Some references (including subsequent publications)

1. R. Clothier *et al*. *A Casualty Risk Analysis For Unmanned Aerial System (UAS) Operations Over Inhabited Areas*. 2007
2. K. Dalamagkidis, K. Valavanis, and L. Piegl. *On Integrating Unmanned Aircraft Systems into the National Airspace System Issues, Challenges, Operational Restrictions, Certification, and Recommendations*. Intelligent Systems, Control and Automation: Science and Engineering, Vol. 36. 2009
3. T. McGeer. *Safety, economy, reliability, and regulatory policy for unmanned aircraft*. Aerovel Corporation, March 2007. (9 pp)
4. T. McGeer. *Aerosonde hazard estimation*. The Insitu Group, 1994. (6 pp)
5. R. Weibel and J. Hansman. *Safety Considerations for Operation of Unmanned Aerial Vehicles in the National Airspace System*. MIT, 2006.

6. Quantitative Risk Management as a Regulatory Approach to Civil UAVs

Tad McGeer, PhD
The Insitu Group
Bingen, Washington
USA

Laurence R. Newcome, Lt. Col. Ret. USAF
Adroit Systems, Inc
Alexandria, Virginia
USA

Juris Vagners, PhD
University of Washington
Seattle, Washington
USA

Abstract

In this paper we argue for quantitative risk management as a regulatory approach to civil Uninhabited Aerial Vehicles (UAVs). The current prescriptive approach to airspace regulation, while appropriate to inhabited aircraft, does not realistically consider the issues of safely operating all UAV's in civil airspace. The wide range of missions for UAVs and the correspondingly diverse physical vehicles to realize these missions makes the prescriptive "one size fits all" approach to regulations inappropriate. Furthermore, there are significant differences between military UAVs and civilian UAVs to warrant separate consideration

Biographies

Tad McGeer – is president of The Insitu Group. He trained as an aeronautical engineer at Princeton and Stanford, an then moved into academic research on walking robots in the 1980s, on the faculty of Simon Fraser University in his native British Columbia. In 1990 he merged his robotic and aeronautical interests for initial studies of the Perseus and Theseus unmanned research aircraft, as Chief Scientist at Aurora Flight Sciences (Manassas, Virginia). He proposed the Aerosonde concept in 1991 and has since managed its development at Insitu.

Laurence R. "Nuke" Newcome - is Director of UAV Activities at Adroit Systems, Inc. of Alexandria, VA, where he is responsible for developing commercial reconnaissance markets for unmanned vehicles. He has worked with Federal Aviation Administration regional offices to obtain Certificates of Authorization for UAV flights in civil airspace on numerous occasions. Prior to joining Adroit, he served in the U.S. Air Force and coordinated deployments of the Predator, Gnat, and Pioneer UAVs into Europe, followed by a tour in the Defense Airborne Reconnaissance Office overseeing development of the Global Hawk and DarkStar high altitude endurance UAVs.

Juris Vagners – is a Professor of Aeronautics and Astronautics and Adjoint Professor of Electrical Engineering at the University of Washington (UW) in Seattle, WA. where he teaches control systems engineering. Prior to joining the UW he worked at Lockheed Missiles and Space Division in satellite guidance, control and navigation. He currently heads a research group at the UW working in conjunction with The Insitu Group and Aerosonde Robotic Aircraft on the development of second generation Aerosonde aircraft.

1 Introduction

The overarching goal of regulation of aerial vehicles is to ensure safe operation. To this end, regulation seeks to minimize the risks to 1) individuals in an aircraft, 2) individuals in other aircraft, and 3) individuals and property over which aircraft fly. Minimization is achieved by promulgating air worthiness, air operation, and air traffic standards, respectively.

Air worthiness standards ensure aircraft are constructed for safe and reliable operation, air operation standards ensure that pilots and mechanics are trained and remain proficient to a common level, and air traffic standards ensure that aircraft are channeled in time, altitude, and geography to reduce the risk of midair collisions and to the risk to individuals on the ground.

In the United States (as well as elsewhere in the world), specific versions of these standards have been developed for air carriers (passenger and freight), general aviation, helicopters, homebuilt aircraft, gliders, and lighter-than-air craft, but not for uninhabited aerial vehicles (UAVs), or in US FAA terminology, remotely operated aircraft (ROAs). An air traffic standard (7610.4) for military (not civil) ROAs exists and Notice N7610.71 issued by the FAA 3/19/99 (Ref 1) makes this standard effective May 1, 1999. The civil version of this standard awaits the manifestation of a need (i.e., a commercial market) for ROA traffic to bring it to life. An air operation standard for ROA pilots/operators is in the discussion phase in the US; it reportedly will not restrict future ROA operations to the exclusive domain of pilots, even if the US Air Force elects to do so. Air worthiness standards for ROAs do not exist, although their manned counterparts are being applied selectively by ROA manufacturers to their products.

In Europe, a comprehensive working group has been set up under the NATO Committee for European Airspace Coordination, which we believe is addressing the issues of Air Traffic Control (ATC). In the United Kingdom, the 970 UAV-Sub-Committee is working on amendments of Def Stan 00-970 to frame a UAV requirements guide. In addition to this group, a new specialist industrial association, the Unmanned Aerial Vehicle Systems Association, has been formed in the UK. This group is specifically looking at Civil Airworthiness requirements. An excellent overview of safety considerations for Uninhabited Combat Air Vehicles (UCAVs) is given in Ref 2. In the reference, the safety relationships within the UCAV environment are highlighted and a comprehensive chart detailing a comparison of the safely issues for UCAVs and inhabited aircraft presented. The chart compares these issues based on operational profiles and common risk factors.

A sense of the broad range of issues covered under flight safety for UAV's intended for military applications can be gained from presentations at the recent 14th Annual UAV Systems Conference held in April, 1999 at Bristol, United Kingdom. The civilian sector has not been addressed with anywhere near the same degree of specificity. While many of the issues are common to both military and civilian UAVs, there are significant differences that should be accounted for in the development of regulations for civil UAVs.

Some risk issues simply do not exist for civilian UAVs. For example, all hazards associated with military stores do not apply, nor those associated with battlefield operations. Others are common to all UAVs, but are not all equally significant for all aircraft, for example, different mission profiles and physical characteristics of the vehicles mitigate some hazards and emphasize others. Thus, a regulatory approach that recognizes such factors should be adopted. One proposed approach is to establish UAVs and UCAVs as a separate class of vehicle much in the same way that helicopters, balloons and gliders are. We argue that this is still too broad to realistically encompass civilian vehicles, and in particular, small civilian vehicles that are intended for use as remote sensing platforms. An approach that recognizes these issues would be consistent with the historical evolution of the air vehicle certification process which recognizes that there is an intimate relationship between the context in which the vehicle is used, the integrity of the vehicle itself and the interactions of the operators of the vehicle and the supporting ground systems such as ATC. In the following sections we examine the implications of such an approach.

2. Hazard Estimation

Principal sources of potential hazards to individuals and property from UAVs arise from 1) the flight integrity of the vehicle itself, 2) midair collisions, and 3) collision with persons or objects on the ground, whether by the entire aircraft or pieces of it following in-flight mishaps. The hazard analysis for inhabited aircraft is driven by the fact that the probability of injury due to potential loss of aircraft is much higher than that due to all other sources. Analysis thus emphasizes probable loss of aircraft due to functional failures, rather than the estimation of ground casualties. Experience has shown that the in-aircraft injuries far outweigh injuries on the ground. Removal of the crew from the aircraft changes this scenario in that now sources 2) and 3) become the drivers of hazard estimation. If airworthiness standards are now also meant to protect those on the ground from frequent falling debris, then the associated probability of occurrence must be examined.

2.1 Military UAVs

Let us first discuss these issues for large UAVs in the military sector. We will address the operations of civilian UAVs separately since these have been much more limited in scope. Since the primary focus of this paper is on civilian UAVs, we will be brief in this discussion. The U.S. military operates some 125 ROAs on a regular basis and, being generous, U.S. manufacturer's flight activities may raise this number to 200, of which perhaps 10 percent, again being generous, may fly on any given day. If all 20 of these ROAs were the largest ROA for which at least 20 flying examples exist, the RQ-1/Predator, which has an area of 200 square feet and a weight of 2300 lbs, the total area at risk below a fleet of 20 falling Predators is some 4000 square feet. The density of the 265 million U.S. citizens spread evenly over the 3.5 million square miles of the U.S. is some 76 people per square mile, or 368,000 square feet per individual, making the odds of an individual being hit by a falling UAV 0.000543, or 1 in 1840. Since its first flight in July 94, six Predators have crashed (not counting two that have been shot down), for an averaged loss rate of one every 9.3 months (as of Feb 99); this interval has actually grown to 18 months for the most recent losses. (Similar mishap data rates are available for Navy RQ-2/Pioneer, Army RQ-5/Hunter, and other large military UAVs). Using the former interval means a falling Predator could be expected to hit a person on the ground once every 1431 years; for the latter, more recent interval, this rate becomes once every 2760 years. This is a conservative estimate because the simplifying assumptions used in reaching it were obviously conservative, such as the U.S. population being uniformly distributed instead of clumped together in metropolitan areas, to name one. But whether the probability of occurrence is once a millennium or once every other millennium per fatality, the risk to people on the ground from the current level of UAV activity does not currently appear to justify air worthiness standards for UAVs.

Air operation standards holding "remote operators" to the same standards as general aviation pilots would pose no hurdle to present and future UAV expansion into the national airspace. Then again, some of the medical and physical limitations placed on pilots would fail the reasonableness test if applied to their UAV counterparts. Example: Sinus blocks should not be disqualifying for that day's UAV flight when its pilot will see no pressure change due to altitude. Indeed, paraplegics could safely navigate all but the few UAVs which replicate rudder pedals in their ground cockpits. The premise that a remotely located pilot would not have the same vested interest in steering his or her disabled aircraft away from populated areas that an onboard pilot would is fallacious on a number of counts. First, the remote pilot is free to focus on salvaging some kind of landing without having to share his or her attention with thoughts of ejection or impact. Second, short of a catastrophic failure, the need to maintain line-of-sight communication with the UAV dictates turning the UAV toward base and/or the highest ground near it in order to maintain control as long as possible. Third, there have already been numerous instances in Bosnia in which the pilot of a "non-returnable" UAV safely and successfully steered the aircraft into controlled crashes in unpopulated areas. The FAA seems to be adopting the general criteria of having "familiarity with operations in the national airspace" as its qualifying standard for future UAV flyers.

Development of air traffic standards for UAVs, however, unlike the two previous regulations, are crucial to their gaining acceptance in the current civil airspace structure. These standards play a large part in minimizing the risk to people and property on the ground, as well as reducing the chances of a midair, by providing aerial highways, altitude floors, and special use bubbles of airspace. Two capabilities, inherent in manned aircraft operations, must be built into UAVs to allow them to operate in this structure with an equivalent level of safety: These are: voice relay, so that the UAV pilot can hear, acknowledge, and respond to directions from air traffic control centers immediately as though he or she were onboard, and "see & avoid," the ability to scan the surrounding airspace for approaching aircraft and take action to avoid a midair. Voice relay is built into, or being retrofitted on, all Predator-size and larger military UAVs, and, as technology reduces costs, will work its way down into the smaller, tactical UAVs. What constitutes the undefined concept of "see & avoid" (there is no formal FAA definition of it) is still being discussed by UAV proponents and the FAA.

Today, regional offices of the FAA issue Certificates of Authorization (COAs) for ROAs to operate within specified areas for periods of up to one year as long as they conform to certain, manned aircraft-like weather, safety, and operating limitations. Although requests for COAs are required at least 60 days prior to flight, FAA responsiveness experienced to date has been from 1 to 10 days. For the near future, Traffic Conflict Alert Systems (TCAS) appear to be part of the answer, perhaps in combination with the reconnaissance UAV's electro-optical or infrared mission sensors. Of course, using these sensors for see & avoid detracts from the mission of the UAV in the first place, scanning the ground vice the surrounding airspace. In the far future, the Global Air Traffic Management (GATM) system is to migrate away from IFF-based, ground-centralized, situational awareness to GPS-based, aircraft-localized situational awareness. In such an environment, the UAV would, capability-wise, be an equal player with its manned counterparts. Bridging the airspace regulatory environment until that time is the challenge.

2.2 Civilian UAVs

To estimate hazards for civilian UAVs we need to rely on standard probability methods as comparatively little data exists on both numbers of aircraft and flight hours. Specifically, we need to estimate the probability of hitting a manned aircraft in flight or an individual or property on the surface. The details of the analysis are given in the appendix and the interested reader can examine these probabilities for any vehicle of interest using the formulae in the appendix or modifying them to fit the scenario of interest.

Since it appears that the most likely civilian UAV applications in the near future will involve small, low speed, low weight vehicles for remote sensing, we focus on this class of vehicles. Of this class, the Aerosonde (see web sites Ref 3,4) has reached the highest level of maturity and accumulated the most flight hours to our knowledge. The Aerosonde was developed by The Insitu Group and Environmental Systems and Services of Australia (responsibility now with Aerosonde Robotic Aircraft, Ltd of Australia) under the sponsorship of the Australian Bureau of Meteorology, US Office of Naval Research, US Department of Energy, Environment Canada, US National Weather Service and the Taiwan Central Weather Bureau. This list of sponsors also serves to support the expectation that the Aerosonde will be the first civilian UAV in significant numbers to enter service on a regular basis for specific meteorological missions. All Aerosonde flights to date have been under mission-specific Special Flight Operations Certificates issued by the appropriate agency in the country (or countries) in which the trials were conducted.

We first wish to emphasize that the Aerosonde development is driven the economics of the application to remote meteorological sensing. This means that aircraft reliability and system complexity (hence cost) are in constant trade-off. This also means that the Aerosonde development has proceeded hand-in-hand with field trials, thus sustaining higher aircraft losses than one would accept in routine service. In 1998, we lost 8 aircraft in about 400 flight hours, with 3 out of 4 aircraft lost in the ultimately successful, historic North Atlantic crossing attempt (Ref 5,6). However, the responsible technical faults are well understood and are being fixed. For example, in the most recent field trials in Hawaii in May of 1999 two Aerosondes flew approximately 70hrs in eight flights with no lost aircraft.

Turning now to the results of our analysis, we first consider the probability of a mid-air collision of an Aerosonde with other traffic (appendix section A.1). Consider first the case of encountering Boeing 747s as an example of large targets. In this case, the applicable collision frequency is 10^{-9} per flight-hour, this being the maximum rate of catastrophic failure considered acceptable by the US FAA. Using 747 parameters of frontal area $\Phi_t = 2x10^{-4}$ km^2, and average speed $V_t =$ 860 km/hr, the corresponding Aerosonde density is $\rho_s = 10^{-8}$ per cubic kilometer.

Next, consider the case of general aviation. Here we set the allowable collision frequency at 10^{-7} per flight-hour, which is closer to the historical rate achieved by the "see-and-avoid" paradigm. The average speed and frontal area parameters are V_t = 200 km/hr and $\Phi_t = 10^{-5}$ km^2. In this case, the corresponding Aerosonde density is $\rho_s = 10^{-4}$ per cubic kilometer. The allowable Aerosonde density in this case is significantly higher, but also significantly higher than anything one would expect in practice, so the probability of collision would still be negligible by current standards, even if nothing were done about avoidance or appropriate Aerosonde distributions. For example, decreased Aerosonde distribution in areas of transiting aircraft such as oceanic tracks can be easily arranged since airways and other busy airspace are well defined over regions of interest to the operators of Aerosondes. Aerosondes can be programmed to avoid such areas both laterally and vertically. Note that this entails no requirement for the avoidance strategy to be perfectly reliable – even if it worked only 90% of the time, it would reduce Aerosonde density tenfold in areas of concern.

Let's next consider crash hazards (appendix section A.2). In this case we first consider the hazard to ships since a highly likely scenario of operations will be over the oceans. For illustration, consider the hazard to ships

arising from Aerosondes doing meteorological reconnaissance over the high seas, e.g. on a transatlantic or transpacific flight. Rough numbers are target density per unit area $\sigma_t = 4x10^{-4}/km^2$ (assuming 10^5 ships randomly distributed over the oceans), target length $l_t = 0.1$ km, averaged over all ship sizes and orientations, target span $b_t = l_t$ with all orientations being equally likely, and average crass frequency $f_i = 10^{-3}$ per flight-hour. Then average collision f_c is about
$4x10^{-9}$ per flight-hour.

Meteorological requirements ultimately may entail about 10^6 annual Aerosonde hours in oceanic reconnaissance. At this rate ships would be hit on the average once every 250 years. Actually, as a hazard estimate this is pessimistic: the probability of seriously *damaging* a ship with a 13 kg Aerosonde would be a good deal smaller. But $4x10^{-9}$ per flight-hour is already small enough to be negligible.

Note that the hazard probability in this case is very much less than the aircraft crash rate – a situation obviously different from that in manned aircraft as we noted earlier! The hazard probability becomes comparable with the crash rate only if the target density is high. Thus as the ship example illustrates, reliability requirements can be substantially relaxed if operations are planned to avoid high density areas.

Let's consider then a flight-plan leg designed to keep the aircraft over reasonably sparsely populated terrain and estimate the probability of hitting a house. There will be some error in tracking the leg and the aircraft may overfly a few bystanders. To illustrate, consider the case for a typical Aerosonde and a typical house. The numbers are: Aerosonde speed $V_s = 80$ km/hr, standard tracking deviation $\sigma_y = 0.05$ km (consistent with flight experience to date), flight path angle $\gamma = 1/20$ (at best L/D, hence conservative for most failures), width of house $b_t = 0.03$ km (typical, not Microsoft, house), height of house $h_t = 0.006$, average crash frequency $f_i = 10^{-3}$ per flight-hour and target cross track position $y_t = 3\sigma_y$. With these numbers the probability of a strike turns out to be about $4x10^{-9}$. On average, one of every 200 million such bystanders passed would be hit.

This result might be questioned on the basis that some failures, such as that of the flight computer, would cause loss of tracking performance. In that case a deadman's switch would kill the engine, but the aircraft would then crash, with equal probability, anywhere within gliding range.

Suppose we want to keep the average strike rate below 10^{-7} per flight-hour, which seems a reasonable guess at the present rate for inhabited aircraft. What restrictions must be imposed on the areas overflown? We use the same numbers as the last example, except with a failure rate of 10^{-4} per flight-hour rather than 10^{-3} per flight-hour, because we are accounting only for events that cause uncontrolled departure from track. The factor of ten is a minimum requirement dictated by economics of Aerosonde operations. In estimating the costs of meteorological reconnaissance by Aerosondes, we presume that most attrition will be caused by adverse weather conditions. If systems failures were to cause attrition at a comparable level, then economics could be improved by making the design more reliable. Hence for minimum cost the systems failure rate must be made small compared to the overall loss rate, i.e. no worse than 10^{-4} per flight-hour. At this rate the maximum allowable σ_t turns out to be about 1 house per square kilometer. Obviously this means that overland operations must be conducted in remote surroundings – but then economical access to such areas is the whole purpose of the Aerosonde project!

3 Conclusions

In this paper we have highlighted some of the features of quantitative risk management as a regulatory approach to UAVs. This approach is most appropriate for civil UAVs in that their operations will, at least in the near future, be strongly mission oriented and the characteristics of these missions are easily defined. As shown by the results for the Aerosonde, for small, light, low speed UAVs operated outside densely populated areas, the hazard probabilities are extremely low.

4 References

1. “Department of Defense (DOD) Remotely Operated Aircraft (ROA) Operations”, Notice N 7610-71, US Dept of Transportation, Federal Aviation Administration, March 19, 1999

2. “Flight Safety Considerations for UCAV Aircraft”, R.C. Wells, D.J. Hamlin, R.B. Smith, Proc. UAVs 14th Int. Conf., Bristol, UK, April 12-14, 1999

3. http://www.insitugroup.com

4. http://www.aerosonde.com

5. “Aerosonde Operations in 1998”. T. McGeer, G. Holland, G. Tyrrell, J. Becker, J. Vagners, P. Ford, Proc. Third Symposium on Integrated Observing Systems: 60-63. American Meteorological Society, Dallas, January 1999

6. “Historic Crossing: An Unmanned Aircraft’s Atlantic Flight”, T. McGeer, J. Vagners, GPS World 10 (2), Feb 1999

Appendix: Aerosonde Hazard Estimation

We are interested in estimating the probability of an Aerosonde hitting a manned aircraft in flight, or an individual or property on the ground. In this appendix we work through the analyses and example calculations for each case. We provide the general specifications for the current generation Aerosonde used to illustrate the calculations.

Aerosonde Mark I General Specification

Mission: Long-range meteorological reconnaissance and environmental monitoring

Wing span: 2.9 m

Wing area: 0.55 sq m

Overall length: 1.7 m

Overall height: 0.6 m

Body depth: 0.19 m

Typical empty weight: 8.2 kg

Maximum launch weight: 13.4 kg

Maximum fuel capacity: 5.0 kg (7 litres)

Powerplant: modified Enya R120

Type: single-cylinder, four-stroke, air-cooled piston engine

Manufacturer: ES&S, Melbourne, Australia

Rating: 0.75 kW

Fuel: Avgas 100LL

Propellor: two-blade, fixed-pitch, 51cm diameter

Electrical power: 10 W typical

Ambient temperature range -10C to +40C

Max level speed: 56 kt

Cruise Speed: 40 kt

Loiter Speed: 40 kt

Max S/L climb @ max TOW: 2 m/s

Service ceiling: 4500 m

Still-air range: 1500 nm

Max endurance: 32 hr in typical cruise profile

Launch: From a car roof, using a cradle that fits most vehicles. Takeoff is normally flown manually by an outside pilot. Automatic takeoff capability has been demonstrated.

Recovery: Belly landing. The aircraft has no undercarriage. Landings are normally flown manually but autoland has been demonstrated.

A1. Midair collision

To formulate the problem, first consider conflict in 2D. The extension to 3D follows readily from the 2D case since aircraft fly at shallow angles so the principal component of converging velocity will remain horizontal. We assume that Aerosondes are operating randomly in a plane. You want to know the probability of collision if you fly across the plane. A collision will occur if there is an overlap of the areas swept out by your aircraft and that swept out by the Aerosonde in time dt. In time dt you sweep through area dA

$$dA(\psi) \approx b_t \sqrt{V_t^2 + V_s^2 + 2V_t V_s \cos\psi}\, dt$$

where V_t is your speed, b_t is your wingspan, V_s is the Aerosondes' speed and ψ is the crossing angle. In this equation we have neglected the Aerosondes' wingspan (small in comparison to anything else up there!). To include it, just replace b_t by (b_t+b_s). Assuming that traffic moves in random directions so that all values of ψ are equally likely, we average over all possible directions to find

$$d\bar{A} \approx \frac{1}{2\pi} \int_0^{2\pi} dA(\psi) d\psi$$

$$d\bar{A} \approx b_t V_t dt \left(\frac{1}{2\pi} \int_0^{2\pi} \sqrt{1 + \left(\frac{V_s}{V_t} \right)^2 + 2\frac{V_s}{V_t} \cos\psi}\, d\psi \right)$$

From this we see that if the ratio of the Aerosonde speed to that of the encountered traffic $\frac{V_s}{V_t}$ increases from zero to 0.5 for example, the bracketed coefficient increases from one to about 1.3. For order of magnitude estimation this change is negligible, so one is justified in using the simpler formula

$$dA \approx b_t V_t dt$$

The probability of collision during the time interval dt for Aerosonde density per unit area of σ_s is

$$P_c(dt) = 1 - e^{-\sigma_s b_t V_t dt} \approx \sigma_s b_t V_t dt$$

with the approximation for $\sigma_s b_t V_t dt << 1$. Hence the probability of collision per unit time is

$$f_c = \frac{dP_c}{dt} = \sigma_s b_t V_t$$

To extend this relationship to the 3D case, we need to replace aerial density σ_s by volumetric density ρ_s and target width b_t by target frontal area Φ_t. Thus f_c becomes

$$f_c = \rho_s \Phi_t V_t$$

We can now evaluate collision frequencies for randomly distributed Aerosondes given assumed densities. Conversely, we can solve this equation for the density of Aerosondes leading to specified collision frequencies, as would be set by regulatory agencies.

Consider first the case of encountering Boeing 747s as an example of large targets. In this case, the applicable collision frequency is 10^{-9} per flight-hour, this being the maximum rate of catastrophic failure considered acceptable by the US FAA. Using 747 parameters of $\Phi_t = 2\text{x}10^{-4}$ km^2, $V_t = 860$ km/hr , the corresponding Aerosonde density is $\rho_s = 10^{-8}$ per cubic kilometer.

Next, consider the case of general aviation. Here we set the allowable collision frequency at 10^{-7} per flight-hour, which is closer to the historical rate achieved by the "see-and-avoid" paradigm. The average speed and frontal area parameters are $V_t = 200$ km/hr and $\Phi_t = 10^{-5}$ km^2. In this case, the corresponding Aerosonde density is $\rho_s = 10^{-4}$ per cubic kilometer. The allowable Aerosonde density in this case is significantly higher, but also significantly higher than anything one would expect in practice, so the probability of collision would still be negligible by current standards, even if nothing were done about avoidance or appropriate Aerosonde distributions. For example, decreased Aerosonde distribution in areas of transiting aircraft such as oceanic tracks can be easily arranged since airways and other busy airspace are well defined over regions of interest to the operators of Aerosondes. Aerosondes can be programmed to avoid such areas both laterally and vertically. Note that this entails no requirement for the avoidance strategy to be perfectly reliable – even if it worked only 90% of the time, it would reduce Aerosonde density tenfold in areas of concern.

A.2 Crash Hazards

Aircraft crash from time to time and thus constitute a hazard to innocent bystanders on the surface. The fundamental principle that the general public must be protected from any form of flying machine has remained the constant driver of safe operation. No matter what the design and purpose of the aircraft, it must be certified and operated with this principle in mind and it is certain that certification and operation of UAVs will follow the same principle.

Suppose that the average frequency of crashes is f_i. Then the probability of a crash in any interval of time dt is

$$P_i(dt) = 1 - e^{-f_i dt} \approx f_i dt \quad where\ f_i dt << 1$$

Thus the probability of crashing in the interval required to cross a “target” of length l_t would be

$$P_i(l_t) = f_i \frac{l_t}{V_s}$$

Meanwhile, the probability of such a target actually being in the flight path during this interval is the product of the combined width of the aircraft and the target ($b_s + b_t$), length along the track $V_s dt$, and the average density of targets on the surface σ_t. The overall probability of a strike, per unit of flight time, thus is

$$f_c = \left(f_i \frac{l_t}{V_s} \right) (b_s + b_t) V_s \sigma_t = f_i l_t (b_s + b_t) \sigma_t$$

For illustration, consider the hazard to ships arising from Aerosondes doing meteorological reconnaissance over the high seas, e.g. on a transatlantic or transpacific flight. Rough numbers are $\sigma_t = 4\text{x}10^{-4}$ per km^2 (assuming 10^5 ships randomly distributed over the oceans), $l_t = 0.1$ km, averaged over all ship sizes and orientations, $b_t = l_t$ with all orientations being equally likely, $b_s << b_t$, and $f_i = 10^{-3}$ per flight-hour. Then f_c is about $4\text{x}10^{-9}$ per flight-hour.

Meteorological requirements ultimately may entail about 10^6 annual Aerosonde hours in oceanic reconnaissance. At this rate ships would be hit on the average once every 250 years. Actually, as a hazard estimate this is pessimistic: the probability of seriously *damaging* a ship with a 13 kg Aerosonde would be a good deal smaller. But $4\text{x}10^{-9}$ per flight-hour is already small enough to be negligible.

Note that the hazard probability in this case is very much less than the aircraft crash rate – a situation obviously different from that in manned aircraft! The hazard probability becomes comparable with the crash rate only if the target density is high. Thus as the ship example illustrates, reliability requirements can be substantially relaxed if operations are planned to avoid high density areas.

Let’s consider then a flight-plan leg designed to keep the aircraft over reasonably sparsely populated terrain. There will be some error in tracking the leg and the aircraft may overfly a few bystanders. We can calculate the associated hazard as follows.

Take the tracking error to be Gaussian with standard deviation σ_y. The probability of crossing a bystander of width b_t at a distance y_t from the track centerline is

$$p(y_t) = \frac{1}{\sigma_y \sqrt{2\pi}} \int_{y_t-(b_t+b_s)/2}^{y_t+(b_t+b_s)/2} e^{-1/2(y/\sigma_y)^2} dy$$

and for $b_t + b_s << \sigma_y$ this becomes

$$p(y_t) \approx \frac{b_t + b_s}{\sigma_y \sqrt{2\pi}} e^{-1/2(y_t/\sigma_y)^2}$$

The probability of a strike is then

$$P_c \approx f_i \frac{l_t}{V_s} \frac{b_t + b_s}{\sigma_y \sqrt{2\pi}} e^{-1/2(y_t/\sigma_y)^2}$$

If the bystanders are high rather than long, so as to be more likely to be hit from the side rather than from above, we approximate l_t by the ratio h_t/γ with γ being the flight path angle.

To illustrate, consider the case for a typical Aerosonde and a typical house. The numbers are: $V_s = 80$ km/hr, $\sigma_y = 0.05$ km (consistent with flight experience to date), $\gamma = 1/20$ (at best L/D, hence conservative for most failures), $b_t = 0.03$ km (typical, not Microsoft, house), $h_t = 0.006$, $f_i = 10^{-3}$ per flight-hour and $y_t = 3\sigma_y$. With these numbers the probability of a strike turns out to be about 4×10^{-9}. On average, one of every 200 million such bystanders passed would be hit.

This result might be questioned on the basis that some failures, such as that of the flight computer, would cause loss of tracking performance. In that case a deadman's switch would kill the engine, but the aircraft would then crash, with equal probability, anywhere within gliding range. Hence the crash radius is $R_i = h/\gamma$. The bystander is at risk if the failure occurs anywhere on a flight segment of length

$$l_i = 2R_i\sqrt{1-\left(\frac{y_t}{R_i}\right)^2}$$

The probability of failure on this segment is given by $P_i(l_i) = f_i\frac{l_t}{V_s}$. If the failure occurs, then the bystander's probability of being struck is just his or her fraction of the affected area i.e.

$$\sigma_t = \frac{(l_t + b_s)(b_t + b_s)}{\pi R_i^2 + 2R_i l_i}$$

Hence the overall probability of striking the bystander is

$$P_c(y_t) = 2f_i\frac{(l_t + b_s)(b_t + b_s)}{V_s R_i}\frac{\sqrt{1-\left(\frac{y_t}{R_i}\right)^2}}{\pi + 4\sqrt{1-\left(\frac{y_t}{R_i}\right)^2}}$$

assuming then that bystanders are randomly distributed across the track, with average areal density σ_t, the average strike probability for all bystanders is

$$\overline{P}_c = 2f_i\frac{(l_t + b_s)(b_t + b_s)}{V_s R_i}\left[\frac{1}{2}\int_{-1}^{1}\frac{\sqrt{1-\bar{y}^2}}{\pi + 4\sqrt{1-\bar{y}^2}}\,d\bar{y}\right]$$

$$\overline{P}_c == 0.24 f_i\frac{(b_t + b_s)}{V_s R_i}$$

The average number of bystanders at risk per unit time is

$$\frac{dN}{dt} = 2V_s R_t \sigma_t$$

and the average rate of bystander strikes is

$$f_x = 0.48 f_i (l_t + b_s)(b_s + b_t)\sigma_t$$

This is essentially the same result as we had earlier. Note that the altitude doesn't appear, except indirectly in the sense that the higher the altitude, the wider the corridor over which the bystander density must be calculated.

Suppose we want to keep the average strike rate below 10^{-7} per light-hour, which seems a reasonable guess at the present rate for inhabited aircraft. What restrictions must be imposed on the areas overflown? We use the same numbers as the last example, except with a failure rate of 10^{-4} per flight-hour rather than 10^{-3} per flight-hour, because we are accounting only for events that cause uncontrolled departure from track. The factor of ten is a minimum requirement dictated by economics of Aerosonde operations. In estimating the costs of meteorological reconnaissance by Aerosondes, we presume that most attrition will be caused by adverse weather conditions. If systems failures were to cause attrition at a comparable level, then economics could be improved by making the design more reliable. Hence for minimum cost the systems failure rate must be made small compared to the overall loss rate, i.e. no worse than 10^{-4} per flight-hour. At this rate the maximum allowable σ_t turns out to be about 1 house per square kilometer.