TIME DELAY ESTIMATION OF ACOUSTIC TONES UNDER LOW SNR CONDITIONS

Stephen Franklin, Joshua Meade and Anthony Finn

Defence and Systems Institute (DASI), School of Engineering
University of South Australia, SA, Australia
Email: stephen.franklin@unisa.edu.au
joshua.meade@mymail.unisa.edu.au
Anthony.finn@unisa.edu.au

Abstract

This paper describes a novel approach for autonomously detecting and tracking aircraft in the vicinity of an unmanned aerial vehicle (UAV). The time difference of arrival (TDOA) of acoustic tones originating from distant aircraft are correlated between spatially distributed microphone pairs located onboard the UAV. The geometry of multiple microphone pairs then allows the elevation and azimuth of the approaching aircraft to be estimated, despite the high levels of narrow- and broadband noise emanating from the engine firing sequence, propeller, airflow over the microphones and mechanical vibration. The technique enables estimation of the signal at levels 20 dB below the broadband noise floor. Potential detection ranges in excess of 1 km are demonstrated when the signal processing is combined with careful suppression of the many noise sources onboard the UAV.

1. Introduction

As lightweight UAVs face strictly limited payload capabilities, the need for a low power, lightweight, high reliability solution for robust collision avoidance from other aircraft is fundamental. Acoustic sensing technologies on UAVs have been an area of significant research for many years, but suffer from short (< 500m) detection ranges [see DOT/FAA/AR-08/41]. Acoustic sensing may be used as a stand-alone sensor system for collision avoidance, or it may be used in conjunction with other sensors to cross-cue higher resolution (e.g. optical) sensors that have very limited field of regard. Although clearly hostage to the propagation of sound in air, acoustic techniques have been demonstrated at effective detection ranges for reasonable aircraft velocities [1]. A major limitation of current techniques, however, are the high levels of narrowband self-noise induced by the firing sequence of the engine and propeller blade rate, the high levels of low frequency broadband noise induced by mechanical vibration and airflow over the microphones, and the impulsive noise induced by the actuators.

Time difference of arrival (TDOA) localisation uses spatially distributed microphones to calculate the propagation time delay between microphone pairs. TDOA values are typically estimated [2] using time-domain techniques such as the generalized cross-correlation (GCC) [3]. These techniques are used to transform the signal into a form whereby the time delays can be estimated [4]. Given accurate time delay information across the microphone sets and the knowledge of the speed of sound through air, the angle of arrival can be obtained for each microphone pair. Ambiguous peaks in the cross-correlation function – caused by the strong harmonic tones in the signal – can be suppressed with the addition of the cross-spectral phase transform pre-filter [5].

TDOA information can also be obtained in the frequency domain and this paper discusses some of the advantages of using these signal processing techniques in terms of improved detection range. Wind noise and mechanical vibration noise mitigation techniques are not the focus of this paper and are discussed in detail elsewhere [6]. A summary of the techniques used is, however, included here.

2. Test Setup

The two key issues identified during the previous trials [1] were: (i) that a microphone mounted on the surface of the UAV inherently causes local air turbulence at the microphone diaphragm, inducing additional broadband noise; and (ii) mechanical vibration, particularly in the UAV wings, induce significant broad band noise which is difficult to remove in the signal processing alone. Commercially available windscreens fail to meet the size constraints associated with smaller UAVs so bespoke small form-factor windscreens were developed in-house. Windshields made of porous material offer low resistance to acoustic waves and high resistance to airflow. The microphones were mounted on/under the surface of the airframe with provision made to offer a low impedance path for the acoustic signals to reach the microphones whilst simultaneously removing broadband noise effects due to the airflow over the microphones.

The effects of mechanical noise were mitigated by mounting the microphones directly to the fuselage of the UAV, where vibration was observed to be minimal but where the acoustic signature was louder. This required changes to the signal processing regime (see later) used in previous experiments to maintain the accuracy and resolution with which oncoming aircraft can be detected. Neoprene rubber and low-density foam was used as interfacing material between the microphone, mounts and the fuselage.

This approach was run in parallel with another that made use of spatially integrating microphone sets designed by Midspar Systems. The Midspar microphones, which are physically larger than the ones described here, are mounted externally and to the wings of a UAV. The technique employs the concept of extended aperture, which allows wind noise of high wave number to be rejected while maintaining response to noise of low wave number within the acoustic space.

Following the success of the laboratory prototype tests [7], the acoustic signatures of approaching aircraft at Parafield Airport in South Australia were observed in the presence of noise generated by an Aerosonde UAV, its propeller and engine. Four PCB-130A-40 ¼ inch ICP array free field low-profile surface pressure 45mV/Pa microphones with integral pre-amplifiers and a 24-bit, 102dB Spurious Free Dynamic Range (SFDR) National Instruments NI-9234 USB Data Acquisition (DAQ) module were used to collect the data. The sampling rate used was 25.6 kHz. This dataset is unique because all of the noise sources are present except for wind noise which can then be added for further analysis when required.

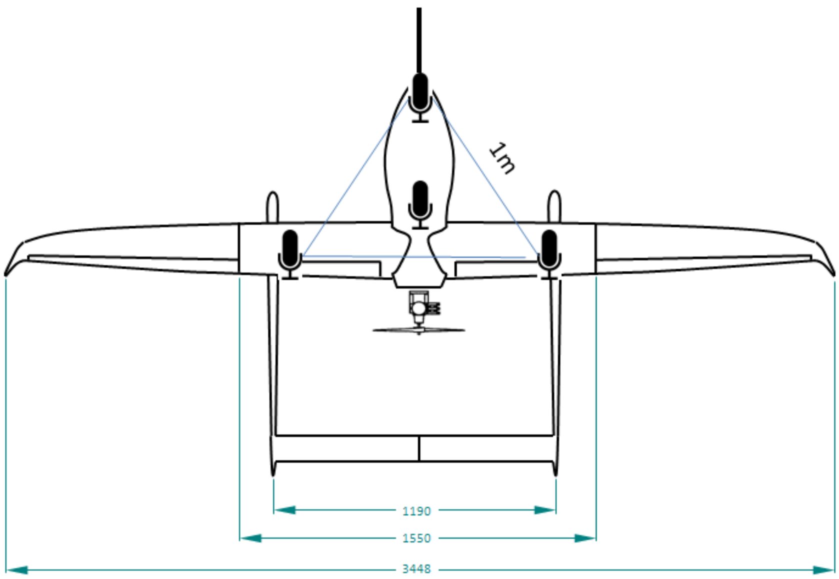

Figure 1. Microphone locations on the Aerosonde UAV

The test setup was placed in an area that is a safe distance away from Parafield Airport in line with the main runway and directly under the landing approach of light aircrafts. A twin piston engine aircraft (DA42 with a noise level of 135dB) was conveniently flying circuits and data was collected on an opportunistic basis with the Aerosonde's engine running at various throttle settings.

3. Methodology

In previous trials [7], correlation techniques were applied in the time domain to estimate the time difference of arrival of signals falling on microphone pairs [1]. Prior to applying such techniques, however, some manipulation in the frequency domain was needed to enhance the signals of interest and to suppress any unwanted signals. This technique has a number of issues that make its performance sub-optimal: it requires significant pre-processing in the frequency domain and works well only if the signal level of the approaching aircraft is significantly higher than the noise level in the frequency band of interest.

To overcome this, a signal processing approach was adopted that can be executed solely within the frequency domain. Based on the Cross-Power Spectral Density (CPSD) of the Fourier transforms of the individual signals arriving at each microphone, time difference of arrival is estimated from the phase shift of signals in each frequency bin. Please refer to [2] comprehensive review of various time delay estimation techniques.

If the signals arriving at two microphones are represented by the x(t) and y(t) and their cross correlation is given by Rxy, their CPSD is then given by

$$S_{xy} = \sum_{k=-\infty}^{\infty} R_{xy}(k) e^{i\omega k} \qquad (1)$$

The above expression represents the discrete-time Fourier transform of Rxy such that Sxy is a complex quantity whose magnitude represents the power density for each frequency bin. The corresponding phase shift represents the time delay. This technique allows us to observe the phase shifts (time delays) at different frequencies. Since the frequency bins are narrow (typically ~1Hz for the sampling frequency and number of samples used although this can be reduced further (to ~0.1 Hz) using decimation) contamination of the time delay estimates due to the broadband noise component is less pronounced when compared to the time domain technique. This improves the detection ranges.

The effectiveness of this technique is illustrated bellow using a simple simulation. Typically, time domain correlation techniques were able to provide reliable detections/estimates only when the narrowband peaks of the approaching aircraft were above the broadband noise floor. As the signals approach the noise floor, the estimates became too noisy to be useful as a means of establishing detection [1]. Using CPSD, however, the signal detection estimates degrade more gradually and statistical estimators (typically, Least Squared Error (LSE) based estimators) can be used to recover useful data from the noisy estimates.

The simulation generates a synthetic test dataset which consists of data from two channels. Channel one has a broadband noise component and a narrowband tone at a given frequency. Channel two also contains a broadband noise component (at the same level) and a narrowband tone at the same frequency as channel one but shifted in phase by 1 radian. Narrowband tones of self-noise are assumed to be excised using notch filters. The CPSD is computed on both channels and the phase information at the frequency of the narrowband tone plotted. This process is repeated for various broadband noise levels, starting from a noise level 20dB less than the signal and progressively increasing in magnitude until it exceeds the signal level by 20dB. The time delay estimates progressively degrade but useful information is still recoverable using LSE estimators.

The trend evident in figure 2 can be replicated for narrowband tones at different frequencies starting from about 50Hz. Also, in this simulation, narrowband tones from the sensing aircraft are assumed to be excised and only one narrowband tone from an approaching aircraft is considered. These assumptions are reasonable because the engine firing rate can be measured and the number of propeller blades known in the sensing aircraft and hence the frequency of the narrowband tones from the sensing aircraft are known precisely. Notch filters can then be used to excise these tones. As for the approaching aircraft, though there are multiple narrowband tones, these tones have a different amplitudes and the propagation loss is dependent on frequency. Often, the strongest tone in the low frequency band (100-400 Hz) will afford the easiest detection in the initial stages.

The next step in the process is to run an LSE estimator to recover the phase shift from the noisy estimates. A first order polynomial fit would readily extract this information in the form

$$y = mx + c \qquad (2)$$

where m is the slope of the line and c gives the y-intercept which in turn directly corresponds to the phase shift. The polynomial fit algorithm works by adjusting the polynomial coefficients by minimizing the LSE. The algorithm provides as output a list of polynomial coefficients along with a covariance matrix that provides information about the goodness of fit. Figure 2 shows the phase shift (blue) at different SNRs and the LSE estimates for the data is also shown in figure 2 (red). Error bars are not indicated in the figure (for clarity) but the length of error bars were found to increase as the SNR worsens.

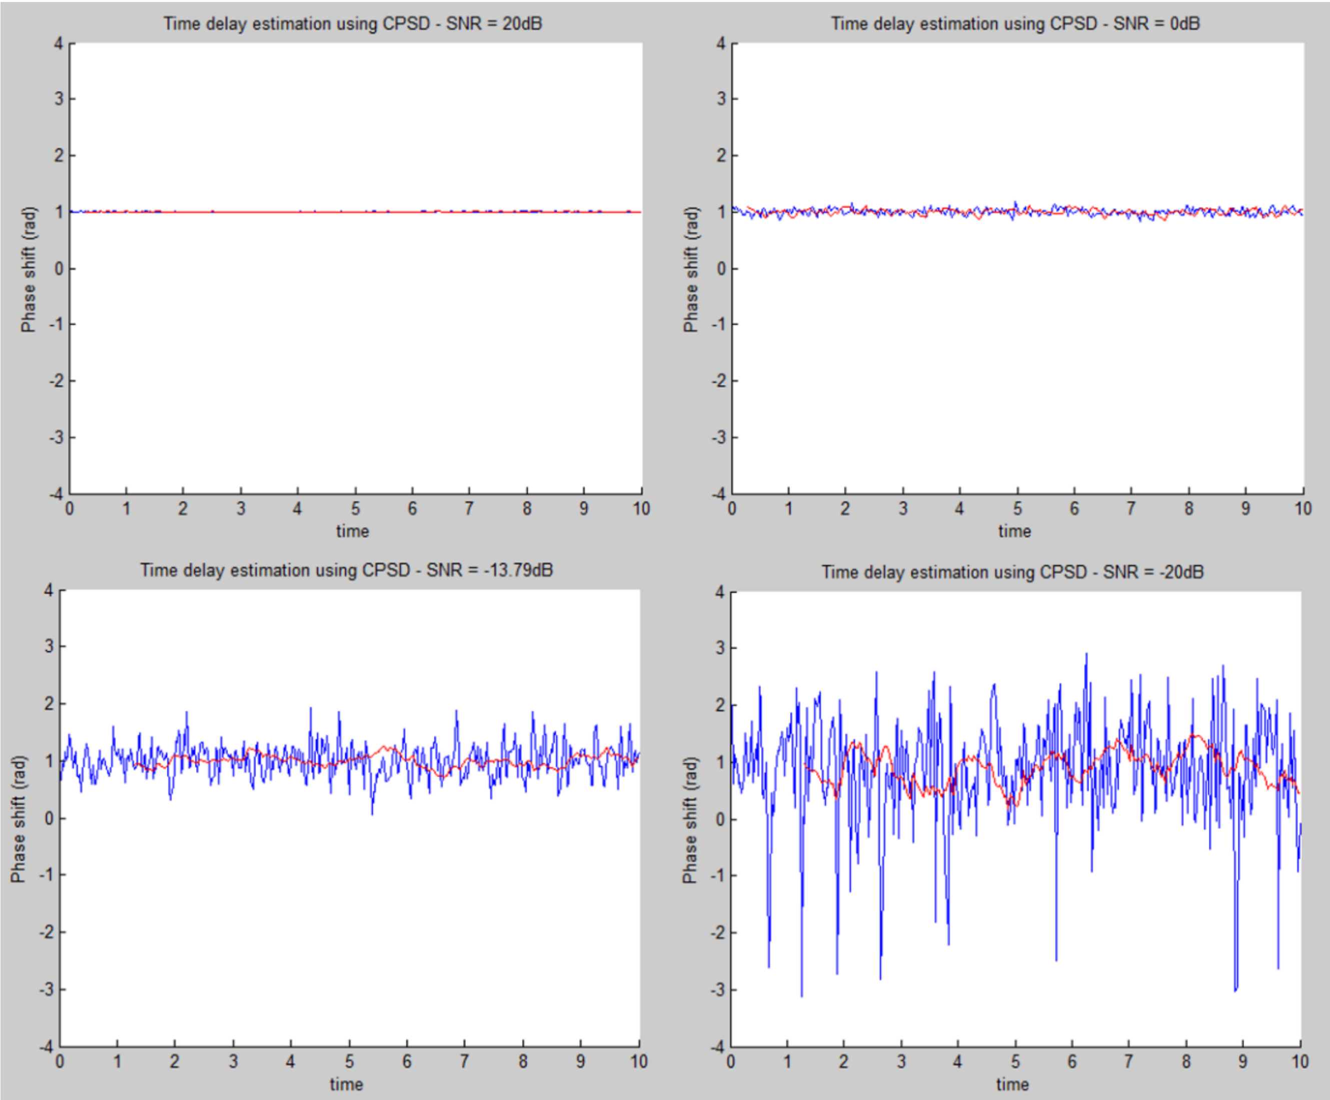

Figure 2. Time delay estimates at different levels of SNR (blue) derived using the CPSD approach after filtering (red) (clockwise from top left +20dB, 0dB, -13.5dB, and -20dB). At +20dB and 0dB convergence is almost instantaneous whereas at -13.5dB and -20dB it takes about 1.25 seconds to converge

4. Drawbacks of the CPSD Approach

The CPSD based approach discussed so far assumes that the frequency of the received signal is known accurately or at least that the frequency does not change during the initial tracking process. In reality, this is not always true. The frequency of the tones are dependent on the throttle settings used on the approaching aircraft and Doppler shift. Doppler shift is a lesser issue because aircrafts on a collision course will not have Doppler shift but we do not have any control over the throttle variations in the approaching aircraft.

5. Testing on Datasets with No Frequency Variation

This technique was therefore applied to previously collected datasets where the approaching aircrafts only had had very small variations in throttle settings as they approached Parafield Airport during the static ground based trials. Using the previous generation of signal processing algorithms the approaching aircraft was tracked for about 15 seconds before the Closest Point of Approach (CPA) [7]. The current generation CPSD based tracking methods are able to perform significantly better and extend the detection range by about 8 seconds (600m for a DA42 twin piston engine aircraft). This makes detection possible 23 seconds or 1900 m before CPA. This is illustrated in figure 4. The problems associated with this technique are also evident in figure 4 (bottom) where there is significant errors in the estimation when the frequency changes.

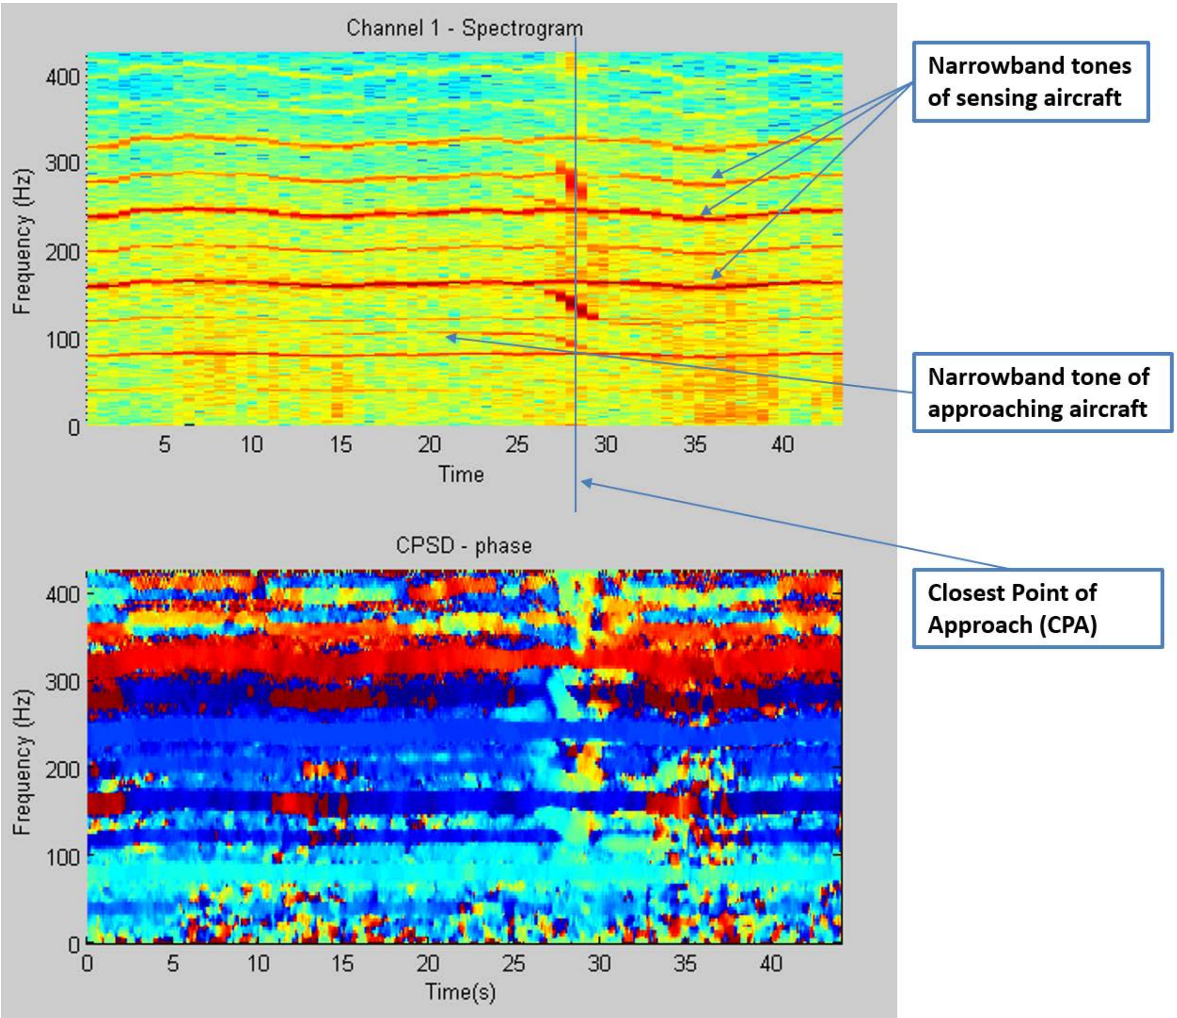

Figure 3. (Top) Spectrogram from microphone 1 showing narrowband tones from sensing and approaching aircraft. Signal strength is indicated using relative colour code varying from blue to red representing weak to strong signals respectively (Bottom) Phase component of CPSD of microphones 1 and 2 showing patches of constant phase shift corresponding to narrowband tones. Phase shift is indicated using a colour code varying from blue to red representing the range (-π to π)

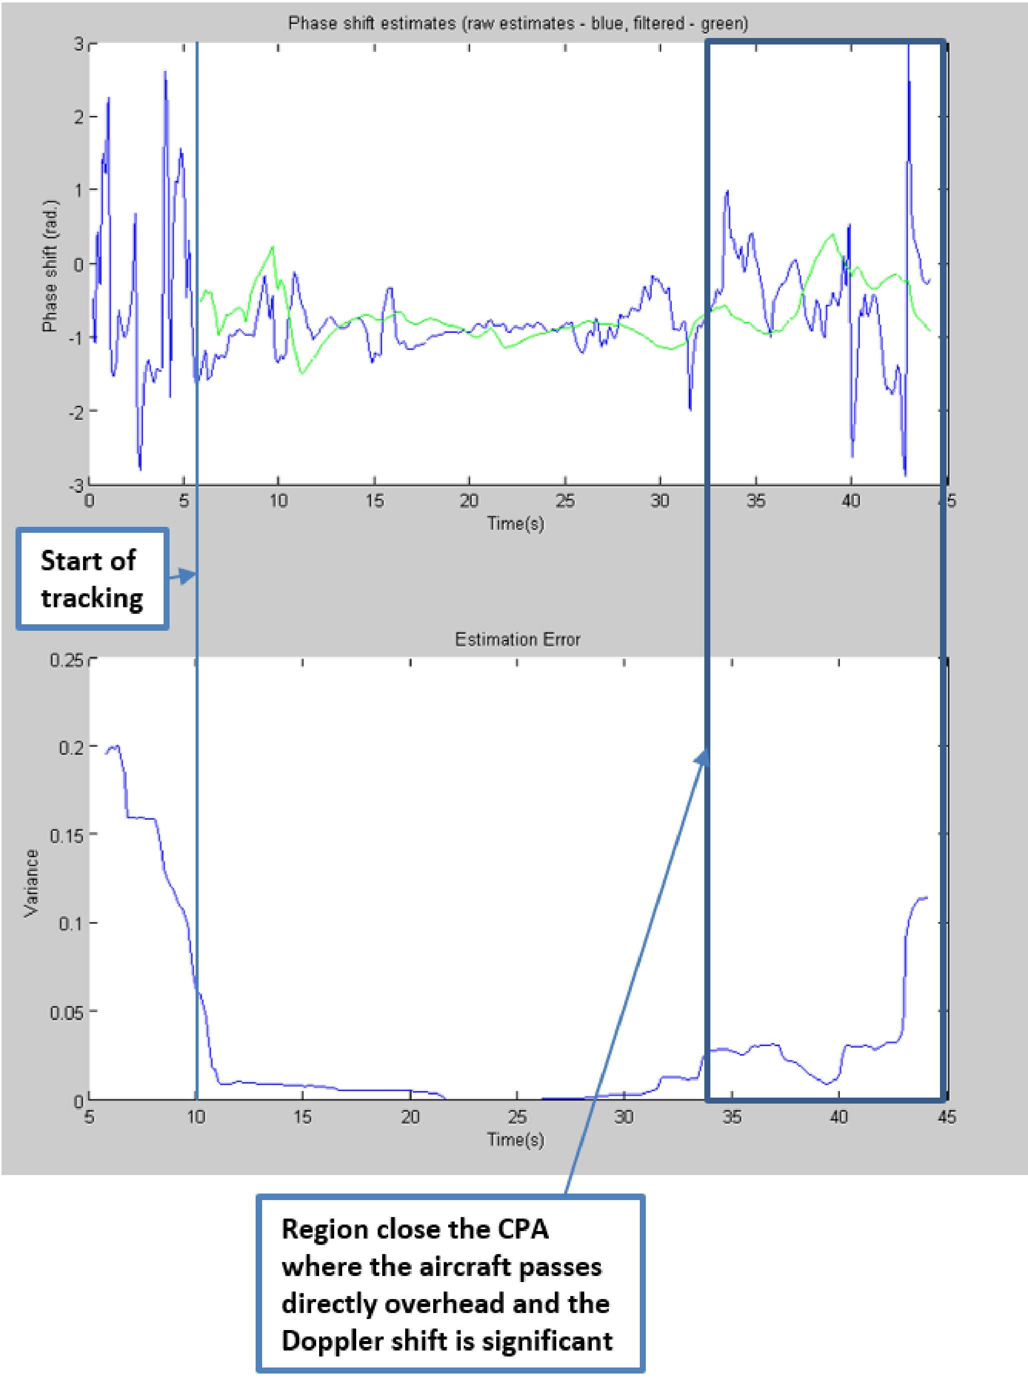

Figure 4. (Top) Phase shift between microphone 1 and 2 at the frequency corresponding to the approaching aircraft (raw estimates – blue, filtered – green). (Bottom) Estimation error (variance)

6. Conclusions

The potential range extension possible with this technique are evident but the sensitivity to changes in frequency is an impediment. Although frequency variations emanating from the approaching aircraft are beyond our control, models can be developed and refined in real-time that would allow us to track the frequency changes even at very low SNRs. Changes in aircraft propeller blade rate that directly affects the frequency of the narrowband tones over time are related and can then be modelled as Markov chains. Moreover, the changes in blade rate are gradual and unlikely to have discontinuities. These properties can be exploited and optimization algorithms can then be used to search the frequency space to find a path of minimal phase shift variations. This is a subject of on-going work currently being undertaken by the authors.

7. Future Work

The work done so far has taken advantage of a range of techniques in the time domain and the frequency domain. The spatial domain however, has not been fully explored. The effect of wind noise was also not fully investigated. Going back to the fundamental properties of the three major signal components, a distinct feature that differentiates the noise sources is the speed at which the signals propagate. Distinguishing between signals propagating at different speeds requires that we have a spatially distributed microphone setup. The current setup being used has four microphones that are spatially distributed but, these microphones are all independent and the way in which the airflow interacts with each diaphragm will be intrinsically different and completely independent of each other. Therefore, the propagations of turbules cannot be tracked spatially. This entails building microphones with a very large aperture (large diaphragms with a minimum diameter of about 50 mm). The pressure variations at different locations on this diaphragm needs to be measured in order to track signals propagating at different speeds. Pressure variations can typically be picked up by PolyVinyliDene Fluoride (PVDF) films with minimal effect on the performance of the microphone, in particular its dynamic response. Once the measurements are made, signals can be enhanced or rejected based on the speed of propagations in the same way that signals with different frequencies can be enhanced or rejected using a single microphone. This would intrinsically improve the SNR and hence have an effect on the signal processing. Currently, prototypes are being built and tested in wind tunnels to evaluate the performance and the potential gain in detection range.

On the signal processing front, an algorithm capable of tracking changes in the frequency of signal from the approaching aircraft is being tested. The potential improvements in detection range in the presence of frequency variations at very low SNR (<10dB) is essential for solving this problem given the various noise sources contaminating the useful signal. Algorithms are currently being tested on synthetic and trials data to validate the performance.

Acknowledgements

This work was funded by “The Sir Ross and Sir Keith Smith Fund” and the authors would like to acknowledge the funding support. We would like to thank Maurice Gonella and Jack Kormas of AAI/Aerosonde Pty Ltd and Dr Brian Ferguson of the Defence Science and Technology Organisation for their on-going support and collaboration on this project.

References

[1] Franklin, S. and Finn, A. *Acoustic Sense and Avoid for UAVs - Trials Results*, DASI, University of South Australia, Sir Ross and Sir Keith Smith Fund Report, 2012.

[2] Jameson, I. *Time delay estimation*, Defence Science and Technology Organisation, 2006.

[3] Knapp, C. and Carter, G. “The generalized correlation method for estimation of time delay,” *IEEE Transactions on Acoustics, Speech and Signal Processing*, **24**(4), 320-327, (1976).

[4] Ward, D.B., Lehmann, E.A. and Williamson, R.C. “Particle filtering algorithms for tracking an acoustic source in a reverberant environment,” *IEEE Transactions on Speech and Audio Processing*, **11**(6), 826-836, (2003).

[5] Ferguson, B.G. and Lo, K.W. “Turbo-prop and rotary-wing aircraft flight parameter estimation using both narrow-band and broadband passive acoustic signal processing methods”, *Journal of the Acoustical Society or America*, **108**(4), 1763-1771, (2000).

[6] Franklin, S., Finn, A. and Wyber, R. *Acoustic Sense & Avoid (Phase II): Real World Validation of Performance Envelope – Final Report*, DASI, UniSA, Sir Ross and Sir Keith Smith Fund Report, 2014.

[7] Finn, A. and Franklin, S. “Acoustic collision avoidance for UAVs”, Proceedings of the Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP) Conference, Adelaide, Australia, 6-9 December 2011.