mykel Kochenderfer at the Stanford Department of Aeronautics and Astronautics for help with reviews and error-finding.

References

[1] Air traffic services: Air traffic control service, flight information service, alerting service. Annex 11 to the Convention on International Civil Aviation, 13th edition, amendment 48, November 2012.

[2] U-space blueprint. Technical report, SESAR Joint Undertaking, 2017. Retrieved from https://www.sesarju.eu/u-space-blueprint

[3] *Introduction to TCAS II, Version 7.1.* U.S. Department of Transportation, Federal Aviation Administration, 28 Feb. 2011. Retrieved from https://www.faa.gov/documentLibrary/media/Advisory_Circular/TCAS%20II%20V7.1%20Intro%20booklet.pdf

[4] Kuchar, James K. and Ann C. Drumm. *The Traffic Alert and Collision Avoidance System.* Lincoln Laboratory Journal 16(2), 2007, pp. 277-96.

[5] Maki, Evan, Andrew Weinert, and Mykel J. Kochenderfer. *Efficiently estimating ambient near mid-air collision risk for unmanned aircraft.* Proc. 10th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, Sep. 2010.