The Technology Partnership
The space to invent.

White Paper

Ready for take-off?

Towards Detect and Avoid for Commercial UAVs

Commercial drones, and the business models they enable, will truly take off when fleets of uncrewed aerial vehicles (UAVs) of many shapes and sizes can operate autonomously in shared airspace. For this to happen, UAVs need to be able to fly beyond the visual line of sight (BVLOS) of remote human pilots, and the technologies that make this possible need to come an at acceptable size, weight, power and cost (SWaP-C).

Here, we focus on the toughest technological challenge on the way to UAV autonomy: last resort tactical airspace deconfliction by means of onboard sensing, also known as Detect and Avoid (DAA). We consider the safety goal posts, what can be achieved with different sensors and data processing techniques, and whether we can meet industry needs for SWaP-C and chart a course towards certification.

Introduction

From delivering vital medical supplies in regions where other infrastructure is lacking to transporting people in air taxis, UAVs are here to stay. But to move beyond the current waiver-based operations and unlock new business models, UAVs need to be able to fly BVLOS in a safe and scalable manner.

We have previously introduced the different technical puzzle pieces needed to make this happen. Ultimately, the ability to mitigate risk hinges on conflict resolution in shared airspace. However, we, the industry, need to solve this problem in a manner that makes adoption of solutions easy. This implies solutions that not only meet performance requirements but do so within the payload limitations of different UAVs and come at an appropriate cost point.

How do we avoid airspace conflict?

The three layers of airspace conflict resolution are:

- Procedural resolution: plan the operation with sufficient separation from other airspace users;
- Air traffic management: use a third-party service whose subscribers share data to improve mutual awareness and communication with other airspace users; and
- **Detect and avoid (DAA): last resort tactical deconfliction.**

Last resort tactical deconfliction is a crucial but challenging technological hurdle consisting of three separate problems – **detection, alerting and avoidance.** In this whitepaper, we focus on **detection** since this remains the bottleneck in terms of meeting SWaP-C and performance.

The market offering for avoidance is relatively mature. Several "avoid" technologies compliant with regulation are available. The alerting function has some interesting challenges for certain use cases such as optionally piloted vehicles (OPVs). What and how should information be displayed to the pilot? There is also interest in learning from heads-up display (HUD) technology in the automotive sector, and how this may transfer to electric vertical take-off and landing (eVTOLs). For the rest of this paper, we focus on detection.

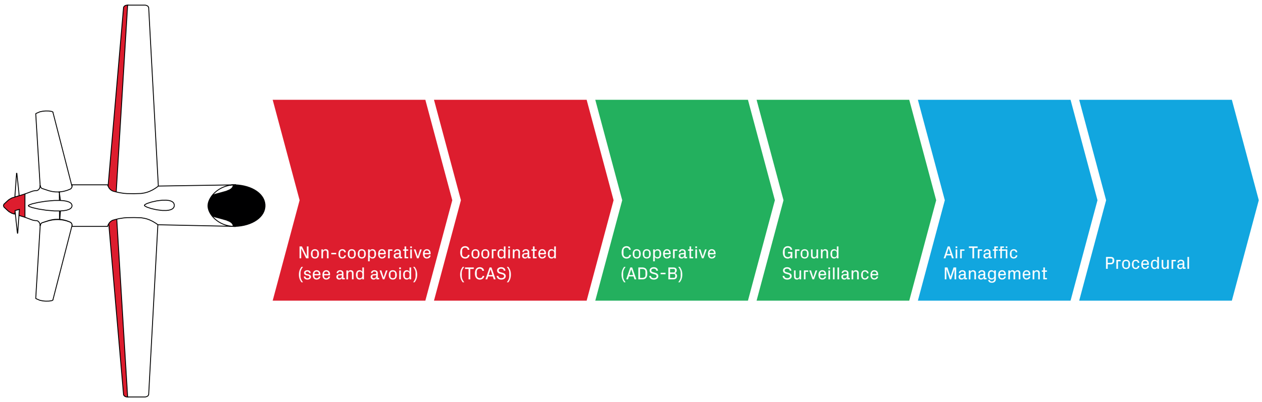

Figure 1 An illustration of the various means to resolve conflict in airspace[1].

First, let's take a step back and set the safety goal posts.

---

1. Adapted from Taxonomy of Conflict Detection and Resolution Approaches for Unmanned Aerial Vehicle in an Integrated Airspace

How is safety defined?

One of the challenges in developing DAA systems is uncertainty about the goal posts. NASA considers statistical metrics (termed as lagging indicators) and performance metrics (termed as leading indicators)[2]. Examples of statistical metrics include safety targets such as number of fatal accidents per flight hour set by the International Civil Aviation Organisation (ICAO), or general aviation mid-air collision statistics. Statistical measures such as these are largely unsuitable as development goal posts because they are not representative of the type of encounters that UAVs will experience and will not scale as the skies fill up with UAVs.

Performance based metrics are more suitable for defining goal posts. They are derived from currently accepted levels of aircraft safety. This is done by (a) defining encounter models of uncrewed-crewed and uncrewed-uncrewed traffic and (b) quantifying the physical and mental limitations of a human pilot. This has now formed the basis of the aircraft collision avoidance standards for small aircraft (ACAS-sXu)[3] published by ASTM International in 2020.

The ACAS sXu specification uses cylindrical boundary boxes around a UAV to define airspace boundaries for 'near mid-air collision' and 'well clear'. The objective of any DAA system is to keep the probability of detecting another airspace user within the cylindrical boundary boxes under a certain threshold. This, in conjunction with the encounter models, provides a basis for determining the detection range vs. the angular rate accuracy required from a DAA system.

Let’s look at some examples to get some tangible numbers.

Take a small cargo drone with a maximum take-off weight (MTOW) of about 40 pounds and cruising speed of about 40 knots. Most UAVs in this class are designed to operate autonomously with minimum pilot intervention – i.e., whilst a collision avoidance operation needs supervision from a remote pilot and can be overridden by them, the pilot wouldn’t need to execute the manoeuvre. Banking angle and turn rates can depend on whether the aircraft is multi-rotor or fixed-wing. Based on the encounter model statistics and risk ratio requirements in ACAS-sXu, the DAA system would need to achieve a detection range of at least 1.5 km with an angular rate accuracy of 0.7 degrees/second.

Note that this is the range needed to detect crewed aircraft and does not consider “unconventional” airspace users like hang gliders and hot air balloons. They are unconventional now, but it is not far-fetched that they might become more common and would need to be detectable (albeit a detection range of less than 1.5 km would be acceptable for slower moving airspace users).

The other question is angular coverage. The Federal Aviation Administration insists on 360-degree azimuthal coverage. In other countries, Mexico for example, a front facing coverage, as for a crewed airplane, is considered sufficient.

---

2. Sense-and-avoid equivalent level of safety definition for unmanned aircraft systems, revision 9, NASA, 2005
3. Standard Specification for Detect and Avoid System Performance Requirements, ASTM International, 2020

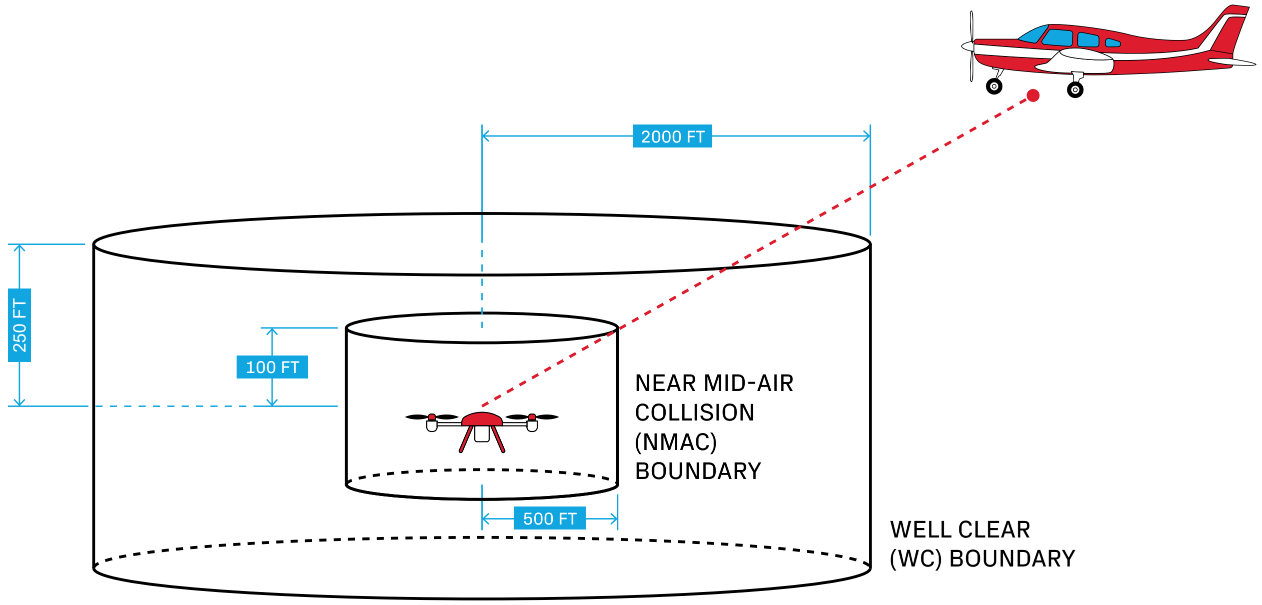

Figure 2 The NMAC and WC boundaries used to benchmark the performance of DAA systems[4]

What does the industry need?

We talked to stakeholders to find out. Our conversations suggest that the requirements for detection systems vary, depending on use case and MTOW.

At the lower end of MTOW – corresponding to UAVs used for inspection and small cargo delivery (both intra-city and long distance) – companies were very sensitive to SWaP-C. A wish list of weight < 1 kg, power consumption < 10-20 W, cost ~ $1000s, and detection range in the order of 1 to 2 km was typical.

Developers of larger, cargo-carrying UAVs and air taxis, whilst also sensitive to cost and size were less sensitive to weight. However, amongst these, those flying in controlled airspace require airborne solutions with much longer range (2 to 3 km) since the closing speeds would be very high.

What is possible?

The toughest technology challenge is undoubtedly an onboard DAA solution for low MTOW aircraft and any solution would need to push the limits of sensing, deep software, SwaP-C, and performance. Based on our conversations with stakeholders, we set ourselves a target of range of detection of greater than 1 km at an aggressive SWaP-C budget of 500g, 10W and BoM $1000.

---

4. Adapted from ACAS-sXu specification

We need to resolve:

- How close to that can we get?
- What is possible in terms of performance for different sensing modalities and SwaP-C?
- Are we also able to demonstrate a route to certification?

To address these questions, we took a ground-up approach by looking at the minimum sensor specification that met performance needs.

Visual techniques

Visual approaches are a good candidate for low MTOW UAVs. We have developed and prototyped several architectures and algorithms to estimate what is possible.

One architecture employs a combination of four wide field of view (FOV) cameras to provide 360-degree coverage and a steerable narrow FOV camera. Output from the wide-angle cameras is used to estimate motion and classify pixels as either background or another airspace user. In the video (Video 1), this classification technique is used to detect hang gliders from a single low-resolution GoPro camera. This approach detects and tracks hang gliders with few false positives.

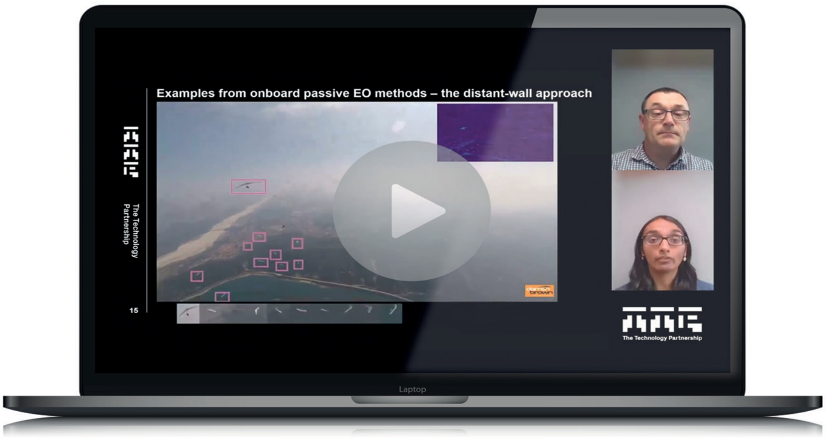

Examples from onboard passive EO methods – the distant-wall approach

Video 1 Demonstration of approach where estimation of motion is used to classify pixels as background or airspace user.

Other, lower complexity architectures for motion estimation could use the drone's IMU data. The airspace user detection stage is used to steer the narrow FOV camera, the output of which is fed to an object classification algorithm. Video 2 shows an object classification machine learning (ML) model trained on a small dataset of helicopters seen from the ground. Inputs to such a classifier could be low resolution patches from the wide-angle system or high-resolution images from the steerable camera.

Video 2 Demonstration of object classification machine learning trained on a relatively small data set of helicopters seen from the ground.

Another architecture, consisting of multiple cameras covering every possible direction, allows stereoscopic distance measurement in a single image frame and removes the uncertainty introduced by the UAV's own motion estimation.

Our analysis suggests that while vision-based systems may not quite yet meet our SWaP -C targets, they can come close to the required performance at about ~ 1kg and ~15 W, at BoM cost below $1000. This is lower than currently available vision-based approaches .

Vision-based approaches also have a role to play in DAA solutions for higher MTOW UAVs. While they would not be the first line of defence, because they cannot provide the range needed, they would serve as a second line of defence to enhance classification of aircraft.

It is worth mentioning that we had also explored LiDAR as a candidate technology for DAA - however, the complexity and cost for the range needed quickly ruled this out.

What about acoustic systems?

To work out what a DAA system could achieve with acoustic detection, we went back to first principles and worked out an audio link budget. Based on this, a conservative estimate of the detection range of a light aircraft is ~700 metres. This calculation assumes an acoustic beamformer (32 mics) to reduce uncorrelated environmental and correlated aero acoustic noise. However, this estimate is conservative since it assumes an SNR of 10 dB for detection, which could be further improved by post-processing.

We tested a simplistic principal component analysis (PCA) based post-processing approach on a relatively small, labelled UAV data set comprising drones, helicopters and background noise. This resulted in several false positives, suggesting that more sophisticated data processing approaches may be required. More sophisticated approaches could look at periodic/harmonic structures in the frequency spectra – for example, the figure below shows spectral features that can be obtained for helicopters – in conjunction with machine learning approaches.

Clearly, getting large enough data sets is key to obtaining these processing gains. This also does not include other noise reduction approaches that can be integrated into the airframe. For example, Zipline, who have launched an acoustic DAA solution with their airframe, have explored modifications to the aircraft wing edges to reduce aero-acoustic noise.

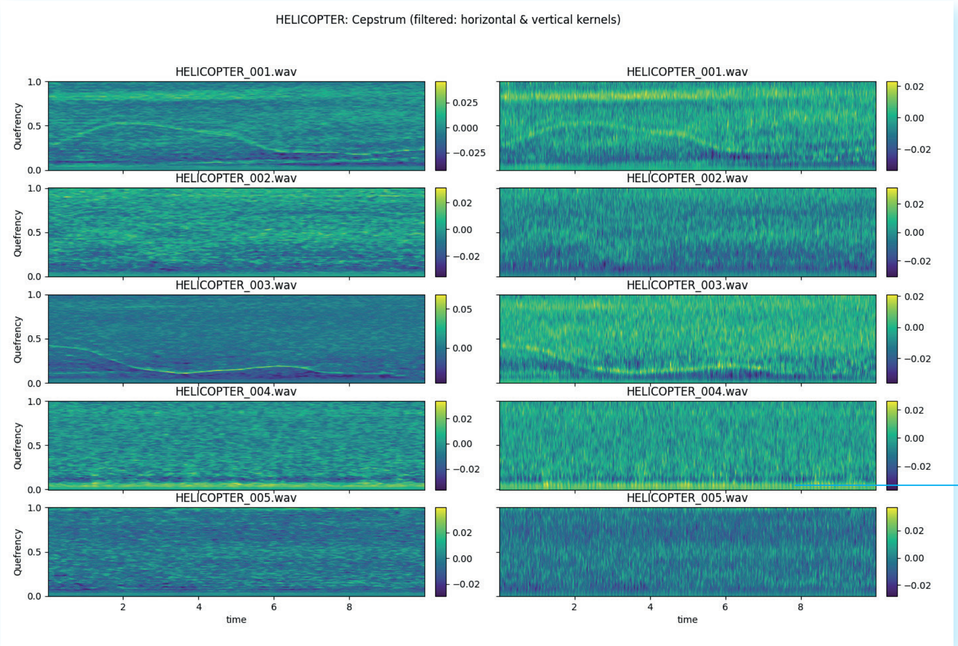

Spectral features indicating chopping or buffetting from the helicopter blades

Figure 3

Analysis of audio from helicopters showing specific spectral features in the frequency spectra.

Even assuming that processing gains will result in an improved detection range of 1-2 km, the downside of the acoustic approach is poor angular accuracy. Plus, acoustic approaches simply cannot detect non-conventional users like hang gliders and hot air balloons. The upside is that the SwaP-C cost of acoustic detection is almost free. We could get away with ~ 1W, $150 and size of 600 mm^2.

Based on our analysis, we think that passive acoustic methods may serve the market in the short term. In the longer term, acoustic approaches may become a second line of defence in a sensor fusion DAA system for smaller MTOW aircraft.

Last but not least: radar systems

The other approach on the table is radar. K-band radar solutions such as those produced by Echodyne and Honeywell are attempting to capture the DAA radar market. In the K-band, to detect a cross-section equivalent to a Cessna at 1 km would require about 30 dBi gain. To keep the form factor small, this implies Effective Isotropic Radiated Power (EIRP) in the range of 80 to 90 dBm, which would drive up the power requirement to quite a way outside our SWAP-C budget for low MTOW UAVs.

An assessment of mm wave radar – which is not currently a regulated air-to-air radar band – suggests we could obtain a 2 km range with a ~30 cm antenna aperture, albeit with a high-power requirement. Overall, for this end of the market, a pure radar-based solution seems unlikely whilst a tip and cue-based sensor fusion approach might be where radar could have a place.

However, radar is certainly the way to go for large UAVs. K-band and X-band radar solutions can provide the range needed, given that the power requirements can be easily met by larger UAVs. The challenge for the market would be to obtain these solutions at a sensible price point which may require pairing up off-the-shelf radar core modules and suitable OTS (OTS) or custom antennas that can meet size, drag and coverage requirements.

Coping with impairments in sensing and algorithms

The discussion so far highlights the strengths and weaknesses of different sensing technologies and their role in DAA for low and high MTOW UAVs. The table below is a summary and flags up other environmental disturbances that can impair the sensing system.

This shows that no single sensing technology will meet the performance needs in all environments, which suggests that sensor fusion may become more important in future. Furthermore, the choice of primary sensing modality may differ depending on MTOW. For example, vision-based approaches may be more suitable for smaller UAVs, while RADAR may be more so for larger UAVs. Secondary sensing systems, beyond improving performance of the primary system, will more importantly provide minimum fallback performance in environmental conditions where the primary sensor fails to perform.

Ultimately, whatever the technology, the detection system needs to drive the probability of “false negatives”, the probability of missing a potential airspace threat, as low as possible. Whilst we would also like to keep the probability of “false positives” low, i.e., misclassifying a non-threat as a threat, this probability can be traded off somewhat against false negatives. Apart from using different sensor systems, another prong of attack to improve detection may be to process the raw data output from any one sensing system with two or more complementary algorithmic approaches.

Sensor technology	**Low MTOW**		**High MTOW**	
	HIGHS	**LOWS**	**HIGHS**	**LOWS**
Vision	▪ Good range (at low closing speeds) ▪ Good accuracy ▪ Reasonable SWaP-C	▪ No night and poor weather operation ▪ Will need to address specular reflections	▪ Can aid airspace user classification	▪ Poor range (at high closing speeds) ▪ No night and poor weather operation
Acoustic	▪ Extremely low SWaP-C ▪ Long range possible (with several caveats around ability to cope with noise)	▪ Poor angular accuracy ▪ Cannot detect “quiet” airspace users (like hang gliders)	NOT SUITABLE	
Radar	▪ Good detection range ▪ Good angular accuracy (traded off against range at this MTOW)	▪ High SWaP-C ▪ Will need to address ground clutter ▪ Will need to address specular reflections	▪ Good range ▪ Good angular accuracy ▪ Suitable SWaP-C	▪ Performance impacted by clutter ▪ Will need to address specular reflections

Table 1 Summary of pros, cons and suitability of various sensor systems for different classes of UAVs

How will we certify these systems?

For any DAA system, especially one using passive methods and artificial intelligence techniques, how will we demonstrate safety? Some existing solutions rely on the guarantee of many flight hours without incident. However, this will never cover all possible corner cases. Aviation regulators are starting to recognise that there is an important place for synthetic data and simulation tools that can extend data coverage and reduce costs incurred in collecting flight data.

With synthetic data come the added burdens of ensuring that simulation adequately matches reality and of developing means to verify the integrity of the data. This is certainly achievable – making it tractable requires an intelligent assessment of which variables are correlated and uncorrelated. Nevertheless, it will never be possible to test a DAA system against all possible scenarios that will be encountered in the air. Thus, it is also important to understand the limitations of the system.

Traditional development assurance frameworks in aviation do not cover machine learning approaches well, i.e., they are designed for situations where “each line of code can be verified”. Now the industry is considering ways to perform rigorous data and inference verification and validation as part of the software certification process. Beyond that, we could ensure our software architecture lends itself more easily to verification and certification. Clearly, an architecture that modifies the neural network “on the fly” based on data observed in flight will face an uphill battle for certification. An architecture that relies on classical methods where possible and neural networks with well-bounded behaviour and training data would be a more balanced approach that could be acceptable by the regulator.

Come and talk to us!

With our first level analysis out of the way, we now have our arms well around the DAA challenge. We are now exploring partnerships to develop the next generation of DAA systems. Depending on your market and what you are looking to achieve, be it lower SWAP, better performance and/or a clearer path to certification, there are various ways TTP can work with you:

- If you are a DAA solutions provider, TTP can assist in the development of your next-generation system. Have you considered sensor fusion, or another ground-up approach using commercial high-volume sensors to minimise SWAP?
- If you are a regulator, TTP – as an independent technology company – can work with you to help develop a staggered technology approach towards certification.
- If you are a UAV operator looking to integrate, customise or develop a suitable DAA solution, have a chat with TTP's autonomous technology experts about your airframe and operational constraints to identify next steps on your DAA journey.

Or is there perhaps a different technical challenge that could benefit from TTP's expertise, scale and agility?

Want to know more about our journey towards Detect and Avoid for Commercial UAVs?

enquiries@ttp.com

TTP is an independent technology company where scientists and engineers collaborate to invent, design and develop new products and technologies.

Working across a wide spectrum of industries including health, life science, industrials and deep technology, TTP creates breakthrough solutions that bring strong commercial value to clients and the benefits of technology to all.

TTP plc
Melbourn Science Park
Melbourn, SG8 6EE, UK
+44 1763 262626
ttp.com/uavs