

Real-time space debris monitoring with EISCAT

J. Markkanen^a, M. Lehtinen^b, and M. Landgraf^c

^a*EISCAT Scientific Association, Tähteläntie 56, 99600 Sodankylä, Finland.
Jussi.Markkanen@sgo.fi.*

^b*Sodankylä Geophysical Observatory, Tähteläntie 62, 99600 Sodankylä, Finland.
Markku.Lehtinen@sgo.fi.*

^c*ESOC, Robert-Bosch-Straße 5, 64293 Darmstadt, Germany.
Markus.Landgraf@esa.int.*

Abstract

Following a feasibility study in 2000–2001 on using the EISCAT ionospheric research radars to detect centimetre-sized space debris in the frame of an ESA contract, we are now finishing a continuation study, aimed at achieving debris detection and parameter estimation in real-time. A requirement is to “piggy-back” space debris measurements on top of EISCAT’s normal ionospheric work, without interfering with that work, and to be able to handle about 500 hours of measurements per year. We use a special digital receiver back-end in parallel with EISCAT’s standard receiver. We sample fast enough to correctly band-pass sample the EISCAT analog frequency band. To increase detection sensitivity, we use coherent pulse-to-pulse integration. The coherent integration is built-in in our method of parameter estimation, which we call the match function (MF) method. The method is derived from Bayesian statistical inversion, but reduces, with standard assumptions about noise and prior, to minimizing the least squares norm $\|z(t) - b\chi(R, v, a; t)\|$, where z is the measured signal and $\{b\chi\}$ is a set of model signals. Because the model signals depend linearly on the amplitude b , it is sufficient to maximize the magnitude of the inner product (cross correlation) between z and χ , the amplitude estimate is then determined by direct computation. The magnitude of the inner product, when properly normalized, is the MF. To construct the set of model signals, we sample the EISCAT transmission, in the same way as we sample the received signal, and apply linearly changing Doppler-shifts to it. Our initial implementation of the MF-method in 2001 was about four orders of magnitude too slow for real-time applications, but we have now gained the required speed factors. A factor of ten comes from using faster computers, another factor of ten comes from coding our key algorithms in C instead of Matlab. The largest factor, typically 100–300, comes from using a special, approximative, but in practice quite sufficient, method of finding the MF maximum. Test measurements show that we get real-time speed already when using a single dual-processor 2 GHz G5 Macintosh to do the detection computations.

1 Introduction

It is estimated that there are approximately 200 000 objects larger than 1 cm currently orbiting the Earth, as an enduring heritage of four decades of space activity. This estimate includes the functioning satellites, but by far the most objects are what is called space debris (SD), man-made orbital objects which no longer serve any useful purpose. Many of the small-sized (less than 10 cm) particles are due to explosions of spacecraft and rocket upper stages, but there are also exhaust particles from solid rocket motors, leaked cooling agents, and particles put into space intentionally for research purposes. The large (> 10 cm) objects have known orbits and are routinely monitored by the U.S. Space Surveillance Network, but information about the smaller particles is fragmentary and mainly statistical. Especially, in Europe there is no radar that is routinely used for monitoring small-size SD.

From 2000, we have been involved in two studies, contracted by ESA (ESA, 1999, 2002), to start utilizing the EISCAT ionospheric research radars also for space debris measurements. Since the early 1980's, the EISCAT mainland radars have been performing ionospheric measurements at least 2000 hours per year; and since the late 1990's, after the EISCAT Svalbard radar became operational, EISCAT has been measuring more than 3000 hours per year. Our aim is to be able to use a substantial amount of these operating hours for simultaneous space debris measurements. In our initial study (Markkanen et al., 2002), we showed that it is feasible, and technically straightforward, to conduct SD measurements in parallel with normal EISCAT ionospheric measurements, without interfering with those measurements.

Our approach is to operate a separate digital receiver back-end, which we call the SD receiver, in parallel with the EISCAT standard digital receiver. This allows us to implement our own, amplitude domain data processing, the match function or matched-filtering (MF) method. The MF-method seeks to increase detection sensitivity by implementing pulse-to-pulse coherent integration. To make the hardware as simple and cheap as possible, the custom-made part of the SD receiver is basically just a fast sampler and digital demodulator; all special processing is done in fast but still cheap general purpose workstations. The SD receiver samples the EISCAT second intermediate frequency (IF2) band fast enough to capture the relevant frequency channels into a single digital stream, without doing the customary channel separation. We need to sample with the rate of about a million complex samples per second continuously. This produces a fairly large amount of data, more than 10 GBytes per

hour. Early on, ESA suggested that we should strive to do the data analysis in real-time, so that the raw data could be quickly disregarded and the data storage requirements kept modest. We describe the main features of the SD receiver in section 3.

A straightforward implementation of the MF-method implies long data vectors, with lengths of hundreds of thousands complex points, to be Fourier-transformed a few thousand times per every second of raw data. At the third European Space Debris Conference, in 2001, we had to concede that with the processing speed that we had achieved at the time, it would take several centuries of CPU time to analyze just one year's quota of EISCAT space debris measurements. However, soon afterwards, M. Lehtinen realized that by accepting some loss of detection sensitivity and a small bias in the velocity estimate, it would be possible to speed up MF computation drastically, typically by more than two orders of magnitude. We use the term fast match function algorithm (FMF) for the resulting computation scheme. We outline the algorithm in section 7 of this paper.

The results of the initial study were encouraging. The achieved detection sensitivity was equivalent to being able to observe a spherical target with diameter of about 2 cm at the range of 1000 km. With the advent of the FMF-algorithm, the processing speed, though still sluggish, was starting to become useful. In 2003, ESA commenced a new study with us to increase the processing power so that large amounts of EISCAT SD measurements could be conveniently handled in real-time. The study has achieved the required processing speed. In addition to the factor of 100 delivered by the FMF-algorithm, we now use computers that are about ten times faster than what we had available in 2001. A final required factor of ten to the speed was obtained by coding the MF- and FMF-algorithms in C, instead of using Matlab as was done in the initial study.

The EISCAT system (Baron, 1984, 1986; Wannberg et al., 1997) consists of three separate radars: monostatic VHF radar, located near Tromsø, Norway, operating at 224 MHz; monostatic but two-antenna EISCAT Svalbard Radar in Longyearbyen, Svalbard, operating at 500 MHz; and tristatic EISCAT UHF radar at 930 MHz, with transmitter in Tromsø and receivers in Tromsø and in Kiruna, Sweden, and Sodankylä, Finland. All the transmitters operate in the megawatt peak power range and routinely utilize high (10–20%) duty cycles.

Even though hard target echoes are routinely detected, standard EISCAT data processing is not optimized for hard targets. The characteristic feature expected from small hard targets is long signal coherence time, typically several hundred milliseconds. By the signal's (phase-) coherence we mean that the signal phase $\phi_0(t)$ obeys a deterministic functional form for some duration of time, called the coherence time.

EISCAT’s normal ionospheric signal has a coherence time less than a millisecond in most parts of the ionosphere. This time is much shorter than the interval between transmitted pulses, which in EISCAT is typically 3–10 ms. Therefore, echoes from individual pulses are uncorrelated, and can only be added usefully in the power domain. This is done by computing, for each pulse m and a set of ranges r , power spectra $G_m(r; f)$, and adding the spectra over the pulses, $G(r; f) = \sum_m G_m(r; f)$. This is called non-coherent pulse-to-pulse integration. It should be noted that, within a single transmission-reception (T/R) cycle, computing the range-gated power-spectrum $G_1(r; f)$ achieves coherent integration of the samples. In fact, for a single uncoded pulse, the MF-method, too, in effect just computes range-gated power spectra.

To achieve coherent integration from pulse to pulse, the MF-method adds the echoes from different T/R cycles in the amplitude domain, taking care that the pulses are added with equal phase. The method, in essence, removes all phase variation from the signal before adding the samples. This is achieved by guessing the phase factor $e^{i\phi_0(t)}$ of the signal, and canceling it by multiplying the signal by the complex conjugate of the guess, $e^{-i\phi_0(t)}$. The guesses in our implementation are generated by brute force. We search through a large set of parametrized model functions, and use the one which achieves best cancellation of the phase, that is, which results in largest integrated amplitude. After the phase variation has been successfully removed, the remaining part of the signal can safely be integrated, both within a single pulse, and from pulse to pulse. Actually, dividing the radar data into T/R cycles is artificial from the MF-method point of view. It is more natural to consider the totality of transmission during an integration period as just a waveform pattern, to be matched against the totality of reception, irrespective of how the patterns are divided into T/R cycles. In particular, there is no need for the T/R cycles to be identical, either in terms of length or transmission content.

As long as the signal stays coherent (i.e. obeys the assumed model), coherent integration suppresses the non-coherent background noise, so that the effective signal-power to noise-power ratio increases in direct proportion to the number of pulses integrated. This increases detection sensitivity. Non-coherent integration, in contrast, does not increase signal-to-noise ratio. The drawback in coherent integration, besides being computationally more demanding due to the long data vectors, is that if the signal model is not accurate, the ensuing phase error can actually start to reduce the integrated amplitude.¹ In our case, coherent integration beyond about 300 ms does not seem to improve detection sensitivity.

¹ We accept that there is a grain of truth in the statement claiming that “most radars utilize non-coherent integration”, because “maintaining coherency [...] is very costly and challenging to achieve.” (Mahafza, 2000)

Part of the reason for the unexpectedly short apparent coherence time is that, although we (in section 5) will derive a model signal that we believe should be quite accurate for small structureless targets, for performance reasons we cannot actually use the ideal model. The approximative model that we actually use, both in the MF- and FMF-algorithms, is suitable for single-frequency-channel transmissions. The different frequency channels in our multi-frequency signal will have slightly different Doppler-shifts because the Doppler-shift depends on transmission frequency. It is impossible to cancel all the Doppler phase factors simultaneously using only the single model phase factor, which is available in the FFT-compatible approximative model.

We derive, in section 4, the MF-method from Bayesian statistical inversion (Lehtinen, 1986). Within the probabilistic Bayesian approach, the estimates for the basic parameters: range, radial velocity, radial acceleration and signal amplitude or signal total energy, are the most probable parameter values, given the measured noisy signal. With our assumptions, this solution is also the one that minimizes the least squares norm between the measured signal and the set of model functions. The solution is also equal to the maximum-likelihood solution.

During several test campaigns over four years, we have collected and analysed about 50 hours of data at the EISCAT UHF radar in Tromsø. These data have been taken mostly for method development and verification purposes. We will not attempt to describe the data in this paper, but only mention that the mean event rate has been about 20 events per hour. The observed height distribution seems to reproduce the well-known main features of SD distribution in the low Earth orbit region. Our measurements cover altitude up to about 4000 km, with gaps that depend on the EISCAT experiment that we have been attached to. It is clear that the EISCAT UHF radar can observe targets down to effective diameter (diameter of a hypothetical spherical target on the antenna optical axis that would give the observed signal strength at the observed range) of about 2.0–2.2 cm at 1000 km range. However, we cannot say much about the actual target cross section, for the EISCAT antennas do not have monopulse feed. At the moment at least, there is no way that would allow pinpointing the actual target direction within the radar beam, and so the target’s radar cross section cannot be deduced from the observed signal strength.² We hope that in the future we can partly alleviate this problem by collecting fairly large amounts of data—perhaps about 500 hours annually—so that the antenna beam pattern can be taken into account statistically, and meaningful comparisons with space debris models made.

² Recently, there has been discussions within EISCAT about emulating the monopulse capability by making the antenna feed slightly asymmetric and rotating the feed rapidly, but the feasibility of this scheme has not yet been verified.

2 The EISCAT UHF radar

So far in our measurements, we have been using mainly the EISCAT UHF radar. The block diagram of the UHF radar's Tromsø site is shown in Fig. 1. The radar's radio-frequency (RF) band is centered at 929 MHz, and there are 14 transmission frequencies available, 300 kHz apart. In the most common, current, EISCAT experiment modes, two frequency channels are used. Recently those have been centered at 929.9 MHz (EISCAT frequency F13) and 930.2 MHz (F14). The RF signal is mixed in two stages to the second intermediate frequency (IF2) band, using local oscillators at 812.0 MHz and 128 MHz, so that F13 maps to 10.1 MHz and F14 to 9.8 MHz. (In the space debris receiver, F13 and F14 ultimately map to -0.1 MHz and +0.2 MHz in the baseband.) The IF2 band is formed by the radar's anti-aliasing filter, which is 7 MHz wide and centered at 11.2 MHz.

In the standard EISCAT data processing, the second IF is digitized by a 14-bit analog-to-digital converter (A/D), which produces a continuous sample stream at the rate of 15 Msamples s^{-1} . The stream of IF2 samples is distributed to the multi-channel, VME-based EISCAT digital receiver, each channel taking one slot in a VME crate. Custom hardware in each digital channel performs quadrature detection, followed by sampling rate reduction appropriate to the typical 10–50 kHz final channel bandwidth. The baseband sample stream is buffered, and further processing to average sample correlation products is done on UNIX-based computers.

The EISCAT UHF transmitter consists of a programmable radar controller that generates the pulse patterns, either uncoded on/off pulses or various classes of binary phase codes; an exciter system that converts the radar controller output to RF around 929 MHz; and a klystron power amplifier that consists of two klystron tubes, in principle able to deliver combined peak power of about 2.5 MW. The power during our space debris measurements has been 1.0–1.5 MW. The maximum transmitter duty cycle is 12.5%, and, in practice, duty cycles near this value are also used in most experiments.

The 32 m UHF antenna has a fully steerable parabolic dish, with Cassegrain optics, and a slewing rate of about $1.3^\circ s^{-1}$ both in azimuth and elevation. The antenna pointing direction is calibrated using celestial radio sources, and is believed to be accurate to better than 0.1° in most directions. The time and frequency references are derived from the GPS system.

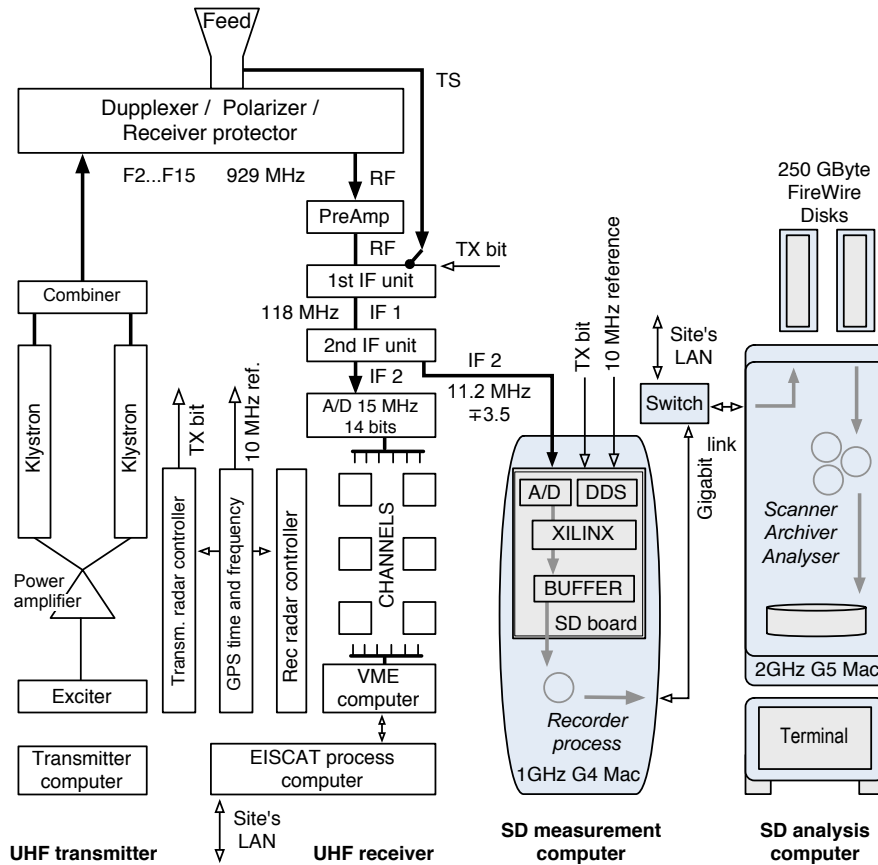


Fig. 1. The space debris (SD) receiver connected to the EISCAT UHF radar. The SD receiver consists of a measurement computer and an analysis computer. The measurement computer hosts a custom signal processing board (SD board). The primary analog input to the SD receiver is the EISCAT second intermediate frequency band (IF2). The input contains, time-multiplexed, both the standard received signal and the transmission sample signal (TS). On the processing board, there is an analog-to-digital converter (A/D) taking 40 megasamples per second; a direct-digital-synthesizer chip (DDS) which provides clock signals on the board, phase-locked to the host radar's 10 MHz frequency reference signal; two Xilinx signal processing chips to perform signal demodulation and sampling rate reduction; and a memory buffer for temporary storage of the complex samples. The recorder program running on the measurement computer moves the samples over the gigabit network link to an external hard disk, mounted on the analysis computer. Target detection is done by the scanner program running on the analysis computer, using the FMF-algorithm. After detection, two other software modules, the archiver and the analyser, perform target parameter estimation.

3 The space debris receiver

To be able to use our own data processing, optimized for hard targets, we use a special digital receiver back-end, the space debris receiver. Signal to the space debris receiver is derived from the EISCAT analog signal path at the

IF2 level. Figure 1 shows the main blocks of the SD receiver, connected to the EISCAT UHF system at the Tromsø site.

EISCAT standard measurement modes normally have more than one frequency channel. EISCAT standard data processing handles this situation in the traditional way, by distributing the IF2 data to multiple hardware channels, each tuned to a particular center frequency. The end result is several sample streams, one on each channel. Our approach in the SD receiver is different. We sample fast enough to capture the relevant part of the analog IF2 band into a single digital stream. We call this type of data multichannel complex data (Lehtinen, 2002). According to the bandpass sampling theorem, if the spread of frequencies is B MHz, one needs to take B million complex samples per second, in the minimum. For its most common measuring modes EISCAT uses the frequencies F13 and F14, which are 300 kHz apart, and then we use 500 kHz sampling rate in the SD receiver. But we have also verified that the SD receiver can handle sampling speeds up to 2.5 Msamples s^{-1} .

In addition to the standard reception, our data processing requires that the transmission waveform be measured. As indicated in Fig. 1, EISCAT provides a transmission sample signal (TS) time-multiplexed into the same data path as the reception. The multiplexer switch is controlled by the receiver protector bit (TX bit in the figure), generated by the EISCAT radar controller microprocessor. We routinely record the receiver protector bit into our data stream to mark out the transmission blocks. The bit is stored into the least significant bit of the imaginary part of the 16 + 16-bit complex integer data words. With this arrangement, the transmission sample signal is automatically sampled with the same sampling rate as the actual reception, though we would like to sample it with a higher rate.

The core of the SD data acquisition system is a custom data acquisition board (SD board), which performs signal sampling, quadrature detection and sampling rate reduction. The board was originally developed for ionospheric tomography by the now defunct Finnish company Invers Ltd.

The analog-to-digital converter on the SD board samples at 40 MHz. The resulting real-valued sample stream is processed by programmable logic chip, from the Xilinx SpartanXL FPGA family, to perform quadrature detection, essentially by doing Hilbert transform. The result of the transform is a complex-valued 10 MHz sample stream, which represent the negative frequency part of the spectral contents of the analog input. The chip then decimates the 10 MHz stream to the final sampling rate. Typical decimation factor M is 20, which yields 500 kHz final sampling rate. The decimation is done by adding samples in blocks of M ; this ensures that the filtering is properly matched to the final sampling rate (see Fig. 3 and the associated discussion in section 4).

It may be noted that there is no separate multiplication (mixing) to baseband in this scheme. Instead, the customary frequency component in baseband is created by the undersampling. With the 40 MHz raw sampling rate, the arrangement requires that the band-limited analog input is centered at 10 MHz. Although it is possible to operate the SD receiver's analog-to-digital converter at other sampling rates, the 40 MHz is a most convenient choice. That the two most often used EISCAT frequencies have been 10.1 MHz and 9.8 MHz, which map as near to the zero frequency as can be hoped for in this processing scheme, is a happy coincidence. The next version of the SD receiver should have a complex mixer built-in.

The SD receiver board is mounted in a Power Macintosh G4 workstation, running under the Mac OS X version of UNIX. We call the Mac G4 the measurement computer. In addition, there is a dual-CPU Macintosh G5 computer for data analysis. The Mac workstations are connected to each other via a gigabit Ethernet link, and are also connected to the host site's local network (LAN). The measurement computer runs software from Invers Ltd to read the samples from a buffer and to write them to hard disk mounted over the gigabit link from the analysis computer. The data accumulation rate to the disk is between 7 and 30 GBytes per hour, depending on the sampling rate. The LAN connection is used to access the EISCAT process computer, to update the time base on the G4 and G5 once every five minutes, using the ntp protocol. This ensures that the time base stays within 20 ms from the time kept in the EISCAT system. This is more than adequate for time stamping space debris events.

Data analysis is done using C and Matlab programs running on the analysis computer. First, a scanner reads a segment (typically 300 ms) of raw data from disk and searches through the segment for hard targets using threshold detection within the framework of the match function method. When the threshold is exceeded, we have a hit. The scanner saves the hit's description to a file and proceeds to next data segment. A second program, the event archiver, inspects the list of hits and combines the hits that correspond to a single target passing through the radar beam into an event. Having determined the time boundaries of the event, the archiver copies the raw data to a separate directory and goes on to look for more events. Finally, the analyser picks data from the event directories and deduces and saves the event parameters.

By far the most time-consuming step in the data analysis is the scanning. Scanning is done by a C program that makes use of the AltiVec vector processor onboard the G5, by calling routines in Apple's DSP library (vdsp), most especially the FFT routine.

The event archiver is also a C program, but it is not performance critical. Most of its time goes to data copying, so the speed is mainly limited by disk

speed. We have saved all raw data from all our test measurements so far—somewhat less than a terabyte—but in routine measurements, only the raw data of events will be saved.

The way to compute the final target parameters is still under development. What the analyser now does, is basically to call the scanner to re-scan the raw data, with maximal time and range resolution but over a narrow range interval, and then make linear or quadratic fits to the range and Doppler-velocity time series. The range and velocity parameters that we normally quote are taken from these fits, for the time instant of maximum signal strength. The analyser is a Matlab program.

The overall processing speed is such that for data taken with 2 MHz sampling rate, it takes 40–45 minutes to scan, archive and analyse one hour of raw data, while keeping the raw data access going on at the same time.

4 The match function method

We want to estimate the parameters of a hard target echo signal $s(t)$ in the presence of wide-band gaussian noise $\gamma(t)$, of variance σ^2 . We denote by $z(t)$ the received signal,

$$z(t) = s(t) + \gamma(t). \quad (1)$$

We denote by $x(t)$ the transmission sample signal (the signal TS in Fig. 1). We ignore here the frequency translations done in the actual receiver, treating z , s , x and γ as complex-valued (detected) signals. The frequency translations affect both the echo signal and the transmission signal by a factor of the form $\exp(i\omega_{LO}t)$, where ω_{LO} is some local oscillator frequency, and hence cancel each other out in correlation products like $s(t)\bar{x}(t)$.

To find an optimal signal estimate (or at least a well-defined estimate), we use the approach of Bayesian statistical inversion. The basic idea is to use a parametrized model for the signal $s(t)$ and find the most probable signal among the model signals, given the received signal $z(t)$.

We specify our set of model signals explicitly in section 5. Here we will make use only of the property that the model depends linearly on one parameter, the complex amplitude b , and in addition depends on a set of other parameters (range, radial velocity and radial acceleration), which we collectively denote by θ , so that

$$s(t) = b \cdot \chi(\theta; t). \quad (2)$$

We sample the signal $z(t)$ using sampling interval τ_s , and get N samples z_n during a time interval T_c , the coherent integration time.

It makes sense that after a specific measurement result z , some parameter values (b, θ) are to be considered more likely than others, in a way that depends on z . That is, the probability of the various imaginable values is describable by some conditional probability density, with z in the condition. In the Bayesian worldview, that density is called the posteriori density, and is denoted here by $D_p(b, \theta|z)$. The inversion problem is to utilize the measurement to find the posteriori density. The posteriori density is the most complete inference that can be made about the parameter values, based on the measurement. Normally, one wants to condense the inference to a few numbers, the parameter estimates. There is no unique way to select “best” estimates, but the standard Bayesian criterion is to use the most probable values:

$$(\hat{b}, \hat{\theta}) = \arg \max_{b, \theta} D_p(b, \theta|z). \quad (3)$$

We now derive the posteriori density. We denote by $D_1(z_n|s_n)$ the conditional probability density function of z_n , given s_n . This is just the probability distribution of the noise $\gamma_n = z_n - s_n$,

$$D_1(z_n|s_n) = \frac{1}{\pi\sigma^2} e^{-\frac{1}{\sigma^2}|z_n - s_n|^2}. \quad (4)$$

For white noise, the conditional joint probability density to produce a particular measured vector z , if the actual signal vector is s , is then

$$D(z|s) = \prod_{n=0}^{N-1} D_1(z_n|s_n) = \frac{1}{(\pi\sigma^2)^N} \cdot e^{-\frac{1}{\sigma^2}\|z-s\|^2}. \quad (5)$$

The density $D(z|s)$ is called the direct theory. Given the direct theory, the Bayesian solution to the inversion problem is

$$D_p(b, \theta|z) = C'(z) \cdot D_{\text{pr}}(b, \theta) \cdot D(z|s). \quad (6)$$

Here $C'(z)$ is normalization factor. The new factor, $D_{\text{pr}}(b, \theta)$, is called the prior density. The prior density is a weight that can be used if it is known a priori—before making the measurement—that some particular signals $s(b, \theta)$ tend to occur more frequently than some others. Using non-trivial D_{pr} might actually make sense when measuring space debris, to throw out detections with highly unlikely parameters. But so far we have used constant prior. For constant prior, it follows from Eq. (6) and Eq. (5) that the posteriori density

is of the form

$$D_p(b, \theta|z) = C(z) \cdot e^{-\frac{1}{\sigma^2} \|z - b \cdot \chi(\theta)\|^2}, \quad (7)$$

where the constant $C(z)$ is determined by the normalization requirement $\int D_p db d\theta = 1$. It follows that finding the most probable parameters amounts to minimizing the least-squares norm,

$$(\hat{b}, \hat{\theta}) = \arg \min_{b, \theta} \|z - b \cdot \chi(\theta)\|. \quad (8)$$

That we should arrive at this most basic technique of parameter estimation, least-squares fitting, is perhaps not surprising. But what we have gained by walking through the Bayesian route is that not only do we have a method for acquiring the parameter estimates, but we have the explicit expression Eq. (7) for the posteriori density. In the future, we intend to make use of the posteriori density in error analysis. Due to the highly non-linear dependence of the model functions χ on the parameters-to-be-fitted, error estimation is not trivial.

A straightforward approach to the minimization problem expressed in Eq. (8) is to discretize the parameter space and perform an exhaustive search. We now show that the search space dimension can be reduced by one by making use of the property that the amplitude b enters the problem linearly. Our result can be confirmed analytically, but will be reasoned here from basic vector geometry. Referring to Fig. 2, the set \mathcal{M} of model vectors $b\chi(\theta)$ consists of 1-dimensional rays \mathbb{C}_χ through the origin of N -dimensional complex vector space \mathbb{C}^N . The rays are generated by the set of vectors $\chi(\theta)$. According to Eq. (8), we need to find the shortest distance between the measured point z and \mathcal{M} . The figure suggests the following strategy. First find the ray $\mathbb{C}_{\hat{\chi}}$ that is as parallel as possible with the vector z . Then the point in \mathcal{M} that is nearest to z is the orthogonal projection \hat{s} of z onto $\mathbb{C}_{\hat{\chi}}$, and is computed in the standard way as

$$\hat{s} = \frac{\langle z, \hat{\chi} \rangle}{\|\hat{\chi}\|^2} \hat{\chi}. \quad (9)$$

So the real problem is to find the maximally parallel ray $\hat{\chi}$. With z fixed, a sufficient measure of parallelism of a ray \mathbb{C}_χ and the vector z is the length of the orthogonal projection of z onto χ ; the longer the projection is, the more parallel (or antiparallel) z and \mathbb{C}_χ are. This measure of parallelism is the match




Fig. 2. Geometric interpretation of the MF-method. The required best estimate of the signal is the point \hat{s} in the set \mathcal{M} of model functions that is nearest to the measured signal z . The “conically-shaped” set \mathcal{M} consists of rays, $\mathbb{C}_\chi = \{a\chi : a \in \mathbb{C}\}$, generated by a set of basic model signals $\chi(\theta)$. The match function value $\text{MF}(\theta)$ is defined as the length of the orthogonal projection of z onto the ray $\mathbb{C}_{\chi(\theta)}$. Maximizing $\text{MF}(\theta)$ gives the ray $\mathbb{C}_{\hat{\chi}}$ that lies as near to the point z as is possible in \mathcal{M} . The estimate \hat{s} is the orthogonal projection of z onto $\mathbb{C}_{\hat{\chi}}$.

function³ MF,

$$\text{MF}(\theta) = \frac{|\langle z, \chi(\theta) \rangle|}{\|\chi(\theta)\|}. \quad (10)$$

The MF is a function of the model parameter θ , but does not depend on the scale of χ : χ and $a\chi$ give the same value of the MF. We maximize the MF to get the maximally parallel model vector $\hat{\chi} = \chi(\hat{\theta})$,

$$\hat{\theta} = \arg \max_{\theta} \text{MF}(\theta). \quad (11)$$

How the maximum, Eq. (11), is computed in practice will be discussed in section 6.

³ Intuitively, the more parallel two signal vectors (functions) are, the more they look alike, which is one reason for our nomenclature. A more serious reason is that MF also stands for matched filter. With velocity and acceleration fixed to zero, $\langle z, \chi(R) \rangle$ can equally well be computed by filtering z with the filter $h(t) = x(t)$ which is matched to the transmitted waveform $x(t)$. The MF is a generalization of this concept to more general pattern matching.

The energy⁴ of a correctly sampled complex-valued band-limited signal $y(t)$ is

$$W_y = \int |y(t)|^2 dt = \tau_s \sum |y_n|^2 = \tau_s \|y\|^2. \quad (12)$$

From Eq. (9)–(12), the energy $W_{\hat{s}}$ of the signal estimate \hat{s} is

$$\frac{W_{\hat{s}}}{\tau_s} = \|\hat{s}\|^2 = \frac{|\langle z, \hat{\chi} \rangle|^2}{\|\hat{\chi}\|^2} = [\text{MF}(\hat{\theta})]^2 = \max \text{MF}^2. \quad (13)$$

We use $W_{\hat{s}}$ as the estimator of W_s , the energy of the signal s . We make no “background subtraction”, even though this means that, as will be discussed at the end of this section, the estimator will be biased.

In summary, the MF-method of parameter estimation has two steps:

- Get the parameters θ by locating the position of MF maximum, Eq. (11).
- Get the signal energy as the square of the value of the MF maximum, Eq. (13).

A noise-free MF is useful for theoretical considerations. Without noise, both factors in the inner product in Eq. (10) are model functions. We will reserve a separate notation, AF, and use the standard name, ambiguity function (Skolnik, 1981), for the noise-free match function,

$$\text{AF}(\theta_0; \theta) = \frac{|\langle \chi(\theta_0), \chi(\theta) \rangle|}{\|\chi(\theta)\|}. \quad (14)$$

In the MF-method, target detection is based on the estimated signal energy exceeding a pre-defined threshold. We set the threshold by visual inspection of the data high enough that there are very few false alarms. We actually need to use a range-dependent threshold, because the lower altitudes, typically up to about 500 km, are often affected by strong clutter from the ionosphere, and require a higher threshold.

We set the detection threshold in terms of the ratio of signal energy to the noise power spectral density (PSD) G_γ . Note that the PSD has dimension of

⁴ We consider the signals $y(t)$ to have the dimension of voltage, and assume unit impedance, so that $|y|^2$ is signal power. The impedance does not matter, since only power ratios, like $|s|^2/|\gamma|^2$, are used when comparing to the physical world. We could also do without explicitly tracking the sampling interval, except that we do not want to change the dimension of energy in the middle of a chain of equations.

energy: $W/\text{Hz} = \text{J}$. We call the dimensionless ratio the energy-to-noise ratio, and denote it by SNR_N ,

$$\text{SNR}_N = \frac{W_s}{G_\gamma}. \quad (15)$$

A special feature of the SD receiver is that the final sampling interval τ_s is precisely matched to the noise-equivalent bandwidth of the receiver. Referring to Fig. 3, the SD receiver bandwidth is determined by the boxcar-in-time finite impulse response filter h , with the response duration equal to the sampling interval. The filter h is implemented digitally, in connection with the sampling rate reduction done in the SD receiver, by adding primary samples in blocks of M , where M is the required sampling rate reduction factor. Figure 3 shows an approximately equivalent analog system corresponding to this behaviour. The input noise to the equivalent filter is $\gamma(t)$. We assume that the noise temperature T_{sys} is defined in such a way that the PSD of the complex-valued wide-band noise can be written as

$$G_\gamma = kT_{\text{sys}}, \quad (16)$$

where k is the Boltzmann constant. During filtering, the PSD is transformed by the squared magnitude of the filter transfer function, so that the filtered noise γ' has PSD

$$G_{\gamma'}(f) = |H(f)|^2 kT_{\text{sys}}. \quad (17)$$

It then follows from the Parseval's theorem, and the special choice of $h(t)$, that noise power (variance) after the filter is

$$P_{\gamma'} = \int G_{\gamma'}(f) df = \left[\int h(t)^2 dt \right] kT_{\text{sys}} = \frac{1}{\tau_s} kT_{\text{sys}}. \quad (18)$$

(This implies that the noise-equivalent bandwidth of the filter is equal to the final sampling frequency $1/\tau_s$.)

The MF-method provides an estimate of the signal energy after the filter $W_{s'} = \tau_s \max \text{MF}^2$. This, and the separately estimated noise power $P_{\gamma'}$, can now be used to estimate the incoming energy-to-noise ratio in front of the filter. In our SD measurements, the EISCAT transmission typically consists of two frequency channels, which appear time multiplexed (also) in the received signal. The frequency channels have narrow bandwidth compared to the width of the filter H . We can take the filter's power gain $|H(f)|^2$ to be constant




Fig. 3. Principal filtering operation in the SD receiver. Shown here is an equivalent analog representation, the actual finite impulse response boxcar-in-time filter with impulse response $h(t)$ and transfer function $H(f)$ is implemented digitally. The duration of the impulse response is designed to be equal to the final sampling interval τ_s . The filter transforms the incoming signal s and noise γ to s' and γ' . The incoming white noise has power spectral density kT_{sys} , which is transformed to $G_{\gamma'}$. The noise samples γ'_n are uncorrelated and have variance σ'^2 .

across a frequency channel. In the filtering, the signal energy W_{ch} in a channel transforms approximately as

$$W'_{ch} = |H(f_{ch})|^2 W_{ch} = \kappa_{ch} |H(f_{ch})|^2 W_s, \quad (19)$$

where f_{ch} is the channel's center frequency. We also introduced in Eq. (19) the factor $\kappa_{ch} = W_{ch}/W_s$ to specify how large a part of the total energy (reception energy, and presumably also transmission energy) goes into the given channel. The filter's effect to the total signal energy is therefore an attenuation by a factor K ,

$$W_{s'} = \left[\sum_{ch} \kappa_{ch} |H(f_{ch})|^2 \right] W_s = K \cdot W_s. \quad (20)$$

Both Eq. (20) and Eq. (18) implicitly contain an unknown, but common, gain and unit-conversion factor. That factor cancels out from the ratio SNR_N , so that we can write

$$\frac{W_s}{kT_{\text{sys}}} = \frac{W_{s'}/K}{\tau_s P_{\gamma'}} = \frac{1}{K} \frac{\max \text{MF}^2}{P_{\gamma'}}. \quad (21)$$

We treat the system temperature as a known radar parameter (110 K for the Tromsø UHF radar), and use Eq. (21) to find the incoming signal energy W_s in physical units. We use that estimate to find a lower limit, RCS_{\min} , for the target's radar cross section (RCS). From the standard radar equation it follows

$$\text{RCS} = \frac{(4\pi)^3 kT_{\text{sys}} \cdot R^4 \cdot W_s}{G(\phi)^2 \cdot \lambda^2 \cdot P_x \cdot \mathcal{D}T_c}, \quad (22)$$

where R is the target range, λ is the radar wavelength, P_x is the transmission power, and \mathcal{D} is the transmission duty cycle such that $\mathcal{D}T_c$ is the actual length of transmission during the integration T_c . The factor $G(\phi)$ is the antenna power gain in the direction of the target within the radar beam, ie: at an angle ϕ offset from the known direction of the antenna optical axes. In the EISCAT systems, it is normally not possible to find the offset angle, and we give RCS_{\min} , which we get from Eq. (22) by setting $\phi = 0$, as a way of cataloguing the observed signal strength.

For the boxcar impulse response h shown in Fig. 3, the modulus of the filter transfer function is $|H(f)| = \left| \frac{\sin \pi f \tau_s}{\pi f \tau_s} \right|$. For a standard dual-frequency SD measurement, with $f_1 = -100$ kHz, $f_2 = +200$ kHz, $\tau_s = 2000$ ns, and $\kappa_1 = \kappa_2 = 0.5$, the overall energy attenuation factor K in Eq. (20) becomes 0.72. For small (spherical) targets the radar cross section, and therefore the received power, varies proportionally to the sixth power of the target diameter, so a 30% underestimate of received power results in 5% underestimate of the size. Considering other problems that we have in determining the target's radar cross section, we have so far ignored the filter effect and used $K=1$. The problems include: not knowing the efficiency of the coherent integration, not knowing the position of the target within the radar beam, and neither knowing accurately the radar's total transmitted power, nor how the power is actually divided into the frequency channels⁵.

For the rest of this paper, whenever we refer to “signal s ” and “noise γ ”, we mean the filtered signal s' and filtered noise γ' of Fig. 3, and we will drop the primes from the notation from now on.

The energy estimate $W_{\hat{s}}$ defined in Eq. (13) is a biased estimate. That is, the expectation value $E W_{\hat{s}}$ over repeated measurements of a given signal s , with different noise, is not equal to W_s , the actual signal energy. Instead, $E W_{\hat{s}}$ is considerably larger than W_s (about $10P_\gamma$ typically). This is seen qualitatively by considering the case when there is no signal present, and thus an unbiased estimate would have zero energy. Because MF^2 is non-negative, $\max \text{MF}^2$ will in any case have expectation value that is larger than zero. But the expectation value will in fact also be larger than the noise power. Taking, for example, the model signal χ having the form of a Doppler-shifted transmission (see Eq. (38))

$$[\chi(\omega)]_n = x_n e^{i\omega\tau_s n}, \quad (23)$$

where the x_n s are the transmission samples and ω is the Doppler-shift. The square of the match function $\text{MF}(\omega) = \frac{|\langle \gamma + 0, \chi(\omega) \rangle|}{\|\chi\|}$ becomes:

⁵ In principle, one should be able to use the measured transmission samples and the known filter gain $|H(f)|^2$ to determine the factors κ_{ch} .

$$[\text{MF}(\omega)]^2 = \frac{|\sum \gamma_n x_n \exp i\omega\tau_s n|^2}{\sum |x_n|^2} \quad (24)$$

$$= \frac{\sum_{nn'} \gamma_n \overline{\gamma_{n'}} x_n \overline{x_{n'}} \exp[i\omega\tau_s(n - n')]}{\sum |x_n|^2}. \quad (25)$$

With the SD receiver's filtering arrangement, the final noise samples are uncorrelated,

$$E(\gamma_n \overline{\gamma_{n'}}) = \delta_{n,n'} \cdot \sigma^2, \quad (26)$$

so we get from Eq. (25)

$$E[\text{MF}(\omega)^2] = \sigma^2, \text{ for all } \omega. \quad (27)$$

This implies that in a typical trial, as a function of ω , $\text{MF}(\omega)^2$ varies around σ^2 , often taking values that are larger than σ^2 . The expectation value of $\max_{\omega} \text{MF}^2$ over multiple trials will thus be larger than σ^2 .

A positive bias in the signal energy estimate in Eq. (22) increases the value of RCS_{\min} , moving it nearer to the actual RCS value. For the weakest signals, RCS_{\min} can conceivably even become larger than the actual RCS.

5 Signal model

We model the radar echo $s(t)$ by assuming that it behaves as if an optical wave were reflected from a mirror moving with constant radial acceleration a_0 . Denoting the delayed time (the time instant of transmission of a short "marker" that is received at time t) by t' with reference to Fig. 4, we express the echo in terms of the transmission sample signal $x(t)$ as

$$s(t) = bx(t'). \quad (28)$$

For the short integration time T_c , a few hundred milliseconds, we neglect all effects due to target structure and internal motions, as well as signal distortions due to the ionosphere, and take the amplitude b to be a complex-valued constant. It should be kept in mind that there are undoubtedly situations where this assumption fails.

For any given target radial motion $r(t)$, the delayed time is determined by the relation

$$t - t' = \frac{2r(\frac{t'+t}{2})}{c}, \quad (29)$$




Fig. 4. Transmitted wave reflected from a point-like target which is moving with constant radial acceleration a_0 . The parabola shows the radial component of the target's position vector in the coordinate frame of a stationary radar antenna, during the few hundred milliseconds of a coherent integration. The full three-dimensional velocity vector is typically very nearly constant in that frame during the integration time. The integration starts at time 0 with the transmission of the first pulse belonging to the integration. At the start of the integration, the target range is R_0 and the radial velocity is v_0 . Note that the diagram is not drawn to scale.

which expresses the time-of-flight in terms of travel distance and propagation speed. With constant radial acceleration, the radial motion is

$$r = r(R_0, v_0, a_0; t) = R_0 + v_0 t + \frac{1}{2} a_0 t^2. \quad (30)$$

For such motion, Eq. (29) is quadratic in t' . The solution of the equation for the propagation time, with appropriate choice of the sign of the square root, is

$$t - t' = \frac{2c}{a_0} \left\{ 1 + \frac{v_0}{c} + \frac{a_0}{c} t - \left[1 + \frac{2v_0}{c} + \left(\frac{v_0}{c} \right)^2 + \frac{2a_0}{c} \left(t - \frac{R_0}{c} \right) \right]^{\frac{1}{2}} \right\}. \quad (31)$$

Equation (31) can be simplified by expanding the square root into a power series. Care must be exercised regarding which terms can be dropped from the expansion. With parameter values that are typical at EISCAT UHF when the antenna is pointed almost vertically,

$$\begin{aligned} R_0 &\approx 10^6 \text{ m}, \\ v_0 &\approx 10^3 \text{ ms}^{-1}, \\ a_0 &\approx 10^2 \text{ ms}^{-2}, \\ \omega_1 &\approx 6 \cdot 10^9 \text{ Hz}, \end{aligned}$$

all terms following the “1” inside the square brackets in Eq. (31) are small compared to unity. Terms can be neglected when the requirement that the corresponding phase angle $\phi_X = \omega_1 \frac{2c}{a_0} X$, where ω_1 is the radar transmission

frequency, stays small during the whole integration time, is met. Taking the first three lowest-order terms of the power-series expansion of $[1 + (\dots)]^{\frac{1}{2}}$, and then disregarding the individual terms which are essentially zero—say all terms for which ϕ_X stays smaller than 0.1 rad when the integration time is less than a second—we are left with

$$t - t' \approx \frac{2}{c} \left[R_0 + v_0 t + \frac{1}{2} a_0 t^2 - (v_0 + a_0 t) \frac{R_0}{c} \right] \quad (32)$$

$$= \frac{2}{c} \left[R_0 + v_0 \left(t - \frac{R_0}{c} \right) + \frac{1}{2} a_0 \left(t - \frac{R_0}{c} \right)^2 \right] \quad (33)$$

$$= \frac{2}{c} r(R_0, v_0, a_0; t - \frac{R_0}{c}). \quad (34)$$

The term $\frac{-R_0}{c}$ is an obvious first order correction to the time instant of pulse reflection; the non-trivial aspect is that this correction is already sufficient. To a good approximation, we can express the model functions $\chi(R, v, a; t)$, used in the MF computation in Eq. (10), in terms of the transmission sample signal $x(t)$ as

$$\chi(R, v, a; t) = x\left(t - \frac{2}{c} r(R, v, a; t - \frac{R}{c})\right). \quad (35)$$

Note that nothing has been assumed about the transmission in this derivation so far. In principle, as long as the transmission can be accurately measured via the transmission sample signal, we do not even need to know what transmission has been used; the MF machinery incorporates the transmission transparently. In principle, this is fine for automated piggy-back measurements, where one does not have any control of the transmission EISCAT might be using at any given time.

The reality, of course, is rather different. A basic problem is that the radar's noise environment is often poorly approximated by our assumption that it consists only of stationary gaussian noise. Distortions occur in practice; for example the ionosphere becomes visible as time- and range-dependent clutter in the data. More or less ad hoc, manual, experiment-specific solutions are used to counter these problems. Also, we do not presently handle the case that the antenna pointing may change during a measurement, even though many EISCAT measurements use cyclical antenna pointing schemes. In practice we need to know beforehand, and even select, the EISCAT measurements we use in the SD work.

6 Computational aspects

In this section we describe how we evaluate the MF in practice. We first derive an approximation for the signal model of Eq. (35). Assume that the transmission can be described by

$$x(t) = \epsilon(t)e^{i\omega_1 t}, \quad (36)$$

where ω_1 is the carrier frequency, and the transmission envelope $\epsilon(t)$ is a slowly varying function, describing, say, a binary phase modulation, as is often the case in EISCAT. This description is normally adequate for a single-frequency transmission. We ignore the correction, R/c , to the pulse reflection time in Eq. (35), and use the special form Eq. (36) to write the model function as

$$\chi(t) = \epsilon\left(t - \frac{2}{c}r(t)\right) e^{i\omega_1\left[t - \frac{2}{c}r(t)\right]}. \quad (37)$$

Inside the slowly varying transmission envelope, we can assume that $r(t)$ stays constant, $r(t) = R$, during the integration time. Then, from Eq. (37) and Eq. (36),

$$\begin{aligned} \chi(t) &= \epsilon\left(t - \frac{2R}{c}\right) e^{i\omega_1\left(t - \frac{2R}{c}\right)} e^{i\left[-\omega_1\frac{2v}{c}t + \left(-\omega_1\frac{a}{c}\right)t^2\right]} \\ &= x\left(t - \frac{2R}{c}\right) e^{i(\omega_D t + \alpha_D t^2)}, \end{aligned} \quad (38)$$

where $\omega_D = -\omega_1\frac{2v}{c}$ and $\alpha_D = -\omega_1\frac{a}{c}$ are the Doppler-frequency and the rate of change of the Doppler-frequency, the ‘‘Doppler-drift’’, respectively. The approximation, Eq. (38), is often used in the literature, usually without the drift term, and is described by stating that the received signal is a delayed-in-time, Doppler-shifted replica of the transmission. With this model, the match function definition in Eq. (10) can be expanded, for continuous-time signals, as

$$\text{MF}(R, v, a) = \frac{\left| \int_0^{T_c} z(t)\bar{x}\left(t - \frac{2R}{c}\right)e^{-i(\omega_D t + \alpha_D t^2)} dt \right|}{\sqrt{W_x}}, \quad (39)$$

where $W_x = \int |x(t)|^2 dt$ is the energy of the transmission sample signal.

For signal vectors, we need to take into account that the transmission samples are only available at times $n\tau_s$. This already forces us to discretize the range variable. With

$$R_j = j\frac{c\tau_s}{2}, \quad (40)$$

the match function becomes

$$\text{MF}(R_j, v, a) = \frac{|\sum_{n=0}^{N-1} z_n \bar{x}_{n-j} e^{-i(\omega_d n + \alpha_d n^2)}|}{\|x\|}, \quad (41)$$

where the normalized Doppler-shift and Doppler-drift are

$$\omega_d = -\omega_1 \tau_s \frac{2v}{c}, \quad (42)$$

$$\alpha_d = -\omega_1 \tau_s \frac{a \tau_s}{c}. \quad (43)$$

At the points

$$v_k = k \frac{2\pi c}{\omega_1 T_c} \quad (44)$$

Eq. (41) can be written as

$$\text{MF}(R_j, v_k, a) = \frac{|\sum_{n=0}^{N-1} (z_n \bar{x}_{n-j} e^{-i\alpha_d n^2}) e^{-i\frac{2\pi k n}{N}}|}{\|x\|}, \quad (45)$$

which shows that at these points the MF can be evaluated using an FFT. The denominator $\|x\|$ is the square root of transmission sample energy, and is (of course) independent of R , v and a .

In most of our data analysis, we have taken the radial acceleration to be a deterministic function of range, $a = a(R)$. We have used the acceleration that corresponds to the target being in circular orbit and the antenna being pointed vertically,

$$a(R) = g_0 \cdot \frac{R_E}{R} \cdot \left(\frac{R_E}{R_E + R}\right)^2, \quad (46)$$

where R_E is the Earth radius 6360 km and g_0 is the acceleration of gravity at zero altitude, 9.8 m s^{-2} . Experimentation with real data has shown that not much sensitivity is lost in practice even if the acceleration is fixed in this way. This is perhaps a bad sign, for we would expect the MF to be a rather sensitive function of acceleration. For instance, inspection of the ambiguity function $a \mapsto \text{AF}(R_0, v_0, a_0; R_0, v_0, a)$ for 0.3 s integration shows that a 5 m s^{-2} error in a causes the integrated signal amplitude to decrease by about 60% (Markkanen and Postila, 2005). But no comparable decrease of the MF maximum seems to occur when we change $a(R)$ by 5 m s^{-2} in the data analysis. On the other hand,

ignoring the acceleration altogether, taking $a = 0$, does reduce the detection sensitivity significantly.

In the routine analysis, we search for the MF maximum only over the two-parameter $[R_j, v_k, a(R_j)]$ -grid. Even then, the detection computations, using full resolution and without any further approximations, become overwhelming. Assuming that we want to cover 1000 km in range and use 0.3 s coherent integration, and that the sampling interval is $0.5 \mu\text{s}$. Then the input data vector is 600 000 points long, and the FFT requires about 60×10^6 floating point operations per range gate. The 13 000 ($1000/0.075$) range gates require about 800×10^9 floating point operations. On a dual-processor 2 GHz G5 workstation, we get about 1 GFlops combined performance for FFTs of this length, so we need about 800 s to handle the 0.3 s of data. A typical EISCAT UHF phase-coded transmission uses baud length of about $20 \mu\text{s}$ or longer. For these, we can relax the range gate separation in the detection phase somewhat, say by a factor of 10. But this still leaves us more than two orders of magnitude short of the required real-time speed.

7 The fast match function algorithm

In this section we describe a fast, although approximative way to evaluate the match function, the FMF-algorithm. The algorithm was developed to compute efficiently a reasonably good approximation of the MF maximum; the FMF is not, in general, a good approximation of the MF over the whole parameter space. We will make use of two special properties of our SD measuring situation in EISCAT.

First, we note that the Doppler-velocity interval that we need to monitor is much narrower than the interval that is actually available with the SD receiver's high sampling rate. For example, for the EISCAT UHF radar with 0.32 m wavelength, our benchmark 2 MHz sampling rate gives unambiguous velocities in the interval $\pm(f_s/2) \cdot (\lambda/2) = \pm 160 \text{ km s}^{-1}$. For near-vertical antenna pointing, it is more than sufficient to monitor the velocity interval $\pm 5 \text{ km s}^{-1}$. Therefore (for each range gate R_j) we can downsample, decimate, the vector w to be Fourier-transformed in the MF-formula Eq. (45),

$$w_n = z_n \bar{x}_{n-j} e^{-i\alpha_d n^2}, \quad (47)$$

by as much as $M_{\text{dec}} = 160/5 = 32$. Typically we use $M_{\text{dec}} = 15$ in this case. We form the decimated vector, w' , by adding w_n s in blocks of M_{dec} . At the same time, we make use of the fact that within a block, the acceleration factor $e^{-i\alpha_d n^2}$ is almost constant. We take it out of the decimation sum, to reduce

both the number of multiplications and the number of complex exponentials that need to be computed.

Second, we make use of the fact that most of the w_n s are zeros. The transmission duty cycle in EISCAT experiments is about 10% at the UHF radar. Therefore, about 90% of the transmission samples x_n are zeros (contain no transmission), in regularly spaced blocks. The products w_n of course need to be computed and decimated only for the non-zero x_n . The (perhaps) surprising step in the algorithm is to ignore the zeros even when collecting the decimated products into the final Fourier-transform input vector, w'' ; we just concatenate the non-zero blocks. The FMF formula is

$$\text{FMF}(R_j, \omega) = \frac{|\sum_{n=0}^{N''-1} w''_n e^{-i\omega n}|}{\|x\|}, \quad (48)$$

which for the points $\omega_k = 2\pi k/N''$ is evaluated using FFT. The vector w'' is typically two orders of magnitude shorter than w . In the benchmark case where w contains 600 000 points, w'' has length $N'' = (1/15) \cdot 0.1 \cdot 600\,000 = 4000$. Due to the much shorter FFT input vector, even allowing for the extra work due to the decimation, in typical cases the FMF is 100–300 times faster than the MF.

Gaining speed by the FMF-algorithm is not in doubt. But what is the price? Decimation, the first step in the algorithm, does not cost us much information. Basically, we are just backtracking from our initial “oversampling”. We can backtrack at this stage, but not earlier, because the vector w is near zero frequency, while the raw data vector z , for multi-frequency transmission, is not. In Eq. (47), each frequency channel in z gets multiplied by the complex conjugate of the corresponding transmission, so the carriers of all the frequency channels (approximately) cancel out simultaneously. Very near the MF maximum, also a possible phase modulation is cancelled out. So the sampling requirement of w is determined by the size of the maximum anticipated Doppler-shift only.

We now inspect the second step. What effect does the removal of periodically repeated blocks of zeros have on the result of the Fourier-transform? We ignore the somewhat trivial decimation step, taking $M_{\text{dec}} = 1$. We want the FMF only in the vicinity of its maximum, so we assume that the correct range and acceleration have already been found and the corresponding phase factors cancelled out from the vector w , and only a Doppler term $e^{i\omega_0 n}$, where ω_0 is the target’s normalized Doppler-shift, still remains. We will also ignore the noise, so we are actually computing what might, for consistency of the nomenclature, be called the fast ambiguity function, FAF.

We assume a single-frequency transmission, consisting of M pulses of L sam-

ples each, transmitted using an interpulse period of P samples. Then, near the maximum, w consists of M pulses of, say, unit amplitude and L samples, each Doppler-shifted by ω_0 , with $P - L$ zeros between each pair of pulses. The non-zero part of w consists of M blocks, and in the m th block, the w_n s take the values

$$w_n^{(m)} = e^{i\omega_0(n+mP)}, n = 0, \dots, L - 1. \quad (49)$$

For computing $\text{FAF}(\omega)$, the blocks $\{w_n^{(m)}\}$ are first concatenated and then multiplied by $e^{-i\omega n}$. The m th block is multiplied by

$$u_n^{(m)} = e^{-i\omega(n+mL)}, n = 0, \dots, L - 1. \quad (50)$$

The contribution $I^{(m)}$ of the m th block to the sum in the denominator of Eq. (48) is

$$I^{(m)} = \sum_{n=0}^{L-1} w_n^{(m)} \cdot u_n^{(m)}. \quad (51)$$

The norm $\|x\|$ in the denominator of Eq. (48) is the sum of ML terms, all equal to unity, so we get from Eq. (48) and (51)–(49)

$$\text{FAF}(\omega) = \frac{1}{ML} \left| \sum_{m=0}^{M-1} I^{(m)} \right| \quad (52)$$

$$= \frac{1}{ML} \left| \sum_m \sum_n e^{i\omega_0(n+mP)} e^{-i\omega(n+mL)} \right| \quad (53)$$

$$= \frac{1}{L} \left| \sum_{n=0}^{L-1} e^{i(\omega_0 - \omega)n} \right| \cdot \frac{1}{M} \left| \sum_{m=0}^{M-1} e^{i(\omega_0 P - \omega L)m} \right| \quad (54)$$

$$= \text{diric}(\omega_0 - \omega, L) \cdot \text{diric}(\omega_0 P - \omega L, M), \quad (55)$$

where we have defined the Dirichlet kernel diric by

$$\text{diric}(x, M) = \left| \frac{\sin xM/2}{M \sin x/2} \right|. \quad (56)$$

The first factor in Eq. (55) represents the Doppler-velocity information available from a single pulse. The factor has an absolute maximum at the target's Doppler-frequency ω_0 , while the zeros nearest to the maximum are at $\omega_0 \pm 2\pi/L$. The second factor in Eq. (55) results from pulse repetition. It has maxima, all equal to unity, at the frequencies

$$\omega_n = \frac{P}{L}\omega_0 + n\frac{2\pi}{L}. \quad (57)$$

In general, none of the maxima ω_n coincides with ω_0 . Therefore, the maximum of $\text{FAF}(\omega)$ is not situated in the expected place ω_0 . A velocity estimate computed with FMF will be biased by an amount that depends on the target velocity. This is not a serious handicap, however, at least not in target detection phase, for the bias is always rather small, less than 0.2 km s^{-1} in the standard measurements.

Potentially more problematic for target detection is the reduction of the FMF maximum value compared to the MF maximum value. Some loss of the integrated signal amplitude is to be expected, for the FMF is no longer the optimal solution to the estimation problem. The amount of reduction depends on the target Doppler-shift ω_0 . When ω_0 is of the form $\frac{2\pi}{P-L}n$, there is no reduction. A more useful property is that the loss in all cases has an upper bound, which is tolerably small. To find the bound, we note that the maximum value of the FAF occurs very near to the ω_n which is nearest to ω_0 . Such an ω_n , according to Eq. (57), is never further away from ω_0 than half the spacing $2\pi/L$ between the ω_n s. Therefore, the FAF maximum value in the worst case is roughly equal to $\text{diric}(\pi/L, L) \approx 2/\pi$, or 64% of the ideal value. This result applies to the single-frequency case. The more realistic case of multiple frequencies is more complicated; but an upper bound still exists.

Often we can observe the target for a few seconds during its beam passage. During that time, its velocity typically varies so much that for some integration, it is near one of the values where the FMF maximum is near the MF maximum. This perhaps explains why we in practice seem to achieve almost the same detection sensitivity with the FMF as with the MF. The main difference is that detection with the FMF proceeds more than a hundred times faster.

8 Acknowledgements

We thank A.P. van Eyken for reading the manuscript and making more than 120 suggestions for improvements, and M. Markkanen for several clarifying discussions.

The EISCAT facility is supported by Finland (SA), France (CNRS), the Federal Republic of Germany (MPG), Japan (NIPR), Norway (NFR), Sweden (VR), and the United Kingdom (PPARC).

References

Baron, M. The EISCAT facility. *J. atmos. terr. Phys.* 46, 469-472, 1984.

- Baron, M. EISCAT progress 1983–1985. *J. atmos. terr. Phys.* 48, 767-772, 1986.
- ESA Directorate of Technical and Operational Support ESOC Ground Segment Engineering Department Mission Analysis Section. Study specification, measurements of small-size debris with backscatter of radio waves. Darmstadt, Germany, 1999.
- ESA Directorate of Technical and Operational Support ESOC Ground Segment Engineering Department Mission Analysis Section. Study specification, real-time space debris detection with EISCAT radar facilities. Darmstadt, Germany, 2002.
- Lehtinen, M. *Statistical Theory of Incoherent Scatter Radar Measurements*. EISCAT Techn. Note 86/45, Eur. Incoherent Scatter Sci. Assoc., Kiruna, Sweden, 1986.
- Lehtinen, M., Markkanen, J., Väänänen, A. et al. A new incoherent scatter technique in the EISCAT Svalbard radar. *Radio Sci.* 37, 3-1 – 3-14, 2002.
- Mahafza, B.R. *Radar System Analysis and Design using MATLAB*. Chapman & Hall/CRC, London, 2000.
- Markkanen, J., Lehtinen, M., Huuskonen A., Väänänen, A. *Measurements of Small-Size Debris with Backscatter of Radio Waves*, Final Report, ESOC Contract No. 13945/99/D/CD, 2002.
- Markkanen, J., Postila, M. *Real-Time Small-Size Space Debris Detection with EISCAT Radar Facilities*, Final Report, ESOC Contract No. 16646/02/D/HK(CS), 2005.
- Skolnik, M.I. *Introduction to Radar Systems*, second ed. McGraw-Hill, Singapore, 1981.
- Wannberg, G., Wolf, I., Vanhainen, L-G. et al. The EISCAT Svalbard radar, a case study in modern incoherent scatter radar system design. *Radio Sci.* 32, 2283-2307, 1997.