

OPTIMISATION OF THE AMBIGUITY FUNCTION OF SHORT-DWELL ARBITRARY PULSED WAVEFORMS

Evan J Hughes and Clive Alabaster

Centre for Electronic Warfare
Department of Informatics and Systems Engineering
Cranfield University
Shrivenham, Swindon
SN6 8LA, UK
e.j.hughes@cranfield.ac.uk
c.m.alabaster@cranfield.ac.uk

Keywords: Pulse-Doppler Radar, Evolutionary Algorithm, Pulse Compression

Abstract

A method based on optimising the ambiguity diagram is described for the design of very short pulsed radar waveforms that still allow both target range and velocity to be measured. The coherent processing interval of the waveform is comprised of a sequence of modulated pulses whose inter-pulse timings and transmission frequencies are chosen independently for optimal performance. When optimised, the waveform structure allows measurement of both target range and velocity whilst maintaining good target visibility at all ranges/velocities. The optimisation process has been conducted using an evolutionary algorithm. The waveforms derived here demonstrate an interesting alternative to conventional low, medium or high pulse repetition frequency processing for situations such as in electronically scanned array antennas where short dwell times are often necessary.

1 Introduction

With modern electronically steered phased array radar systems, there is often a significant restriction on the available dwell period for each waveform as a very large volume of space is being scanned using a narrow pencil-beam antenna pattern. The requirement to detect high velocity targets at long ranges leads to a conflict of design choice for the waveform timings as a High Pulse Repetition Frequency (HPRF) is needed in order to measure velocity unambiguously, whereas a Low Pulse Repetition Frequency (LPRF) is needed in order to determine the unambiguous range. There is no single PRF choice that allows both target range and velocity to be measured unambiguously within a single dwell period. The compromise choice is to transmit a sequence of shorter pulse bursts, each at a different PRF in order to allow the ambiguities to be resolved. As multiple bursts must be transmitted within the total allowable dwell period, each of the bursts is therefore compromised on the maximum possible Coherent Processing Interval (CPI) and the number of pulses that can be integrated to improve the total energy on the target and therefore detection range.

In the field of pulsed radar waveforms, the pulse repetition frequency (PRF) falls into three regimes [1, 2]:

Low PRF. These are sufficiently low to avoid range ambiguities, however, velocity data is likely to be highly ambiguous. Low PRF affords very little ability to reject clutter, particularly from fast moving platforms since main beam clutter (MBC) and its ambiguous repetitions in the Doppler domain overwhelm returns at all Doppler frequencies.

High PRF. A high PRF waveform is one which is sufficiently high so as to avoid Doppler ambiguities, however range data is likely to be highly ambiguous. HPRF can reject MBC well but suffers from side lobe clutter (SLC) which becomes spread between $\pm V_P$, where V_P is the ground speed of the platform. Consequently, HPRF is ideally suited to velocity measurement applications. Since also a large number of pulse returns may be received in each coherent processing interval (CPI), large processing gains are possible leading to good long range detection performance.

Medium PRF. Medium PRF is ambiguous in both range and Doppler. MPRF is seen as a compromise between high and low PRF since it avoids the severe problems of MBC experienced by LPRF and the SLC related problems of HPRF. Medium PRF subdivides the beam dwell time into several CPIs (typically 6 to 8); each transmitting a different PRF, in order to decode the true range and velocity of targets and to ensure all round target visibility in conditions of clutter. In order to accommodate multiple CPIs, each CPI duration must be limited which, in turn, leads to more modest processing gains and detection ranges.

One possible solution to the problem of measuring both range and velocity unambiguously in a single dwell is to allow the spacing between each of the pulses to be chosen independently, allowing far more flexibility in the properties of the final waveform. The velocity processing can be performed using a sparse Discrete Fourier Transform (DFT) process which can be made tolerant of the arbitrarily chosen pulse timings. For an ideal waveform, the ambiguity diagram [3, Chapter 6] structure would approach a single central spike, surrounded by a low plateau.

Previous work on studying the benefits of using arbitrarily spaced pulse trains for performing measurement in both range and velocity [4] demonstrated that, for small num-

bers of pulses (< 10) it is possible to use optimisation procedures that use an ambiguity diagram of the waveform in order to produce waveform timings that have a useful degree of unambiguous visibility in both range and velocity simultaneously. The previous work studied a simple waveform structure where small numbers of pulses were used and every pulse had the same modulation characteristics and a consistent transmission frequency. Alternative ‘classical’ design approaches were compared to waveforms that were optimised and it was determined that the use of optimisation tools can indeed produce useful waveform timings. The previous work indicated that the issue of pulse eclipsing is often in conflict with the requirements of low-ambiguity range and velocity performance however, and if many more pulses are to be packed into a short dwell, the issues of minimising eclipsing whilst also optimising the ambiguity function need to be addressed.

This paper extends the previous technique significantly by considering the transmission of Linear Frequency Modulated (LFM) pulses that are offset in frequency to allow the effects of eclipsing to be partially decoupled from the range and velocity ambiguity diagram. The dwell times considered are also much shorter and the number of pulses employed much larger. The resulting optimisation problem is far more difficult and so evolutionary optimisation methods have again been applied to attempt to produce waveforms with a ‘thumb-tack’ like behaviour.

Section 2 describes the structure of the waveforms that are being optimised and the methods of processing for extracting target returns, Section 3 describes the optimisation process and Section 4 shows example performance results of an optimised waveform. Finally Section 5 concludes.




Figure 1: Waveform timings for envelope of transmitted pulses

2 Waveform Structure

The design aim is to derive waveforms which yield an ambiguity diagram as close to the ideal as possible, whilst trading ambiguities against eclipsing losses. A general pulse sequence of N pulses is illustrated in Figure 1, in which the pulse widths are $\tau_1, \tau_2, \dots, \tau_N$, the Pulse Intervals (PI) between the pulses are T_1, T_2, \dots, T_N and the start times of the rising edges of the pulses are t_1, t_2, \dots, t_N . The first pulse is always taken to start at zero time ($t_1 = 0$). The optimisation process also allows the total dwell time to be fixed as T .

The individual PIs may be chosen to be ambiguous in range or velocity, however to guarantee that the furthest target of interest has been illuminated by all of the pulses, the last PI of the pulse sequence determines the maximum unambigu-

ous range as shown in (1), where $c = 3 \times 10^8 m^{-s}$ is the speed of propagation.

$$R_{mu} = \frac{c(\tau_N + T_N)}{2} \quad (1)$$

The pulse start time of pulse k may be determined as in (2), where the timing of the first pulse is $t_1 = 0$.

$$t_k = \sum_{i=1}^{k-1} (\tau_i + T_i) \quad (2)$$

For this study, each of the pulses is modulated with a section of a long Linear Frequency Modulated (LFM) Chirp signal. The long chirp will generally be equivalent to 3 to 10 times the length of the longest transmitted pulse and is therefore never transmitted in its entirety. The section of the long chirp used to modulate each pulse is intended to be different to allow extra degrees of freedom to the optimisation process. The optimisation process can therefore select the start time and essentially the centre frequency of each chirped pulse. On reception, the received short pulses are compressed (i.e. cross-correlated) using the full long LFM chirp characteristics, rather than each of the individual short chirp sweeps. The result is a cross-correlation of the full long chirp sweep with each of the individual short pulse sweep returns. The resulting cross correlation profiles have interesting properties in that the central compression peak may appear offset in time and phase when compared to the actual transmission time (and phase) of each pulse. The pulse compression can be performed easily using a padded Fast Fourier Transform (FFT) process. If the zero padding of the reference long chirp and received signal is chosen carefully, an apparent shift in the pulse positioning can be created by inducing artificial range-Doppler coupling effects from the transmission of a range of LFM chirp sections.

$$\Delta R = \frac{c}{2B} \quad (3)$$

Figure 2 shows an enlarged view of an example pulse after compression. As the central section of the reference chirp was not used, the resulting compression peak is offset in time. The time-sidelobes are visible clearly and have been reduced from the usual sinc pattern by the application of a Hamming window envelope to the transmitted chirps.

For example if the long reference chirp has a total bandwidth of $B=20\text{MHz}$ and pulses of $\tau = 10\mu s$ are transmitted with a bandwidth of 7.5MHz (using (3) gives a range resolution $\Delta R = 20m$), the reference chirp has the same sweep rate as the shorter transmitted chirp and so would have an equivalent length of $T_c = 26.7\mu s$ if it was ever transmitted. After cross correlation, the correlation profile will have a total width of $(10 + 26.7)\mu s$ and the location of the central spike will be dependent on what the centre frequency of the transmitted chirp was. Figure 3 shows a diagram of the zero padding needed in order to allow FFT based pulse compression using the received signal of length T and the theoretical reference long chirp with the same sweep rate as the transmitted pulses which would have an equivalent length of T_c . The compression process is performed via FFT as

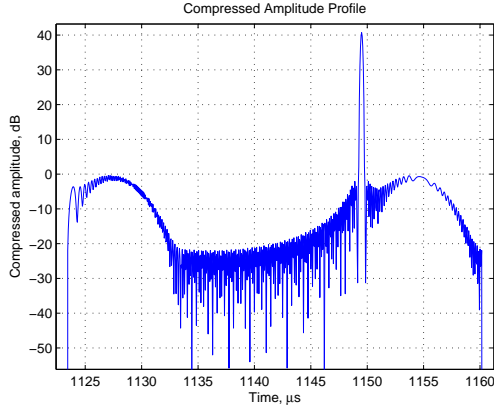


Figure 2: Enlarged view of a pulse after pulse compression. The shift in the location of the pulse compression peak is clear, along with the pulse compression artefacts which now occupy a time period of $\tau_i + T_c$.

detailed in (4), where $S_{Rx}(t)$ denotes the received complex samples at time t (down converted to baseband), $C_{ref}^*(t)$ denotes the complex conjugate of the long reference chirp waveform (baseband definition prior to multiplication with carrier) defined in (5) and $S_c(t)$ denotes the complex pulse-compressed received waveform.




Figure 3: Zero padding of waveforms to be cross correlated using a FFT. The upper trace shows the reference chirp of theoretical length T_c which is post-padded with zeros of length T (the waveform length in Fig. 1). The lower trace shows the received signal of length T that is pre-padded with a section of zeros of length T_c . Note that the shading of the pulses in the lower trace indicate that they are from different sub-sections of the larger chirp in the upper trace.

$$S_c(t) = \text{IFFT}\{\text{FFT}(C_{ref}^*(t)) \times \text{FFT}(S_{Rx}(t))\} \quad (4)$$

$$C_{ref}(t) = \exp\left(j2\pi t \left(\frac{Bt}{2T_c} + \frac{B}{2}\right)\right) \quad (5)$$

In practice the pulse compression process is performed on sampled data. The received signal must be sampled at a rate commensurate with the total bandwidth of the reference chirp (signal is complex so sampling at the bandwidth provides sufficient total samples to satisfy the Nyquist criterion). Thus in the previous example, the received data would be sampled at $B = 20\text{MHz}$, i.e. the sampling will lead to the reference chirp consisting of $N_c = B \times T_c$ samples. After pulse compression, if the zero padding strategy in Figure 3 is used, the sample indices i_k of the peak of the k^{th} pulse compressed return is given by (6), where δ_k is the time offset from the start of the reference chirp for the section of the reference chirp that was actually transmitted.




Figure 4: Behaviour of a target return when observed with an arbitrary waveform. At each time offset t_r , the DFT process only needs to consider the N received cells corresponding to the transmitted pulses.

In practice, δ_k is chosen in the range $\delta_k \in [0, T_c - \tau_k]$ to ensure that an entire chirp segment is captured.

$$i_k = B(t_k + (T_c - \delta_k)) \quad (6)$$

2.1 Sparse DFT processing

Target extraction is performed by correlating the transmitted pulse train with the received echo data at each possible range offset of interest as shown in Figure 4. As pulse compression has been applied first, the targets are well isolated, reducing the processing required for the pulse-train correlation step as only the range-bins where pulses are expected need to be included in the correlation process. For each range offset, all Doppler phase profiles for the velocity span of interest are correlated, allowing a range-velocity map to be constructed.

The process segments copies of the received data by offsetting by each of the pulse start times (modified by the compression process) and stacks them, allowing all the range cells to be identified and processed as a block (similar to figure 5). The DFT process of pulse-compressed returns may be described as in (7), where $r_{R,v}$ is the magnitude of the complex DFT correlation output at range R and velocity V . Note that (7) copes with unequal time sampling, as opposed to standard DFT techniques which assume a fixed pulse repetition interval.

$$r_{R,v} = \left| \sum_{k=1}^N S_c\left(t_k + (T_c - \delta_k) + \frac{2R}{c}\right) \exp\left(-\frac{j4\pi v}{c}\right) \right| \quad (7)$$

The segmentation of the pulse compression waveform introduces an effective time offset into the location of the compression peak which persists into the processing of the DFT. Thus the location of the compression peak is controlled not only by the pulse start time, but also by the choice of chirp segment offset time.

2.2 Eclipsing

The number of pulses eclipsed at any range R may be calculated as in (8), where n_R is the number of pulses eclipsed at range R , t_m and t_k are pulse transmission start times of the N pulses, τ_k is the length of pulse k and c is the speed of propagation. The summations are a summation of Boolean values, given the logical outcome of the inner expression. The summation process is shown graphically in

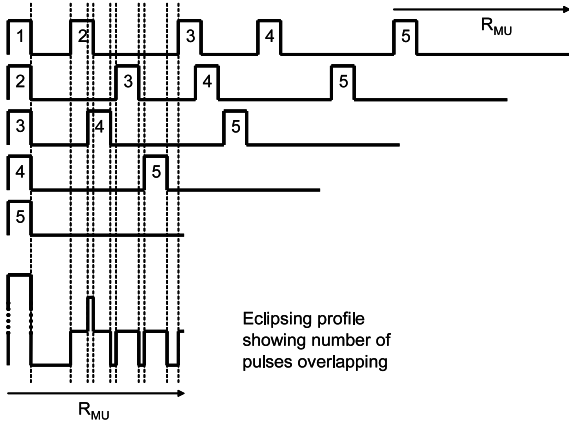


Figure 5: Separating pulse train into individual fragments for each transmitted pulse envelope in order to establish overall eclipsing profile.

Figure 5. The locations of regions of eclipsing are governed only by the start times and lengths of pulses; by allowing the choice of chirp fragment frequency in transmission, the performance of the DFT process may be almost completely decoupled from the issues of eclipsing.

$$n_R = \sum_{k=1}^{N-1} \sum_{m=k}^N \left((t_m - t_k) \leq \frac{2R}{c} \leq (t_m - t_k + \tau_k) \right) \quad (8)$$

3 Optimisation Procedure

Evolutionary Algorithms (EAs) are optimisers which use an approach inspired by the natural phenomenon of biological evolution [5] and directly exploit the Darwinian concept of ‘survival of the fittest’ where the best specimens of a species live long and produce many offspring while the weaker members of the population die young and have few or no offspring. Successive generations tend to become dominated by the best features from previous generations. In the context of mathematical optimisation, each member of a population is a potential solution to the optimisation problem, which in this paper is a waveform structure. The ‘fitness’ (quality) of each population member is measured by an ‘objective function’, which in this paper is to minimise the velocity and range ambiguity of the waveform whilst also limiting the effects of range eclipsing.

EA optimisation is an iterative process with each iteration representing one ‘generation’. Each generation should comprise members that are generally fitter than their predecessors as a result of selective breeding and replacement. As the population becomes fitter as a whole, the individual members will begin to appear very similar to one another until the population eventually converges to an optimal solution. We found that the best results were obtained using an ‘Evolutionary Programme’ (EP) [5]. In the terminology of the EP, each population member is called a ‘chromosome’ which comprises a number of individual ‘genes’. In this paper, a chromosome represents the PI sequence and chirp offsets of a transmitted waveform. The optimisation process performed by the EP is:

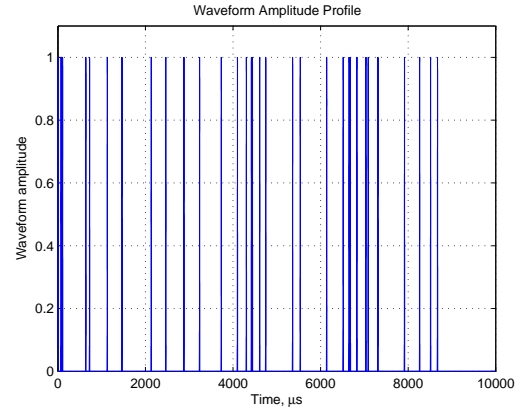


Figure 6: Amplitude profile of optimised schedule.

1. Generate an initial population of random chromosomes,
2. evaluate the fitness of each member of the population,
3. select a subset of the population to be allowed to ‘reproduce’,
4. combine the selected subset of ‘parents’ into ‘offspring’,
5. introduce some ‘mutations’ (changes) into the new offspring,
6. select the best P solutions from the sorted set of the parents and offspring.
7. repeat from step 2 for a given number of generations.

For simplicity, the transmitted pulse lengths, τ , have been kept constant and therefore the chromosome structure used in this paper consists of two sections. The first section has $N-1$ timing genes, each gene being in the range $g_t(i) \in [0, 1]$ and the pulse start times are established using (9) and (2). The last interval, T_N , is set by the required unambiguous range and so is not encoded in the chromosome.

$$T_i = \frac{g_t(i)}{\sum_{m=1}^{N-1} g_t(m)} \left(T - \frac{2R_{mu}}{c} - \sum_{k=1}^N \tau_k \right) \quad (9)$$

The second section of the chromosome consists of chirp genes, $g_c(i)$, and the chirp section used for transmission, $C_i(t)$, at pulse i is given in (10) and $t = 0 \dots \tau_i$ is used to create the chirp.

$$C_i(t) = \exp \left(j2\pi t \left(\frac{Bt}{2T_c} + g_c(i)(T_c - \tau_i) \right) \right) \quad (10)$$

The total chromosome structure is therefore $2N-1$ real-valued genes. The pulse start timings and chirp offset timings are quantised to the sample resolution defined by the chirp bandwidth B before the waveform is analysed. To reduce the effects of spectral spreading during pulse compression, each of the transmitted chirps is amplitude weighted using a Hamming window [3].

The evolutionary programme used a working population of 20 chromosomes and ran for 500 generations. The offspring were created by creating a copy of each parent chromosome in turn and then crossing over 70% of the genes

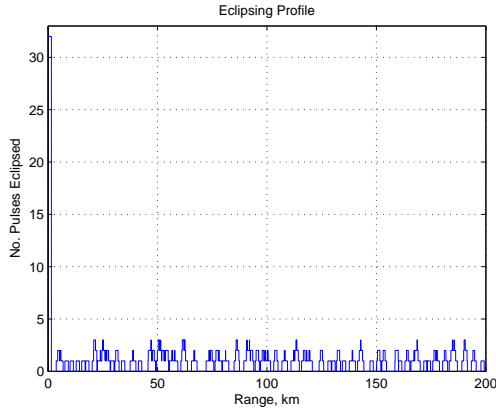


Figure 7: Eclipsing profile of optimised waveform.

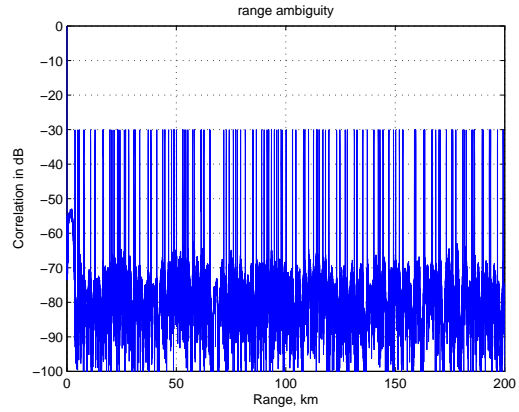


Figure 8: Range ambiguity profile of optimised waveform (vertical scale in dB).

with genes from another chromosome chosen at random from the parent population; which genes were swapped was chosen at random. The offspring chromosome was then altered through mutation where 90% of the gene values were perturbed by a small random value which was generated from a Gaussian distribution. The EP was run multiple times and although on each run the solutions differed (suggesting the results are good local optima and not necessarily the global-optimum solution), they were of similar performance.

4 Example Optimised Waveform

4.1 10GHz 10ms Dwell, 32 Pulses

The evolutionary process was set to search for waveforms for a radar with a carrier of 10GHz, $R_{mu}=200\text{km}$, range resolution of 20m, 32 transmitted pulses all with lengths of $\tau = 10\mu\text{s}$, total dwell of 10ms and a velocity space of interest of $\pm 1500\text{ms}^{-1}$. The structure of the waveform has an average duty cycle of 3.2%.

The waveforms have all been assessed based on the pattern of eclipsing in range and the overall range-velocity ambiguity diagram. If DFT processing is applied as in (7), the ambiguity diagram may be formed as the product of the magnitude of the DFT profile in range for a zero velocity target, multiplied by the magnitude of the DFT profile in velocity for a zero range target. The resulting matrix may now be scaled by the maximum peak value in the matrix (i.e. the central peak) and provides the ambiguity diagram. For clarity, the range and velocity slices are presented.

Figure 6 shows the amplitude profile of an example optimised waveform. It is clear that the pulse intervals are non-uniform, helping to prevent ambiguities in the range and Doppler processing. Figure 7 shows the eclipsing profile of the optimised waveform. At time zero, all of the pulses have a presence hence there are all 32 pulses visible as overlapping to the left of the figure. The number of pulses eclipsed at any one range varies with time over the full R_{mu} of interest but the maximum level has been minimised by the evolutionary algorithm. Figure 8 shows the range profile of the ambiguity function, plotted on a logarithmic vertical

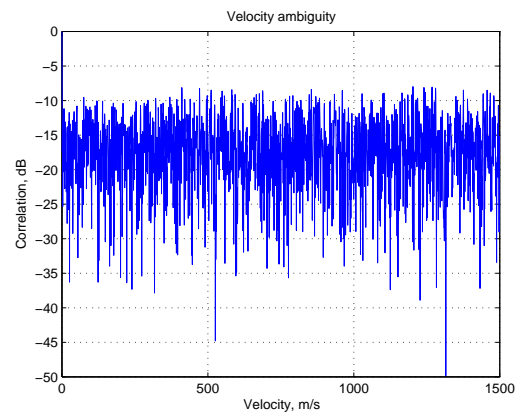


Figure 9: Velocity ambiguity profile of optimised waveform (vertical scale in dB)

axis. The profile has the main correlation peak at 0dB on the far left of the figure, falling rapidly as range advances. Overall the ambiguity profile shows a good separation between the target detections and the time-sidelobe artefacts of the pulses in the waveform.

Figure 9 shows one side of the velocity ambiguity function profile (it is symmetric), plotted on a logarithmic vertical axis. The profile has the main correlation peak at 0dB on the far left of the figure. The correlation behaviour falls rapidly and is very variable with velocity out to the maximum velocity of interest. The Doppler sidelobe artefacts are higher in the velocity domain than in range however the combined sidelobe attenuation levels will provide useful detection of high-flying targets in a multi-function radar system. As the number of pulses used, N , is increased, the level of the time and Doppler sidelobe artefacts reduces, enhancing performance further. The evolutionary method is capable of optimising waveform structures with hundreds of pulses if necessary.

As the number of pulses is increased for a given waveform length T , the eclipsing profile and range ambiguity function become more difficult to optimise, however the velocity ambiguity profile generally becomes simpler to opti-

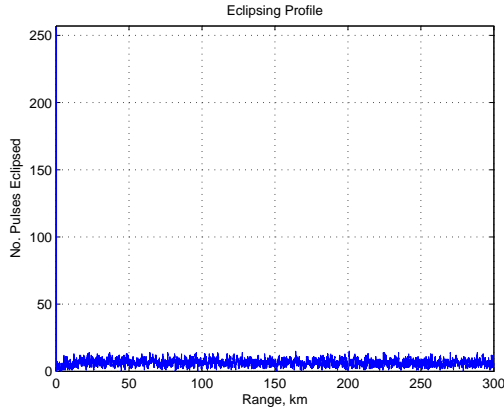


Figure 10: Eclipsing profile of optimised waveform.

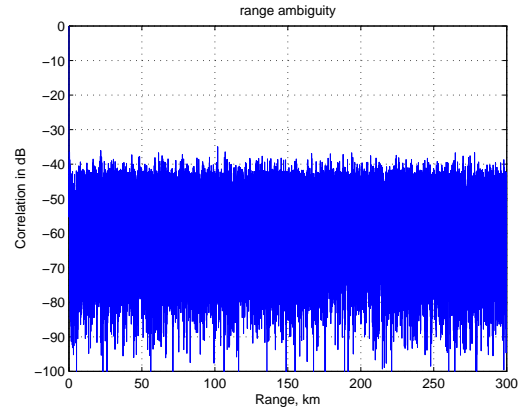


Figure 11: Range ambiguity profile of optimised waveform (vertical scale in dB).

mise as the individual pulse unambiguous velocity regions in Doppler due to the shrinking pulse timing intervals become larger.

4.2 1GHz 22ms Dwell, 256 pulses

The evolutionary process was reconfigured to search for waveforms for a radar with a carrier of 1GHz, $R_{mu}=300\text{km}$, range resolution of 20m, 256 transmitted pulses all with lengths of $\tau=2\mu\text{s}$, total dwell of 22ms and a velocity space of interest of $\pm 1500\text{m/s}^{-1}$. The structure of the waveform has an average duty cycle of 2.3% and represents a waveform that would be appropriate for a surveillance radar with a 2° beamwidth and a 4 second scan.

Figure 10 shows the eclipsing profile of the optimised waveform. At time zero, all of the pulses have a presence hence there are all 256 pulses visible as overlapping to the left of the figure. The evolutionary algorithm has again worked well to minimise the range eclipsing. % Figure 11 shows the range profile of the ambiguity function, plotted on a logarithmic vertical axis. The profile has the main correlation peak at 0dB on the far left of the figure, falling rapidly as range advances. Overall the ambiguity profile again shows good time-sidelobe performance.

Figure 12 shows one side of the velocity ambiguity function profile. The profile has the main correlation peak at 0dB on the far left of the figure. The Doppler sidelobe artefacts are lower than in the 32-pulse example and allow the velocity of targets to be quantified well. The waveform is thus suitable for use in a classic medium to long-range surveillance radar application, yet easily provides unambiguous measurements for both range and velocity.

5 Conclusions

The optimisation process has yielded range-velocity capable waveform designs that are practical to transmit and process, yet are temporally very short. The eclipsing losses have been optimised to a low level, aided by allowing the transmission of chirp sections which help to decouple the behaviour of the ambiguity function from the eclipsing profile. In comparison to classic HPRF or MPRF wave-

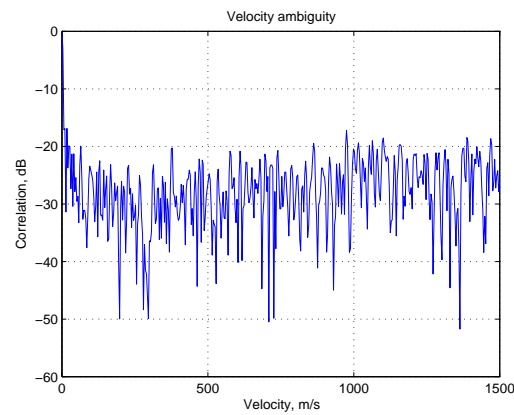


Figure 12: Velocity ambiguity profile of optimised waveform (vertical scale in dB)

forms, the short-dwell arbitrary waveforms would allow the surveillance volume to be scanned far more frequently, allowing tracking performance to be improved and more flexible time-management systems to be implemented.

References

- [1] C.M.Alabaster, *Pulse Doppler Radar: Principles, Technology, Applications*. SciTech Publishing, 2012.
- [2] C. M. Alabaster, E. J. Hughes, and J. H. Matthew, "Medium PRF radar PRF selection using evolutionary algorithms," in *IEEE Trans. Aerospace and Electronic Systems*, vol. 39, no. 3, July 2003, pp. 990–1001.
- [3] B. R. Mahafza, *Radar Systems Analysis and Design Using MATLAB*. Chapman & Hall / CRC Press, 2000.
- [4] C. M. Alabaster and E. J. Hughes, "Arbitrary pulsed radar waveform," in *Waveform Diversity and Design*, Lihue, Hawaii, USA, Jan. 2012.
- [5] K. Deb, *Multi-objective optimization using evolutionary algorithms*. John Wiley & Sons, 2001, ISBN 0-471-87339-X.