Radar Waveform Optimisation as a Many-Objective Application Benchmark

Evan J. Hughes

Department of Aerospace, Power and Sensors,
Cranfield University, Shrivenham, Swindon,
Wiltshire, England. SN6 8LA
e.j.hughes@cranfield.ac.uk

Abstract. This paper introduces a real, unmodified *Many-Objective* optimisation problem for use in optimisation algorithm benchmarking. The radar waveform design problem has 9 objectives and an integer decision space that can be scaled from 4 to 12 decision variables. Proprietary radar waveform design software has been encapsulated in a fast and portable form to facilitate research groups in studying high-order optimisation of real engineering problems.

1 Introduction

Real engineering problems are often characterised by many objectives. Pareto ranking has been exploited in recent years to develop a large number of excellent multi-objective optimisation algorithms which can solve bi-objective optimisation problems effectively and reliably, for example, NSGA-II [2]. However, it is known that Pareto ranking alone does not scale well to problems with large numbers of objectives (4+ typically cause problems) [9, 5]. Currently, there are few algorithms that are designed specifically to tackle many-objective problems.

This paper describes a real, unmodified engineering problem and the software is provided to allow optimisation with many-objectives to be studied, and hopefully efficient optimisation algorithms developed.

The problem has 9 objectives, and from 4 to 12 integer decision variables, each in the range [500,1500] inclusive, giving 1001 alleles per decision variable. It is known that some of the objectives are not totally independent, and it is suspected that the Pareto set is concave in places, with regions of low density.

The problem is the design of a waveform for a *Pulsed Doppler Radar*, typical of many airborne fighter radar systems. The radar system is required to measure both range and velocity of targets. Unfortunately, with the very long ranges (100 nautical miles typical) and very high velocities (Mach 5 possible), with a simple waveform it is only possible to measure either: range unambiguously but ambiguous velocity; velocity unambiguously but with the range ambiguous; or with both range and velocity ambiguous. For example, if velocity is measured, then target range may only be known modulo by say 5 kilometres, i.e. a target at 108km would appear at 3km.

To allow full unambiguous measurements, a set of simple waveforms is transmitted, each subtly different from the last. The results of the multiple waveforms are then combined in order to resolve the ambiguities. The problem is how to choose the set of simple

waveforms. Previous work in this area has led to the development of an evolutionary algorithm capable of designing practical waveforms [1].

This radar waveform design problem is interesting in that in a practical radar system, an entire set of non-dominated solutions would need to be created prior to each mission. While the radar is active, it will choose a general location on the non-dominated surface, based on current radar operating conditions, then select a waveform randomly which is local to this chosen location. The random choice helps prevent 3rd-party interception of the waveform as it is changing constantly, yet the waveform is biased towards an optimal radar configuration. Thus the radar chooses its operating point on the non-dominated surface dynamically on-line, from a non-dominated set that is likely to remain fixed for each mission.

An initial analysis of the properties of the objective surface has been performed and a demonstration of the typical behaviour of two different optimisation algorithms, NSGA-II and MSOPS, on the function is presented.

Section 2 details the radar design problem and section 3 describes the format of the objective function software. Section 4 introduces initial results from analysing the non-dominated surface and section 5 describes the results of comparing the performance of two example optimisation algorithms. Finally section 6 concludes.

2 Radar Waveform Design

2.1 Introduction

Radar systems are categorised by the rate at which they transmit pulses of energy toward the target, called the Pulse Repetition Frequency, or *PRF* [10]. There are three broad categories: Low PRF with few pulses (20 typical) and big gaps between them (1 millisecond typical); High PRF with many pulses (thousands) and short gaps (few micro seconds); and Medium PRF where there are a moderate number of pulses (64 typical) and moderate gaps (100 μS typical).

Low PRF radar systems can measure range exactly, but velocity measurements are ambiguous for any velocities greater than the maximum unambiguous velocitiy V_{mu}, given in (1) where F_{prf} is the pulse repetition frequency in Hertz and λ is the wavelength of the transmitted pulses.

$$V_{\text{mu}} = \frac{F_{\text{prf}}\lambda}{2} \tag{1}$$

The maximum unambiguous range of the radar is given by (2), where $c \approx 3 \times 10^8\,\text{ms}^{-1}$ is the speed of propagation of the pulse.

$$R_{\text{mu}} = \frac{c}{2F_{\text{prf}}} \tag{2}$$

A typical Low PRF radar may have a PRF of 1kHz, yielding a maximum unambiguous range of $R_{\text{mu}} = 150\text{km}$ and a maximum unambiguous velocity of $V_{\text{mu}} = 15\text{ms}^{-1}$, assuming a 10GHz transmission frequency ($\lambda = c/F_{\text{TX}}$, therefore $\lambda = 0.03\text{m}$).

The main advantage of low-PRF radar is the ability to measure target range directly using simple pulse delay ranging. However, low-PRF radar suffers from a lack of velocity visibility, since unwanted ground returns and undesired slow moving targets get

folded over and over into the small velocity window. Low-PRF radar is best suited to operation in the absence of ground clutter returns, for example where a radar is looking up at high-flying targets, rather than looking down and the radar beam is striking the ground.

A typical High PRF radar may have a PRF of 100kHz, yielding a maximum unambiguous range of $R_{\text{mu}} = 1.5$km and a maximum unambiguous velocity of $V_{\text{mu}} = 1500\text{ms}^{-1}$, assuming a 10GHz transmission frequency.

The principle advantage of high-PRF radar, is the ability to detect high closing-rate targets, in what is essentially a noise-limited environment. However, detection performance is poor in tail aspect (low closing-rate) engagements, where targets compete directly with the velocity spectrum of the sidelobe clutter, where transmissions out of the side of the antenna beam strike the ground and provide echoes back. Furthermore, the highly ambiguous range response causes the sidelobe clutter to fold within the ambiguous range interval. Consequently, sidelobe clutter can only be discarded by resolving in velocity. High PRF radar is very good where small relative velocities are not often seen, or when exceptionally good antennas are available that have very little spurious radiation out of the side of the beam.

Medium-PRF radar is a compromise solution designed to overcome some of the limitations of both low and high-PRF radar. By operating above the low-PRF region, the ambiguous repetitions of the ground clutter spectrum may be sufficiently separated without incurring unreasonable range ambiguities. Consequently, the radar is better able to reject mainbeam clutter when in a look-down scenario through velocity filtering without rejecting too many targets. By operating below the high-PRF region, the radar's ability to contend with sidelobe clutter in tail-chase engagements is improved. Targets may now be extracted from sidelobe clutter using a combination of velocity filtering and range gating.

For example, the mainbeam clutter may be 20ms^{-1} wide, so can be 'notched' out as long as $V_{\text{mu}} > 20\text{ms}^{-1}$. However, if targets are folded in to the notched region, they cannot be detected and the region is said to be *blind*. When the pulse is transmitted, the receiver is turned off to protect it and the ranges at multiples of Rmu are now *eclipsed* and no targets may be detected here either. A second eclipsing region may also be applied to help reject the effects of the sidelobe clutter from the ground. The region will extend from a range that is just shorter than the aircrafts altitude (i.e. the first range at which an echo from the ground could occur) to often a few kilometres ahead of the altitude return.

A typical Medium PRF radar may have a PRF of 10kHz, yielding a maximum unambiguous range of $R_{\text{mu}} = 15$km and a maximum unambiguous velocity of $V_{\text{mu}} = 150\text{ms}^{-1}$, assuming a 10GHz transmission frequency. However, in many applications, both range and velocity will now be ambiguous.

Medium PRF radars possess excellent clutter rejection characteristics which render them an attractive proposition for airborne intercept (AI), fire control systems, ground based air surveillance, weapon locating radar and a variety of other applications.

2.2 PRF Selection

Each PRF is characterised by regions of blind velocities and ranges associated with the velocity filtering of mainbeam clutter and time gating of sidelobe clutter and associated eclipsing losses. These blind zones are depicted in black on a blind zone map, as in figures 1 & 2.

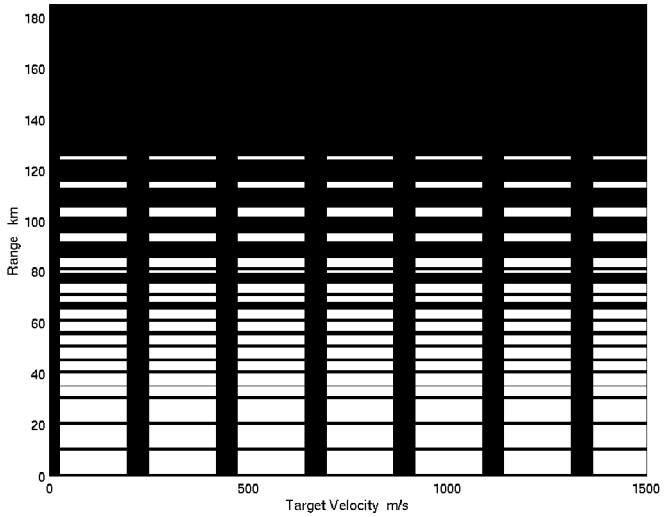

Fig. 1. Blind zones for a single, clutter limited, medium PRF waveform with PRF 14.9kHz

Multiple bursts of pulses are required in order to perform target detection and to resolve range and Doppler ambiguities. This is achieved by transmitting burst of pulses at a number of PRFs within the dwell time on target and sequentially measuring and comparing the ambiguous information received from every PRF. For example, eight different PRFs may be used but must all be able to be transmitted sequentially within the dwell time on the target, with each PRF burst having 64 processed pulses (64-point Fast Fourier Transform (FFT)) and a short period of time in which to change over PRFs. In practice, the change-over time is to allow the first pulse to reach, and return, from the furthest possible target of interest. Thus extra pulses are transmitted in a process termed *Space Charging*. For example, if the maximum range was 185km, and $R_{\mathrm{mu}} = 15km$, 13 extra pulses would be sent giving a total of 77, but only 64 would be processed, making sure that 64 returns from both the closest and furthest targets were contained within the processing window.

The positions of blind zones vary with PRF, therefore, by applying suitable PRFs in a multiple-PRF detection scheme, not only may range and Doppler ambiguities be resolved, but also the blind zones may be staggered to improve target visibility. Ground clutter returns received through the antenna sidelobes may be strong enough to over-

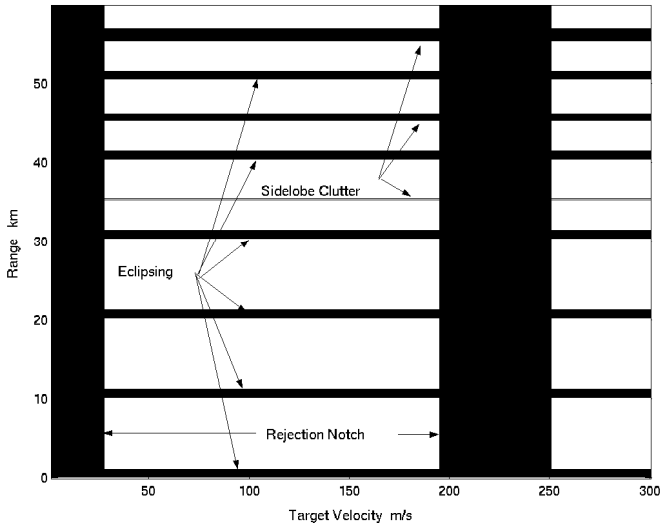

Fig. 2. Expanded view of Blind zones of Fig. 1

whelm weak target signals, consequently blind ranges tend to worsen with increasing range, as shown in figure 3.

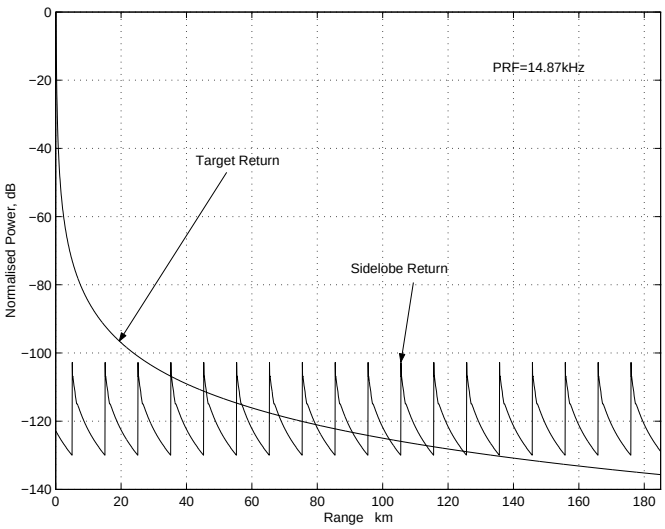

Fig. 3. Comparison of target return and sidelobe clutter for a single, noise limited, medium PRF waveform with PRF 14.9kHz

Conventionally, three PRFs are required to be clear in range and Doppler in order to resolve range and Doppler ambiguities and to declare a target detection.

The number of PRFs within a schedule must be selected carefully; too few and the ability to overcome range-velocity blind zones will be hindered. With too many PRFs, then, depending on the average PRF, there may be insufficient time to transmit the entire PRF schedule within the dwell time on target. Typically, eight PRFs are employed spanning about an octave.

If significant harmonic relationships exist between any of the PRFs chosen, then it may not be possible to resolve all of the ambiguities and the schedule is not *decodable*. In reality, targets have a physical size too and extend outside of individual range or velocity cells. It is desirable to make sure that a schedule is not only technically decodable, but also decodable in the presence of range or velocity extended targets. If the decodability criteria is broken by a large target, then false targets or *ghosts* will be observed. Unfortunately there is no means of distinguishing a ghost target from a real target and so all must be processed as true detections, leading to false alarms.

Because of the relatively wide size of the rejection notches, the possibility remains for a PRF schedule to be decodable and still have some rejection notch overlap; this is found to be a particular problem at the first repetitions of the ambiguous velocity intervals. The consequences of such occurrences are bands of velocities in which the radar is blind, or nearly blind (three PRFs clear only), at *all* ranges, thereby allowing a target to approach at a particular velocity with minimum risk of detection. Nothing can be done about the rejection notches, centred on zero Hz, which blind the radar to crossing targets.

After the pulses have been received by the radar, they are decoded using the coincidence algorithm [8]. The coincidence algorithm operates by taking the target returns and for each range bin, performing an FFT across all pulses in the PRF. Thus a map of range-velocity is produced. The regions of heavy clutter are then notched out and a detection algorithm is then used to identify potential targets within this ‘folded’ (ambiguous) range-velocity map. The process is repeated for each of the PRFs.

The next stage is to decode the targets and resolve the ambiguities. This is performed by taking each range-velocity map and repeating them until the maximum range-velocity extent of interest has been covered. For a single PRF, this will give many ranges and velocities at which a target may be present. The process is repeated for all the PRFs and the results overlaid. If a true target is present, it will appear in the same position in all PRIs (yet may not be detected, or may have been eclipsed or notched out). Any region of the range-velocity map that has 3 or more coincident detections is declared as being a true target. The process works well but issues can arise where very fast targets have moved between range cells between the first and last PRF being transmitted and do not necessarily align in the coincidence process. The problem is known as *range-walk* and is accounted for in the software.

In the radar problem encapsulated for this paper, the length of the transmitted pulse is directly proportional to the delay before the next pulse. This keeps the duty cycle of the transmitter constant. The radar is also frequency-hopping in that the transmission frequency changes for each PRF. The result is that the wavelength will also change and so the order of transmission for the PRFs is important.

The selection of PRFs in a medium PRF set is therefore based on the following constraints:

1. A spread of values which enable the resolution of range and velocity ambiguities to ensure basic decodability,
2. Removal of totally blind ranges and velocities,
3. The total time required for transmission of the waveform must be within the target dwell time.

The objectives then become:

1. Maximise the size of the target in range that can be decoded without ghosting,
2. Maximise the size of the target in velocity that can be decoded without ghosting,
3. Maximise the size of the clutter patch in range that can be tolerated before blind ranges occur,
4. Maximise the size of the clutter patch in velocity that can be tolerated before blind velocities occur,
5. Minimise the total time required for transmission of the waveform.

Objectives 1 to 4 may be calculated using the process outlined in [7] by calculating what is the target extent at any range/ velocity that can be tolerated before problems arise. If any of the objectives are zero or negative, i.e. a negative target size is the maximum that can be tolerated, then it implies that one of the constraints has been violated.

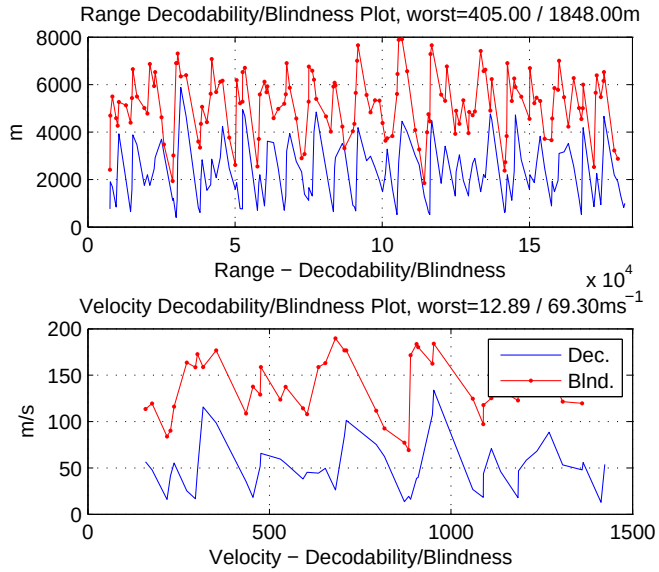

Fig. 4. Minimum target sizes for PRI set $[50.8\ 50.0\ 55.6\ 64.0\ 86.2\ 70.1\ 67.4\ 96.8]\mu s$

For the practical radar design in section 2.3, as each range and velocity has an associated minimum target size as demonstrated in figure 4, the worst case and median performance are actually of interest. The result is that for objectives 1 to 4, both the median overall performance and minimum overall performance need to be maximised, yielding 9 objectives in total. The constraints can be applied easily in objective space after the optimisation process. For the PRI set $[50.8 50.0 55.6 64.0 86.2 70.1 67.4 96.8]\mu s$ results shown figure 4, the corresponding objectives are:

1. Median range decodability: 2205.0 metres,
2. Median velocity decodability: 47.9 ms^{-1},
3. Median range blindness: 5310.0 metres,
4. Median velocity blindness: 135.5ms^{-1},
5. Minimum range decodability: 405.0 metres,
6. Minimum velocity decodability: 12.9ms^{-1},
7. Minimum range blindness: 1848.0 metres,
8. Minimum velocity blindness: 69.3ms^{-1},
9. Dwell time: 44.8ms.

As objectives 1 to 8 are positive, and objective 9 is less than 50ms, the PRI set shown will form a viable radar waveform.

2.3 The Radar Model

A radar model based on an airborne fire control type application was derived to trial the fitness of PRF sets. The model assumes approximately 10GHz operation, 64-point FFT processing, 10% fixed duty ratio of pulse length to Pulse-Repetition Interval (PRI), linear FM pulse compression achieving a variable compression ratio with the PRF and that platform motion compensation is applied (i.e. the location of the ground is shifted in velocity back to zero, rather than being left at the platforms forward speed). The maximum target velocity with respect to the ground was taken as 1500 m/s and the maximum range was taken to be 185 km (100 nmi). These and other operational characteristics are summarised in Table 1. It is intended that the model should be representative of the types currently in service or about to enter service.

3 Software Structure

The software for the radar design problem is available for download from [3]. As the exact design analysis algorithms are proprietary, the software source is not provided, rather a compiled but portable binary file. In the interest of maximising portability between platform types, and simultaneously protecting the proprietary algorithms, Matlab *P*-code format has been used. Matlab *P*-code is a platform independent pre-parsed binary format used by the Matlab engine, helping to reduce the options for de-compilation that 'C', Fortran or Java would present.

The function `testpris.p` takes a $1 \times N$ vector as an input, where $N \in [4, 12]$ is the number of decision variables, and outputs a 1×9 vector of metrics. Each of the

Table 1. Summary of the radar model's characteristics

Parameter	**Value**
Carrier frequency	9.97 GHz for PRF_1, each following PRF -30MHz
Minimum PRI	50 μs
Maximum PRI	150 μs
Compressed pulsewidth	0.5 μs
Receiver recovery time	1.0 μs
Range resolution	75m
FFT size	64 bins
Duty cycle	10% fixed
Maximum target dwell time	50ms
Maximum target Velocity	$\pm 1500 ms^{-1}$
Maximum detection range	185.2 km (100 nmi)
Number of PRFs/PRIs	4 to 12
Number of PRFs for coincidence	3

decision variables is an integer in the range [500,1500] inclusive and represents the set of Pulse Repetition Interval values between $[50.0\mu s, 150.0\mu s]$ in steps of $0.1\mu s$.

The 9 metrics that are output represent:

1. Median range extent of target before schedule is not decodable (in metres),
2. Median velocity extent of target before schedule is not decodable (in ms^{-1}),
3. Median range extent of target before schedule has blind regions (in metres),
4. Median velocity extent of target before schedule has blind regions (in ms^{-1}),
5. Minimum range extent of target before schedule is not decodable (in metres),
6. Minimum velocity extent of target before schedule is not decodable (in ms^{-1}),
7. Minimum range extent of target before schedule has blind regions (in metres),
8. Minimum velocity extent of target before schedule has blind regions (in ms^{-1}),
9. Time required to transmit total waveform (in milliseconds, to be **minimised**)

The metrics 1 to 8 are to be maximised, while metric 9 is to be minimised. There are 9 corresponding constraints: the first eight metrics must all be greater than zero, and the 9^{th} metric must be less than 50 ms. In order to simplify the conversion of the objectives all to minimisation, and to simplify the constraint process, a wrapper function has been provided `objpri.m` that will allow a $P \times N$ matrix to be provided, and a $P \times 9$ matrix is returned with all of the metrics arranged for minimisation, and aligned so that if any are greater than zero (the maximisation is converted to a minimisation by negating), then the solution can be considered not feasible as a practical waveform. The function also allows an entire population (size P in the example above) to be passed as one matrix.

The current version of the MSOPS [4] optimisation algorithm code used to generate the results found later in this paper is also provided as an example of how the objective function may be implemented.

The run-time of the objective function is quite short, considering it is an un-modified engineering application. Under Matlab and on a 1.8GHz Pentium 4 processor, Microsoft Windows XP, table 2 indicates the observed processing times for 10000 evaluations, and therefore times for single objective vector calculations.

Table 2. Example processing times for objective vector calculation

N	Time 10000 eval (sec)	Time 1 eval (ms)
4	21.6	2.16ms
8	33.6	3.36ms
12	46.4	4.64ms

4 Initial Objective Surface Analysis

The objective surface consists of 9 objectives and from 4 to 12 decision variables. Some relationships are known to exist between pairs of objectives, and also between the chromosomes and objectives.

The first main relationship is that if the number of decision variables is less than 9, then the objective region must be a projection of the lower-dimensional decision space manifold into the 9-objective space: thus not every possible objective vector is defined. With greater than 9 decision variables, the converse is true and there is likely to exist extensive many-to-one mappings between decision space and objective space. At present, it is not clear if one particular choice of decision space dimensionality provides the entire Pareto surface of the problem. It is hypothesised however that this is not the case and that the full Pareto surface will be comprised of sections where the decision space dimensionality changes. From a radar design perspective, the number of PRFs used and therefore the decision space dimensionality are not critical, as long as the schedule is valid and performs well. Short schedules are often attractive as they tend to require less processing time, although this processing aspect is a design preference rather than an objective, and is useful for refining the choice of PRF schedule from the full Pareto set.

The second relationship is that the first 4 objectives are median values, and objectives 5-to-8 are the minimum values, of the same 4 data sets. Thus objective 1 will always be better than (or equal to) objective 5 etc. Objectives 1 & 3, 5 & 7 are metrics associated with the performance in range and tend to have a degree of correlation, i.e. they may not conflict strongly. Similarly, Objectives 2 & 4, 6 & 7 are metrics associated with the performance in velocity and again do not tend to conflict strongly with each other either, however they do conflict with the objectives associated with range.

The third relationship is between the decision variables and the dwell time (objective 9). The objective is calculated from equation 3, where $F_9(\mathbf{x})$ is objective 9, $\mathbf{x}$ is the decision vector (which is integers in units of $0.1\mu s$) and $\lceil \cdot \rceil$ is a rounding-up operation. The first part of the sum accounts for the time to transmit 64 pulses for the FFT. The second part of the sum calculates how many pulses are required to space-charge to the maximum range of interest R_{max}. The objective has a nearly linear relationship, apart from the quantised space-charge offset. Interestingly, the order of the decision variables also has no influence on the objective, implying that it is therefore approximately unimodal (i.e. no local optima), but multi-global (i.e. more than one global optima exist). At first sight, the minimum value appears to occur when all of the decision variables are at their minimum, i.e. for an 8 PRI system, the minimum total dwell time would be 35.6ms, giving an objective value of -14.4. However, if unambiguous range of the minimum possible PRI is not an integer fraction of the maximum range, then the constant space-charge term may be rounded up. It may therefore be possible that under some

conditions, the global optima may not quite occur when all the decision variables are at their minimum values.

$$F_9(\mathbf{x}) = 1000 \sum_{i=1}^{N} \left(64 \frac{x_i}{10^7} + \left\lceil \frac{2 \times 10^7 R_{\max}}{x_i c} \right\rceil \frac{x_i}{10^7} \right) - 50 \tag{3}$$

Similarly a fourth observation can be made that the velocity-related objectives **are** dependent on the order of the decision variables (i.e. if the elements of a decision vector are re-ordered, the objective values may change). This coupling is due to the frequency-hopping radar design: the target velocity produces a *Doppler shift* of the carrier frequency; the shift amount is dependent on the carrier frequency itself. As the first PRI described by the first decision variable is transmitted at 9.97GHz, the second at 9.94Ghz etc., the order of the decision variables will change the effective transmission frequency, and therefore the velocity performance of the waveform. The objectives that are associated with range however are only very weakly correlated to the decision variable ordering. The modification of the objective value that occurs with a re-order is due to the effects of target range-walk. Given that a relatively large range resolution of 75 metres is used in this design, in a dwell of 50 milliseconds, a target must be travelling at $1500ms^{-1}$ or faster in order to move range cells during the dwell. The effect is thus only very small in this example as only targets at the limit of the velocity of interest will be affected. Objectives 1 & 3 are the most likely to undergo any change as these are calculated based on the medians. Objectives 5 & 7 are calculated using minimum and although possible, it is unlikely that any order-dependence will be observed.

The number of decision variables influences both performance, and also ultimately which constraints are most difficult to satisfy.

5 Algorithm Comparison

An initial examination of the ability of multi-objective optimisers to explore the objective surface was performed. Two primary optimisers: NSGA-II and MSOPS were used, along with a 3rd which is a steady-state derivative of the MSOPS algorithm and is currently under development (unpublished prototype which is run in a 'Pareto ranking mode' to aid confirmation of NSGA-II results). The experiments were to generate non-dominated surfaces for the application problem using 8 decision variables. Although NSGA-II is known to be less suitable for many-objective problems when compared to bi-objective problems, it has been included as a useful reference algorithm.

Each optimiser was run 30 times for 20,000 function evaluations. In each run of each optimiser, all 20,000 points that were generated were collected and the non-dominated surface of these points established, rather than relying on the contents of the final population alone e.g. as is common in NSGA-II.

The non-dominated surfaces of the 90 independent experiments were collated into a single group consisting of 775,140 points. Initially, an attempt was made to create the composite non-dominated surface to establish the contribution from each algorithm. However after over 24 hours of processing, the non-dominated set was still incomplete, but it was clear that all of the runs of all of the algorithms made a significant contribution

to the surface – in the high dimensionality of the objective space, the proportion of non-dominated solutions is very large. The group of 775,140 points was used to establish a lower-bound reference point, and a range for each of the objectives for scaling purposes.

The lower reference used was: [-7035.0 -81.3 -27150.0 -296.9 -2700.0 -22.3 -7660.5 -100.5 -14.4]. The range of each objective was calculated as: [7110.0 84.7 23130.0 249.5 2775.0 27.0 10511.0 100.3 78.0]. Figure 5 shows a plot of 5000 example non-dominated points, and figure 6 shows the same points, but normalised. It is clear that the relationship between the objectives is non-trivial

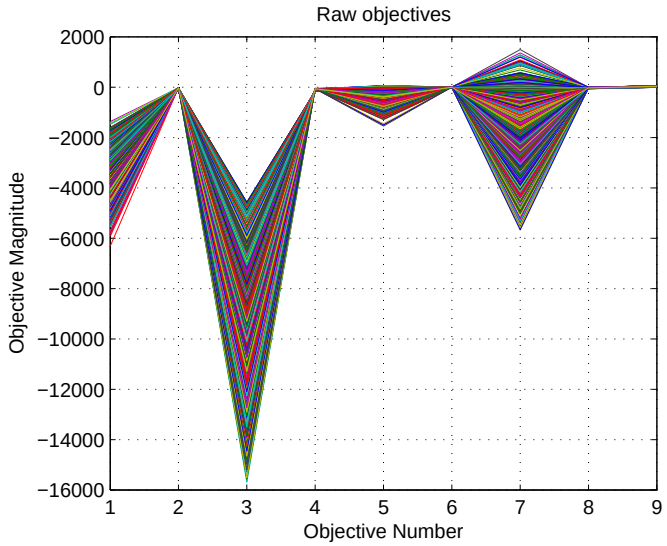

Fig. 5. Plot of 5000 non-dominated objective vectors. Objectives are un-normalised

A 10 million-point random search was performed of the objective function in order to help establish the relative performance of each of the algorithms. Unfortunately, every one of the 90 EA runs entirely outperformed the 10 million-point random search, preventing useful normalisation by exploiting the cumulative density function of the aggregated objectives [6].

Figure 7 shows an approximation of the distribution of the median attainment surface of the 3 algorithms over the 30 runs. Each of the 3 lines of the graph is calculated by:

1. Generate 200 approximately uniformly distributed unit length target vectors over the entire objective space.
2. For each of the 200 target vectors in turn, calculate the weighted min-max result for all the points in each of the 30 sets of repeated experiments. The minimum value in each of the 30 sets is taken, yielding a 30 by 200 matrix for each of the 3 algorithms.

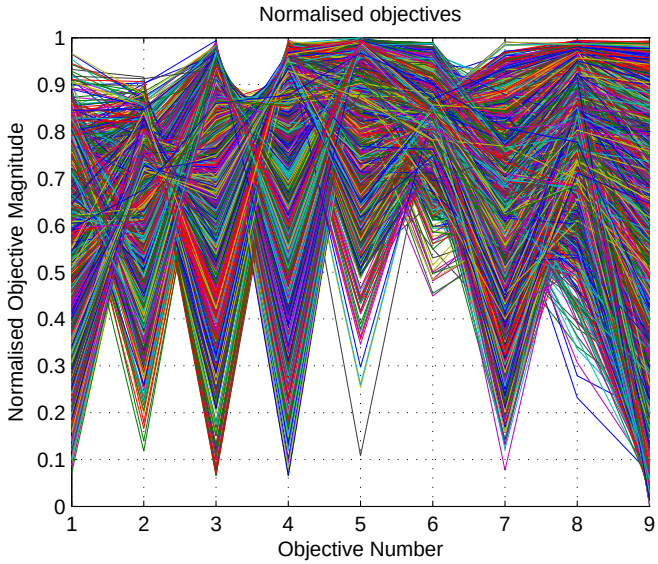

Fig. 6. Plot of 5000 non-dominated objective vectors. Objectives are normalised so that full 775,140 point vector set lies in range [0,1].

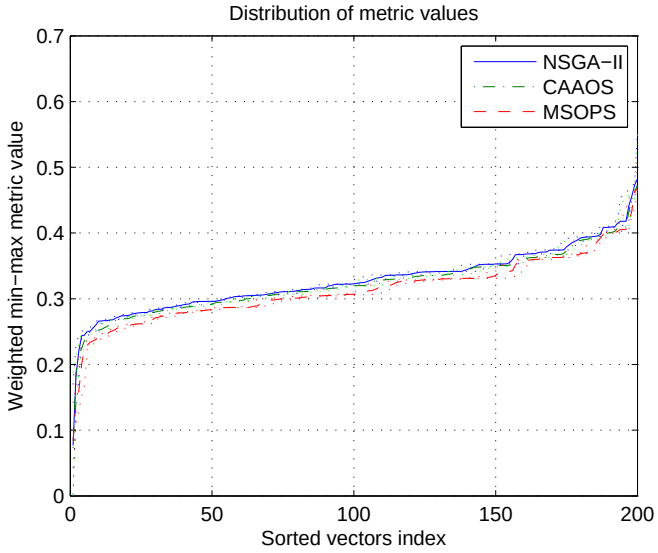

Fig. 7. Sampled median attainment surface distributions for NSGA-II (solid), MSOPS (dashed) and prototype algorithm (labelled CAAOS, dash-dot). Lower values are superior. Dotted lines show 95% confidence intervals obtained through 10 independent choices of the 200 target vectors.

3. The median of each column of this matrix is taken, creating a 200 element vector for each algorithm tested: the vector is a sampled approximation of the median attainment surface.
4. Each of the 200 element vectors is sorted to build a cumulative distribution, and then plotted.

As a 200 point sample of a 9-dimensional objective space is very sparse, the process was repeated 10 times, each with a different set of 200 points. The results were used to create 95% confidence interval bounds and plotted as dotted lines on figure 7, with the median of the 10 results as the thick-line. It is clear that despite 200 points being a sparse sample, the results are reassuringly accurate.

The performance of the three algorithms is very similar, with MSOPS leading slightly everywhere, as anticipated from previous studies [5]. As all 3 algorithms have produced similar results, despite entirely independent trials and algorithms, it is hypothesised that the obtained non-dominated sets are quite close to the Pareto optimal solution, but as the difference between MSOPS and NSGA-II shows, the set can only be considered non-dominated rather than Pareto. Additionally, as this is a 9-objective problem and NSGA-II has performed reasonably well with 20,000 evaluations, it is suspected that the overall density over the majority of the Pareto surface is high, allowing the problem to be approximated very well within 20,000 evaluations. However some of the results from the MSOPS trials suggest that there are regions of low density (the weighted min-max metric can converge well, however the Vector-Angle Distance Scaling metric is poor, suggesting a low density of points). Thus it is anticipated that the objective may be viable for study where only very few function evaluations are available (typical for on-line optimisation within a radar system).

It has been observed that there are concavities in certain dimensions (e.g. between objectives 1 & 2), and it may be the edges of the set that are sparse. Early indications from analysis using MSOPS also suggests that there may be regions of discontinuities and possibly sections of disconnected objective space.

6 Conclusions

This paper has presented and described a real engineering application of many-objective optimisation, and also provided access to software that allows the application to be studied by other researchers in the field.

The objectives can be calculated quickly enough to allow for practical optimisation algorithm development, and the surface appears complex enough to test the performance of visualisation and surface analysis tools.

At times, industrial acceptance of multi and many-objective optimisation algorithms has been slow. It is hoped that by providing an un-simplified real-world problem to use as an empirical benchmark, others can be encouraged to do the same for other problems, allowing more credibility to be attached to optimisation algorithm performance.

I also hope that as researchers in the field develop better algorithms for many-objective optimisation, the results can be collated and the true Pareto set for this optimisation problem approached. This collected set would be made available to extend the non-dominated data already provided from the production of this paper.

References

1. C. M. Alabaster, E. J. Hughes, and J. H. Matthew. Medium PRF radar PRF selection using evolutionary algorithms. In *IEEE Transactions on Aerospace and Electronic Systems*, volume 39, pages 990–1001, July 2003.
2. K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. In M. Schoenauer, K. Deb, G. Rudolph, X. Yao, E. Lutton, J. J. Merelo, and H.-P. Schwefel, editors, *Parallel Problem Solving from Nature – PPSN VI*, pages 849–858, Berlin, 2000. Springer.
3. E. J. Hughes. Many-objective radar design software. http://code.evanhughes.org.
4. E. J. Hughes. Multiple single objective pareto sampling. In *Congress on Evolutionary Computation 2003*, pages 2678–2684, Canberra, Australia, 8-12 December 2003. IEEE.
5. E. J. Hughes. Evolutionary many-objective optimisation: Many once or one many? In *IEEE Congress on Evolutionary Computation, 2005*, volume 1, pages 222–227, Sept. 2005.
6. E. J. Hughes. Multi-objective equivalent random search. In *Parallel Problem Solving From Nature, PPSN IX*, Reykjavik, Iceland, September 2006. Springer LNCS-4193.
7. A. Kinghorn and N. Williams. The decodability of multiple-prf radar waveforms. In *Radar 97 (Conf. Publ. No. 449)*, pages 544 – 547, Oct. 1997.
8. W. H. Long and K. A. Harringer. Medium PRF for the AN/APG-66 radar. *Proceedings of the IEEE*, 73(2):301–311, Feb. 1985.
9. R. C. Purshouse. Evolutionary many-objective optimisation: An exploratory analysis. In *The 2003 Congress on Evolutionary Computation (CEC 2003)*, volume 3, pages 2066–2073, Canberra, Australia, 8–12 December 2003. IEEE.
10. M. I. Skolnik, editor. *Radar Handbook*. McGraw-Hill, 2nd edition, 1990. ISBN 0-07-057913-X.