Passive Geolocation with 3D TDOA

A CRFS White Paper

Prepared by:

Dr Michael Knott
Dr Alasdair Edge

CRFS Limited
Building 7200,
Cambridge Research Park
Beach Drive, Cambridge,
CB25 9TL, UK
Tel: +44 1223 859 500

CRFS Inc.
4230-D Lafayette Center Drive
Chantilly
VA 20151
USA
Tel: +1 571 321 5470

Contents

1 INTRO TO RF GEOLOCATION ____ 3

2 TDOA BASIC OPERATING PRINCIPLES (IN 2D) ____ 4

2.1 CROSS-CORRELATION FUNCTION ____ 4

2.2 POSITION DETERMINATION ____ 4

3 FACTORS AFFECTING ACCURACY ____ 7

3.1 ACCURACY OF TIMING ____ 7

3.2 SAMPLE RATE AND BANDWIDTH ____ 8

3.3 SIGNAL PERIODICITY ____ 8

3.4 NETWORK GEOMETRY AND DILUTION OF PRECISION ____ 8

3.5 OBSTACLES AND MULTIPATH ____ 9

3.6 MOVING EMITTERS AND THE DOPPLER EFFECT ____ 10

4 3D TDOA GEOMETRY ____ 11

5 EXAMPLE APPLICATIONS ____ 11

5.1 ADS-B SPOOFING DETECTION ____ 11

5.2 SEARCH AND RESCUE WITH UAVS ____ 12

5.3 DRONE DETECTION AND GEOLOCATION ____ 13

6 PROOF OF CONCEPT ____ 14

6.1 BUFFALO, NEW YORK NETWORK ____ 14

6.2 NETWORK SETUP ____ 14

6.3 TEST PROCEDURE ____ 14

6.4 RESULTS ____ 14

7 CONCLUSION ____ 16

1 Intro to RF Geolocation

The concept of using radio frequency (RF) waves to locate objects is one that will be familiar to anyone acquainted with radar. Radar systems send short pulses of radiation out in a known direction, and then wait for any return signal that has bounced off an object in that direction. Measuring the time difference between the sent and received signals allows the distance to the object to be calculated. Radar is an example of an active geolocation technique, because it requires a signal to be generated by the user if objects are to be located. If the object that we wish to locate generates its own RF signal, then we can use passive geolocation techniques that rely on detecting the RF signal being produced. These techniques use a network of spatially-separated receivers to pinpoint the location of RF transmissions.

Passive techniques have a number of significant advantages when compared to active techniques like radar. Because they do not require an RF signal to be generated, they do not interfere with other uses of the spectrum in the surrounding area. This also means that their use cannot be detected, allowing the user to “see without being seen.” Their effectiveness does not depend on the radar cross-section of the object being located, allowing for the detection of even very small RF emitters. In addition to locating the target, analysis of the nature of the RF signal received can allow identification and classification of its source. The infrastructure required is also usually significantly smaller and less expensive, making a much wider range of deployment scenarios possible. There is only one significant drawback of passive systems, which is their reliance on the target producing an RF signal. In theory, this means that they can be evaded by avoiding sending out RF signals. However, this will likely be virtually impossible in practice – for a plane, it would mean avoiding all communication or use of radar, and, for a UAV, its flightpath would need to be entirely pre-programmed.

The range of passive geolocation techniques includes Angle of Arrival, Power of Arrival and Time Difference of Arrival. Angle of Arrival (AoA) uses at least two directional antenna arrays, each of which detects the direction from which the transmission originated. By extending lines of bearing (LOBs or LOB) out in those directions from each array, we can roughly locate the source at the intersection of those lines. Power of Arrival (POA) uses the fact that the power of signals drops off with distance traveled. By measuring the differences in the power received at two different locations, we can determine how much farther the source is from one location than the other. If we have at least three receivers, we can use this information to determine the source location. In this whitepaper, we will focus on the passive technique known as Time Difference of Arrival (TDOA), which uses the time difference between receipt of signals at spatially separated receivers to determine the location of the source.

2 TDOA Basic Operating Principles (in 2D)

2.1 Cross-Correlation Function

The first thing that we need to determine when using TDOA is whether, and at what time, the different receivers picked up the same signals. To do this, we collect the received spectra from each receiver, and then calculate the cross-correlation function (CCF) between each pair of spectra. The cross-correlation function is a way to determine mathematically how similar two signals are. One function is multiplied by the complex conjugate of the other and then integrated to give a value indicative of the similarity. This procedure is then repeated for all possible values of the time delay, τ, between the two signals. If the two signals are denoted f(t) and g(t), then

$$(f * g)(\tau) \triangleq \int_{-\infty}^{\infty} \overline{f(t)}\, g(t+\tau) dt$$

By progressively moving the functions relative to one another, we can find the time delay where the two signals are most similar, which will correspond to the maximum value of the CCF, and will indicate the actual time delay, T, between receipt of the signals at the two Nodes. We can then convert this time difference into a distance (by multiplying by the speed of light, c), which will tell us how much further one of the receivers is from the source than the other.

2.2 Position Determination

So how do we determine the position of the source, given the difference in distance from each of the receivers? If we take two receivers, R_1 and R_2, and have a signal arriving at R_1 a time T later than at R_2, we can say that the source is a distance δ_1 further from R_1 than R_2, where

$$\delta_1 = cT$$

(Again, assuming that the RF signal travels at the speed of light). What can we say about the location of the source based on this information? If we have two receivers separated by a distance 2D, and an emitter at an arbitrary point (x, y), then we will have a setup as in Figure 1 below.

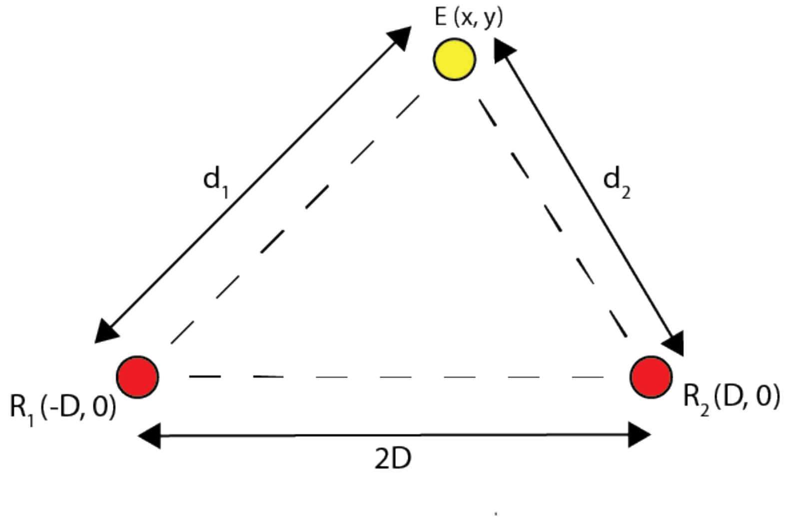

Figure 1: Distances from two receivers to an arbitrary point

The distances d_1 and d_2, from R_1 and R_2 to the emitter are

$$d_1 = \sqrt{(x+D)^2 + y^2}$$

$$d_2 = \sqrt{(x-D)^2 + y^2}$$

If we know that the emitter is δ_1 further from R_1 than R_2, then we have

$$d_1 - d_2 = \delta_1$$

and therefore

$$\sqrt{(x+D)^2 + y^2} - \sqrt{(x-D)^2 + y^2} = \delta_1$$

This can be simplified to

$$\frac{4x^2}{{\delta_1}^2} - \frac{4y^2}{(4D^2 - {\delta_1}^2)} = 1$$

For any given value of δ_1, this equation can be satisfied by a continuous line of points constituting a hyperbola. Figure 2 illustrates this, with hyperbolas plotting out given differences in distance from the two receivers.

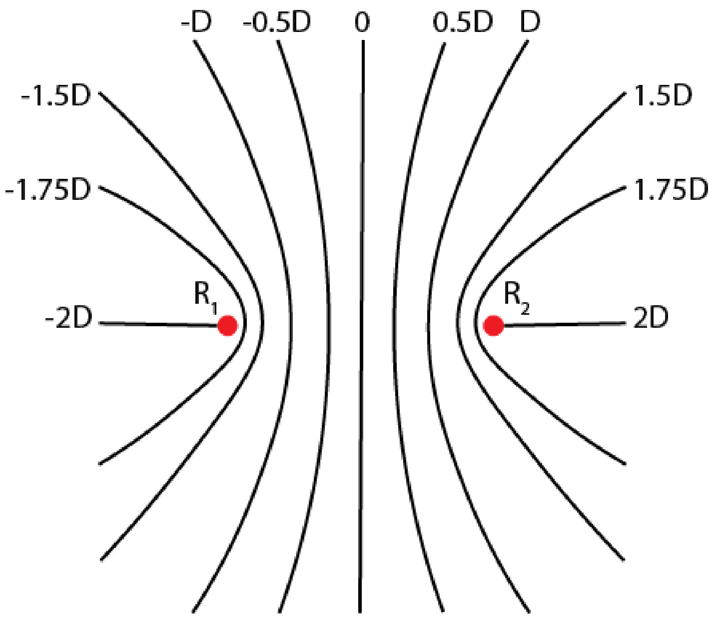

Figure 2: Curves indicating possible transmitter locations that correspond to various values of δ_1

So if we have just two receivers, we can only say that the source lies somewhere along the relevant hyperbola. A third receiver, R_3, is required, with a time difference T' between receipt of the signal at R_3 and R_2. We then know that the source is a distance

$$\delta_2 = cT'$$

closer to R_3 than R_2. The set of points located δ_2 farther from R_3 than R_2 also traces out a hyperbola, and the source will be located at the point where the two hyperbolas intersect. The time difference between arrival at R_3 and R_1 can also be used in the same way, and will result in a third hyperbola that intersects the other two at the same point.

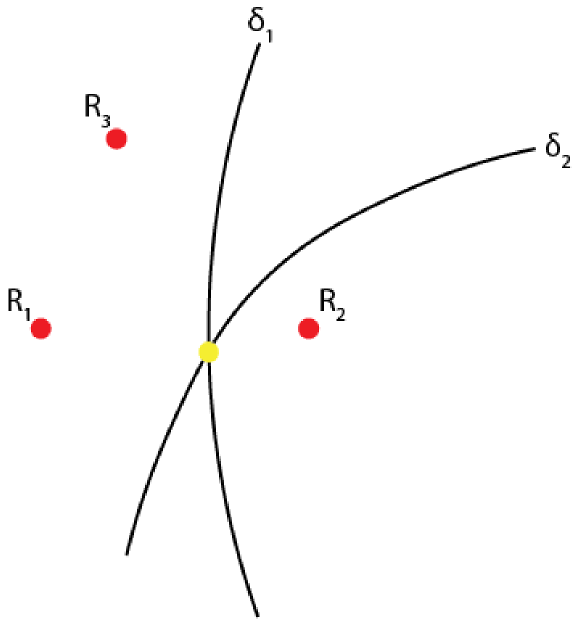

Figure 3: Two hyperbolas intersect at the source location

3 Factors Affecting Accuracy

3.1 Accuracy of Timing

When determining the accuracy of TDOA measurements, the first thing to consider is the accuracy of the timing. Errors in the time recorded will translate directly into errors in the distance measurements. These errors can come in two parts: firstly, the uncertainty in the time recorded at each receiver, which will depend on the clock used in each receiver, and secondly, the accuracy of the synchronization between the receivers. This synchronization will generally be carried out using GPS signals (or signals from another positioning network, such as GLONASS or BeiDou), which will mean the accuracy will be dependent on the reception. In very good GPS conditions, the synchronization error will be around 30 ns RMS, corresponding to a spatial accuracy of around 9m, but under typical conditions that would likely increase to tens of meters. This also means that the system operation is dependent upon being able to receive a GPS signal. In situations where this is not possible (due to, say, jamming or solar flares), an alternative means of synchronization will be required. If the outage is sufficiently short, it may be possible to extend the timing from the last signal received, but this will lead to a gradually increasing error over time. In the case of longer outages, a physical synchronization link between the receivers, such as CRFS's optical SyncLinc, can be used to maintain an accuracy equivalent to very good GPS conditions.

3.2 Sample Rate and Bandwidth

The time resolution that we can obtain will also depend on the rate at which we sample the incoming signal. If we were to sample at a rate of 10 MHz, we would obtain a time resolution of 1/frequency = 0.1 µs, and a spatial resolution of (0.1 µs x c) = 30m. It might seem as though this error could be reduced to virtually zero by just increasing the sample rate, but this will only work for sample rates up to the modulation rate or bandwidth of the signal. Beyond this point, differences in successive samples will be primarily due to noise, rather than changes in the signal. The effect of this is that TDOA is best suited to locating wideband signals, with an accuracy of 10m at bandwidth of 10 MHz, reducing to 100m at 1 MHz, and 1km at 100 kHz.

3.3 Signal Periodicity

Another potential issue for TDOA analysis occurs when the signal being received is periodic. For a signal of period P, the signal will be the same for any time $t_n = t_0 + nP$, where n is any integer. Thus, when we calculate the CCF, if there is a peak at a value T_0, there will be identical peaks at ($T_0 + nP$), and we will be unable to determine which of these values is the actual time difference between receipt of the signal at the two locations. In some cases, we can use physical insight to rule out all but one possibility (if the period is sufficiently long, and therefore the difference between possible distances is sufficiently large). In practice, this will mean a period of milliseconds or greater, giving possible locations which are hundreds of kilometers apart.

3.4 Network Geometry and Dilution of Precision

Two further factors that affect the accuracy of TDOA are the location of the source relative to the receiver network, and the geometry of the receiver network itself. In very general terms, the accuracy is highest when the source is located in the triangle formed by the three receivers. As we have seen in sections 3.1 and 3.2, there will be an uncertainty in the measurement of the time difference of arrival, which leads to a corresponding uncertainty in the determination of the difference between the distance of the source from each receiver. Let's say that the uncertainty corresponds to a value ±D in the units of Figure 2, and that the source is calculated to be equidistant from A and B. In the situation where the source is located directly between the two receivers (y=0), we can see that this leads to a smaller uncertainty in the x-coordinate then when the source is located at the top of the Figure (in practice, the uncertainty in the position of the latter would not be a straight line as shown, but the effect of the increased spacing between the adjacent hyperbolas is easiest to see when presented like this).

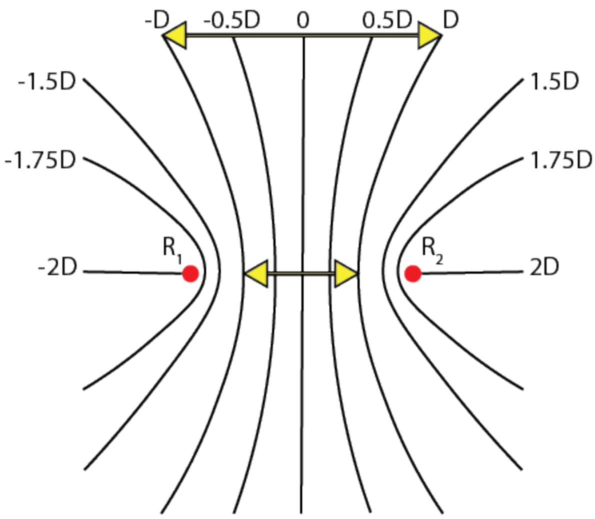

Figure 4: Effects of emitter position on x-coordinate uncertainty

3.5 Obstacles and Multipath

The accuracy of measurements that can be obtained can also be affected by the presence of obstacles. In the simplest cases, obstacles located between emitter and receiver may absorb the RF signal before it reaches the receiver. Over longer distances, the curvature of the Earth itself may block signals.

TDOA is, in general, unaffected by multipath effects. Signals that arrive at receivers via indirect routes will almost invariably be weaker than those travelling directly, and the CCF will pick out the strongest signal as the one corresponding to the actual time delay of arrival. However, there are possible situations where multipath could affect the accuracy of TDOA results. Consider a situation combining a blocked direct line-of-sight with a strong secondary path to the receiver, as in Figure 5.

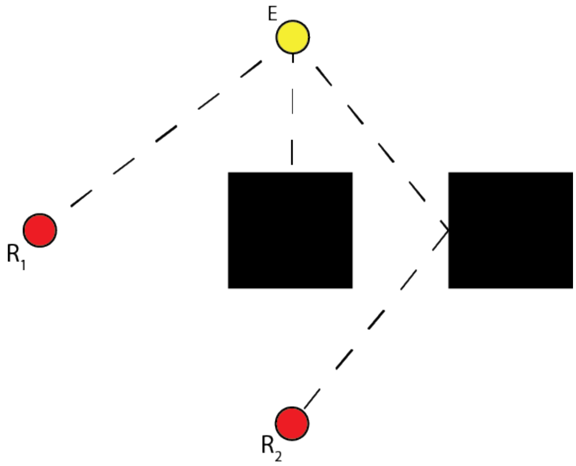

Figure 5: Multipath Effects in TDOA

In this case, the strongest signal received at R_2 will be the reflected signal, which has to travel an additional distance compared to the direct path. This will cause the emitter to appear to be further from R_2 than it actually is.

The most practical way to avoid the effects of obstacles is to raise the receivers above ground level, as the majority of obstacles will be located on the ground. This can be done by mounting receivers on top of buildings, on poles, or even on UAVs (see section 5.2). It may also be advantageous to model the effects of different receiver placements in order to optimize geolocation performance before deploying a network for real. CRFS offers simulation and modeling tools allowing users to do so.

3.6 Moving Emitters and the Doppler Effect

As described in section 2.1, TDOA relies on being able to recognize when the same signal is being received at spatially-separated receivers. The CCF

However, in situations where the emitter is moving, we need to take into account the Doppler Effect. The frequency of the signal, f, as measured at the receiver will be altered, and will be related to the original emitted frequency, f_0, by

$$f = f_0\left(1 + \frac{v}{c}\right)$$

where v is the relative velocity between receiver and emitter (positive values corresponding to movement towards each other).

The velocity of the source will, in general, not be the same relative to each of the receivers, and so the measured frequency of the signal will no longer be the same at each. The peak in the CCF at the value corresponding to the true time delay will be reduced accordingly. In extreme cases, the reduction may be so great as to render the peak impossible to differentiate in the CCF, making the true time difference unknown. In situations where more than four receivers are present in the network, it may be possible to select pairs of receivers where the relative velocities are very similar, and the frequency shift is correspondingly smaller. However, the effect is likely to be greatest when emitters are located within the network boundary (as that is when there is the highest chance of the emitter moving towards one receiver and away from another simultaneously). Given the discussion of network geometry in section 3.4, we can see that this will result in a reduction of obtainable accuracy.

4 3D TDOA Geometry

Extending TDOA geometry into three dimensions necessarily complicates the geolocation procedure somewhat. The set of points in 3D that are a given difference further from one receiver than another is not a curve, but is instead a surface – the three-dimensional extension of a hyperbola known as a hyperboloid. Adding a third receiver will define a second hyperboloid, which will overlap with the first one along a curve. That means that we need a fourth receiver, and the third hyperboloid that it defines, if we are to determine a point location for the source.

5 Example Applications

5.1 ADS-B Spoofing Detection

Automatic dependent surveillance-broadcast (ADS-B) is a surveillance technology currently being introduced in the US (Next Generation Air Transportation System), Europe (Single European Sky ATM Research) and across the globe as a replacement for secondary surveillance radar, and possibly in future to replace primary radar, for air traffic monitoring and control. It requires planes to determine their current position using GPS, and then to broadcast the details of this position (ADS-B OUT). This information can then be used by air traffic control to coordinate the movement of planes through their airspace. Planes may also carry an onboard receiver (ADS-B IN), which allows them to receive the ADS-B OUT transmissions of other planes in the vicinity.

ADS-B is an attractive solution for air traffic control management for a number of reasons. Receiving stations can be built in almost any location, which means that ADS-B does not suffer from the problem of black spots in remote locations where building radar stations is unfeasible. Increased accuracy of location data (from more regular broadcasts of position) allows for smaller minimum aircraft separations, and therefore higher capacity. And as already described, it allows for plane-to-plane communication of position. However, it also has a number of potential drawbacks that need to be considered.

Because ADS-B is unencrypted, anybody receiving a transmission can determine the location of the plane sending that transmission. This possibility is used in a harmless (and indeed, useful) way by websites such as FlightAware, which plot (almost) live position updates of commercial aircraft. A more worrying possibility is that because ADS-B transmissions are unauthenticated, spoof ADS-B signals can be generated. By creating a signal that matches the required messaging protocol, details of a non-existent ‘ghost’ aircraft can be sent to either ATC ground stations or to other planes. This could cause major disruption if flights need to be diverted to avoid the ghost aircraft. Moreover, it is possible to create multiple ghost signals simultaneously, greatly increasing the possible disruption.

3D TDOA offers a means to mitigate the harm caused by ADS-B spoofing. By analyzing the signal time of arrival at TDOA stations, we can determine whether the signal actually originates from the location it purports to, based on the ADS-B encoded information. In cases where there is a disparity between the two location determination methods, air traffic controllers can be alerted, and investigation of the signal source can be carried out.

5.2 Search and Rescue with UAVs

Applications of 3D TDOA are not limited to situations with an airborne target and ground-based receivers. Consider the situation where a search and rescue team is trying to locate a person trapped by an avalanche in mountainous terrain. There are some existing locator technologies (such as the RECCO radar-based system, which uses radar pulses that can be reflected by a passive reflector that can be carried), but they almost all require a prior awareness of danger, so that the necessary steps can be taken (such as ensuring that the person has a RECCO reflector in their clothing/on their person). Even in cases where someone is unprepared for the possibility of avalanche, though, it is highly likely that they will be carrying a mobile phone. This raises the possibility of using TDOA to locate them using RF signals from the phone. However, in the kind of terrain where avalanches are likely to occur, there are also likely to be significant physical obstacles to signal propagation at ground level. These obstacles can be avoided by mounting receivers on unmanned aerial vehicles (UAVs), which can be flown high enough to ensure their line-of-sight is not blocked. In this situation, with airborne receivers and a ground-based target, 3D TDOA is necessary, especially as small errors in the estimated position could lead to significant (and life-threatening) delays in the search and rescue operation.

5.3 Drone Detection and Geolocation

Of course, UAVs/drones are not always deployed in such a constructive manner. Recent incidents at airports around the world have demonstrated the need to be able to detect and locate drones that threaten the safety of air traffic, either through negligence or deliberate ill intent on the part of their operators. The economic impacts of these incidents can also be substantial, with the two-day closure of Gatwick estimated to have cost in excess of $64m.[1] Existing radar systems have some utility in tracking drones, but have a number of significant drawbacks, including problems differentiating drones from birds (which have a comparable radar cross-section).

3D TDOA geolocation of drones can be carried out by receiving RF communication sent from drones to their operators (normally at 2.4 or 5.8 GHz). This has the advantage of avoiding any confusion with birds, which are unlikely to be sending out any RF signals. The one drawback of using TDOA for drone location is that it depends on drones sending out signals, so it could potentially be bypassed by a drone with a pre-programmed flight path that does not need to communicate with its operator while in flight. However, there are very few scenarios where such pre-programming would be operationally useful.

[1] http://fortune.com/2019/01/22/gatwick-drone-closure-cost/

6 Proof of Concept

6.1 Buffalo, New York Network

The unencrypted nature of ADS-B transmissions described in section 5.1 also makes them an ideal testbed for establishing the practical performance of a 3D TDOA system, by comparing its geolocation results with the GPS position of planes broadcast through ADS-B. As a proof of concept exercise, a network of 4 RFeye Nodes was used to geolocate planes in the area surrounding Buffalo, New York, and these results were compared to the locations broadcast over the 1090 MHz ADS-B channel. The system was able to accurately geolocate and track aircraft with speeds of up to 580 knots over hundreds of km of operational airspace (in third party testing, tracking aircraft at supersonic speeds has been shown to be possible).

6.2 Network Setup

Results of a 48-hour period monitoring air traffic in Buffalo, New York can be seen below. In this case, the RFEye Nodes were arranged with a 70km baseline, to monitor an area around 900km across. The triangle in Figure 6 has Nodes at the three vertices, and the fourth Node can be seen as a small square within the triangle boundaries.

6.3 Test Procedure

Testing took place over a 48-hour period. Nodes were controlled by CRFS's real-time RFeye Site software. A 3D TDOA mission was used with the Nodes participating set up to monitor the 1090 MHz aircraft ADS-B transponder frequency. To determine accuracy, a module was created that decoded the ADS-B transmissions to provide true GPS location and altitude. This was compared for each point with the 3D TDOA derived result.

6.4 Results

ADS-B transmitted locations were compared with TDOA-calculated locations, and the difference between the two was determined. Figure 6 below shows the results, with the Nodes forming the central green triangle, each green point representing an error of less than 500m, yellow points errors of <5km and red points errors of <10km. The black squares within each of those regions show what that level of uncertainty will mean in practice.

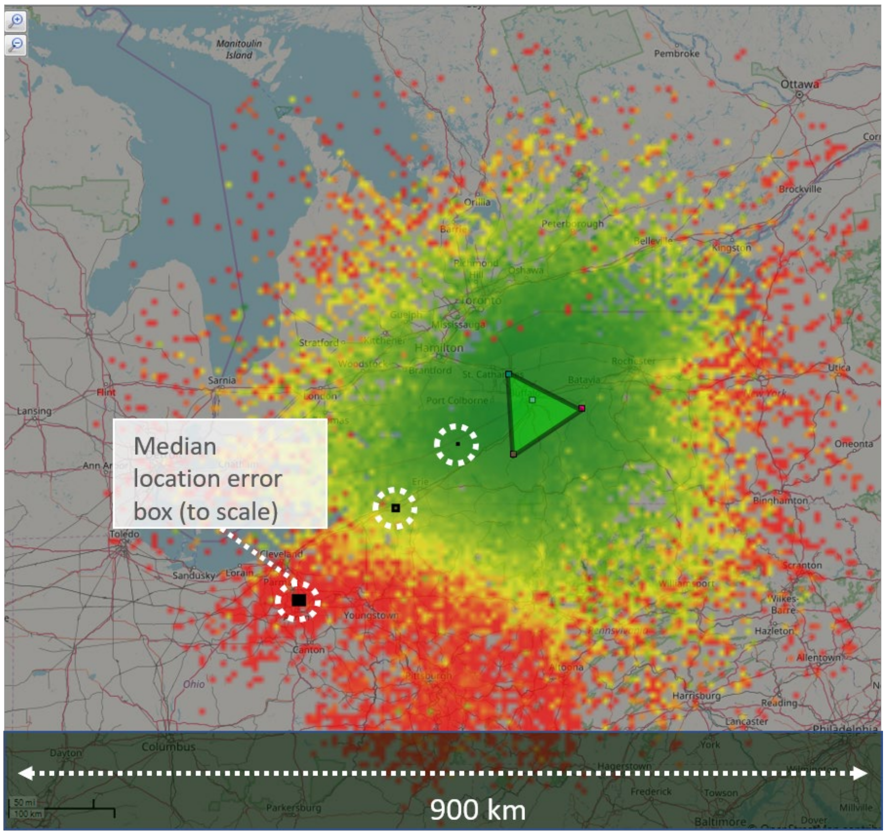

Figure 6: Comparison of 3D TDOA with ADS-B locations

Accuracy increases as planes approach the network, and within the network itself median errors are less than 200m.

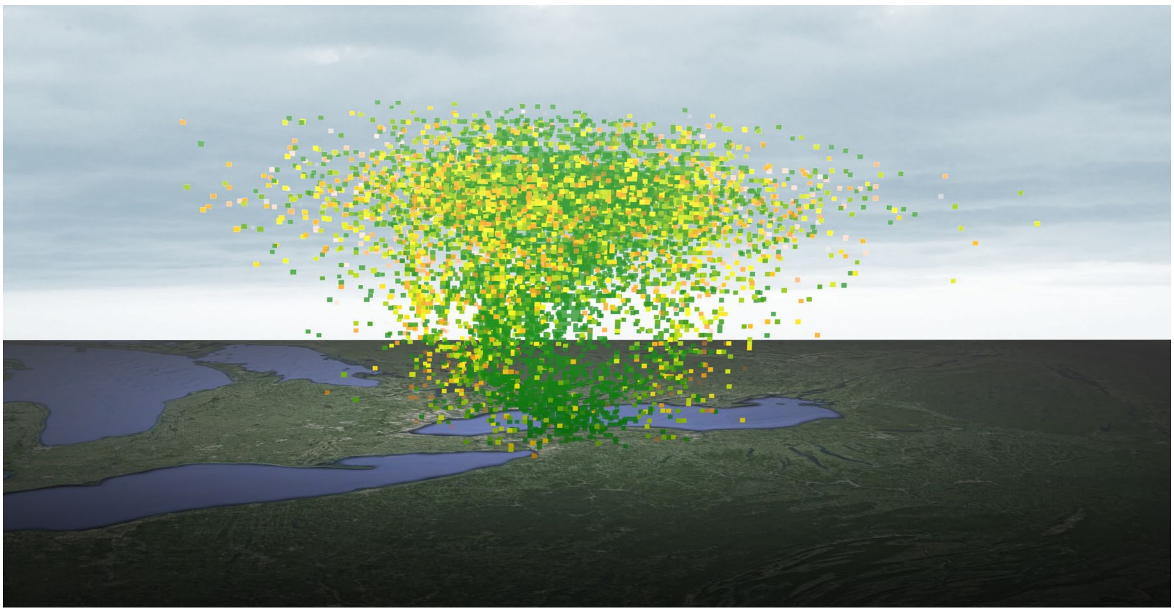

Figure 7: 3D plot of observed plane positions around the Buffalo Node network

7 Conclusion

In this whitepaper, we have looked at the underlying theory and operating principles of 3D TDOA, discussed some of its applications, and seen an example of the technique being used in practice. Advances in technology are likely to lead to further deployment scenarios that cannot currently be anticipated – consider the way that UAVs have rapidly moved from being a rarity to being widely commercially available to the general public. The flexibility to respond to these developments is a significant strength of 3D TDOA, as any object producing RF signals can be geolocated. The applications covered here are far from an exhaustive list of even current uses, though, so if you have a specific 3D TDOA deployment scenario in mind, then contact us on enquiries@crfs.com to discuss how RFeye Node technology can meet your requirements.

Legal Information

Disclaimer