

FM radio based bistatic radar

P.E. Howland, D. Maksimiuk and G. Reitsma

Abstract: An experimental bistatic radar system is described that detects and tracks targets to ranges in excess of 150 km from the receiver, using echoes from a non-cooperative FM radio transmitter. The system concept and limitations on performance are described, followed by details of the processing used to implement the system. An adaptive filter algorithm is described that is used to efficiently remove interference and strong clutter signals from the receiver channels. A computationally efficient algorithm for target detection using Doppler-sensitive cross-correlation techniques is described. A simple constant false alarm rate algorithm for target detection is described, together with a description of a Kalman filter based target association algorithm. Representative results from the system are provided and compared to truth data derived from air traffic control data.

1 Introduction

In this paper we describe an experimental radar system developed over the past eighteen months that detects and tracks aircraft by receiving and processing echoes from a single non-cooperative frequency modulated (FM) commercial radio station.

Of all the transmitters of opportunity available in the environment, broadcast transmitters represent some of the most attractive for surveillance purposes, owing to their high powers and excellent coverage.

At first glance, analogue television transmitters seem the obvious choice of illuminator, as they have very high equivalent radiated powers. However, despite the instant appeal of their pulse-like waveform structure, it is quickly found that the waveform is far from suited for radar usage when used in a conventional radar matched filtering approach [1]. Howland showed, however, that it is possible to exploit the Doppler and bearing information in echoes of the television video carrier signal to track aircraft at ranges of up to 260 km from the receiver and 150 km from the transmitter [2]. In this processing approach the receiver bandwidth is only a few kilohertz to capture the range of target Doppler shifts and is thus only a small fraction of the 5.5 MHz television waveform bandwidth. It is therefore referred to as 'narrowband processing'.

The disadvantages of that approach, however, all stem from the relatively low information content in the Doppler measurements. In order to locate the target at all, the system must observe the target's Doppler history for an extended time before there is sufficient information to locate the target. Even then, the use of nonlinear estimation techniques to calculate the target's trajectory mean that a good initial estimate of the target's location must be available. In the restricted case of a forward-scatter radar then the unique

geometry can be exploited to derive an analytical expression for the target's location [3] but for the more general bistatic problem it is necessary to resort to elaborate global optimisation schemes [2]. The former approach is limited in its operational applications and the latter is neither robust nor computationally efficient.

2 Wideband processing

By contrast, 'wideband processing' is defined in passive radar systems as the use of a receiver bandwidth that is comparable to the bandwidth of the waveform being exploited. For example, a typical FM radio broadcast occupies a bandwidth of about 100 kHz. Radar range resolution is approximately equal to $c/2B$ and hence we see that the FM broadcast offers a potential range resolution of up to 1500 m. It is clear that the signal offers useful target ranging information. Although we focus on the use of FM radio signals in this paper, it should be noted that the approach described here is generic and applicable to any transmission of opportunity with a reasonable ambiguity function, such as cell-phone transmissions or digital radio or television waveforms.

The suitability of a signal for target location is governed by more than its bandwidth, however, and of more importance is the ability of the radar receiver to unambiguously locate the target. It turns out that the noise-like characteristics of the FM radio signal are well suited for this, and it has an ambiguity function that approaches the idealised thumb-tack surface [4].

In order to achieve the processing gain necessary to detect weak target echoes in a background of noise and interference it is necessary to achieve an equivalent to the optimal matched filter processing used in conventional radar systems. Passing an echo through a matched filter is equivalent to the correlation of the radar echo with a delayed replica of the transmitted signal, and it is this approach that must be used directly in a passive bistatic radar system. As the transmitted signal is not available directly to the receiver, a dedicated receiver is required to collect the transmitted signal.

The greatest limitation on system performance is the interference received from the transmitter being used to detect aircraft. This unwanted direct signal correlates perfectly with the reference signal and produces range

© IEE, 2005

IEE Proceedings online no. 20045077

doi: 10.1049/ip-rsn:20045077

Paper first received 10th September 2004 and in revised form 13th January 2005

The authors are with NATO C3 Agency, Oude Waalsdorperweg 61, 2597AK The Hague, The Netherlands

E-mail: paul.howland@nc3a.nato.int

and Doppler sidelobes that are several orders of magnitude greater than the echoes that are sought. To detect anything but the closest of targets it is necessary to remove this signal, by both angular nulling with the antenna and adaptive echo cancellation in the receiver. However, eventually the dynamic range of the receiver limits the cancellation and so the principal limitation on system performance lies with the analogue-to-digital converter technology.

After correlation processing the radar has measurements of the target that are very similar to those of conventional radars: bistatic-range, bearing and Doppler. The accuracy of each of these measurements, however, is quite different. Range and bearing are a factor of ten or so worse than a conventional microwave radar owing to the lower bandwidth of the FM radio signal, whereas Doppler is two or three orders of magnitude more accurate (owing to the extended integration times possible with passive radar). The radar can exploit the excellent Doppler information to provide a resolution comparable to conventional radars and by simultaneously using multiple transmitters the system can achieve target location accuracies that may be even better.

Another feature of FM radio based bistatic radars is that they are simultaneously unambiguous in both range and Doppler. This useful property makes them ideal for detecting long-range high-speed targets, such as auroral disturbances in the ionosphere [5] or even man-made space objects.

3 Expected system performance

In our system we are exploiting a single, vertically-polarised FM radio transmitter located at Lopik, some 50 km behind the receiver. The transmitter has a mean ERP of 50 kW and frequency of 96.8 MHz. The transmitter is located on a 375 m mast and provides excellent long-range low-level illumination.

The receiver is located on the roof of the NATO C3 Agency in The Hague, and is approximately 20 m above ground level. The Agency is on the edge of some sand dunes which lead to the North Sea approximately 2 km away. The receiver comprises two vertically-polarised half-wave dipoles over a wire-mesh backplane 1.5 wavelengths by 1.5 wavelengths in size. The receiver antenna is steered so as to try to place the transmitter in a null in the antenna pattern to reduce the unwanted direct signal. The radar surveys a sector approximately 120° in azimuth and is steered at an angle 45° west of true north, looking over the North Sea towards the United Kingdom.




Fig. 1 Predicted system coverage




Fig. 2 Predicted signal-to-interference ratio

Tx ERP = 50 dBW, Tx height = 375 m, Tx-Rx distance = 50 km, Rx height = 75 m

Assuming a target with a radar cross-section of 10 m^2 and a cross-correlation processing gain of 47 dB, we predict surveillance over the region shown in Fig. 1. The upper illustration shows contours of coverage above 15 dB signal-to-noise ratio, whilst the lower illustration shows a slice in elevation along boresight. Both Figures include a simple model for the elevation lobing effects of both the transmitter and receiver, which result in the break-up of the coverage at longer ranges. This modelling seems reasonably accurate; we reliably observe aircraft at ranges of up to 150 km from the receiver, the range of the first deep null in the coverage in the Figure.

As mentioned above, the limiting factor in the performance of our passive coherent location (PCL) system is the direct-path interference received from the transmitter. This interference is up to 90 dB greater than the echo we would expect to see from a 10 m^2 aircraft at a range of 150 km, and illustrates the need for good interference rejection. The signal-to-interference ratio, after taking account of the antenna suppression, but before correlation or filtering, is shown in Fig. 2.

4 System overview

The system described in this paper was built on a low budget and is one of the simplest architectures that can be used to explore this technology. A block diagram of the system hardware is shown in Fig. 3 and of the processing algorithm in Fig. 4.

Reading from left to right, the signal is collected by a digital receiver system comprising of at least three channels. This allows for one reference channel and two surveillance channels for direction finding. An adaptive filter is applied to the two surveillance channels to reject the unwanted transmitter signal and then the digital data from the three channels are fed to the cross-correlator that outputs two amplitude-range-Doppler (ARD) surfaces. To achieve the necessary processing power, we use a small cluster of Intel Pentium-4 computers running Linux to perform the processing in real-time, in parallel. The system updates the complete surveillance region once every five seconds.

A conventional constant false alarm rate (CFAR) detection scheme is then applied to each ARD surface to determine the range and Doppler of each target.




Fig. 3 System hardware block diagram




Fig. 4 Signal processing block diagram

The complex amplitude of a target's echo received by each surveillance channel is then fed to the direction finding processor. With only two surveillance channels the direction finding system uses phase-interferometry to estimate the target bearing [2].

At this stage in the processing the system has determined the range, bearing and Doppler of a number of targets. In order to further process the data it is necessary to associate this plot data with individual targets and this is performed using a conventional Kalman filter tracker in the box marked 'Plot-to-target association (line tracking)'.

Finally, having associated plots-to-targets, the range/Doppler/bearing data for each target are processed by a nonlinear estimator to determine the target's location, speed and heading. Use of a nonlinear estimator allows optimum use of the Doppler information in this tracking process.

The following Sections now describe the processing in greater detail and conclude with some representative results from the experimental system.

5 Data collection

Data is collected on three antennas, one reference channel and two surveillance channels. These antennas are connected to Cubic Communications VXI-3570A digital receivers, which sample the data in quadrature in a bandwidth of 110 kHz. These data are then transmitted in real-time over a gigabit local area network (LAN) to the processing cluster. In the current configuration, one second's worth of data from all three receivers is collected and then processed every five seconds, the current bottle-

neck in the processing being a combination of the performance of the LAN and the adaptive signal cancellation (which is not parallelisable).

The first step in the processing is to adaptively cancel any unwanted direct signal from the surveillance channels.

6 Adaptive removal of the direct signal and surface clutter

Although the cross-correlation processing between the reference and surveillance channels causes any unwanted reference signal in the surveillance channel to be confined to the zero-Doppler and zero-range bin, the range and Doppler sidelobes of this autocorrelation function remain significant. At best, with a 1 second integration time and 50 kHz effective bandwidth, these will be 47 dB below the main autocorrelation peak. However, given that the direct signal may be 80–90 dB greater than the echoes themselves, this means that the sidelobes remain some 30–40 dB higher than the echoes we are seeking. This is compounded by strong surface clutter returns from the sea surface to a bistatic range of around 50 km.

It is therefore critical that the direct signal and clutter is removed from the surveillance channels before cross-correlation processing is attempted. An adaptive noise canceller like that in Fig. 5 is used. The goal of the canceller is to estimate the desired signal $d(n)$ from a noisy observation

$$x(n) = d(n) + w_1(n) \quad (1)$$

recorded by the surveillance antenna, where $w_1(n)$ is the unwanted interference. The signal from the reference




Fig. 5 Adaptive noise canceller structure

antenna, $w_2(n)$, is used to estimate interference. The task of the adaptive filter is to estimate $\hat{w}_1(n)$ from $w_2(n)$. Then, this estimate is subtracted from the signal from the primary sensor leaving only an estimate of the true echo signal

$$e(n) = x(n) - \hat{w}_1(n) \quad (2)$$

To implement the adaptive filter the joint process estimator algorithm is used. The structure of the filter is presented in Fig. 6 and consists of two parts:

- an adaptive M -stage lattice predictor
- an adaptive tapped delay line

The structure of the M -stage lattice predictor is shown in Fig. 7. The output signals at the m th stage are

$$f_m(n) = f_{m-1}(n) + k_m^* b_{m-1}(n-1), \quad m = 1, 2, \dots, M \quad (3)$$

$$b_m(n) = k_m f_{m-1}(n) + b_{m-1}(n-1), \quad m = 1, 2, \dots, M \quad (4)$$

where M is the predictor order. The variables $f_m(n)$ and $b_m(n)$ are the m th forward prediction error and the m th backward prediction error, respectively. The coefficient κ_m is the m th reflection coefficient.

From each stage of the filter, the backward prediction error $b_m(n)$, $m = 0, 1, \dots, M$ is connected to the input of a finite impulse response (FIR) filter.

The M -stage lattice predictor transforms the sequence of the correlated input samples $x(n)$, $x(n-1)$, $\dots$, $x(n-M)$ into a sequence of uncorrelated prediction errors $b_0(n)$, $b_1(n)$, $\dots$, $b_M(n)$. The second part of the filter uses the




Fig. 6 Joint process estimator




Fig. 7 Lattice predictor structure

backward prediction to estimate the desired signal $y(n)$. The first part of the filter is equivalent to the Gram-Schmidt algorithm, the second part of the filter is equivalent to a multiple regression filter.

Mathematically, the joint process estimator is described by two algorithms. The first algorithm, called the gradient adaptive lattice (GAL), is used to adjust the coefficients κ_m in the lattice predictor structure. In the second filter, the coefficients $h(n)$ are updated using the normalised LMS algorithm (NLMS). These algorithms are described in [6] and are not repeated here.

Although this filter structure may appear excessively complex, we have found experimentally that the eigenvalue spread in the correlation matrix prevents the direct use of simpler algorithms, because the convergence rate is too slow. The lattice predictor decorrelates the input data vector, thus the spread of the eigenvalues is smaller, the convergence rate of the filter is faster, and the reference signal is removed from echoes properly.

Our selection of the prediction order M for the filter does not use any formal method, like Akaike's information-theoretic criterion or minimum-description-length criterion, but instead relies on empirical observation. A value of $M = 50$ has been found to be optimum.

7 Target detection by cross-correlation

Having adaptively filtered out the direct path signal, it is necessary to search for the Doppler-shifted and time-delayed echoes of the targets. This processing step serves two distinct purposes within the radar:

- to act as a matched filter for the radar system and provide the necessary signal processing gain to allow detection of the target echo
- to estimate the bistatic range and Doppler shift of the target

This results in a range resolution (with one transmitter) of approximately 2–3 km (depending on the instantaneous

modulation of the signal, which depends on the radio programme content). The Doppler resolution is the reciprocal of the coherent integration time, and so typically 1 Hz, corresponding to a velocity resolution of around 1.5 ms^{-1} . For all intents and purposes there is no minimum or maximum unambiguous range or Doppler. The maximum range is set by the integration time (1 second gives a maximum range of 150000 km) and the maximum Doppler by half the sample rate of the signal, thus typically $\pm 150 \text{ kHz}$ or so, or about 750 times the speed of sound. When implementing the radar receiver, the designer can select the subset of ranges and Doppler shifts of interest.

In practice, the coherent integration time is limited by migration of the target out of the Doppler, and sometimes range, cell of interest. An integration time of around 1 second is optimal for most civilian air traffic and provides a processing gain of around 47 dB.

7.1 Algorithm concept

The range–Doppler estimation concept is illustrated in Fig. 8. The algorithm operates on a 1 second sample of data and generates Doppler-shifted copies of the reference signal that act as a bank of filters, each matched to a different target velocity.

The processing in Fig. 8 is analogous to the calculation of the ambiguity function and can be written in discrete time notation as

$$|\Psi(\tau, v)| = \left| \sum_{n=0}^{N-1} e(n) d^*(n - \tau) e^{j2\pi v n / N} \right| \quad (5)$$

where Ψ denotes the amplitude–range–Doppler (ARD) surface that we are seeking to calculate, $e(n)$ denotes the filtered echo signal and $d(n)$ represents the reference signal. The variable τ denotes the time delay corresponding to the bistatic time difference of arrival (TDOA) of interest and v denotes the Doppler shift of interest. The reference signal $d(n)$ can be weighted using a standard weighting function (such as those defined in [7]) before calculation of the ARD surface in order to reduce the range and Doppler sidelobes, at the expense of a broadened main peak and slight loss in processing gain.

Simplistically, the most obvious way to implement this processing would be to calculate the discrete Fourier transform of $e(n)d^*(n - \tau)$ for each range of interest. To compute the ARD surface, for each range of interest we must:

- rotate the elements of $d(n)$ and conjugate to obtain the required time delay, $d^*(n - \tau)$

- multiply the rotated $d(n)$ and $e(n)$
- calculate the FFT of $e(n)d^*(n - \tau)$
- discard data from Doppler bins not of interest

Note that this algorithm allows the calculation of a limited number of ranges, but all possible Doppler shifts (limited only by the sample rate).

8 An efficient implementation

The major drawback of the approach presented above is the excessive processing load owing to calculations of the Fast Fourier Transforms for long input signals. We resolve this issue by applying a decimation technique that allows us to discard data at Doppler frequencies we know targets do not exist, before calculating the Fourier transform. This modified integration algorithm utilises some extra processing steps to decimate the signal but greatly reduces the overall computation complexity with almost no loss in signal processing gain.

The algorithm can be summarised as follows:

Parameters:

$d(n)$ - a reference signal

$e(n)$ - an echo signal

p - an initial time delay, which plays an important role in the distributed version of the algorithm. For single CPU

$p = 0$

N_i - a number of range bins

R - a decimation factor

Computation:

$d_m(n) = \text{conjugate}(\text{delay_by}(d(n), p))$

$N_F = \lfloor N/R \rfloor$, a number of points in the FFT after decimation

for $k = 1, 2, \dots, N_b$

- $s(n) = d_m(n)e(n)$ - product of a conjugated, time delayed reference signal and an echo signal
 - $s_d(n) = \text{CIC}(s(n), R)$ - decimation by factor R
 - $s_d(n) = \text{LPF}(s_d(n))$ - out-of-band filtering
 - $S = \text{FFT}(s_d(n), N_F)$ - computation of the Doppler velocities
 - $\text{delay_by}(d_m(n), 1)$ - delay reference signal in time domain by one sample
 - $R_d(k) = S$ - building up a range–Doppler surface
- end for

The main difference between the basic algorithm and this modification is the presence of two additional functions in the processing path, a cascaded integrator-comb (CIC) filter and a low-pass FIR filter (LPF). The CIC is a very efficient implementation of a decimation filter and is described in [8] and [9].




Fig. 8 Cross-correlation concept

The algorithm works as follows. First, the product of the conjugated, time delayed reference signal with the echo signal is calculated. The signal, $s(n)$, enters the CIC structure where it is integrated using a one stage integrator, decimated by factor R , and then differentiated in a one stage comb section.

The integrator operates at the original sampling rate f_s . After decimation, the comb section operates at the reduced sampling rate of $f_{ns} = f_s/R$. Therefore, the length of the output vector from the filter is a factor of R smaller. Next, the decimated signal $s_d(n)$ is low-pass filtered in order to remove out-of-band frequencies.

Finally, calculating the FFT algorithm on the filtered data vector $s_d(n)$ results in all the Doppler velocities for all targets from a specific range of frequencies at that range.

In our system the following parameters are typical:

- receiver bandwidth, $BW = 110$ kHz
- input signals are quadrature sampled at $f_s = 195.313$ kHz
- decimation factor, $R = 128$
- sampling frequency after decimation $f_{ns} = 1525.88$ Hz
- cut-off frequency for the 5th order low-pass (symmetric) FIR filter, $f_c = 300$ Hz

The difference in computational complexity between the original and modified versions of the algorithm are as follows:

- $O(\text{CIC})$ - N complex additions for the one stage integrator, $N_d = N/R$ additions for the one stage comb integrator
- $O(\text{FFT})$ - $N_d \log_2(N_d)$ complex operations
- $O(\text{LPF})$ - $5N_d$ complex multiplications and $5N_d$ additions

The computation cost of the extra steps is

$$O_t = O(\text{CIC}) + O(\text{LPF}) + O(\text{FFT}) \quad (6)$$

For $N = 2^{18} = 262144$ samples, $R = 128$, $N_d = 2048$ the value of O_t is $2^{18} + 2048 \log_2(2048) + 10 \times 2048 \simeq 305123$ operations. The complexity of the algorithm described in Section 7.1 is $O_{\text{FFT}} = N \log_2(N) = 2^{18} \log_2(2^{18}) = 4718592$ operations. Therefore, the speed-up factor S_f is $S_f = (O_{\text{FFT}})/(O_t) = 15.46$.

To further increase the processing speed, we parallelise the algorithm by spreading the computation of different sets of time delays, p , amongst different computers. We are able to achieve a 1 second coherent integration in just under 1 second by using this algorithm and six 2.6 GHz Pentium-4 machines in parallel.




Fig. 9 Example results of cross-correlation processing

Figure 9 shows a typical amplitude–range–Doppler surface created using the adaptive interference and clutter processing and cross-correlation processing described above. Target echoes are visible in this Figure as the bright returns at different bistatic ranges and Doppler shifts.

Having generated the ARD surface it is then necessary to automatically detect the range and Doppler bins in which valid targets lie. This is performed using a constant false alarm rate (CFAR) algorithm.

9 CFAR detection algorithm

In order to maintain a constant probability of false alarm, the detection threshold changes according to an estimate of the noise variance. The conventional cell-averaging constant false alarm rate with guarded cells algorithm (CA-CFAR) is used [10]. The algorithm operates on the full ARD surface, first in the range domain, then in the Doppler. The optimum parameters were found empirically to be:

- number of cells for averaging $M = 10$, $M/2$ cells at each side of the cell under test
- threshold level $K_o = 3$ dB

10 Direction of arrival estimation

For our initial system we are implementing a simple angle estimation process using phase interferometry, as described in earlier work [2]. The angle of arrival of a target echo, Φ , is related to the phase difference of the received signal at the two surveillance antennas, Θ , by:

$$\Phi = \frac{2\pi d}{\lambda} \sin(\Theta) \quad (7)$$

where d is a distance between the dipoles, λ is the wavelength. In order to minimise any angular ambiguities the antennas are mounted half a wavelength apart. This gives a 180° ambiguity, targets behind the antenna and in front of the antenna cannot be distinguished, but in practise this is acceptable owing to the reasonable front-to-back ratio of the antenna gain pattern, which means targets behind the antenna are rarely detected. The phase of each echo on the ARD surface is calculated using the argument of their complex value. Any phase mismatch between the two channels is removed in software using a simple calibration coefficient.

11 Target association

Although many conventional air surveillance radars output raw detection data for tracking by an external system, it is better for a passive radar to internally track aircraft detections from each transmitter. This tracking is performed in the range–Doppler-bearing domain.

By using an internal tracker, the system is then able to forward the associated plot data for

- association of returns from different transmitters of the same target (in a multistatic system)
- target state estimation (described below)

In our experimental system we use the basic Kalman filter described in Section 1.5 of [11]. The measurement vector (8) comprises measurements of range, R_k , and Doppler, F_k and bearing, Φ_k , from the ambiguity surface, and the state vector (9) comprises range, range-rate, Doppler, Doppler-rate, bearing and bearing-rate.

$$\mathbf{z}(k) = (R_k F_k \Phi_k)' \quad (8)$$

$$\mathbf{x}(k) = (r(k)\dot{r}(k)f(k)\dot{f}(k)\phi(k)\dot{\phi}(k))' \quad (9)$$

We exploit the fact that the measurements of Doppler are proportional to the rate of change of range (10) and use a modified form of the state prediction (11)

$$\mathbf{F}(k) = \begin{pmatrix} 1 & 0 & -\lambda\tau & 0 & 0 & 0 \\ 0 & 0 & -\lambda & -\lambda\tau & 0 & 0 \\ 0 & 0 & 1 & \tau & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \tau \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \quad (10)$$

$$\mathbf{x}(k+1|k) = \mathbf{F}(k)\hat{\mathbf{x}}(k|k) \quad (11)$$

where τ is the update rate and λ is the wavelength. The fact that the Doppler is proportional to rate of change of range means that it is particularly easy for even the basic Kalman filter to track targets in the range–Doppler space.

For updating the state prediction covariance matrix (13), we use the standard definition of the state transition matrix (12) as follows:

$$\mathbf{F}(k) = \begin{pmatrix} 1 & \tau & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & \tau & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \tau \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \quad (12)$$

$$\mathbf{P}(k+1|k) = \mathbf{F}(k)\mathbf{P}(k|k)\mathbf{F}(k)' + \mathbf{Q}(k) \quad (13)$$

The tracker uses a standard 4-out-of-5 track initiation logic. A track is ‘preliminary’ until it has satisfied this condition, after which it is ‘confirmed’.

The association gate is defined according to (14)

$$[\mathbf{z} - \hat{\mathbf{z}}(x+1|k)]' \mathbf{S}(k+1)^{-1} [\mathbf{z} - \hat{\mathbf{z}}(x+1|k)] \leq \gamma \quad (14)$$

The gate threshold γ is set as 11.4 corresponding to a probability of 0.99 with three degrees of freedom. When maintaining preliminary tracks, we increase the gate size by a factor of 1.5 to increase the probability of association. To reduce the computational load associated with (14) we

first apply an ‘early gate’ to all the plots under consideration, to reject distance outliers. This simple gate simply rejects any plot more than 10.0 Hz, 3.0 km or 1 radian away from $\hat{\mathbf{z}}(k+1|k)$.

The basic logic of the tracker is:

- update all confirmed tracks with the closest plot to $\hat{\mathbf{z}}(k+1|k)$ falling within the association gate defined in (14). If no plots are present, rate-aid the track.
- using any remaining plots, update all preliminary tracks with the closest plot to $\hat{\mathbf{z}}(k+1|k)$ falling within the association gate defined in (14). If no plots are present, rate-aid the track.
- using any remaining plots, initiate new tracks

A graphical display of the filter’s prediction and covariance estimates was used to tune the filter parameters. Using this, the initial value of the covariance matrix was set as

$$\mathbf{P}_{\text{init}}(0|0) = \begin{pmatrix} 5.0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.0225 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.04 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.8 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.06 \end{pmatrix} \quad (15)$$

and the covariance matrix modelling process errors was set as

$$\mathbf{Q}(k) = \begin{pmatrix} 3.0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.02 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.05 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.8 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.6 \end{pmatrix} \quad (16)$$

where the basic units of measurement are km, Hz and radians. Note that the bearing parameters are still preliminary and will undoubtedly benefit from additional tuning.

Example results of this tracking process are shown in Fig. 10. In this display, an error ellipse is displayed around each track estimate, together with an arrow indicating the track prediction. Preliminary tracks are shown in light grey, whilst confirmed tracks are darker. The majority of preliminary tracks are owing to false alarms and are never promoted into confirmed tracks.




Fig. 10 Example target tracks in range–Doppler space

Even without any special logic or modifications of the tracker to cope with manoeuvring targets, the basic filter has been found to be extremely robust and original plans to implement an interacting multiple model (IMM) filter have not proved necessary.

12 Target state estimation

At the time of writing, the integration of bearing into the system is very new. We therefore use the simple trigonometrical transformation from bistatic range and bearing into a local Cartesian coordinate scheme for displaying the target on a geographical map. Example results are shown in Fig. 11. Whilst this serves to quickly show the detections on a map, it effectively discards all the excellent Doppler information provided by the system. We are therefore currently implementing a proper estimation of target location from range, bearing and Doppler by extending the earlier work of Howland [2] to modify the extended Kalman filter to include range (as well as Doppler and bearing) in the state equations. It will not be necessary to use the elaborate genetic algorithm or Levenberg–Marquardt approaches to initiate each track, initial Doppler and bearing measurements from the target association stage will be sufficient.

13 Verification of system performance

To verify the detection performance of the system we use a live civil air traffic control data feed. We read these data into a display program, filter the data geographically to cover only the region in which our system detects targets and then display them on a map (Fig. 12).

We also take these data and calculate the bistatic range and Doppler shifts that each target would present to our system. We display these data on a separate range–Doppler display (Fig. 13) with their Mode 3/A identifiers. Selecting an aircraft’s range/Doppler plot highlights the same aircraft on the map. Selected tracks on both displays are shown as squares, and these are visible on Figs. 12 and 13. We can also overlay the truth data directly on the system’s range–Doppler display (Fig. 14) as a quick check of which aircraft have been detected.




Fig. 11 Example target tracks overlaid on map




Fig. 12 ATC truth data with aircraft detected by passive radar (boxes)

These techniques provide a real-time verification of the system’s performance. In the examples provided here, it is possible to compare the detections in Fig. 14 with the truth tracks in Fig. 13. The latter tracks can then be related directly to the aircraft’s locations by comparing Mode 3/A codes between Fig. 12 and 13, or simply by looking for tracks highlighted with the squares. Furthermore, it is possible to compare the output of the passive radar’s target association tracker, in Fig. 10, with the truth tracks, in Fig. 13. In both cases we feel the results are encouraging.

Using this approach we have verified that the radar is able to reliably detect and track aircraft at bistatic ranges of beyond 300 km, corresponding to ranges of up to 150 km from the receiver. Track initiation is reliable and false target tracks are rarely observed.

At the time of writing we are still verifying the final track accuracy against Eurocontrol data. Accuracy is limited by the poor bearing accuracy, but by inspection it seems that most tracks are within 10 km of their true location, and often much closer. Once we have calibrated the bearing measurement system, further tuned the line tracker and implemented the extended Kalman filter estimation of target state we expect this to be much improved.




Fig. 13 Civil ATC truth data converted to bistatic range/Doppler space for comparison




Fig. 14 Example detections showing overlaid civil ATC data

14 Future work

We are currently extending the system to use an eight-channel circular array. This will offer 360° azimuthal coverage and the opportunity to implement adaptive beamforming techniques to assist in interference suppression. The use of adaptive beamforming will also bring the benefit of flexibility and the possibility to exploit multiple FM radio stations. We plan to do this sequentially, rather than simultaneously, and update the target-state estimation processing accordingly. The greatest challenge in the move to multiple transmitters will be the development of robust target-association logic, allowing echoes from different transmitters but the same target to be associated. However, the approach promises to offer far greater target location accuracy and more robust detection.

15 Conclusions

We have described a passive radar system that detects and tracks aircraft at ranges beyond 150 km in real time using simple computer hardware, a dipole antenna and a single FM radio station as the source of illumination.

The principal challenges in this work were the cancellation of the unwanted direct signal and surface clutter returns in the surveillance channels, and the development of a cross-correlation algorithm capable of processing the data in real time. We described both processes here.

We have shown how a standard CA-CFAR algorithm can be used to detect aircraft in the range-Doppler surface. We have shown how a simple four-state Kalman filter can be used to reliably associate target echoes in the range-Doppler space, before the subsequent estimation of the target's location, heading and speed using a nonlinear estimation technique, such as the extended Kalman filter or particle filter. We have verified the performance of our system using live air traffic control data.

16 Acknowledgments

The authors gratefully acknowledge the support of the Supreme Headquarters Allied Powers Europe (SHAPE) and Allied Command Transformation (ACT) for providing the funding under their Scientific Programmers of Work for this research. They also acknowledge the support of Christian Lievers and Albert Huizing of the TNO Fysisch en Elektronisch Laboratorium (TNO-FEL) in The Hague for the development of the receiver control software used in this project. They thank Dr Krzysztof Kulpa of the Warsaw University of Technology and RADWAR S.A. for suggesting the use of efficient cross-correlation processing techniques.

17 References

- 1 Griffiths, H.D., and Long, N.R.W.: 'Television based bistatic radar', *IEE Proc. F, Commun. Radar Signal Process.*, 1986, **133**, (7), pp. 649–657
- 2 Howland, P.E.: 'Target tracking using television-based bistatic radar', *IEE Proc., Radar Sonar Navig.*, 1999, **146**, (3), pp. 166–174
- 3 Blyakman, A.B., Ryndyk, A.G. and Sidorov, S.B.: 'Forward scattering radar radar moving object coordinate measurement'. IEEE Int. Radar Conf., May 2000, pp. 678–682
- 4 Griffiths, H.D., Baker, C.J., Ghaleb, H., Ramakrishnan, R., and Willman, E.: 'Measurement and analysis of ambiguity functions of off-air signals for passive coherent location', *Electron. Lett.*, 2003, **39**, pp. 1005–1007
- 5 Sahr, J.D., and Lind, F.D.: 'Passive radio remote sensing of the atmosphere using transmitters of opportunity', *Radio Sci.*, 1998, **284**, pp. 4–7
- 6 Haykin, S.: 'Adaptive filter theory' (Prentice Hall, Upper Saddle River, NJ, USA, 2002, 4th edn.)
- 7 Harris, F.J.: 'On the use of windows for harmonic analysis with the discrete fourier transform', *Proc. IEEE*, 1978, **66**, pp. 51–83
- 8 Hogenauer, E.B.: 'An economical class of digital filters for decimation and interpolation', *IEEE Trans. Acoust. Speech Signal Process.*, 1981, **ASSP-29**, (2), pp. 155–162
- 9 Willson, Jr., A.N., Kwentus, A.Y., and Jiang, Z.: 'Application of filter sharpening to cascaded integrator-comb decimation filter', *IEEE Trans. Signal Process.*, 1997, **45**, (2), pp. 457–467
- 10 Nitzberg, R.: 'Radar signal processing and adaptive systems' (Artech House, Boston, London, 1999)
- 11 Li, X., and Bar-Shalom, Y.: 'Multitarget-multisensor tracking: principles and techniques' (YBS, 1995, 2nd edn.)