MEASUREMENTS OF SMALL-SIZE DEBRIS WITH BACKSCATTER OF RADIO WAVES

J Markkanen, M Lehtinen, A Huuskonen, A Väänänen

Final Report
of
ESOC Contract No. 13945/99/D/CS
with
Sodankylä Geophysical Observatory

ESA/ESOC Technical Management
M Landgraf, R Jehn

March 2002

Abstract

Our main objective in this ESA contract study was to define a concept how an existing European radar can be used to detect and characterize small-size space debris (SD) by applying low cost hardware and software upgrades. We concentrate on the EISCAT incoherent scatter ionospheric radars, located in northern Scandinavia and on Svalbard. We show performance estimates which indicate that the EISCAT UHF (930 MHz) radar is the most sensitive and accurate of the EISCAT radars for SD work. The expected performance is the same as the German FGAN radar performance.

For parameter estimation we apply Bayesian statistical inversion. We start from a four-parameter model (amplitude, target range, velocity, acceleration) of the SD signal and find a probabilistic solution for the parameters, given a sampled noisy signal m. The parameter estimation leads to maximization of a quantity we call the Generalized Match Function (GMF). The GMF matches the measured signal against a time-shifted, frequency-shifted transmission $T_{t,\omega}$, GMF $= |<m, T_{t,\omega}>|$. The GMF method automatically makes use of the signal coherence to enhance detection sensitivity. Also, since the transmission $T_{0,0}$ is directly measured, very little needs to be known about what the host radar is doing, the full information we can gain from the measurement is contained in m and $T_{0,0}$. We develop the GMF formalism from first principles in this report.

For the SD measurements we want coherent integration, which EISCAT neither needs nor provides for. To allow atmospheric and SD measurements to proceed simultaneously, we use a separate receiver backend of our own in parallel with the EISCAT receiver. The SD receiver has a multi-channel sampler, a memory buffer, and programmable control. The samples flow to the buffer which is visible to a workstation. The workstation stores the data to disk at the sampling rate, 2–3 Msamples/s. In this study, we analysed our data off-line; the final aim is real-time detection. We estimate that real-time detection in the GMF scheme requires at least 1 Gflop/s computing speed. We have developed the necessary programs to scan, archive and analyse off-line the large amounts of data, 20–30 GB/h, that our approach produces. We describe the main features of this software.

A test campaign in February 2001 at the EISCAT UHF radar showed that objects with diameter about 2 cm are readily observable when in the centre of the radar beam. In 4.5 h, we found 56 clearly detectable hard targets between 490 and 1480 km altitude, with effective diameters between 1.9 cm and 52 cm. We elaborate on these results. We also compare the strongest event with catalogue predictions.

We conclude that piggy-backed SD measurements in EISCAT are efficient, technically straightforward, and (almost) ready to go.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background and overview

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 includes the functioning satellites, but by far most of the 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 of the smaller particles is fragmentary and mainly statistical.

Our aim in this work has been to determine the feasibility of using the European ionospheric radar system EISCAT in northern Scandinavia for the study of small-size SD. Since the early 1980s, the EISCAT radars have been performing ionospheric measurements about 2000 hours per year. In this feasibility study we have shown that it is technically straightforward to piggyback the SD measurement on top of the normal EISCAT ionospheric measurements, without interfering with those measurements. This could open up an extensive and, in terms of sensitivity and accuracy, competitive new radar resource for SD measurements.

The EISCAT system [1, 2, 3] consists of three separate radars: monostatic VHF radar, located near Tromsø, Norway, and operating at 224 MHz; monostatic but two-antenna EISCAT Svalbard Radar in Longyerbyen, 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. Except for a few special campaigns related to meteor measurements [4, 5], EISCAT has not been used to measure hard targets. The normal ionospheric EISCAT experiments, especially on the VHF radar which customarily measures at an altitude around 1000 km, see often, typically once in 5–10 minutes, a point-like, Doppler-shifted target, which stays visible from a few seconds to half a minute. These strong echoes are called "satellites" by the EISCAT community, and are a considerable nuisance in the ionospheric work. Slightly less frequently, the satellite echoes are actually observed by all EISCAT radars, but to our knowledge, no study of them has ever been made.

Even though capable of seeing hard targets, the normal EISCAT experiment setup is not very suitable for detailed measurements of SD. The time resolution is rather poor, typically 5–10 s, and the range resolution has traditionally been several kilometers in the altitude region above 600 km which is of interest in SD measurements. Most problematic is that the traditional EISCAT experiments operate in the "power domain", by computing

for each transmission-reception cycle, the interpulse period (IPP), various signal autocorrelation estimates, called the lagged products, and only saving the (averaged) products. Power domain is fine for the ionospheric plasma where the signal correlation times are shorter than the IPP, and thus the signal is incoherent from IPP to IPP. However, operating in power domain means unnecessary neglect of phase information for signals with phase coherence time up to several hundred milliseconds, covering tens or hundreds of IPPs. For signals with long phase coherence time it makes good sense to work in the "amplitude domain", by adding together individual complex samples in such a way that they have nearly identical phase. This *coherent integration* enhances the signal with respect to the incoherent background noise, making it easier to detect weak signals.

Because of the problems in the standard EISCAT experiment arrangements from the SD measurement point of view, and because of our principle of non-interference with EISCAT operation, we have not used EISCAT digital signal processing facilities in this work. Instead, we have received the analog signal in parallel with the EISCAT receiver, with a separate receiver backend unit of our own. We have hooked to the the EISCAT analog signal path on a suitable coupling point, provided by EISCAT. After that point, we have sampled and processed the signal with our own data acquisition system.

The use of a separate signal processing system, fully under our own control, has allowed us to implement dedicated amplitude domain signal processing. The basis for our SD detection and data analysis is Bayesian statistical inversion [6]. The starting point is a parametrized model of the measurement m as

$$m(t) = s(A, R, \omega, \alpha; t) + \gamma(\sigma; t). \qquad (1.1)$$

The signal s, the target echo, depends on the target range R; the Doppler-shift ω, which corresponds to the radial velocity; and the Doppler-drift α, which corresponds to the radial acceleration of the target. The fourth signal model parameter is the amplitude A, which depends on the target cross section, target range, and radar parameters through the radar equation. In our model, the noise γ is parametrized by its variance σ^2 alone.

The presence of noise introduces a random element into the measurement, which, for weak signals, can be the dominating effect. The statistical inversion approach acknowledges the inherent probabilistic nature of a noise-dominated measurement (or of any measurement), and provides a well-defined procedure to find the fundamental end result of such a measurement. The end result cannot be just a set of four numbers A, R, ω, α. For a noise-dominated situation such a set would be meaningless at the best and misleading at the worst. A more complete answer is the *posteriori density* $D_{\mathrm{p}}(A, R, \omega, \alpha|m)$. It is the conditional probability density for the echo pa-

rameters actually having values A, R, ω, and α, given the measurement m and, possibly, some external information about what to expect.

Though desirable, it is difficult to present usefully the four-dimensional density D_{p}. When listing measurement results in this report, we give only the four numbers A, R, ω, α, and quantities like radial velocity and effective diameter, directly derivable from them, but no error bars. Our parameter values are the Bayesian estimates. These are the most probable values of the parameters, given the measurement. We do not forget about the errors completely, though. We have used the statistical inversion theory to estimate the achievable parameter accuracy, mainly to compare the performance of several radars. In this work the error estimates are based on the linearised model of the measurement. The linearised model uses only the lowest-order terms in the power-series expansion of the model signal $s(A, R, \omega, \alpha)$. We discuss the regime of validity of these error estimates in Section 4.2.4 of this report.

The Bayesian parameter estimation leads to a pattern matching problem. The task is to find the particular waveform, from a set of model waveforms $\{\epsilon^{R,\omega,\alpha}(t)\}$, which gives the best fit to the measured noisy waveform $m(t)$. We call the quantity indicating the degree of match the *generalized match function*, GMF. For sampled data (sample vectors), we define it by

$$\mathrm{GMF} \;=\; \left| < m, \frac{\epsilon^{R,\omega,\alpha}}{\|\epsilon^{R,\omega,\alpha}\|} > \right|, \qquad (1.2)$$

where $< x, y >$ means the inner product of complex vectors and $\|x\|$ is the associated norm. Geometrically, this definition means that the GMF is the length of the orthogonal projection of the measured vector to the model vector. The projection is a suitable quantity to describe how "parallel" or similar two waveforms are. Parameter estimation means adjusting the model parameters so that the model vector gets as parallel to the measurement vector as possible. This is equivalent to finding the GMF maximum. The GMF is the central tool in our analysis of SD data.

It turns out that as a function of R, with fixed ω and α, Eq. (1.2) can be viewed as representing a filtering operation, with m the filter input and the model waveform the filter. In this sense, the acronym stands also for "generalized matched filter". The other parameters besides R then are the "generalization". The final result of our GMF-based parameter estimation method is not new. It reproduces the well-known result of matched filtering that says that the time-reversed replica of any given waveform is the matched filter for the original waveform [7, page 156]. However, we will not need the filtering aspects of the GMF in this work.

Neither will we dwell in this work on the connection between our GMF and what has been called the (radar) ambiguity function [8, pages 411–420]. If there is no noise and no acceleration, our GMF essentially reduces to the

radar ambiguity function. For our purposes we find the name GMF more descriptive, and also less ambiguous, than the original name—ambiguity functions of several different kinds are used extensively in incoherent scatter radar data analysis.

The GMF-based data handling is tailored for efficient processing of the coherent echoes of small, smoothly moving hard targets. In particular, the method provides coherent integration of the signal. The importance of a coherent integrator as a signal detector is that its peak output is proportional to the signal energy W_s rather than just the signal power. Also the noise will manifest itself in the output of the detector; there will always be some possibility of confusing pure noise with the signal. For the coherent integrator, the ratio W_s/P_γ might therefore appear as a good measure to characterize how probable a correct detection is. When the ratio is "large", we expect high probability for correct detection. There is a technical problem, however. The noise power depends on the receiver bandwidth B via the relation $P_\gamma = kT_\gamma B$, where k is Boltzmann constant and T_γ is the noise temperature. For the inherently very narrowband SD signals, there is no natural receiver bandwidth that we could use as a reference when deciding which ratio is large and which is small. We force uniqueness by agreeing always to assume unit noise bandwidth. Thus we define our basic quantity with which to measure signal detectability, the *energy-to-noise ratio*, as

$$\text{SNR}_\text{N} = \frac{W_\text{s}}{kT_\gamma}. \tag{1.3}$$

This is a dimensionless quantity. For correctly sampled signals the sampling interval τ is the reciprocal of the receiver bandwidth. If the signal lasts for N samples, and P_s denotes the mean signal power, we have

$$\text{SNR}_\text{N} = \frac{N\tau P_\text{s}}{kT_\gamma} = \frac{NP_\text{s}}{kBT_\gamma} = N \cdot \text{SNR}, \tag{1.4}$$

where

$$\text{SNR} = P_\text{s}/P_\gamma = |A|^2/\sigma^2 \tag{1.5}$$

is the usual signal-to-noise ratio, computed from the complex signal amplitude A and the variance σ^2 of the Gaussian noise. The directly measurable quantities SNR and N depend on the receiver bandwidth used in the measurement, their product SNR_N does not.

In addition to justifying our notation for the energy-to-noise ratio, (1.4) makes explicit the main benefit of coherent integration: the integration amplifies the signal-to-noise ratio, directly proportional to the integration time. In the GMF method, the length of integration is in principle limited only by the coherence time of the signal. We can expect coherent integration to have considerable effect on detection sensitivity and parameter estimation accuracy.

1.2   Study objectives

According to our contract agreement with ESOC, the objective of this study is "to define a concept how existing European radar or radio astronomy facilities can be used to detect and characterize small size debris by applying low cost hardware and software upgrades. In this study a suitable and cost-efficient concept shall be developed and its feasibility shall be demonstrated in a test campaign. The data recording and analysis chain shall be described in full detail" [9].

1.3   Structure of the study and organization of this report

This study was formally divided into four work packages

1. definition of a concept to detect small-size debris,
2. specification of data collection and analysis,
3. software and hardware upgrades,
4. test campaign.

In *work package 1*, we defined our measurement model. We applied statistical inversion to linearised version of the model, to compare the EISCAT radars as well a few others in terms of accuracy and sensitivity. The main purpose of the comparison was to find the EISCAT radar that would be most suitable for the test campaign of WP4. This work is described in Chapter 2 of this report.

In *work package 2*, we presented our data collection system. We had actually used a preliminary but functional version of the data collection system already in November 1999, a few months before entering the ESA contract in January 2000, to do ionospheric measurements at EISCAT Svalbard Radar. The data collected during that campaign turned out to be very useful test data when developing the GMF-based analysis procedures during WP2. These developments are described in Chapter 3.

In *work package 3*, we continued developing the analysis software. Also, a new version of the data collection system was being developed by a Finnish company, Invers Ltd, which had actually also lent us the hardware we had used in the November ESR measurements. These developments, both in software and in hardware, were low level and are not interesting enough to be detailed separately in this report.

We carried out the test campaign of *work package 4* in February 2001 at EISCAT UHF radar in Tromsø, in connection with a two-week Finnish EISCAT campaign. Analysis results from the test campaign are detailed and discussed in Chapter 4.

The structure of this report follows closely the order of the work packages. After this introductory chapter, in *chapter 2* we introduce the theory on which our measurements and analysis are based. We start with a simple sensitivity comparison of several radars, based on the radar equation. Then we specify our signal model and give a relatively self-contained derivation of the posteriori density D_p. The signal model assumes single-frequency transmission and hence misses an interference that turned up in actual analysis; discussion of this complication is in Appendix A. We derive Bayesian estimation formula for the basic parameters range, velocity, acceleration, and amplitude, showing the crucial role of the GMF. We use linearised theory to make accuracy comparisons between the various EISCAT radars. Next we describe the discretization that we have used for the actual detection and parameter estimation, and summarize the estimation procedure. In the last part of chapter 2 we handle the computational aspects of the GMF method. In particular, we estimate the computing power needed to perform SD detection in real time.

In *chapter 3* we describe our measuring and analysis system, both the hardware and the software. First we give an introduction to the EISCAT UHF radar, which is the radar we used during the test campaign. Then we describe our SD receiver and explain how it was used during the test campaign. On the software side, we explain the whole processing chain, from the sampling done under control of firmware, through scanning, detection, and archiving, up to the plotting of summaries of analysis results.

In *chapter 4* we describe the test campaign and the analysis results. We summarize the campaign execution. Then we display the analysis results in various ways. We use the analysis results to discuss the achievable sensitivity and accuracy, and to evaluate the usefulness of a fast but approximate method for GMF evaluation which we call the FastGMF algorithm. We defer discussion of the details of the underlying EISCAT experiments and the SD analysis setup to experiment-specific Appendixes B and C.

In *chapter 5* we summarize what we have learned from this study, but also discuss a few problems, related to measurement accuracy and signal coherence. We state our conclusions from the test campaign. Finally, we suggest steps for further development.

Chapter 2

Theory

2.1 Analysis by the radar equation

2.1.1 Radar equation for point targets

The power observed at the radar receiver is computed from the radar equation, which for point targets reads

$$P_{\mathrm{r}} = \frac{P_{\mathrm{t}}G}{4\pi R^2}\ \sigma_* \ \frac{A_{\mathrm{eff}}}{4\pi R^2},\tag{2.1}$$

where P_{r} is the power at the receiver input, P_{t} is the transmitter power, G is the antenna gain, R is the range to the scattering point, A_{eff} is the effective area of the antenna, and σ_* is the radar scattering cross section of the target [8, page 4]. The effective area is related to the antenna gain by

$$A_{\mathrm{eff}} = \frac{G\lambda^2}{4\pi},\tag{2.2}$$

where λ is the wavelength of the transmission.

2.1.2 Scattering cross section of a perfectly conducting sphere

The radar scattering cross section depends on the form and the size of the reflecting particle in a complicated manner. For the modelling purposes we assume that the debris particle is a sphere. Then, for diameters d small compared to the wavelength, the scattering cross section is proportional to d^6, and for large diameters it is equal to the geometrical cross section $\frac{1}{4}\pi d^2$. The first region is called the Rayleigh region and the latter the optical region. In between there is the resonance or Mie region, in which the scattering cross section shows an oscillatory behaviour. A simple rule says that the Rayleigh approximation is valid if the diameter is less than one fifth of the wavelength and the optical approximation when the diameter is larger than 10 times the wavelength. The cross section is shown in Fig. 2.1.

The thick line in Fig. 2.1 gives a simplifying approximation, in which the first part follows the Rayleigh approximation, the latter part is the optical approximation, and the resonance region is neglected altogether. This is reasonable because in this study the radars operate mostly in the Rayleigh region. The high-frequency radars operate in the resonance region, but there is no point in using the curve in full, as the true debris particles are not perfect spheres, and the result is approximate in any case. Hence the scattering cross section in our model calculations obeys the formula [10, page 95]

$$\frac{\sigma_*}{\frac{1}{4}\pi d^2} = \begin{cases} 9(\frac{\pi d}{\lambda})^4, & \text{when } d < \lambda/(\pi\sqrt{3}), \\ 1, & \text{when } d > \lambda/(\pi\sqrt{3}). \end{cases}\tag{2.3}$$
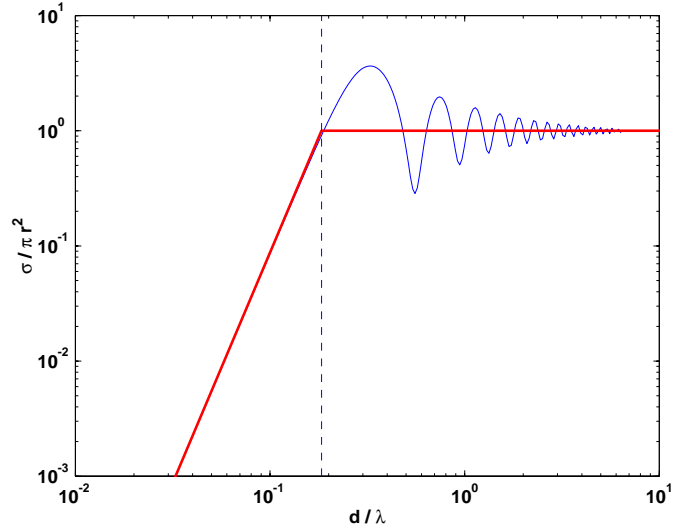
Figure 2.1: The scattering cross section divided by the geometrical cross section for a perfectly conducting sphere. The thick line follows first the Rayleigh approximation and then the optical approximation. The dashed line indicates the crossover point.

2.1.3   A simple sensitivity comparison of six radars

Table 2.1 lists three EISCAT radars and three other radars included for comparison. The non-EISCAT radars are the TIRA tracking radar of Forschungs-gesellschaft für Angewandte Naturwissenschaften (FGAN) at Wachtberg in Germany, a new weather radar of the Finnish Meteorological Institute at Luosto (LUO) in Finland, and a hypothetical radar which we call the U5G. The U5G radar combines the 32 m antenna of the EISCAT UHF system at Sodankylä with a 3 kW transmitter to operate at about 5 GHz. The relevant properties of the radars are given in Table 2.1.

    We want to compare the sensitivity of the radars for small-size SD detection, that is, the radars' ability to detect weak, but coherent, signals. We assume that all the radars can perform coherent integration. Then we can use the energy-to-noise ratio SNR_N as the measure of sensitivity. We have computed in Table 2.1 two benchmark numbers related to SNR_N. First, we predict the SNR_N that the radars should observe when measuring a "standard target", a 3 cm sphere at 1000 km range. We assume 0.1 s coherent integration. Second, we compute the diameter that a target at 1000 km range should have in order to produce $SNR_N = 1$, again with 0.1 s coherent integration. We call this diameter the reference diameter.

    The result of the comparison is that the EISCAT UHF and the FGAN radars are the most sensitive of the existing radars. The other radars are weaker for detection of small SD particles. It appears that the suggested

Table 2.1: Sensitivity comparison of six radars. The ESR system refers to the 42 m antenna.

Radar properties	VHF	ESR	UHF	FGAN	U5G	LUO
Frequency [MHz]	224	500	933	1330	5000	5625
Wavelength [m]	1.338	0.600	0.321	0.225	0.060	0.053
Transmission [MW]	3.0	1.0	1.5	1.6	0.003	0.3
Antenna gain [dB]	43.1	45.3	48.1	49.7	62.5	47.5
Max duty cycle [%]	12.5	25.0	12.5	3.7	25.0	0.12
Noise temperature [K]	100	60	80	209	60	300
Energy-to-noise ratio for 3 cm target						
Cross section σ_* [mm^2]	0.16	3.9	47	190	710	710
Received power [aW]	0.18	0.81	15.3	69.3	12.1	0.96
$\sqrt{\mathrm{SNR_N}}$	1.3	4.9	13.2	9.4	19.1	0.2
Ref. diameter [cm]	2.5	1.6	1.1	1.3	0.5	12.7

special U5G system, which combines a high frequency, low power transmitter with a large antenna, would be by far the most sensitive system.

The reference diameter must not be thought as the minimal detectable size. What can actually be detected depends on how the detection is actually done, and can be done. It turns out that the condition $\mathrm{SNR_N} = 1$ is quite too weak to ensure reliable detection. We found during the test campaign that we actually need $\sqrt{\mathrm{SNR_N}}$ to be about 5 for the target to be clearly detectable. Therefore, the actually detectable size is considerable larger than the reference diameter, see Fig. 4.3.

2.2 Statistical inversion

2.2.1 Signal model

We denote the radar transmission by

$$\mathcal{T}(t) = \epsilon^{\mathrm{T}}(t) \cdot \exp(i\omega_{\mathrm{rad}}t), \qquad (2.4)$$

where the slowly varying $\epsilon^{\mathrm{T}}(t)$ is called the transmission envelope, and ω_{rad} is the radar radio frequency (RF). We assume a pointlike target, at range R_0. If the target is stationary, the received backscattered echo $s(t)$ is just the transmission, delayed by pulse propagation delay $2R_0/c$ and scaled by some scale factor A_0. In the receiver, the RF signal $A_0 \mathcal{T}(t - 2R_0/c)$ is first translated to an intermediate frequency (IF), $\mathcal{T} \to \mathcal{T}'$, and then averaged by the receiver impulse response $p(t)$. In our measuring system, we actually

form the impulse response digitally by adding samples, and the samples are integer number of wavelengths apart. The complex exponentials can be taken out from the convolution $p*\mathcal{T}'$, and we have

$$
\begin{aligned}
s(t) &= A_0\, p*\mathcal{T}'(t - \frac{2R_0}{c}) \\
&= A_0\, p*\epsilon^{\mathrm{T}}(t - \frac{2R_0}{c}) \cdot \exp(i\omega_{\mathrm{IF}}t) \cdot \exp(-i\,\omega_{\mathrm{rad}}\frac{2R_0}{c}). \quad (2.5)
\end{aligned}
$$

We assume that during the beam passage the target has constant radial acceleration, $a_r = \mathrm{const} = a_0$. Then the target range is

$$
R(t) = R_0 + v_0 t + \frac{1}{2}a_0 t^2, \quad (2.6)
$$

where v_0 and R_0 are the initial velocity and initial range, at time $t = 0$. We use Eq. (2.5) also for a slowly moving target, with R_0 replaced by $R(t)$. Then the echo model becomes

$$
s(t) = A_0 \cdot \epsilon^{R_0,\omega_0,\alpha_0}(t) \cdot \exp(i\,\omega_{\mathrm{IF}}t) \cdot \exp(-i\omega_{\mathrm{rad}}\frac{2R_0}{c}), \quad (2.7)
$$

where

$$
\begin{aligned}
\epsilon^{R_0,\omega_0,\alpha_0}(t) &= \epsilon^0(t - \frac{2R(t)}{c}) \cdot \exp(-i\omega_0 t) \cdot \exp(-i\alpha_0 t^2), \quad (2.8) \\
\epsilon^0(t) &= p*\epsilon^{\mathrm{T}}(t), \quad (2.9) \\
\omega_0 &= 2v_0\omega_{\mathrm{rad}}/c, \quad (2.10) \\
\alpha_0 &= a_0\omega_{\mathrm{rad}}/c. \quad (2.11)
\end{aligned}
$$

We assume that the transmission $\mathcal{T}(\mathrm{t})$ is also directly available as a "transmission sample signal", processed similarly to the actual reception, so that it sees the same receiver impulse response. The factor $\epsilon^0(t)$ in (2.8) is the envelope of this IF signal. Note that in Eq. (2.8) we have taken the complex exponentials outside the convolution. This is a good approximation as long as the phase drift due to Doppler-frequency, $\Delta\phi_\omega = \omega_0 t_{\mathrm{p}}$, and the phase drift due to the acceleration, $\Delta\phi_\alpha = \alpha_0 t_{\mathrm{p}}^2$, are small during the duration t_{p} of the impulse response. The impulse response duration is about equal to the sampling interval τ, $\tau \sim 1$ μs. Typical Doppler-velocity is 1.5 km s^{-1}, which corresponds to $\omega_0 = 2\pi \cdot 9$ kHz. Typical Doppler-acceleration is 100 m s^{-2}, which corresponds to $\alpha_0 t_p = 2\pi \cdot 0.3$ mHz. Clearly, both $\Delta\phi_\omega$ and $\Delta\phi_\alpha \ll 1$.

The last factor in Eq. (2.7) is constant when there is only a single frequency in the transmission, and can be dropped. In the case of multiple frequencies the term must be taken into account, this case is handled in Appendix A. The factor $\exp(i\,\omega_{\mathrm{IF}}t)$ will either be paired with its complex conjugate in an inner product, or be taken the magnitude of. In both cases, it will yield unity. We will drop it from our expressions from now on.

We denote the sampled waveforms with bold typeface, so that for example $\mathbf{e}^{R,\omega,\alpha}$ is a vector with elements

$$\epsilon_n^{R,\omega,\alpha} = \epsilon^{R,\omega,\alpha}(n\tau). \tag{2.12}$$

2.2.2 Measurement in the presence of noise

The measured, complex-valued (detected) signal m contains noise γ on top of the target echo s,

$$m(t) = s(t) + \gamma(t). \tag{2.13}$$

We assume complex Gaussian noise, that is,

$$\gamma = \gamma_{\text{re}} + i \cdot \gamma_{\text{im}} \longleftrightarrow (\gamma_{\text{re}}, \gamma_{\text{im}}), \tag{2.14}$$

where γ_{re} and γ_{im} are independent Gaussians, with zero mean and variance $\sigma^2/2$. Then, for any given moment of time, the joint probability density function (PDF) of the two-dimensional vector $(\gamma_{\text{re}}, \gamma_{\text{im}})$ is

$$D(\gamma) = \frac{1}{\pi\sigma^2} \exp(-\frac{1}{\sigma^2}|\gamma|^2). \tag{2.15}$$

With this normalization, σ^2 equals the mean noise power

$$P_\gamma = \text{E}|\gamma|^2 = \text{E}(\gamma_{\text{re}}^2 + \gamma_{\text{im}}^2) = \sigma^2. \tag{2.16}$$

We assume that the noise samples $\gamma_n = \gamma(n\tau)$ are uncorrelated. This requires that the duration of the receiver impulse response is not longer than the sampling interval τ. Then the samples $m_n = m(n\tau)$ are statistically independent. We will also assume that the noise is stationary, even though this is a rather poor approximation for the actual radar data.

We denote by $D(m_n|s_n)$ the conditional PDF of m_n, given s_n. This is just the PDF of the noise,

$$D(m_n|s_n) = \frac{1}{\pi\sigma^2} \exp(-\frac{1}{\sigma^2}|m_n - s_n|^2). \tag{2.17}$$

Assume that we take M samples, and denote by $\mathbf{m}$, $\mathbf{s}$ and $\mathbf{\Gamma}$ the M-dimensional sample vectors. For future reference, assume also that we have N signal samples, so that the signal vector $\mathbf{s}$ is padded to length M by zeros. Because of the assumed independence of the samples, the conditional joint density of all the M samples ("measurements") is

$$D(\mathbf{m}|\mathbf{s}) = \prod_{k=0}^{M-1} D(m_k|s_k) = \frac{1}{(\pi\sigma^2)^M} \cdot \exp(-\frac{1}{\sigma^2}\|\mathbf{m} - \mathbf{s}\|^2). \tag{2.18}$$

We used the standard notation for complex vectors $\mathbf{x} = (x_0 \ldots x_{M-1})$, with the inner product of $\mathbf{x}$ and $\mathbf{y}$ as

$$< \mathbf{x}, \mathbf{y} >= \sum_{k=0}^{M-1} x_k \overline{y_k} \qquad (2.19)$$

and the norm as

$$\|\mathbf{x}\| = \sqrt{< \mathbf{x}, \mathbf{x} >}. \qquad (2.20)$$

2.2.3   Parameter estimation, inversion solution

Based on a M-point measurement $\mathbf{m}$, we want parameters A, R, ω, and α of the signal. In general, it is not possible to solve the parameters uniquely, and the best we can then do is to find a joint probability density for the various conceivable values. Equation (2.18) gives the so called *direct theory* of the estimation problem, the conditional probability density of the measurement $\mathbf{m}$, given signal $\mathbf{s}$. The *inverse problem* is to find the PDF of the signal parameters, A, R, ω, α in this case, given a measurement $\mathbf{m}$. This sought-for PDF we denote by $D_{\mathrm{p}}(A, R, \omega, \alpha | \mathbf{m})$ or by $D_{\mathrm{p}}(\mathbf{s} | \mathbf{m})$. It is called the *posteriori density.*

For brevity, we rewrite the direct theory (2.18) as

$$D(\mathbf{m}|\mathbf{s}) = D(\mathbf{m}|A, R, \omega, \alpha) = \frac{1}{(\pi \sigma^2)^M} \cdot e^{\mathcal{L}}, \qquad (2.21)$$

where the "likelihood function" $\mathcal{L} = \mathcal{L}(A, R, \omega, \alpha, \mathbf{m})$ is

$$\mathcal{L} = \frac{-1}{\sigma^2} \|\mathbf{m} - \mathbf{s}(A, R, \omega, \alpha)\|^2, \qquad (2.22)$$

and $\mathbf{s}(A, R, \omega, \alpha)$ means the signal model vector that corresponds to the indicated parameters.

All aspects of the measuring situation would be known if we would know the joint probability density $D_{\mathrm{sm}}(\mathbf{s}, \mathbf{m})$, for all $\mathbf{m}$ and $\mathbf{s}$. According to the basic property of conditional probability, the joint density D_{sm} can be expressed in two ways,

$$D(\mathbf{s}|\mathbf{m})D_{\mathrm{s}}(\mathbf{m}) = D_{\mathrm{sm}}(\mathbf{s}, \mathbf{m}) = D(\mathbf{m}|\mathbf{s})D_{\mathrm{m}}(\mathbf{s}), \qquad (2.23)$$

where $D_{\mathrm{m}}(\mathbf{s})$ and $D_{\mathrm{s}}(\mathbf{m})$ are

$$D_{\mathrm{s}}(\mathbf{m}) = \int d\mathbf{s} \; D_{\mathrm{sm}}(\mathbf{s}, \mathbf{m}) \qquad (2.24)$$

and

$$D_{\mathrm{m}}(\mathbf{s}) = \int d\mathbf{m} \; D_{\mathrm{sm}}(\mathbf{s}, \mathbf{m}). \qquad (2.25)$$

Solving Eq. (2.23) for the sought-for density $D(\mathbf{s}|\mathbf{m})$ gives the Bayes inversion formula

$$D_{\mathrm{p}}(\mathbf{s}|\mathbf{m}) = \frac{1}{D_{\mathrm{s}}(\mathbf{m})} \cdot D_{\mathrm{m}}(\mathbf{s}) \cdot D(\mathbf{m}|\mathbf{s}). \qquad (2.26)$$

In the context of statistical inversion theory, the marginal density $D_{\mathrm{m}}(\mathbf{s})$ is called the *prior density* or the *a priori density*, and is denoted by $D_{\mathrm{pr}}(\mathbf{s})$. According to its definition (2.25), the prior density describes the relative frequency of the parameters in the target population. The joint density D_{sm} is, of course, not known (if it where, there would be nothing left to do), and so the a priori density cannot be actually computed from Eq. (2.25). Instead, one *assumes* some form for D_{pr}, and inserts it to the right-hand side of Eq. (2.26). We incorporate a priori information into the inversion problem in this way. For lack of anything better, or for simplicity, it is common to assume a uniform a priori density in some interval, and zero outside. In the SD measuring problem it might make good sense to use non-flat a priori distributions for all parameters. However, in this study we assume flat a priori distribution.

Once D_{pr} and the direct theory density $D(\mathbf{m}|\mathbf{s})$ have been fixed, the factor $D_{\mathrm{s}}(\mathbf{m})$ in (2.26) can be determined, at least in principle, by requiring that $D_{\mathrm{p}}(\mathbf{s}|\mathbf{m})$ is a properly normalized probability density. We have

$$\begin{aligned} D_{\mathrm{p}}(A, R, \omega, \alpha|\mathbf{m}) &= C(\mathbf{m}) \, \mathrm{e}^{\mathcal{L}(A,R,\omega,\alpha,\mathbf{m})} \\ &= C(\mathbf{m}) \, \mathrm{e}^{-\frac{1}{\sigma^2}\|\mathbf{m}-\mathbf{s}(A,R,\omega,\alpha)\|^2}, \qquad (2.27) \end{aligned}$$

where $C(\mathbf{m})$ is a normalization factor that does not depend on A, R, ω, α.

The posteriori density D_{p} contains everything there is to be known about the target parameters on the basis of the measurement $\mathbf{m}$ and the prior information; it is the complete answer. If the actual target parameters are $A_0, R_0, \omega_0, \alpha_0$, we hope that the posteriori density would be well concentrated around the point $(A_0, R_0, \omega_0, \alpha_0)$. If we are pressed to give, instead of the full distribution, single numbers as the solution of estimation problem, the *Bayesian estimate* is the vector $(\widehat{A}, \widehat{R}, \widehat{\omega}, \widehat{\alpha})$ that maximizes the posteriori density, with $\mathbf{m}$ fixed.

Equation (2.27) shows that we get the signal parameter estimates by any of three equivalent ways, namely,

    ▷ from the maximum of the posteriori density D_{p},
    ▷ from the maximum of the likelihood function $\mathcal{L}$,
    ▷ from the minimum of the norm $K = \frac{1}{\sigma^2}\|\mathbf{m} - \mathbf{s}(A, R, \omega, \alpha)\|^2$.

To find the minimum of the norm, we write K as

$$\sigma^2 K = \|\mathbf{m}\|^2 + \|\mathbf{s}\|^2 - 2\mathrm{Re}(<\mathbf{m}, \mathbf{s}>)$$

$$
\begin{aligned}
&= \ \|\mathbf{m}\|^2 + \left| \|\mathbf{s}\| - \frac{<\mathbf{m},\mathbf{s}>}{\|\mathbf{s}\|} \right|^2 - \left| \frac{<\mathbf{m},\mathbf{s}>}{\|\mathbf{s}\|} \right|^2 \\
&= \ \|\mathbf{m}\|^2 + \left| \|\mathbf{s}\| - \frac{<\mathbf{m},\mathbf{s}>}{\|\mathbf{s}\|} \right|^2 - \mathrm{GMF}(R,\omega,\alpha)^2, \qquad (2.28)
\end{aligned}
$$

where we observed that the GMF, as defined in Eq. (1.2), does not depend on the signal amplitude A,

$$
\begin{aligned}
\mathrm{GMF} \ &= \ \frac{\left| <\mathbf{m}, \mathbf{e}^{R,\omega,\alpha}> \right|}{\|\mathbf{e}^{R,\omega,\alpha}\|} \qquad\qquad (2.29) \\
&= \ \frac{\left| <\mathbf{m}, A\mathbf{e}^{R,\omega,\alpha}> \right|}{\|A\mathbf{e}^{R,\omega,\alpha}\|} \\
&= \ \left| \frac{<\mathbf{m},\mathbf{s}>}{\|\mathbf{s}\|} \right|.
\end{aligned}
$$

Equation (2.28) shows that the parameter estimation task divides into two jobs: minimizing the second term (to its absolute minimum, zero), and maximizing the GMF term. We have the following two-step estimation recipe.

1. Maximize the term GMF^2 in Eq. (2.28). This gives the estimates $\widehat{R}$, $\widehat{\omega}$ and $\widehat{\alpha}$,

$$
(\widehat{R},\widehat{\omega},\widehat{\alpha}) = \arg \max_{R,\omega,\alpha} \mathrm{GMF}. \qquad\qquad (2.30)
$$

2. Select the A in $\mathbf{s} = A\mathbf{e}^{\widehat{R},\widehat{\omega},\widehat{\alpha}}$ in such a way that the second term in (2.28) becomes zero. This requires

$$
\widehat{A} = \frac{<\mathbf{m}, \mathbf{e}^{\widehat{R},\widehat{\omega},\widehat{\alpha}}>}{\|\mathbf{e}^{\widehat{R},\widehat{\omega},\widehat{\alpha}}\|^2}. \qquad\qquad (2.31)
$$

From (2.31) we get

$$
\|\widehat{\mathbf{s}}\| = \|\widehat{A}\mathbf{e}^{\widehat{R},\widehat{\omega},\widehat{\alpha}}\| = \mathrm{GMF}(\widehat{R},\widehat{\omega},\widehat{\alpha}). \qquad\qquad (2.32)
$$

The signal energy W_s is the sum of the squares of the signal samples, times the sampling interval,

$$
W_s = \sum \tau |s_n|^2 = \tau \|\mathbf{s}\|^2, \qquad\qquad (2.33)
$$

so that from (2.32) we get an estimate for the signal-to-noise ratio,

$$
\widehat{\mathrm{SNR}} = \frac{\langle P_{\widehat{s}} \rangle}{P_\gamma} = \frac{W_{\widehat{s}}/(N\tau)}{\sigma^2} = \frac{1}{N} \frac{\mathrm{GMF}(\widehat{R},\widehat{\omega},\widehat{\alpha})^2}{\sigma^2}, \qquad\qquad (2.34)
$$

where N is the number of signal samples.

Under very good SNR, $\mathbf{m} \approx \mathbf{s}_0 = A_0 \mathbf{e}^{R_0,\omega_0,\alpha_0}$, and we have

$$
\begin{aligned}
\text{GMF}(R_0,\omega_0,\alpha_0) &\approx \frac{|<A_0\mathbf{e}^{R_0,\omega_0,\alpha_0}, e^{R_0,\omega_0,\alpha_0}>|}{\|e^{R_0,\omega_0,\alpha_0}\|} \\
&= \|A_0\mathbf{e}^{R_0,\omega_0,\alpha_0}\|.
\end{aligned}
\tag{2.35}
$$

On the other hand, still under high SNR, it follows from the Schwartz inequality that

$$
\text{GMF}(R,\omega,\alpha) \leq \|A_0\mathbf{e}^{R_0,\omega_0,\alpha_0}\|, \quad \text{for all } R,\omega,\alpha.
\tag{2.36}
$$

Equations (2.36) and (2.35) imply that (R_0,ω_0,α_0) is the GMF maximum point, and we have $\text{GMF}_{\max} = \|s_0\|$. Thus we have the desirable result that our estimation method is able to recover the actual signal parameters in the case of very high SNR.

The behaviour of the parameter estimates under poor SNR is much less clear. For instance, the estimate $\|\hat{\mathbf{s}}\|$ will not go to zero when the actual signal vanishes. When the signal is very weak, $\mathbf{m} \approx \mathbf{\Gamma}$, and the estimate (2.32) becomes

$$
\|\hat{\mathbf{s}}\| \approx \max_{R,\omega,\alpha} \left| \sum_{n=j}^{N+j-1} f_n \gamma_n \right|,
\tag{2.37}
$$

where the coefficients f_n depend on (R,ω,α) and the shape of the transmission envelope, and are normalized by $\sum |f_n|^2 = 1$. The index j relates to the range parameter R. Obviously, whatever value the $\|\hat{\mathbf{s}}\|$ in (2.37) evaluates to, it will not evaluate to zero, and, for small signal amplitude, it is also essentially independent of the amplitude. This means that the amplitude of weak signals tends to be overestimated (see Fig. 2.10).

Even when we take in (2.37) the max only over the ω and α parameters, at a fixed range, the result is non-zero, and is almost independent of the signal for weak signals. It is also independent of the range. This prompts us to define for target detection purpose a dimensionless quantity, which we call simply the "Ratio" and denote by $\mathcal{R}$, as

$$
\mathcal{R} = \max_{\omega,\alpha} \frac{\text{GMF}(R,\omega,\alpha)}{\sigma}.
\tag{2.38}
$$

The Ratio maximum, $\mathcal{R}_{\max}$, is just the square root of the energy-to-noise ratio,

$$
\mathcal{R}_{\max} = \max_{R} \mathcal{R} = \max_{R,\omega,\alpha} \frac{\text{GMF}}{\sigma} = \frac{\|\hat{s}\|}{\sigma} = \sqrt{\frac{W_{\hat{s}}}{P_\gamma}} = \sqrt{\text{SNR}_{\text{N}}}.
\tag{2.39}
$$

When there is a target in the measurement, $\mathcal{R}(R)$ should exhibit a peak near the target range R_0, and, away from R_0, should relax to a more or less
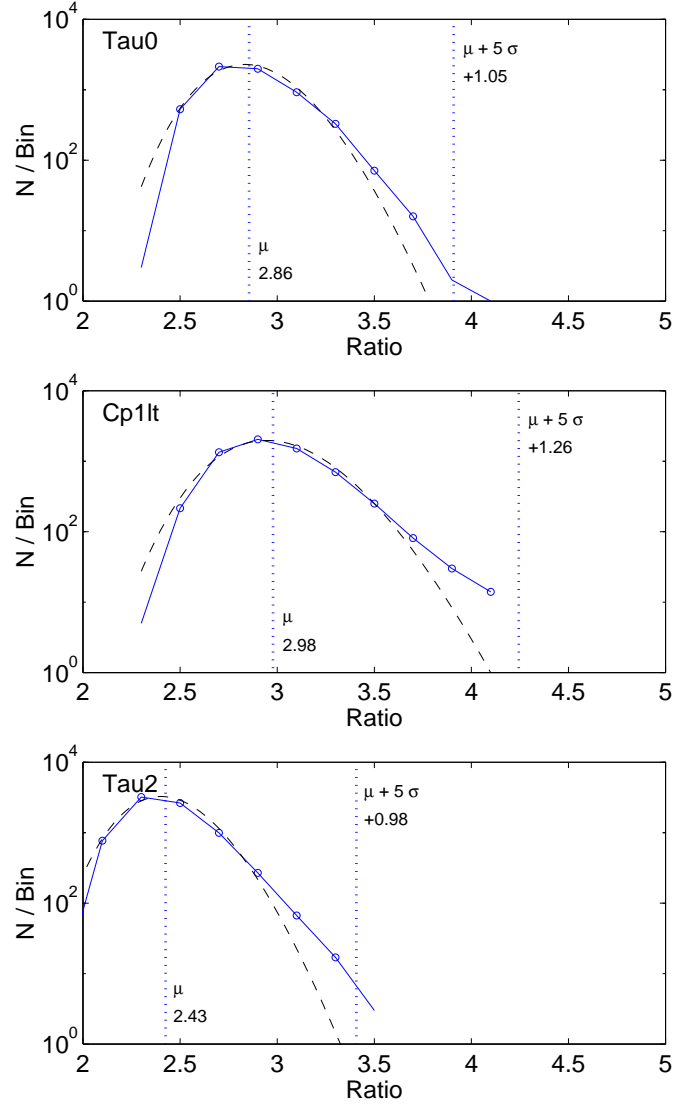
Figure 2.2: Distribution of values of "background" $\mathcal{R}$ in three experiments. Tau0 data are simulated. The integration time was 0.3 s in all cases. The data in each panel are from one full-resolution scan. The dashed line is log-normal distribution with the same mean and variance (μ, σ^2) as the measured data.

noisy "background" level, in principle computable from expressions of the type (2.37). See Fig. 3.4 for an example of actual behaviour of the ratio. Thus, $\mathcal{R}$ is a suitable tool to be used in a threshold-based detection scheme

$$\text{Detection} \longleftrightarrow \max_R \mathcal{R} > \text{Threshold.} \tag{2.40}$$

The ratio $\mathcal{R}$, as a quantity to be kept in computer storage, represents a sensible compromise between the full three-dimensional, but very large, $\text{GMF}(R, \omega, \alpha)$ on one hand, and the single number $\text{GMF}_{\max}$, which no more indicates the data variability at all, on the other.

The detection threshold should be set so that we get at most a prescribed percentage of false detections (detections due to noise only). To do this would require knowledge about the distribution of the $\mathcal{R}$ values in the absence of signal. Figure 2.2 shows a preliminary, numerical inspection of the matter. However, we have not based our detection criterion on the false-alarm rate in this work. Instead, by experimenting with several threshold values, we have set the threshold so high that all our detections actually represent real targets.

2.3 Error analysis in linearised model

The posteriori covariance matrix Σ_p of the model parameters A, R, ω, α is calculated from the equation

$$\Sigma_p = (\mathbf{D^T}\Sigma_m^{-1}\mathbf{D})^{-1}, \tag{2.41}$$

where $\mathbf{D}$ is a matrix formed by the partial derivatives of the model signal, $s(t)$ in Eq. (2.7), with respect to the model parameters [11, pages 74–75]. Assuming the partial derivates constant amounts to using a linearised model of the signal. $\mathbf{D^T}$ is the (conjugate-)transpose of $\mathbf{D}$, and Σ_m is the error covariance matrix of the measurements. In our case the measurements are independent, and hence Σ_m is diagonal. Moreover, we assume that all the measurements, the N individual samples, have equal error σ caused by system noise. In this case, Eq. (2.41) simplifies to

$$\Sigma_p = \sigma^2(\mathbf{D^T}\mathbf{D})^{-1}, \tag{2.42}$$

where σ^2 is the noise power. The errors of the model parameters are obtained as square roots of the diagonal elements of Σ_p.

2.3.1 Results as a function of diameter

Figures 2.3 and 2.4 show the relative error of the signal model parameters as a function of the sphere diameter for four radars. Two radars shown in Table 2.1 were omitted from the model calculations. These were the Luosto
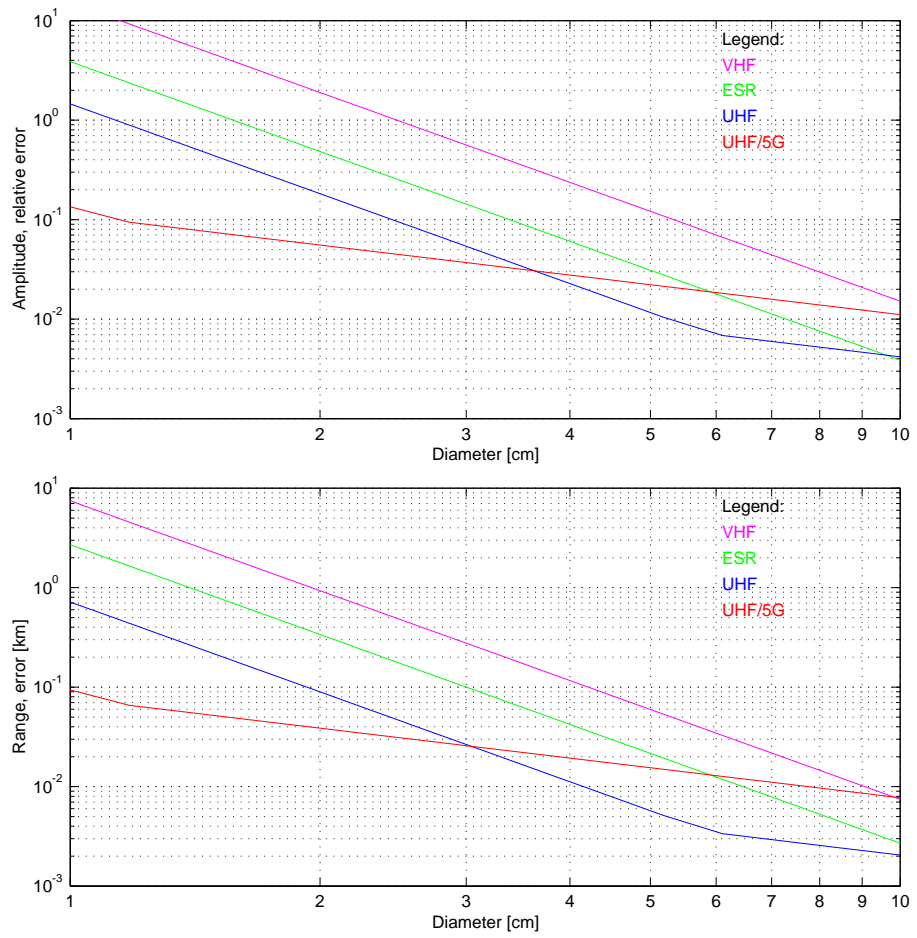
Figure 2.3: Relative error of the amplitude and absolute error of the range as a function of the sphere diameter.
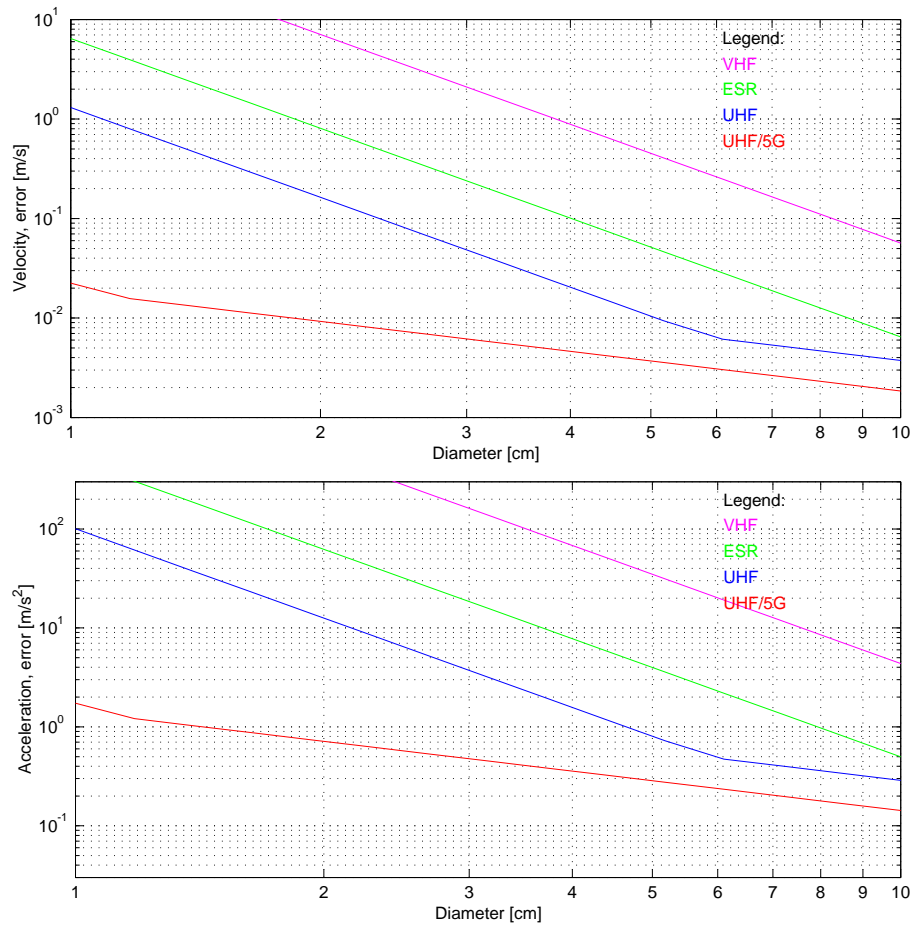
Figure 2.4: Absolute error of the velocity and the acceleration as a function of the sphere diameter.

Table 2.2: The absolute error of amplitude A and the relative error of range R, velocity v, and acceleration a for the diameters given on the left for the EISCAT UHF radar. The values are for 1000 km range and 0.1 s coherent integration.

diameter [cm]	A	R [m]	v [m s^{-1}]	a [m s^{-2}]
1.0	1.46	717	1.31	101
1.2	0.84	415	0.76	58
1.4	0.53	261	0.48	37
1.6	0.36	175	0.32	25
1.8	0.25	123	0.22	17
2.0	0.18	90	0.16	13

weather radar, because its performance is clearly inferior to the others, and the FGAN radar, because it is very close to the EISCAT UHF in its properties.

Most of the experiments with the EISCAT radars use phase-coded pulses, in which the transmission phase is flipped at selected intervals. The error estimates show that the accuracy of the amplitude, velocity, and acceleration estimates is not affected by the phase coding.The range determination, on the other hand, benefits substantially if the phase is flipped, or gaps are added to the transmission, because the edges are what give the accuracy for the range determination. Hence all results will be given for phase-coded pulses only.

Warning

Table 2.2 gives the errors for some selected diameters. We give the error limits down to very small diameters. In view of our preliminary results in Section 4.2.4 concerning the validity of the linearised model, this might not be entirely justified.

2.3.2   Results as a function of velocity and acceleration

The value of the radial velocity has a small effect on the error. Within the range 0–2000 m s^{-1}, the error of any of the model parameters varies by less than a factor of 0.05. Thus we may neglect the velocity in the error model studies, and present the results for a single velocity value, which is 1000 m s^{-1}.

The values of the acceleration within the range 0–100 m s^{-2} has a small effect on the results, but the effect is even smaller than that of the veloc-
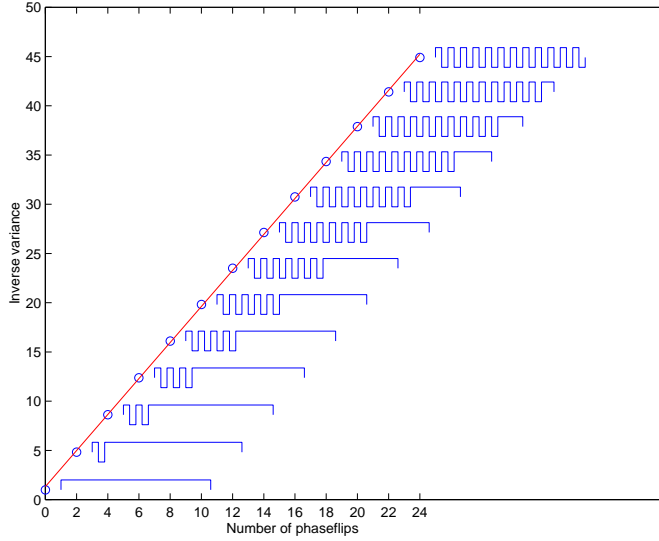
Figure 2.5: Inverse of the variance of the range measurement as a function of the number of phase flips. Corresponding data points are shown with open circles. The fitted line (red) has slope 1.9. The phase codes used in the computations are also shown.

ity. Thus all results shown are valid for any reasonable value of the radial acceleration the debris object may have.

2.3.3 Effect of the phase flips

The results for range in Fig. 2.3 were given for a phase-coded pulse, because the phase coding increases the accuracy of the determination of the range. The effect is illustrated in detail in Fig. 2.5 which shows $1/\sigma_R^2$ as function of the number of phase flips. The variance decreases inversely proportional to the number of phase flips. This result is useful. The variance σ_R^2 decreases as the number of phase flips increases, but the other parameters are not affected by this. The mean number of phase flips in the codes used for calculating the results in Fig. 2.3 was slightly less than 12. As the maximum number of phase flips would be 24 with the assumptions used, the variance could be halved, leading to a $\sqrt{2}$ decrease in the error of the range determination.

2.4 Computational aspects

2.4.1 Discretization of the model parameters

The Bayesian parameter estimation method in Section 2.2.3 requires that we have a continuously parametrized set of model shape vectors $\mathbf{e}^{R,\omega,\alpha}$, Eq. (2.12), with the waveform $\epsilon^{R,\omega,\alpha}(t)$ defined in Eq. (2.8). The problem

here is that it is not clear how we should in practice perform the maximization of the GMF in the continuous case. In this work, we have solved the problem in a straightforward manner, by discretizing the parameter space into uniform, finite, three-dimensional grid $\{(R_j, \omega_k, \alpha_l)\}$ and performing an exhaustive search over the grid.

We introduce also another approximation, replacing in Eq. (2.8) the time-dependent propagation delay $2R(t)/c$ by the constant delay $2R_0/c$. Typically a single integration lasts for 0.1–0.3 s. During that time most SD targets move less than 0.5 km in the radial direction if the antenna is pointed nearly vertical. Due to the constant-range approximation, the model waveform can become 3–4 μs stretched or contracted with respect to the actually received waveform. The "mismatch" in GMF due to this approximation is small, and on the positive side, we gain the following three benefits, one of which is essential.

$\triangleright$ The norm $\|\mathbf{e}^{R,\omega,\alpha}\|$ becomes independent of the model parameters R, ω, α. It only depends on the transmission shape ϵ^{T} and the integration time, and can be taken out of the GMF when computing the GMF maximum. This is basically a convenience only.

$\triangleright$ We can use the sampled transmission vector $\mathbf{e}^0$ directly as the first factor in the signal model, Eq. (2.12), see below. This is a rather necessary feature if we cannot assume anything from the transmission shape *a priori*, or do not want to sample the transmission faster than the signal itself.

$\triangleright$ Because the only velocity dependence in Eq. (2.12) now is in the exponential $\exp(i\omega t)$, it becomes possible to evaluate the GMF with FFT. This is essential to keep the computations manageable.

We acquire the echo shape directly by sampling the transmission $p*\epsilon^{\mathrm{T}}(t)$, with the same sampling interval τ that we use for the signal sampling. The result is the basic signal shape vector $\mathbf{e}^0$, with elements

$$\epsilon_n^0 = p * \epsilon^{\mathrm{T}}(\delta\tau + n*\tau), \; \mathrm{n} = 0,\ldots \mathrm{N}-1, \qquad (2.43)$$

where $\delta\tau$ is an unknown offset, $\delta\tau < \tau/2$. We discretize the range with resolution δR corresponding to the sampling interval,

$$R_j = j \cdot \delta R, \; j = 0, 1, \ldots, \qquad (2.44)$$

$$\delta R = \frac{\tau c}{2}. \qquad (2.45)$$

Typically the sampling interval is 0.5 μs, corresponding to $\delta R = 75$ m. Reception from range R_j, $\mathbf{e}^{R_j}$, is modelled by the transmission $\mathbf{e}^0$, shifted j steps to the right,

$$\mathbf{e}^{R_j} = (0\ldots 0 \; \epsilon_0^0 \ldots \epsilon_{N-1}^0 \; 0\ldots 0). \qquad (2.46)$$

Note that even if the radar target would happen to be exactly at some of the ranges R_j, the model $\mathbf{e}^{R_j}$ could still be a slightly wrong approximation of the actual reception, due to the sampling offset $\delta\tau$.

The discretization in frequency is determined by the discrete Fourier transform that is used in evaluating the GMF,

$$\omega_k = k \cdot \delta\omega, \ \ \mathrm{k} = 0, \pm 1, \dots \tag{2.47}$$

$$\delta\omega = \frac{2\pi}{T_c} = \frac{2\pi}{M\tau}, \tag{2.48}$$

where M is the number of measured points (the length of the vector $\mathbf{m}$) and the typical integration time $T_c \approx 0.3 \ s$, so $\delta\omega \approx 2\pi \cdot 3$ Hz. The corresponding velocity step δv_{D},

$$\delta v_{\mathrm{D}} = \frac{\lambda}{2} \cdot \frac{1}{T_c}, \tag{2.49}$$

is 0.5 m s^{-1} for the EISCAT UHF radar with radar wavelength $\lambda \approx 30$ cm.

Finally, the discretization of the acceleration parameter can be set at will and need not even be uniform. If the acceleration is not accounted for correctly, the phase of the model signal will drift away from the phase of the actual signal during the integration. This leads to mismatch which reduces the value of $\mathcal{R}_{\mathrm{max}}$. The same thing happens if the velocity is matched incorrectly. Due to finite velocity grid spacing, this can always happen to some degree. It follows from Eq. (2.49) that the maximum velocity quantization error, equal to half the velocity grid spacing δv, causes maximum phase mismatch

$$\Delta\Phi_v = 2\pi \cdot \frac{(\delta v/2)\, T_c}{\lambda} = \pi/2 \tag{2.50}$$

during the integration. It appears sensible to select the acceleration grid spacing δa so that the phase error

$$\Delta\Phi_a = 2\pi \cdot \frac{(1/2)(\delta a/2)\, T_c^2}{\lambda} \tag{2.51}$$

due to the maximal acceleration quantization error $\delta a/2$ also is $\pi/2$. This gives acceleration grid spacing

$$\delta a = \lambda/T_c^2, \tag{2.52}$$

which is 3.5 m s^{-2} for 0.3 s integration.

An example of a two-dimensional GMF (we ignore acceleration) is shown in Fig. 2.6, plotted on the (v, R)-plain. The data are from our November 1999 test campaign at ESR, where the active EISCAT experiment was the so-called tau0 experiment. The data have high SNR, thus Fig. 2.6 serves to illustrate the form of the GMF in the absence of noise. Two features are conspicuous. First, the GMF is well localized both in range and in velocity

around the maximum point. Second, the GMF oscillates strongly both in range and in velocity direction.

In the range direction, the GMF is strongly dependent on the transmitted codes. In particular, the short-period (600 m) oscillation apparent in Fig. 2.6 is due to interference between the two frequency channels of the tau0 experiment, and would be absent in a single-frequency case. See Appendix A for details. We believe that the interference phenomenon is a manifestation of the fact that the presence of multiple frequencies actually brings in more detailed information about the combination of the range and velocity parameters. This additional information is seen as the fine structure in the GMF and, therefore, also in the posteriori distribution. There are radar applications where two radar frequencies are used to find narrow layers or small objects by applying interferometry between the frequencies. The fine structure seen in the GMF suggests that such interferometry is a built-in feature in the GMF method of analysis.

The $\text{GMF}(v = v_{\text{D}}, r)$, as a function of range, is essentially the autocorrelation of the code envelope ϵ^{T}. Thus, it has a finite length in the range direction. Moreover, the main part of the autocorrelation has width of the order of the constant phase element of the phase code, 9 km in the case of tau0 experiment.

In the velocity (frequency) direction, the GMF behaviour is not sensitive to the code details, instead the shape corresponds to the various time scales inherent in the measurement, such as pulse length, interpulse period, and integration time. For example, the main mass of the GMF is concentrated on a frequency interval of width $2/L$, where L is the length of transmission in an IPP, about 2 ms in this case, corresponding to 0.3 km s^{-1} on the velocity axis.

We now summarize the parameter estimation scheme in the discretized case. We denote by $G_{j,k,l}$ the generalized match function on the parameter grid point $(R_j, \omega_k, \alpha_l)$,

$$G_{j,k,l} = G(R_j, \omega_k, \alpha_l) = \frac{\left| <\mathbf{m}, \mathbf{e}^{R_j,\omega_k,\alpha_l} > \right|}{\sqrt{W_{\text{T}}}}, \qquad (2.53)$$

where W_{T} is the energy of the sampled transmission $\mathbf{e}^0$,

$$W_T = \|\mathbf{e}^0\|^2 = \|\mathbf{e}^{R_j,\omega_k,\alpha_l}\|^2 \quad \text{for all j, k, l.} \qquad (2.54)$$

We locate the $G_{j,k,l}$ maximum position $(\widehat{j},\widehat{k},\widehat{l})$ using the ratio $\mathcal{R}$ as an intermediate.

Step 1. For each range gate j, we find and store the maximum value of $G_{j,k,l}$ with respect to k and l, together with the location (K_j, L_j) of the maximum,
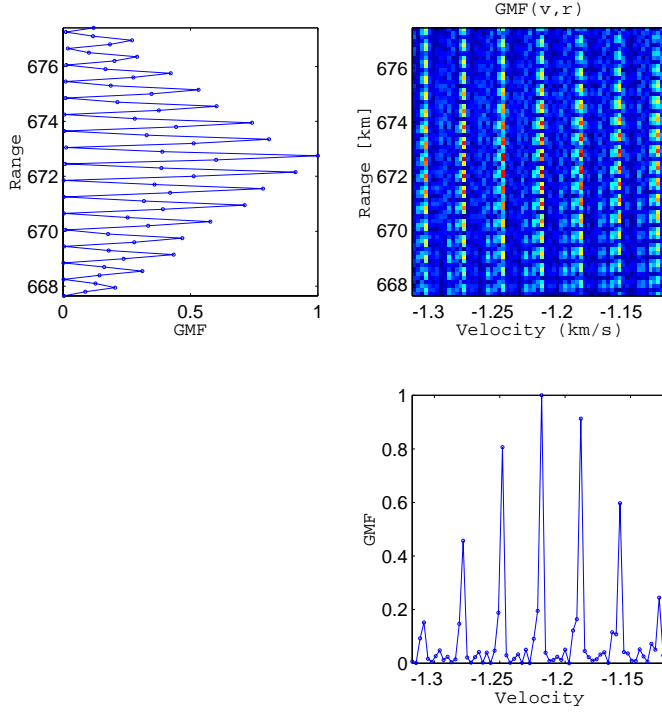
Figure 2.6: Measured generalized match function. The data are from our November 1999 campaign at ESR and have the standard experiment tau0 as transmission. The figure shows the region around the maximum position of the GMF. The surface $\mathrm{GMF}(v, r)$ is shown in the top-right panel as a colour-coded image, while the other two panels show cuts of the surface through the maximum position.

$$\mathcal{R}_j = \max_{k,l} G_{j,k,l}/\sigma \tag{2.55}$$

$$(K_j, L_j) = \arg\max_{k,l} G_{j,k,l}. \tag{2.56}$$

Step 2. We find the maximum $\mathcal{R}_{\mathrm{max}}$ of $\mathcal{R}_j$. Its location gives $(\widehat{j},\widehat{k},\widehat{l})$

$$\widehat{j} = \arg\max_j \mathcal{R}_j, \tag{2.57}$$

$$(\widehat{k},\widehat{l}) = (K_{\widehat{j}}, L_{\widehat{j}}). \tag{2.58}$$

Step 3. We read the shape parameter estimates from the location of the maximizing point, and the SNR estimate from the GMF value at that point,

$$(\widehat{R},\widehat{\omega},\widehat{\alpha}) = (R_{\widehat{j}}, \omega_{\widehat{k}}, \alpha_{\widehat{l}}), \tag{2.59}$$

$$\widehat{\mathrm{SNR}} = \frac{1}{N}(\mathcal{R}_{\mathrm{max}})^2. \tag{2.60}$$

This straightforward discrete-parameter implementation is an approximation to the full continuous-parameter inversion solution which we formulated in Section 2.2.3. The estimates (2.59) and (2.60) retain the inherent errors described by the posteriori density $D_{\mathrm{p}}(A, R, \omega, \alpha|\mathbf{m})$. But, in addition, they incorporate quantization errors because only a discrete set of parameter values are considered, and they also depend on the related incomplete knowledge of the transmission shape. Under sufficiently good SNR, the quantization errors will dominate over the errors due to noise [though the required SNR might be *very* high, see Eq. (4.4) below].

The actual computation of the $G_{j,k,l}$ makes use of FFT. We can expand the definition (2.53) as

$$G_{j,k,l} = \left| \sum_{n=0}^{M-1} u_n^{(j,l)} \exp(-2\pi i \frac{kn}{M}) \right|, \qquad (2.61)$$

where

$$u_n^{(j,l)} = m_n \overline{\epsilon_{n-j}^0} \exp\left(-i\,\alpha_l \tau^2 n^2\right)/\sqrt{W_T}. \qquad (2.62)$$

This shows that for each range gate j and acceleration parameter l, $G_{j,k,l}$ can be computed for all Doppler-shifts k by a single M-point FFT.

However, we note immediately that the computation task involved in detection, using full resolution and without any further approximation, becomes overwhelmingly large. Assume we want to cover 1000 km in range and use 0.3 s integration. Assume that the sampling interval is 0.5 μs. Then $M = 600\,000$, and the FFT requires about 60 Mflops. The $1000/0.075 \approx 13\,000$ range gates require about 800 Gflops. If we also want to cover, say, a ± 20 m s^{-2} zone of acceleration values around a range-dependent first-guess value, with grid spacing 3.5 m s^{-2}, we need about 8 Tflops. For a workstation one can expect about 1 Gflop/s computing speed, so we would need about 8000 s to handle the 0.3 s of data, or three years to handle one hour. Fortunately, it is possible to make several further approximations, so in reality the prospects for real-time detection by affordable means are much better.

2.4.2   Fast GMF algorithm

The radar duty cycle in EISCAT experiments is about 10% in the UHF and about 20% at ESR. Therefore, in the GMF computation the FFT is done on a vector that has 80–90% of its length zeros, in regularly placed blocks. Also, the 10–20 kHz bandwidth needed to cover the SD signal is narrow compared to our sampling frequency of a few MHz. This means that not much information will be lost if we re-sample our data with reduced rate. When it is possible to accept some loss of accuracy, these considerations make it possible to reduce the number of arithmetic operations drastically,
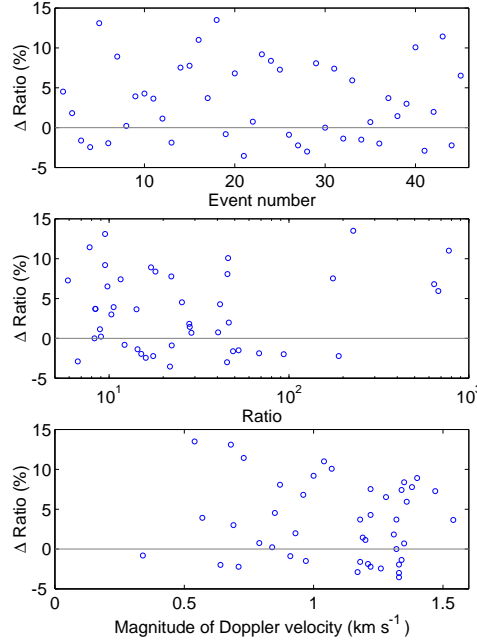
Figure 2.7: Effect of GMF algorithm on detection sensitivity. The panels show difference in $\mathcal{R}_{\mathrm{max}}$ when estimated with the full GMF and the FastGMF algorithms, respectively.

typically by a factor of about one hundred. We call the resulting method for fast but approximative GMF evaluation the FastGMF algorithm.

2.4.3 Detection in real time

The computation speed required for real-time detection depends

  ▷ on the used GMF algorithm,
  ▷ on the length of integration, which together with the sampling interval determines the input data size M,
  ▷ on the number of range gates to handle,
  ▷ on the number of acceleration points to use,
  ▷ on how soon after a scan we make the next scan.

We analysed our February 2001 test campaign both with the straight-forward "full" GMF algorithm and the hundred times faster approximative FastGMF algorithm, and found the difference in detection sensitivity to be surprisingly small. For example, Fig. 2.7 shows the difference $\Delta\mathcal{R}_{\mathrm{max}}$,

$$\Delta\mathcal{R}_{\mathrm{max}} = \frac{\mathcal{R}_{\mathrm{max}}(\mathrm{GMF}) - \mathcal{R}_{\mathrm{max}}(\mathrm{FastGMF})}{\mathcal{R}_{\mathrm{max}}(\mathrm{GMF})},$$

for all analysed events, plotted as function of various parameters: event number, signal strength, and also target Doppler-velocity. The $\mathcal{R}_{\mathrm{max}}$ difference

is only of the order of 10%, mostly in favour of the full GMF as expected,
but sometimes the FastGMF has yielded the higher value. This means that
other factors than just the algorithm are affecting the value of $\mathcal{R}_{\max}$. What-
ever these factors are, they do not appear to be much related to the signal
strength, because there is no obvious trend in the second panel of Fig. 2.7.
This assures us that there are not any strong selection effects in the play,
even though all the data were originally detected using the FastGMF algo-
rithm. We therefore conclude that it is safe to use the FastGMF algorithm,
for target detection at least.

It seems that it does not help much to increase integration beyond 0.3 s.
This is illustrated in Fig. 2.8 and Fig. 2.9. It is not clear why we do not
gain—and can even lose—from increased integration. The signal model,
rather than our software, must be failing some way or another, because we
do get the expected $\sqrt{N}$ behaviour [Eqs. (2.39) and (1.4)] when we feed
simulated data to our analysis programs, see Fig. 2.10. Factors contributing
to the model failure may include the limited target coherence time itself,
ionospheric effects on the signal coherence, the varying signal strength when
the target moves through the radar beam, and also the fact that the current
implementation of the GMF-method does not account for the motion of the
echo envelope during the integration. In the estimate of the computing load
below, we will use 0.3 s as the length of the integration.

The required range gate spacing depends on the transmitted code. If
there is only a single frequency in the transmission, GMF in the range
direction is a smoothly varying function, the code autocorrelation function,
centred near the actual target range. With multiple frequencies—which
almost always is the case in EISCAT—on top of this smooth variation, the
GMF exhibits a much faster oscillation in the range direction, as shown
in Fig. 2.6. At first sight, one might therefore conclude that instead of
the baud length, the scale determining the range grid spacing is the much
shorter period of the interference pattern. This would suggest the use of the
maximum range resolution available.

It turns out that this is not necessary, the relevant scale even in the
multi-frequency case is the code baud length. The positions of the peaks
and zeros in the interference pattern depend also on the frequency ω. Ba-
sically, the pattern in range direction shifts slowly to larger ranges when
ω is increased. The detection function at point r, $\mathcal{R}(r)$, is formed from
$\text{GMF}(r,\omega,\alpha)$ by picking the maximum over ω and α. This causes the min-
ima of the interference pattern in the GMF to be partially filled up in $\mathcal{R}(r)$.
We conclude that it suffices to use range grid spacing that is some sizeable
fraction of the code baud length, a third, say. In fact, when analysing cp1lt
data we have used range grid spacing 9 μs, which is 43% of the baud length
21 μs. This might appear a rather sparse grid, but the test campaign showed
that in practise the loss of detection sensitivity is tolerable. For example,
Fig. 2.11 shows that in 90% of cp1lt events, we lose less than 27% of $\mathcal{R}_{\max}$
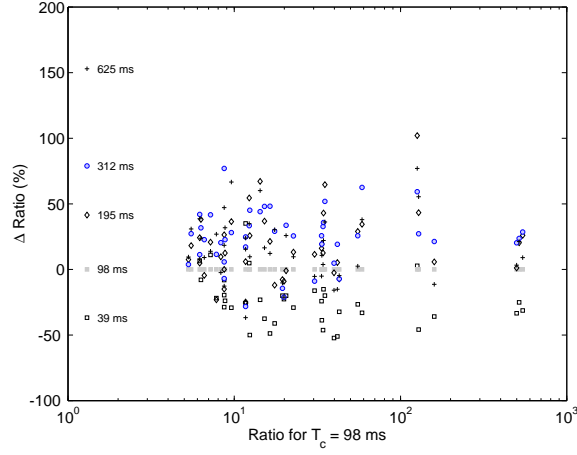
Figure 2.8: Dependence of $\mathcal{R}_{max}$ on integration time. The cp1lt data from the test campaign was analysed with five different integration times T_c. The ratio $\mathcal{R}_{max}^T$ is compared with the reference ratio, $\mathcal{R}_{max}^{98}$. The markers in the left indicate the expected difference $\Delta\mathcal{R}_{max}$ if the coherent integration were fully successful. X-axis is $\mathcal{R}_{max}^{98}$, Y-axis is $(\mathcal{R}_{max}^T - \mathcal{R}_{max}^{98})/\mathcal{R}_{max}^{98}$.
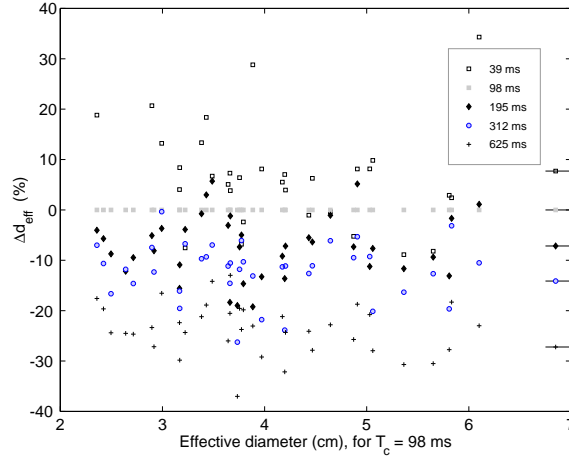
Figure 2.9: Dependence of effective diameter on integration time. Data are cp1lt data from the test campaign. The effective diameter for the indicated integration times is compared with the diameter obtained with 98 ms integration. If the coherent integration would be successful, the effective diameter would not dependent on the integration time. Markers to the right of the data points show the mean change from the reference value, for the different integration times.
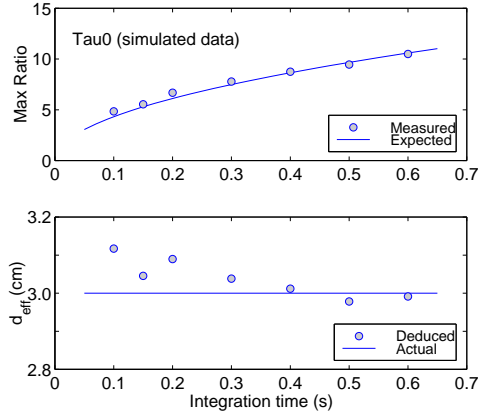
Figure 2.10: Testing coherent integration with simulated data. The simulation is based on ESR tau0 experiment. The simulated 3 cm target is at 1000 km range.
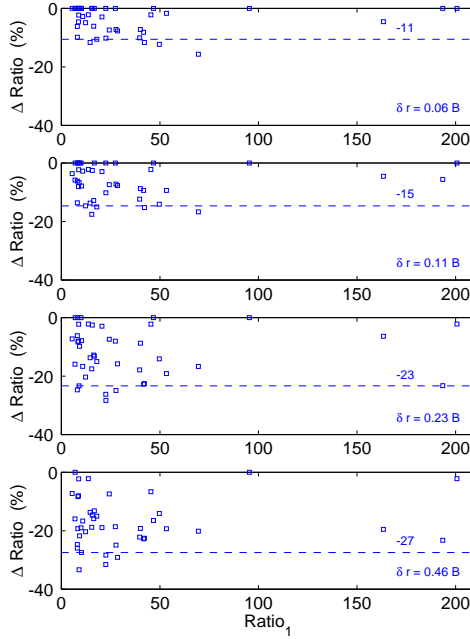
Figure 2.11: Effect of range grid spacing on detection sensitivity. The panels show the reduction of $\mathcal{R}_{\mathrm{max}}$ from the full resolution value, $\mathcal{R}_{\mathrm{max,1}}$, for different grid spacings δr. The horizontal axis is the full-resolution value. The grid spacing is given in terms of the cp1lt phase code baud length B (21 μs). In this data set, the sampling interval was 0.6 μs, so $\delta r = 0.06B$ corresponds to using every other gate. In each panel, 90% of data points are above the horizontal dashed line. The analysis used the FastGMF algorithm. Data are the cp1lt events from the test campaign.

compared to using the maximal range resolution.

The test campaign showed (see Fig. 4.1) that in practice we can cover about 1000 km in range, say from 500 km to 1500 km, with a 100–200 km gap somewhere inside.

Selecting the acceleration grid for detection has turned out to be very simple: no grid is needed. A single, range-dependent value is good enough as long as the antenna is pointed nearly vertically. The value a_O we have used is computed assuming that the target is in a circular orbit,

$$a_O = g_0 \cdot \frac{R_E}{h} \cdot (\frac{R_E}{R_E + h})^2, \qquad (2.63)$$

where R_E is the Earth radius 6360 km, g_0 is acceleration of gravity at zero altitude, 9.8 m s^{-2}, and h is the target altitude. We do not know how circular the orbits actually are, and in view of our criterion (2.52) for the acceleration grid spacing, which gives 3.5 m s^{-2} for 0.3 s integration, more than the single value might appear desirable. However, both simulations and the data from the test campaign show that not much sensitivity is lost in practice even if the acceleration is not varied.

Figure 2.12 shows, with simulated data, the effect of an error in the model acceleration to the amplitude estimate, thus, to the detection sensitivity. We simulated populations of strong (SNR = 100%) targets. The populations had progressively larger spread in the uniformly distributed acceleration near the circular-orbit value, and uniformly distributed velocity in the interval 900–1100 m s^{-1}. We assumed the circular-orbit acceleration in the detection. Two observations may be made from Fig. 2.12. First, there is a rather large spread in the distribution of the amplitude estimates even when the correct acceleration is used in detection. This initial spread is probably due to the velocity quantization error, which can cause the amplitude to be underestimated by about 60% (see Fig. 4.9). Second, even with acceleration errors up to ± 20 m s^{-2}, the typical loss of sensitivity is only about 40%. These results are essentially independent of the used GMF evaluation method.

Figure 2.13 shows for the test campaign data the reduction of $\mathcal{R}_{\max}$ if $\alpha = 0$ is used in the GMF. The mean reduction is 33%. However, only one of the events would have been left undetected even in this case. We conclude that for target detection with vertical radar beam it is acceptable to use only the single, range-dependent value of acceleration.

If we further assume that it is sufficient to sample with 2 MHz rate, and accept that with 0.3 s integration it is possible to tolerate 0.2 s blank time between successive scans, we can determine the number of flops per seconds needed to handle the detection in real time. This we have done in Table 2.3.

Our detection and analysis software has been programmed using the MATLAB high level programming language and programming environment.
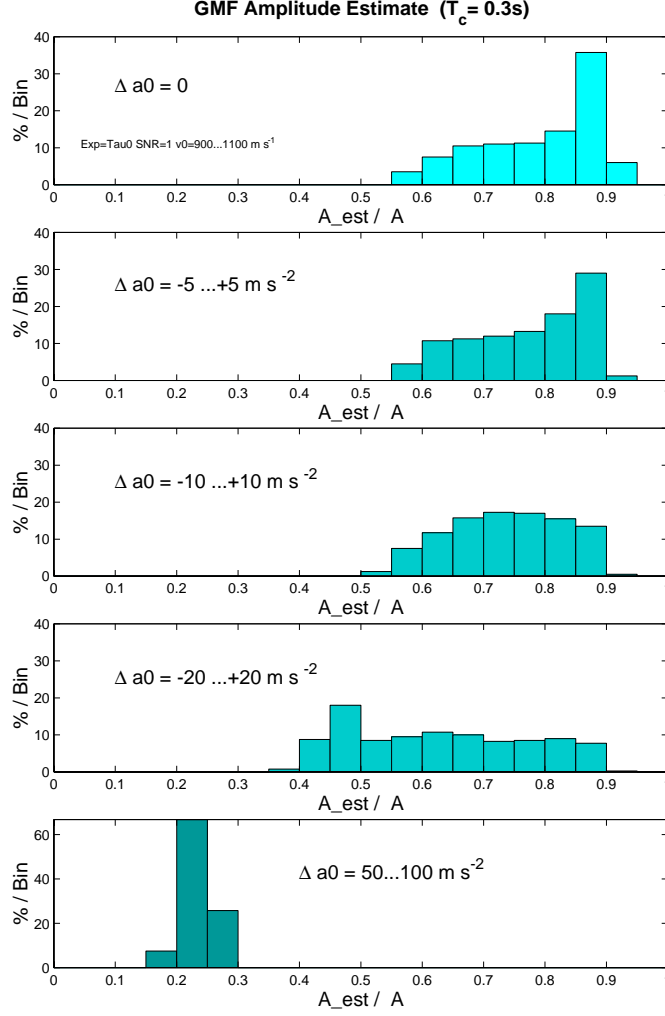
Figure 2.12: Distribution of amplitude estimates, $\widehat{A}/A_0$, for populations of simulated targets. The spread Δa_0 of target accelerations in the population increases from top panel towards the bottom panel as indicated, the velocity spread is unchanged. In the topmost panel, all targets have the same, circular-orbit, acceleration. All targets are at the same range, and have the same strength A_0, large compared to noise. In all panels, only the single, circular-orbit value of the acceleration is used in detection.
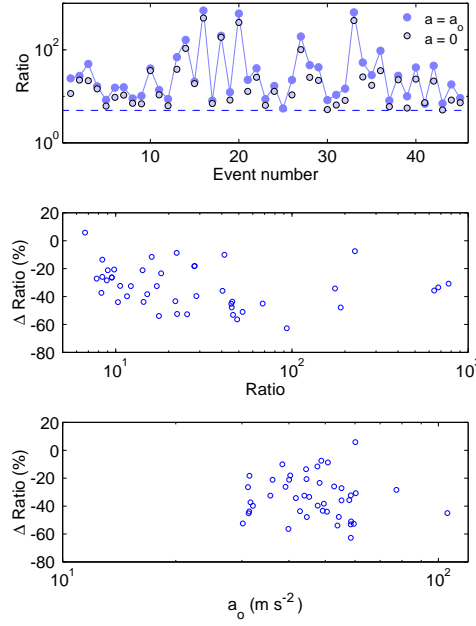
Figure 2.13: Effect of acceleration correction on $\mathcal{R}_{max}$. The cp1lt data was analysed with FastGMF, with 0.3 s integration and full spatial resolution. The top panel shows $\mathcal{R}_{max}$ both without acceleration correction and with the circular-orbit value for the correction. The dashed line is the detection threshold (=5). Only event number 25 was not detected when no acceleration correction was in use. The middle panel shows reduction of $\mathcal{R}_{max}$ as function of $\mathcal{R}_{max}$. Mean drop is 33%. The bottom panel shows reduction of $\mathcal{R}_{max}$ as function of the circular-orbit acceleration value.

MATLAB itself is an interpreted environment, but it allows user-defined, compiled c-language extensions. The GMF evaluation has been implemented as such an extension. We used MATLAB's flops counter to get the number of arithmetic operations needed for a single range gate in the FastGMF algorithm. This number is shown in the second line of Table 2.3 for the three experiments that we have been analysing in this work, namely, cp1lt and tau2 for UHF and tau0 for ESR. For example, the cp1lt experiment requires 0.71 Mflops per gate. We emphasize that this is not a guess but an actually measured number in the real job. Note that we also do acceleration correction here. We only handle one acceleration value per gate, but at least for nearly vertical pointing, more values are not needed. Next we get the number of range gates to compute. For example, in cp1lt we need 740 range gates, to cover 1000 km with 1.35 km (9 μs) resolution. This requires 525 Mflops. To perform two scans per second thus requires about 1.1 Gflops. We get a similar speed requirement for tau2. ESR is easier, basically because of the lower frequency that is used there.

Table 2.3: SD detection computing power requirements in some EISCAT experiments, in the FastGMF algorithm. We assume 0.3 s coherent integration and scan repetition period 0.5 s. We require range coverage 1000 km, and range grid spacing as used in analysis of test campaign data. The Mflops/gate values are taken from MATLAB flops counter during actual execution of the detection algorithm.

	cp1lt	tau2	tau0
Sampling rate (μs)	0.5	0.5	1.0
Mflops in FFT/gate	0.27	0.27	0.27
Mflops in other ops/gate	0.44	0.53	0.56
Total Mflops/gate	0.71	0.80	0.83
Baud length (μs)	21	36	64
Range resolution (μs)	9	10	15
Number of gates/scan	740	666	444
Number of gates/s	1480	1330	890
Req. reading rate (MBytes/s)	2.6	2.6	1.3
Req. computing rate (Gflops/s)	1.1	1.1	0.7

We estimate that four G4-level power-PCs should be able to handle this task in real time, when the algorithms are carefully coded in a compiled language. It is encouraging that in preliminary FFT benchmarking, we have got almost 3 Gflops/s performance on a 500 MHz G4 laptop.

Chapter 3

Measuring System

3.1   Hardware

3.1.1   EISCAT UHF radar

The main blocks of the EISCAT UHF radar are shown in Fig. 3.1. The EISCAT receivers at all three UHF sites, Tromsø, Kiruna, and Sodankylä, are almost identical, the main difference being that at the receiving-only sites Kiruna and Sodankylä there is no need for the duplexer and the receiver protector. Also the polarizer arrangements are somewhat different. The Tromsø UHF receiver has a cooled preamplifier, which takes in the radio-frequency (RF) signal around 929 MHz from the antenna through the duplexer-polarizer system. The RF signal is mixed by the first local oscillator sinusoidal signal and the high-frequency part of the mixing result filtered out. These operations produce the first intermediate-frequency signal (IF1) around 117 MHz. To minimize cable losses, this conversion takes physically place very near the preamplifier, up in the antenna.

The IF1 is then taken down from the antenna to the receiver hall. There it is downconverted to the second intermediate frequency (IF2) around 11 MHz by mixing with second local oscillator signal and again disregarding the higher frequency mixing product. The analog signal path is bandlimited to around 7 MHz width by a filter centred at 11.25 MHz. This filter is the antialiasing filter of the system.

The second IF is digitized by a 14-bit A/D converter, which produces a continuous sample stream at the rate of 15 Msamples/s. 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. Each of the digital channels performs quadrature detection by multiplying the incoming real-valued sample stream by a complex harmonic sequence $e^{-i2\pi f_K t_n}$, where f_K is a frequency near 11 MHz, and filtering out the high-frequency part. The channel hardware also reduces the sampling rate (decimates) from the 15 Msamples/s to a value typically 10–50 kHz. The continuous, decimated streams of baseband samples are then time-gated to the channels' buffer memories, under the control of a special radar controller device, which controls the gating with 100 ns time resolution.

The complex-valued, four-byte-precision sample vectors in the channel buffer memories are the input to the subsequent signal processing. The next step normally is to compute various lagged products of the samples and average them over a few seconds. The processing takes place in a general purpose workstation, which is mounted directly on the receiver VME crate and runs the Sun Solaris flavour of UNIX. The lagged product computation proceeds at about 200 Mflops/s. The data transfer from the channel buffer memories to the VME computer memory, a typical place of a bottleneck in VME-based systems, proceeds at about 33 Mbytes/s in block-DMA mode.

The time-averaged lagged products are transferred over a 100 MHz ded-
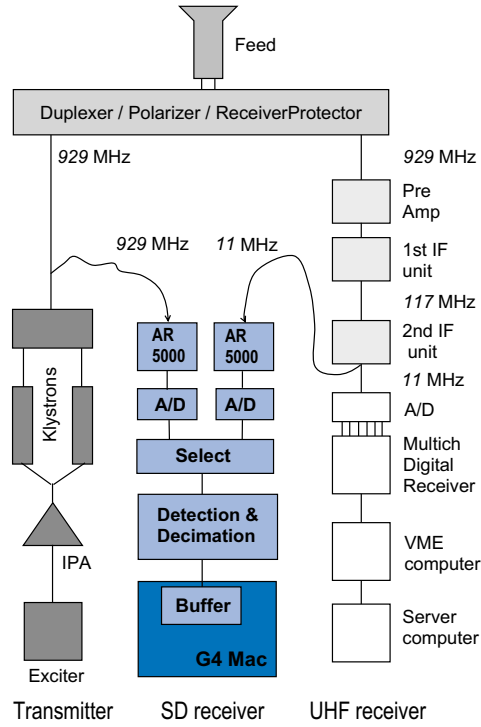
Figure 3.1: EISCAT Tromsø UHF radar and our SD receiver

icated Ethernet link (future update to a Gigabit link is planned) to a 14-processor Sun Enterprise 4500 server, where further processing then takes place—speeds up to about 1.5 Gflops/s have been achieved [12]. The processed data are saved to the server disks (there is about TByte of disk space available), and finally transferred to tape archive.

The EISCAT UHF transmitter consists of a programmable radar controller that generates the pulse patterns at DC level, either uncoded on/off pulses or various classes of binary phase codes; an exciter system that converts the radar controller output to RF around 928 MHz; and a klystron power amplifier that consists of two klystron tubes, delivering combined peak power of about 2.5 MW. The maximum power which can actually be transmitted is limited by the waveguide system to a value just below 2 MW (as of beginning of 2002). The maximum transmitter duty cycle is 12.5%, and duty cycles near this value are also used in most experiments in practice. A set of 14 transmission frequencies is available, spaced by 300 kHz and going from 926.0 MHz upwards.

The 32 m UHF antenna has a fully steerable parabolic dish, has Cassegrain optics, and has rotation rate of about 80°/min both in azimuth and elevation. The four 30 kW antenna motors, two on both axes, are controlled by a rather simple-minded servo system. The servo is quite sufficient

for the occasional change of pointing required by ionospheric work, where the path taken from point to point is not of much interest. The lack of built-in synchronization and explicit speed control between the axes would make difficult to achieve any movement which requires fast controlled motion of both axis. If SD tracking would ever be required, this part of the antenna control system would probably need to be redesigned.

The time and frequency base at all EISCAT sites is from the GPS system.

3.1.2   The SD receiver

To be able to measure in parallel with EISCAT, an additional data collecting system is needed. Coherent integration requires that the raw samples are available, so an independent system from the analog level onwards is needed. Our SD receiver allows us to hook on any analog frequency, and as a matter of principle, we would prefer to take the signal as near the antenna (as unprocessed) as possible. After discussion with EISCAT staff, for the February 2001 test campaign in Tromsø, we decided to use the EISCAT second IF as the analog input. In this section we describe the SD receiver. Figure 3.1 shows the main blocks of the SD receiver and its main connection to the EISCAT UHF system.

EISCAT experiments normally contain more than one frequency. EISCAT itself handles this situation in the traditional way, by having multiple hardware channels, each tuned to a particular frequency[1]. The end result is several sample streams, one for each channel. Our approach in the SD receiver is different. We sample fast enough to always capture the whole relevant IF2 band into a single digital stream. If the spread of frequencies is B MHz, we take slightly more than B million complex samples per second. We call this kind of multi-frequency, complex-valued, digital data the IQ data. The largest frequency spread during the test campaign was 2.1 MHz. We handled it by taking 2.5 million complex samples per second. The individual EISCAT frequency channels have rather narrow bandwidth, a few hundred kHz at most, so the IF2 frequency band has large gaps without useful information, see the bottom panel of Fig. 3.2. It would make sense to compress the data by cutting off the empty segments by suitable FFT processing, but we have not done this in this study. Preferably the compression should happen on-line.

In addition to sampling the target echo, we also need to sample during the transmission periods, to get the code patterns. At its ESR site, EISCAT provides the transmission sample signal on the same data path as the reception, so no special arrangements are needed there. In Tromsø, this service is planned but was not available during the test campaign, so we had to resort to ad hoc arrangements. We sampled the transmission with the same

---

[1]Under program control, the tuning can be changed in microsecond time scale.

sampling rate as the reception.

Our data acquisition system has originally been developed for ionospheric tomography by a Finnish company, Invers Ltd. The basic system consists of a sampling section, a detection and decimation section, and a computer interface section which has a large data buffer and a control interface. The system supports three analog input channels. We used two channels, one for transmission sampling, the other for reception sampling. The transmission channel was level-detected, and when found active, was selected as the data source. For all other times, the receiver channel was used as the data source.

The raw sampling rate was 40 MHz. The resulting real-valued sample stream was processed by programmable logic circuitry to perform detection, essentially by doing Hilbert transform as follows. If the 40 MHz stream has samples $x_1, x_2 \ldots$, the first complex sample is $z_1 = (x_1, x_2)$, the next complex sample is $z_2 = (x_5, x_6)$, and so on. The stream of the z_n-samples thus has sampling interval 100 ns. For narrowband analog input around 10 MHz, the z_n-sequence represents the positive-frequency part of the double-sided analog spectrum, the negative-frequency part being zeroed in the sampling.

The complex 10 MHz sample stream is decimated as required and written into the output buffer. The decimation is done by adding an appropriate number of consecutive complex samples, thus ensuring proper filtering.

It may be noted that no separate conversion to baseband is done in this scheme. What one has in the baseband is some Nyquist replica of the positive-frequency part (or the negative-frequency part) of the original, real-valued analog signal. To work, this arrangement requires the band-limited analog input to be centred at a suitable frequency. The necessary frequency translation is done by the AR5000 broadband communication receiver, placed in front of the analog-to-digital converter.

The output buffer and the control circuitry are mounted on a PCI slot of a fast workstation. We are using Power Macintosh G4 workstations, running under the Mac OS X version of UNIX. Software from Invers Ltd is used to read the data from the buffer and write them to hard disk. Typical data accumulation rate is between 20-30 Gbytes/hour, depending on the sampling interval. For final storage, data are copied to 60–70 Gbyte FireWire disks, which is a rather slow process.

Figure 3.2 shows a snapshot of the raw data produced by the SD receiver. In this study, all subsequent signal processing has been done off-line, using raw data on disks.

3.2   Software

Figure 3.3 shows a conceptual block diagram of our SD software, as it would be used in real-time operation. All the main tasks were implemented already for the test campaign, but the software was operated off-line, using data
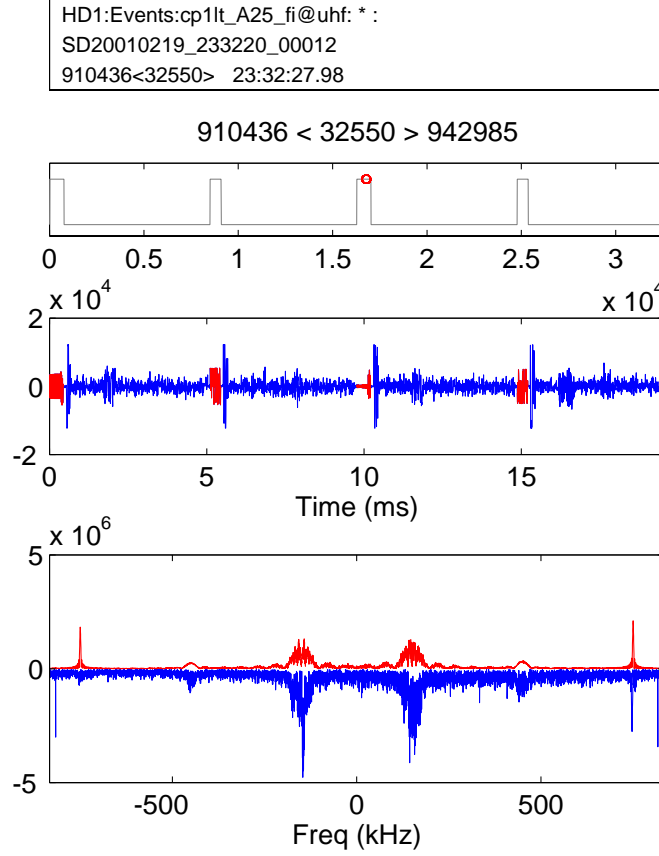
Figure 3.2: Four IPPs, 19.53 ms, 32550 samples of cp1lt raw data from the test campaign. The top panel gives location of the data in the data stream. The indicated time 23:32:27.98 gives approximate UT at the beginning of the displayed data. The x-axis is samples. The second panel shows transmission ON-bit (line) and GPS full-second mark-bit (circle). The full-second mark occurs here at 23:32:27.99 with respect to the uncorrected UT time axis. This shows that the uncorrected time is lagging the GPS time only about 10 ms. The third panel shows the real part of the complex samples, starting from long pulse (LP) transmission, followed by alternating code (AC) transmission. Transmission is in red, reception in blue. The transmissions are LP, AC, LP, AC. The strong echoes immediately after the transmission blocks are clutter. The increased noise at about 7 ms and 17 ms is due to calibration noise injection. In each IPP, there is strong target echo from 960 km range which shows up 6.4 ms from the leading edge of the respective transmission. The bottom panel shows power spectra of transmission (red) and reception (blue, sign inverted). The frequency axis is from -833 kHz to +833 kHz, corresponding to 0.6 μs sampling interval.

recorded on disks. In real-time operation, only a large buffer of raw data would be kept on disk. Whether operating real-time or off-line, conceptually we view the data as a continuous stream.

The software to process the data stream is a chain of four modules: the scanner, the detector, the archiver, and the analyser. For this study, we have implemented the processing chain as two MATLAB$^{\mathrm{TM}}$ programs, SPDSCAN and SPDVIEW. MATLAB is a commercial, high level interpreted language and a programming environment, well suited for graphics and rapid prototyping of signal processing applications. However, a major inconvenience during this study has been that MATLAB is not adequately supported on the Macintosh platform.

3.2.1 Data stream

The digital SD data stream consists of primary data and auxiliary data. The primary data, produced by the SD receiver under the control of the GURSIP$^{\mathrm{TM}}$ software of Invers Ltd, consist of 32-bit words. Of a word, 15+15 bits are used by a complex integer which represents a measured sample, either a transmission sample or a reception sample. The remaining two bits contain two flags. One flag tells whether the data sample is a transmission sample or a reception sample. The other flag is a GPS-based full-second time mark, accurate to a few μs.

The auxiliary data consist of a time stamp and radar state information. The GURSIP software inserts a fully qualified time stamp, based on the recording computer's timekeeping, into a log file at the start of the recording. Starting from the time stamp and counting samples in the continuous sample stream, we get the actual UT time of any data segment with about 10 ms accuracy, mainly limited by the accuracy of the workstation clock, see Fig. 3.2. For routine SD measurements, antenna pointing direction should be recorded. In EISCAT, pointing data are only available via the radar's process computer. We would need EISCAT to provide a server program, capable of returning the azimuth and elevation when queried. It is useful to have the transmission power included into the data stream also, although the relative changes of the power can be monitored via the transmission samples. Neither the antenna pointing direction nor the transmission power were available on-line during the test campaign. On the other hand, only fixed antenna pointing was used, so we could easily keep track of the pointing direction during data analysis .

3.2.2 Scanner

The most time consuming task of the data processing, GMF computation, is done by the program SPDSCAN. It implements step 1 of the GMF method, Eqs. (2.55) and (2.56), saving to files $\mathcal{R}_j$, K_j and L_j for each scan, together
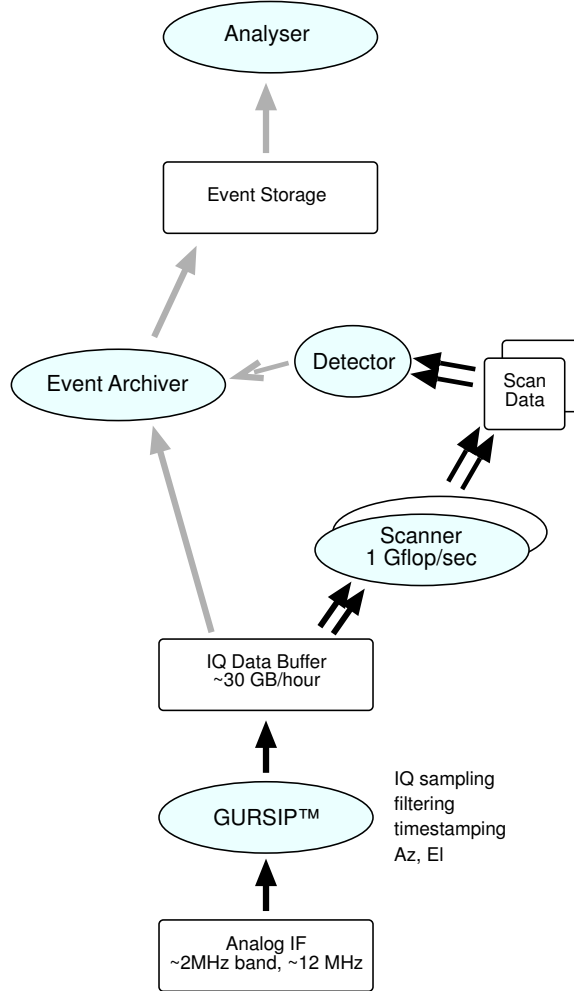
Figure 3.3: Main software modules for SD measurements. The GURSIP software by Invers Ltd controls the low-level signal processing, which takes the analog input and produces complex samples around DC, at the rate of a few Msamples/s. In addition, this software inserts the necessary auxiliary information into the data stream. The scanner computes the GMF. In a real-time system, the scanner is the bottleneck and might need to run on several networked workstations in parallel. The detector uses the scanner data for the detection decision and notifies the event archiver to save or release the corresponding data in the IQ-data buffer. The analyser performs high resolution re-scans and makes fits to time series to get event parameters.
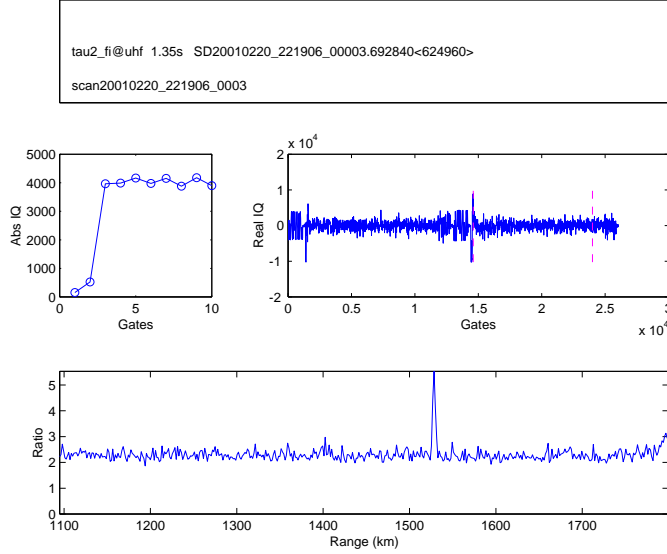
Figure 3.4: Standard screen display of the SD scanner program, SPDSCAN, showing a target hit. The top panel shows the location of the scan in the input data stream. The middle panels show the first two IPPs, about 13 ms, of the raw data in a 0.3 s integration. The bottom panel is the ratio $\mathcal{R}$. The scan UT time is 22:19:08, 20 February 2001.

with a set of supplementary information. For real-time detection, it might become necessary to run the scanner on several networked computers. The scanning task is straightforward to parallelise on scan-by-scan basis, so we do not expect the parallelisation to present a big problem in practice. Parallelisation on this level, including the distribution of the tasks and collecting the results, is a standard feature of the GURSIP software. Figure 3.4 illustrates the appearance of SPDSCAN on computer screen. The time instant shown is near the beginning of the strong event which we use for catalogue comparison in Section 4.4.

The files produced by SPDSCAN are used as the input by the detection, archiving, and analysis program, SPDVIEW.

3.2.3 Target detection and initial parameter estimates

Before actually attempting detection, it is possible to use SPDVIEW to inspect interactively the ratio curves produced by SPDSCAN. It is also possible to "clean" the curves, for instance by removing systematic distortions. For example, the ESR radar used to have a pronounced variation of receiver gain during the reception window, and the UHF radar cp1l experiment has calibration noise injection in the middle of the reception window. These trends and distortions are averaged and removed, or at least reduced, using data

that do not contain manifestly large signal contribution, see Section B.3 for more details. The result is a set of fairly flat versions of the ratio $\mathcal{R}$ profiles. Detection is done on the cleaned data, by finding $\mathcal{R}_{\max}$—step 2 in the GMF method—and comparing it against a fixed threshold. By inspection of the test campaign data, we somewhat arbitrarily adopted $\mathcal{R}_{\max} > 5$ as the detection criterion, see Section 4.2.2.

When the detection threshold is exceeded, SPDVIEW performs step 3 of the GMF method [Eqs. (2.59) and (2.60)] to get initial estimates for target range, Doppler-velocity, and the signal-to-noise ratio. If also the acceleration has actually been deduced by the scanner, by varying the acceleration parameter in the GMF, also the acceleration is returned. However, in almost all cases, we have fixed the acceleration to the circular-orbit value $a_{\mathrm{O}}(r)$ [Eq. (2.63)].

The initial parameter estimates obtained as side product of detection are deduced directly from the profiles collected by the scanner, with the particular range, time and spectral resolution used by the scanner. These preliminary values are used to group the individual scans into events, corresponding to a single target when it moves though the radar beam. SPDVIEW provides automatically a preliminary grouping, based on finding a detectable spike in the $\mathcal{R}$ profiles in nearly the same range in nearby scans. The grouping can be improved by interactive inspection.

3.2.4   Event archiver

After an event has been identified, SPDVIEW is used as a rudimentary archiver to save the raw data and the auxiliary information to an event-specific directory. The amount of raw data to be archived is large but not unmanageably large. In the test campaign, about 15 events per hour were found. To store 10 seconds of raw data around each event would amount to storing 150 seconds/hour. Assuming 2 Msamples/s sampling rate and four-byte complex samples, 150 seconds/hour corresponds to about 900 measuring hours per terabyte of available storage, if all ranges are kept. As the largest available tape units and FireWire disks presently take about 100 Gbytes of data, storing one year's measurements as uncompressed raw echoes would require about 10 or a dozen such tapes or disks. This seems a quite reasonable cost. A planned FFT-based data compression algorithm, for storing only the necessary limited frequency bands, could result in a further compression by a factor of about 5 in the average. Storing only limited range intervals in an experiment which has 12.5% radar duty cycle, would reduce the storage needs by a further factor of 8. If both compression modes are used, the required storage space for one year's measurements goes down to about 25–30 Gbytes.

3.2.5 Analysis

We get initial estimates of target range and velocity as side product of the detection, but the final determination of the target parameters is done separately with SPDVIEW, using raw data from the event archive. The analysis consists of two phases. First the event's raw data are re-scanned with full time and range resolution, and optionally using the full GMF algorithm. This becomes possible because after detection, we can restrict to a few kilometers the range where the GMF is evaluated. In the second phase, we plot the parameters from individual scans as function of time, and, when possible, fit low-order polynomials to the range and velocity data. Figure 3.5 is an example of such an event summary plot.

Panels (2)–(4) of the event summary plots show the maximum ratio $\mathcal{R}_{\mathrm{max}}$, range R, and Doppler-velocity v_{D} (positive away from the radar). The header panel gives the event time t_0, which we define as the time of maximum signal strength. The event time is used as the origin of the time axes in the data panels. The header also gives $R(t_0)$ and $v_{\mathrm{D}}(t_0)$.

If there are enough points where $\mathcal{R}_{\mathrm{max}}$ exceeds the chosen threshold, also range rate $\mathrm{RR} = \mathrm{d}R/\mathrm{d}t$ and radial acceleration $a_{\mathrm{D}} = \mathrm{d}v_{\mathrm{D}}/\mathrm{d}t$ are estimated for time t_0, by fitting either a straight line or a parabola to the R and v_{D} points, respectively. We normally use smaller threshold value in the analysis than in the detection. The range rate RR is redundant, and in general less accurate than the directly measured Doppler-velocity. When the fits are available, the header panel gives also $\mathrm{RR}(t_0)$ and $a_{\mathrm{D}}(t_0)$. In this case, also the quoted range $R(t_0)$ and velocity $v_{\mathrm{D}}(t_0)$ are taken from the fits. If there are too few good points to allow fitting, the quoted range and velocity are read directly from the R and v_{D} data.

The parameter line in the event summary plot header gives also the effective target diameter d. It is the diameter of a conducting sphere that at the observed range would give the observed SNR, if the target would be in the centre of the radar beam at the moment of observation. Normally, the effective diameter underestimates the true size; for very weak events, the effective diameter can also be an overestimate, see Fig. 2.8. For several events in the test campaign, antenna side lobes were visible in the $\mathcal{R}_{\mathrm{max}}(t)$ plot. It might be possible to get the actual off-centre angular distance in these cases, but except for a single occasion (see Section 4.4), we have not attempted to determine it.
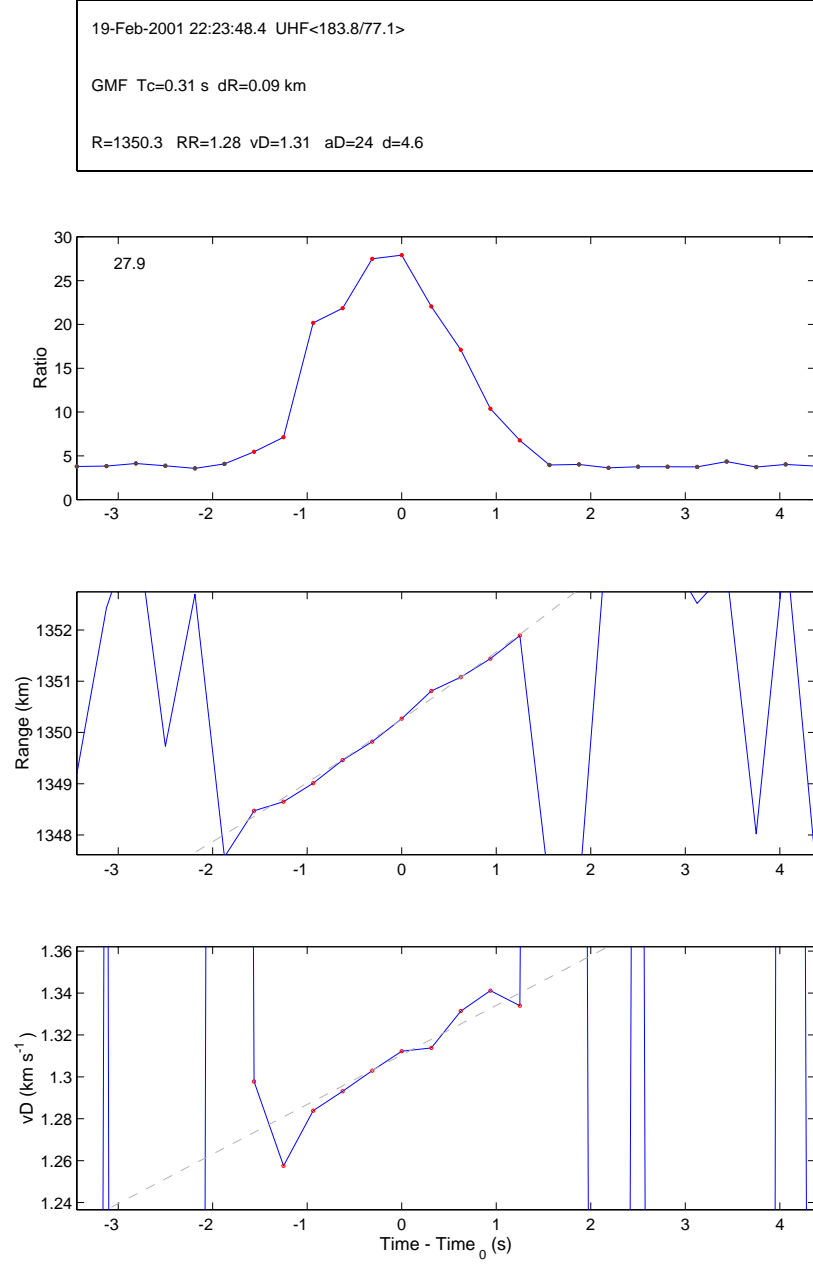
Figure 3.5: Event summary plot for an event in experiment cp1lt, 19 February 2001. The event was analysed with the standard GMF algorithm. The header panel shows the deduced event parameters, which include time of maximum signal strength, and, for that instant of time, target range, range-rate as fitted from range data, Doppler-velocity, and minimum diameter.

Chapter 4

Test Measurements

4.1   Test campaign

Our ESA contract stipulates a series of test measurements to be performed and analysed. We did the measurements at EISCAT UHF radar during a Finnish experiment campaign which lasted from 11 February to 23 February 2001, using our SD receiver as shown in Fig. 3.1.

Two standard EISCAT experiments were in use during the campaign, the cp1lt experiment (Appendix B) and the tau2 experiment (Appendix C). We took 16.5 hours of cp1lt data, with sampling rates from 2.5 MHz to 1.1 MHz, and 1.7 hours of tau2 data, with 2 MHz sampling. This gave about 0.5 Tbytes of data, which were stored on 13 FireWire disks. We have analysed 2.8 hours of cp1lt and 1.2 hours of tau2. Both our group and EISCAT had various difficulties especially in the beginning of the campaign, and so the analysed data is from late in the campaign.

4.2   Analysis results

4.2.1   Detection rate and altitude coverage

There were 45 (cp1lt)+11 (tau2) clear hard target events in the 4 hours of data. "Clear" here means that the event had more than one scan exceeding the detection threshold, and that the deduced target parameters were sensibly located in (R, v_D)-space from scan to scan. The mean event rate was 13 events per hour. The combined SD search region in the two experiments was from 400 km to 1750 km in altitude. Figure 4.1 shows the altitude distribution of the events and also indicates the blind zone of each experiment. The blind zone is an altitude interval from which no target echoes can be received, due to an ongoing transmission. The blind zone in tau2 is in the middle of the highest event density in cp1lt. This might partially explain the somewhat lower event rate in tau2, nine events per hour, compared to 16 events per hour in cp1lt.

4.2.2   Detection sensitivity and effective diameters

After experimentation with various values, $\mathcal{R}_{\max} = 5$ was used throughout as the detection threshold. The weakest event of each experiment is shown in Fig. 4.2. The background value of $\mathcal{R}$ is about 3 in both experiments, and it might appear from the figure that threshold 5 is a little too high (see also Fig. 2.2). Indeed, no false alarms in the higher search window (the region on the high altitude side of the blind zone) were triggered in the four hours of data. However, in the low-altitude search windows the situation was more complicated, as the data were a lot more disturbed there. But even in the lower window, after subtracting the stationary distortions, the threshold value 5 appeared to be suitable. On the other hand, it is clear
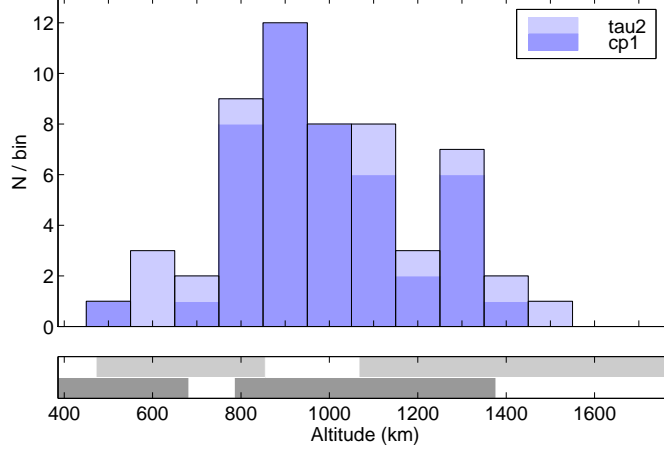
Figure 4.1: Altitude distribution of events in the test campaign, February 2001. The cp1lt data are from 19 February, 22:23–01:11 UT (45 events), the tau2 data are from 20 February, 21:26–22:40 UT (11 events). The top panel shows the altitude distribution of all the 56 events, the bottom panel shows the SD search regions in the two experiments. In tau2, the first blind zone is in 855–1067 km, in cp1lt, in 682–785 km altitude.
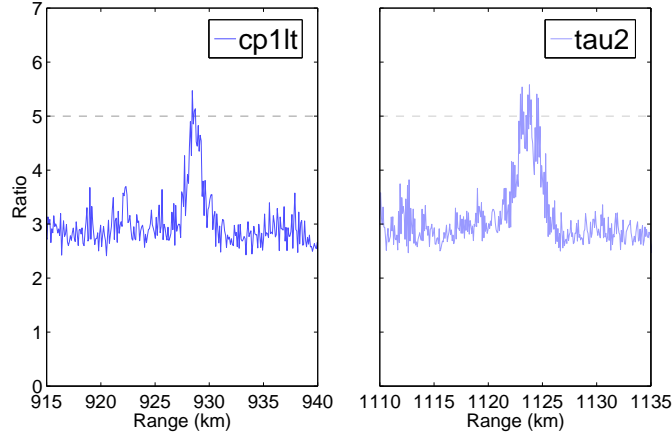
Figure 4.2: Weakest analysed events in the test campaign. The cp1lt event corresponds to effective diameter 2.1 cm, the tau2 event to effective diameter 2.3 cm. The detection threshold 5 is indicated by the dashed lines. Data were scanned with FastGMF, using full spatial resolution.

from Fig. 4.2 that one cannot use much lower threshold than 5. For example, already value 4 gives a false alarm in almost every scan that goes through the full altitude range (observe that Fig. 4.2 shows only about 3% of the full range).

Given the length of integration, duty cycle, system temperature, transmission power, and target range, $\mathcal{R}_{\max}$ can be converted to effective target diameter with the help of the radar equation. Figure 4.3 shows the resulting threshold diameter as a function of range for cp1lt and tau2. For example, at 1000 km range, $\mathcal{R}_{\max} = 5$ corresponds to 2.1 cm diameter in cp1lt and 1.9 cm in tau2. The greater sensitivity of tau2 is due to the higher duty cycle, 8.8% in tau2 contrary to 7.0% in cp1lt. The assumed integration time is 0.31 s for both experiments. The assumed transmission power is 1 MW which was a typical value during the test campaign. Part of the time, the power was lower, but we have not made any corrections due to this. The sensitivity estimate in Fig. 4.3 also assumes that the coherent integration has been successful and that the GMF method estimates the peak signal amplitude correctly. In practice, several factors cause the signal amplitude to be underestimated, up to several tens per cent. This reduces the actual sensitivity. On the other hand, the assumed transmission power in Fig. 4.3 is only half of the 2 MW that the refurbished EISCAT UHF transmitter is ultimately expected to deliver. Therefore, the curves in Fig. 4.3 give a rather good idea about the detection sensitivity that we can realistically expect from the UHF system.

Figure 4.3 shows the small-diameter events from the test campaign. The full effective-size distribution is plotted in the top panel of Fig. 4.5. Because we do not usually know how far off-axis the target passed through the radar beam, its effective diameter gives only a lower bound for the actual size. One probably should not make strong inferences from a data set this small, but in the higher ranges, the events appear to lie rather far from the threshold curve. Range-integrated size distribution, containing all the analysed events, is shown in Fig. 4.4. Also it shows some kind of deficiency of events near the detection threshold. This might indicate that we do not detect all events that are near the detection threshold. Such an effect is to be expected, but a proper statistical study should be done about how large the effect should be.

4.2.3   Velocity estimates

The remaining basic parameters, the Doppler-velocity and the acceleration, are shown in the middle and bottom panels of Fig. 4.5. The velocity data looks sensible. There is even the expected small asymmetry between the positive and negative velocities. There are 22 events with positive v_r, with mean 1.03 m s^{-1}, and 23 events with negative velocity, with mean -1.16 m s^{-1}. Two factors cause this 13% bias towards negative velocities. First, the an-
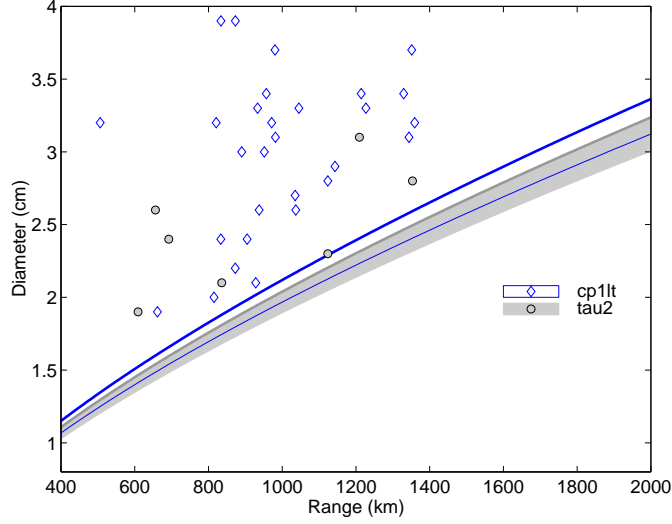
Figure 4.3: EISCAT UHF detection sensitivity limit with 0.3 s integration and 1 MW transmission power. The minimum detectable size is shown as function of range. There are two curves for both experiments of the test campaign. The upper curve corresponds to the actually used detection threshold $\mathcal{R}_{\max} = 5$, the lower curve corresponds to $\mathcal{R}_{\max} = 4$. The markers show measured effective diameters, estimated using standard GMF algorithm.
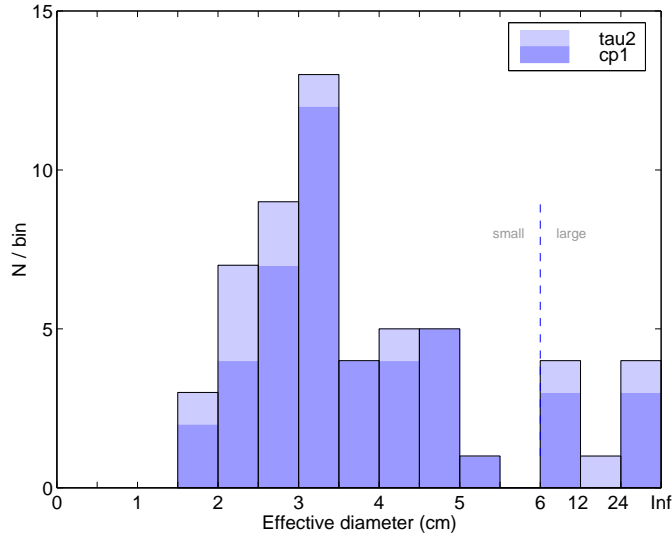
Figure 4.4: Size distribution of events in the test campaign. The dashed vertical line indicates the value $d_{\mathrm{eff}} \approx \lambda/5$. Note that the figure's x-axis is non-linear to the right of this point.
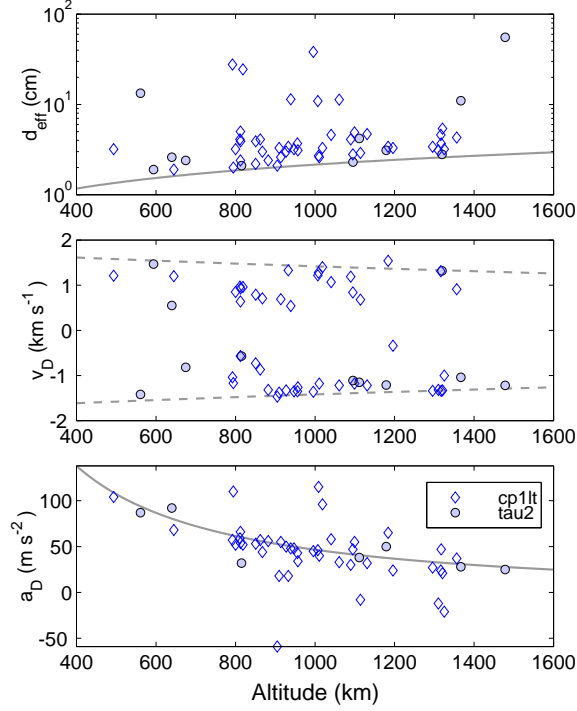
Figure 4.5: Event parameter estimates in the test campaign. The data was analysed with the full GMF, with 0.3 s integration. The top panel shows the effective diameter. The solid line in the top panel indicates the threshold diameter (for cp1lt). The middle panel shows the Doppler-velocity v_D. The dashed lines give the maximum and minimum radial velocity for targets in circular orbits, for the beam elevation 77°. The bottom panel shows the acceleration estimate, computed by a linear fit to $v_D(t)$ during beam passage. The solid line is the vertical-beam, circular-orbit acceleration a_O.

tenna is pointed slightly, 3.8°, to the west (at elevation 77.1°), and second, the maximum of the low-orbit debris inclination distribution is around 65°. Consequently, the SD objects approach mainly from the west, and thus have a small bias of flying towards the station. Flying towards the radar corresponds to negative radial velocity [13, Fig. 2.4-3, p. 2.4-3]. Another curiosity is that the scatter of the positive and negative velocity points is clearly different.

The target radial velocity shown in Figs. 4.5, 4.6, and 4.10 refers to velocity at the time of maximum signal strength. Typically the target could be followed from one to several seconds, and in these cases the given Doppler-velocity is obtained by a linear fit to the v_D points from consecutive 0.3 s integrations. We have not attempted any quantitative error analysis here. Nevertheless, the analysis summary plots, such as Figs. 3.5 and 4.12, reveal
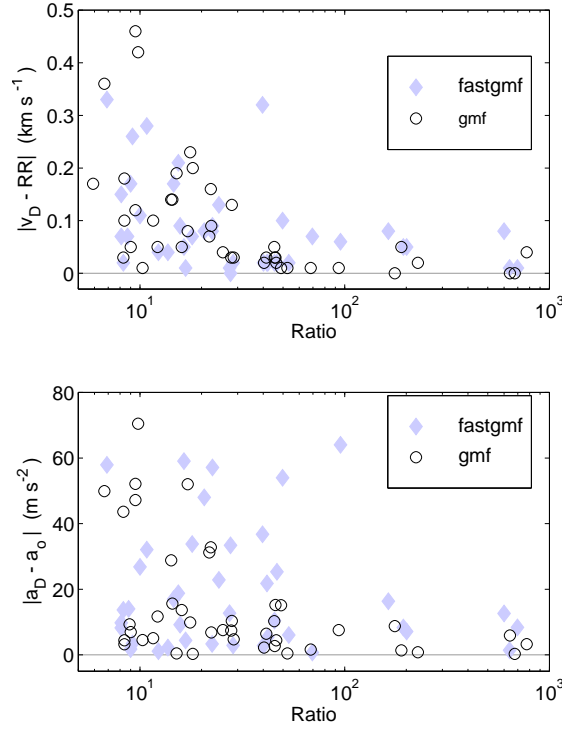
Figure 4.6: Effect of signal strength on the velocity and acceleration estimation. The top panel shows the difference of directly measured Doppler-velocity and the velocity computed by fitting range data. The bottom panel shows the difference of the acceleration from the circular-orbit value. The acceleration was determined via fits to the Doppler-velocity data.

that in the $v_D(t)$ curves there are "semi-random" jumps, up to a few tens of m s^{-1} in GMF analysis, and up to 100 m s^{-1} in FastGMF analysis. The jumps are much larger than the theoretical velocity error estimates, based on linearised signal model (Section 2.3), would suggest. We discuss this problem in Section 4.2.4, below.

The top panel of Fig. 4.6 compares the Doppler-velocity v_D with the velocity RR obtained by linear or quadratic fit to the range data. The weakest targets have the largest difference between the two estimates. The difference can be several hundred m s^{-1}, but is seems probable that the inherent scatter in the RR points is larger than in the v_D points, as hinted by the top and middle panels of Fig. 4.10. Therefore, a large part of the $|v_D - RR|$ difference is probably due to the inaccuracy in the RR data rather than in the Doppler-data.

4.2.4   Velocity accuracy estimates

We reserve a more thorough study of the measurement accuracy for future work; here we offer preliminary remarks about the Doppler-velocity (Doppler-frequency) error. We need to inspect the behaviour of the full posteriori probability density $D_{\mathrm{p}}(A, R, f | \mathbf{m}) \sim \exp(-\|\mathbf{m} - \mathbf{s}\|^2 / \sigma^2)$ in Eq. (2.28). We evaluate D_{p} in the frequency direction, through the D_{p} maximum position (A_0, R_0, f_0). We neglect noise in the measurement $\mathbf{m}$ and assume simple, single-frequency transmission. Expanding the definitions involved gives

$$D_{\mathrm{p}}(f) = D_{\mathrm{p}}(f_0) \cdot e^{-\mathrm{SNR_N} \cdot 2 \left\{ 1 - \cos[\pi(f - f_o)(KP - P + L)] \cdot \frac{\sin \pi(f - f_0)L}{\pi(f - f_0)L} \cdot \frac{\sin \pi(f - f_0)KP}{K \sin \pi(f - f_0)P} \right\}},$$
(4.1)

where $\mathrm{SNR_N}$ is the energy-to-noise ratio (1.3), P is interpulse period, L is length of a single transmission, and K is the number of IPPs in the coherent integration, so that the integration time is $T_{\mathrm{c}} = KP$. For the GMF we similarly find from Eq. (2.29)

$$\frac{\mathrm{GMF}}{\sigma} = \sqrt{\mathrm{SNR_N}} \cdot \left| \frac{\sin \pi(f - f_0)L}{\pi(f - f_0)L} \right| \left| \frac{\sin \pi(f - f_0)KP}{K \sin \pi(f - f_0)P} \right|.$$
(4.2)

We expand the exponent in Eq. (4.1) into a power series and use typical values for K, P, and L to simplify the expression. We get

$$D_{\mathrm{p}}(f) = D_p(f_0) \cdot e^{-\frac{1}{2\sigma_f^2}(f - f_0)^2 \, + \, \mathcal{O}\left[(T_{\mathrm{c}}(f - f_0))^4\right]},$$
(4.3)

where

$$\sigma_f^2 = \frac{3}{8\pi^2} \frac{1}{T_{\mathrm{c}}^2 \, \mathrm{SNR_N}}.$$
(4.4)

Thus we get a Gaussian distribution, with variance σ_f^2, provided that $|f - f_0| \ll 1/T_{\mathrm{c}}$. Note incidentally that $1/T_{\mathrm{c}}$ is also the frequency step for the discretized GMF. Using the Gaussian distribution is equivalent to using the linearised signal model of section 2.3.[1] Therefore, for example with 0.1 s coherent integration, using linearised theory for error estimation is legal—or at least self-consistent—when the posteriori density is so narrow that most of the probability mass is well within 10 Hz (1.6 m s^{-1}) from f_0. This should happen when the signal is strong enough.

We can use the full posteriori density from Eq. (4.1) to find how strong the signal needs to be for the linearised theory to be valid. The centre column of Fig. 4.7 shows the posteriori density for a few targets of increasing size, the rightmost column shows a blow-up around the maximum at $v = v_0$. We assume EISCAT UHF radar, transmission duty cycle 10%, target at 1000 km

---

[1] Actually, there is a small (about 25%) discrepancy between the velocity error we have shown in Fig. 2.4 and the error σ_f computed from Eq. (4.4) and shown in Fig. 4.8.
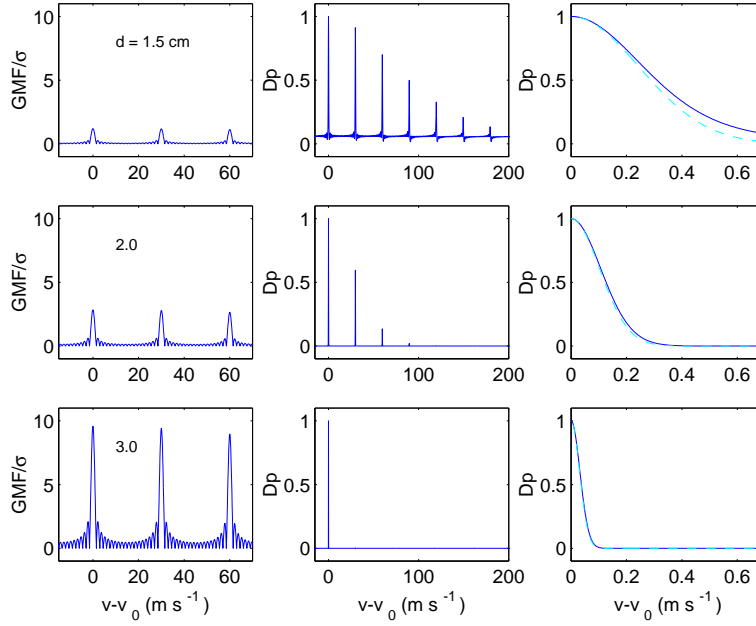
Figure 4.7: Shape of the posteriori density as a function of signal strength. The leftmost column of panels show the GMF, the centre column gives the (non-normalized) D_{p} in the velocity direction, the rightmost column is a blow-up of the centre column near the main peak. The dashed line in the rightmost panels is the Gaussian approximation of D_{p}.

range, and coherent integration $T_{\mathrm{c}} = 0.1$ s. The leftmost column shows the GMF, in the $\sqrt{\mathrm{SNR_N}}$ units that we normally use in target detection. The dashed line in the rightmost column shows the posteriori density for the linearised theory, computed from the Gaussian approximation (4.3).

Figure 4.7 shows that for a weak signal (d = 1.5 cm, $\sqrt{\mathrm{SNR_N}} \ll 5$), the posteriori probability is spread over a wide velocity interval. From such a measurement we cannot single out a preferred velocity for the target, in other words, we cannot "detect" the target (in the velocity space). The error estimate cannot be computed from the Gaussian approximation, the estimate would be quite too small. On the other hand, for a strong signal (d = 3.0 cm, $\sqrt{\mathrm{SNR_N}} \gg 5$), the posteriori density is well concentrated around the single point v_0. The linearised theory now gives a good error estimate. It hardly can be a coincidence that the localization of the posteriori density occurs when $\sqrt{\mathrm{SNR_N}}$ reaches the value about 5, which we experimentally, in the test campaign, found to be a good detection threshold.

We may attempt to find the region of validity of the Gaussian approximation on the basis of the localization property of the posteriori density. An easy-to-compute indicator of D_{p} localization is the size of the first side
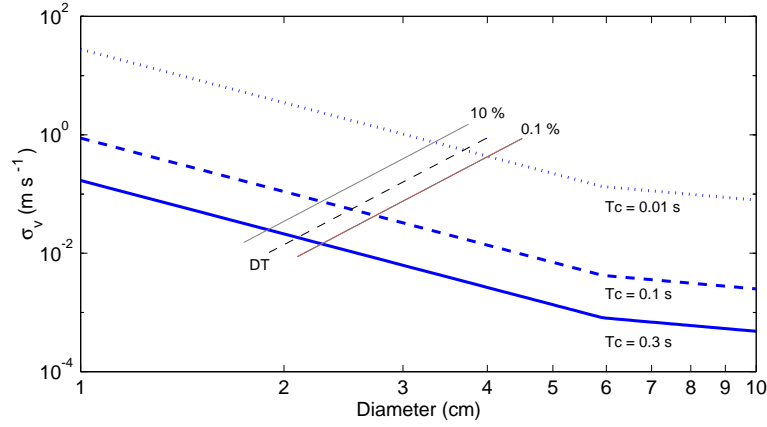
Figure 4.8: Validity regime of the linearised error model. The thick lines give the velocity error σ_{v} as a function of target size, computed from the Gaussian approximation of the posteriori density, for three integration times T_{c}. The thin solid lines (10%) and (0.1%) delimit the transition region where the posteriori density localizes when we move from left to right along the error curves. The thin dashed line (DT) corresponds to the detection threshold that was used in the test campaign. The linearised theory should be applicable to the right of the transition zone.

maximum, at $f_1 = f_0 + 1/P$, compared to the main maximum at f_0. We will consider the density localized when the ratio $D_{\mathrm{p}}(f_1)/D_{\mathrm{p}}(f_0)$ is 0.1% or smaller. We take the velocity error curves of the linearised theory, and mark on each of them the target size that corresponds to the localization limit, to get the line labelled "0.1%" in Fig. 4.8. The line labelled "10%" is found in a similar way. The velocity error curves in Fig. 4.8 assume target at 1000 km range. The $T_{\mathrm{c}} = 0.1$ curve is compatible with the curve we have shown earlier, in Fig. 2.4. The region of validity of the Gaussian error estimates begins somewhere between the two lines marked by 10% and 0.1%, and extends to the right in the figure.

The dashed line labelled "DT" in Fig. 4.8 represents our standard detection threshold $\sqrt{\mathrm{SNR_N}} = 5$, and corresponds to a high degree of D_{p} localization, $D_{\mathrm{p}}(f_1)/D_{\mathrm{p}}(f_0) = 2\%$. Therefore, we seem to have got the—problematic—result that if the target can be detected at all, it is safe to use the (too) small velocity error bars of the linearised theory for it.

Conceivably, a loophole in the above argument could be that we have been using the continuous-parameter GMF. Could the reason for the larger-than-predicted velocity fluctuations be that in our actual data analysis, we maximize the GMF only on a discrete grid of velocity values? The answer appears to be "no". Figure 4.9 shows, as function of velocity through the GMF maximum position, both the continuous-parameter $\mathrm{GMF}(v)$ and the
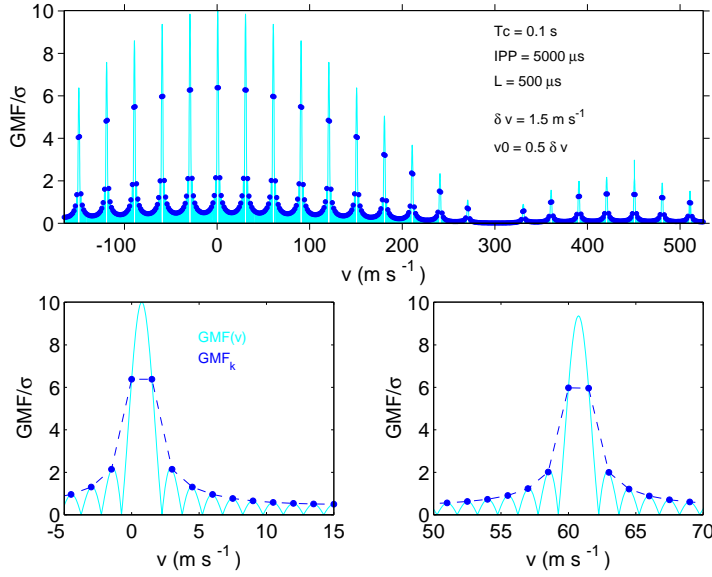
Figure 4.9: Continuous-parameter GMF and discretized GMF. The figure shows the theoretical, noise-free GMF as a function of velocity, for a simple experiment as specified in the top panel. The solid line is the continuous-parameter GMF(v), the dots give the discrete-parameter GMF(v_k). The two bottom panels are blow-ups around the main maximum and the second side maximum. The target velocity $v_0 = 0.75$ m s^{-1} is exactly half of the velocity grid spacing. The continuous-parameter GMF has its main maximum at v_0, the (now ambiguous) position of the main maximum of the discrete-parameter GMF differs from it only by the maximum velocity quantization error, $\delta v/2 = 0.75$ m s^{-1}.

discrete-parameter GMF$_k$ = GMF(v_k). Their relation depends on the actual radial velocity v_0 of the target. If the velocity is equal to some available grid point (in the figure, the point $v_0 = 0$), the maximum of GMF(v_k) is at $v = v_0$. Even the value of the GMF$_k$ maximum equals to the value of the continuous-parameter GMF. The worst case is when the target velocity falls precisely halfway between two velocity grid points, this is the situation shown in Fig. 4.9. The GMF$_k$ values of the maxima are about 60% lower than the GMF values, but the main GMF$_k$ maximum is inside the main maximum of the continuous-parameter curve. Thus, the maximum possible velocity error is just the velocity quantization error, 0.75 m s^{-1} in this example. It does not look probable that the discretization could be the main source of the error size discrepancy.

It seems that we must claim the large velocity errors to noise. It seems that the noise causes so much variation in the GMF that, every now and then, a wrong peak will be picked by the maximization algorithm; this

will then immediately give velocity error that is of the observed magnitude, tens of meters per second. We suspect that we are seeing here a manifestation of an old problem of pulse radar engineering, the so called range-Doppler dilemma. It is an old maxim of radar science that if one tries to increase the unambiguous range by decreasing the pulse repetition frequency (PRF = 1/IPP), one runs into the danger of the velocity becoming ambiguous. In Fig. 4.9, the separation of the local GMF maxima is equal to the PRF. The envelope curve of the maxima is determined by the pulse length, L. If L is fixed and we decrease the PRF, the level of the GMF side maxima increases. This means that the noise has better chance to cause a wrong peak to be selected, in other words, the velocity has become more ambiguous.

If the large errors really are related to the range-Doppler dilemma, the errors can, possibly, be mitigated. Contrary to what is sometimes stated [8, page 139], the range-Doppler dilemma is not an unavoidable limitation of radar operation, but can actually be circumvented. One of us (Lehtinen) actually owns a patent for a method to do precisely that [14]. The Lehtinen method is based on varying the radar's pulse repetition rate in a carefully chosen manner, and is therefore called the SMPRF-method (Simultaneous Multiple Pulse Repetition Frequencies).

Digging the skeleton of the range-Doppler dilemma from its hard-earned grave is not the quantitative solution to the error size discrepancy we would like to have. The error analysis, when based on the full posteriori density, should be able to predict the errors correctly. We did assume zero noise when deriving the form of the posteriori density; this might not be permissible in the non-linear problem. It might also be essential to take the range direction into consideration when computing the error estimates in the full theory.

4.2.5 Acceleration estimates

The acceleration estimate is plotted at the bottom panel of Fig. 4.5. This estimate is not given for the weakest targets at all, but when available, it comes from the same linear fit to the $v_D(t)$ data as the Doppler-velocity. The accuracy of the slope of the fitted line depends strongly on how long the target stays visible. On the basis of the event summary plots, it is clear that Fig. 4.5 includes acceleration points that deviate significantly from the expected value. Most of the targets should have almost circular orbits, and their radial accelerations are expected to fall near the solid line given in Fig. 4.5. The large deviations from this expected value a_O strengthen the conclusion that the acceleration cannot be estimated in this way for most of our events. The bottom panel of Fig. 4.6 shows that increasing the signal strength in itself does not necessarily improve the estimate much. However, in the few cases where there is a strong signal which can be followed for several seconds, we normally get an acceleration estimate that is within a

few m s^{-2} from a_O.

Our initial idea was to obtain the acceleration directly for each individual scan by varying the acceleration parameter in the GMF. It now seems that the GMF does not depend strongly enough on the acceleration, at least when we use only 0.3 s integration. Figure 2.13 shows the change of $\mathcal{R}_{\max}$ (for cp1lt data) when the acceleration parameter in FastGMF is changed from the circular-orbit value to zero. The mean reduction is 33%, there is large scatter, but not the expected trend. If a_O is approximately the correct acceleration, the larger it is, the larger is the error we make when changing to $a = 0$. Thus the drop $\Delta\mathcal{R}_{\max}$ of $\mathcal{R}_{\max}$ should increase with increasing a_O. No such trend is visible in the bottom panel of Fig. 2.13. Only the scatter in $\Delta\mathcal{R}_{\max}$ increases somewhat with increasing a_O.

Thus, the question of how to best determine the acceleration is still open. For the time being, as long as our data are taken with almost vertical beam, the best estimate for acceleration is just the circular-orbit, vertical-beam value. This strategy will need to be reconsidered if handling data from lower elevations becomes relevant.

4.3 FastGMF versus GMF

In MATLAB c-implementation, the FastGMF algorithm is typically about 100 times faster than the full GMF. For speed reasons, it is mandatory to use the FastGMF in the detection phase, but it would be good to be able to use it for the actual parameter estimation also. It took 10 days of CPU time to scan the 45 cp1lt events with the GMF, even though we used as short range interval as possible. The integration was 0.3 s, and full spatial and temporal resolution were used in the scans.

In Section 2.4.3 we pointed out that both simulation and data from the test campaign have shown that very little can be gained from the use of the full GMF algorithm instead of the much faster FastGMF. To understand why FastGMF can be this good in comparison to the full GMF, consider the following. In the optimal case, both GMF and FastGMF are able to recover the signal amplitude exactly. The actual value of $\mathcal{R}_{\max}$ estimate depends more significantly on how far the actual target parameters happen to be from the discrete grid of range and velocity values than the computational algorithm. For a target that we can follow a reasonable length of time, both range and velocity change appreciably, and the phase space trajectory has a good chance to get near some grid point, irrespective whether we are using GMF or FastGMF. To rephrase this argument a bit more formally, recall that in the parameter estimation we are maximizing, with respect to j and k, an expression of type

$$G_{j,k} = \frac{1}{T_c} \left| \int_0^{T_c} dt\, e^{i(\omega_k - \omega_0)t} \times \epsilon^{R_j}(t)\overline{\epsilon^{R_0}(t)} \right|, \tag{4.5}$$

where $\epsilon^R(t)$ means the transmission envelope translated to position R and normalized to unit amplitude, R_j and ω_k are points on a fixed grid, R_0 and ω_0 are the actual target parameters, and T_c is the integration time. Then the crucial factor is, when the point (R_0, ω_0) moves, how near will it approach some of the grid points.

This interpretation is consistent with the data in the bottom panel of Fig. 2.7. The figure shows that the fastest moving targets tend to benefit the least from the use of the full GMF. The other side of the coin is that even the full GMF algorithm tends to underestimate $\mathcal{R}_{\max}$, and, moreover, it is not immediately obvious how to best improve the estimate. Tightening the grid in range direction improves the amplitude estimate $\widehat{A}$, but tightening the grid in frequency direction does not improve the amplitude estimate at all. The latter result comes about because, due to the FFT, the frequency grid spacing δf is coupled, via $\delta f = 1/T_c$, to the shape of GMF: when the grid becomes tighter, T_c becomes larger, and then the GMF oscillates more wildly. To see this quantitatively, note from (4.5) that if $R_0 = R_{j_0}$ for some j_0, and A_0 is the actual signal amplitude, we have

$$\frac{\widehat{A}}{A_0} = \max_k G_{j_0,k} = \max_k \left| \frac{\sin \pi (f_k - f_0) T_c}{\pi (f_k - f_0) T_c} \right|, \tag{4.6}$$

which always attains the same minimum value, $\frac{2}{\pi} \approx 0.64$ (at $|f_k - f_0| = 1/2T_c$), no matter how small the frequency grid spacing is [see also Eq. (4.2) and Fig. 4.9]. If the target Doppler-frequency f_0 is picked at random, the expected value of $\widehat{A}/A_0$ is

$$\left\langle \frac{\widehat{A}}{A_0} \right\rangle = \frac{1}{0.5} \int_0^{0.5} df \, \frac{\sin \pi f}{\pi f} \approx 0.87, \tag{4.7}$$

independently of how tight the frequency grid is.

One possibility to improve the amplitude estimate, while still using FFT for the GMF computation, is to multiply the measured signal $m(t)$ by an additional phase factor

$$g^{(\delta k)}(t) = \exp(i \, \frac{2\pi \, \delta k}{T_c} \, t), \quad \delta k \leq 0.5, \tag{4.8}$$

at the same time when $m(t)$ is multiplied by $\exp(i\alpha t^2)$ to make the acceleration correction. The benefit is that then the coupling $\delta f = 1/T_c$ no more holds; in addition to the frequencies $f_k = k/T_c$, we are allowing also the frequencies $(k + \delta k)/T_c$ in the signal model. The simplest strategy is to take $\delta k = 0.5$, in addition to $\delta k = 0$. This improves the worst-case amplitude estimate from $\widehat{A}/A_0 = 0.64$ to $\widehat{A}/A_0 = \sin \frac{\pi}{4}/\frac{\pi}{4} \approx 0.90$, at the price of roughly doubling the computing time. The δk-correction could in principle be worthwhile in the final analysis when computing target diameter.
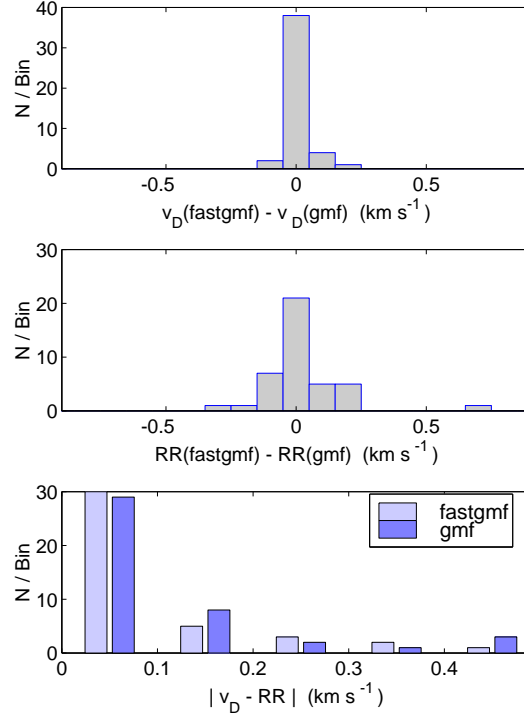
Figure 4.10: Effect of GMF algorithm on velocity estimates. The top panel is a comparison of Doppler-velocity v_D in GMF and FastGMF. The middle panel is a comparison of the range rate $RR = dR(t)/dt$. The bottom panel is a comparison of the consistency between v_D and RR.

However, we have not applied the correction during this study. The dominant source of error in target size determination with the EISCAT radars is the uncertainly about where the target went through the radar beam. The δk-correction becomes more relevant if this bigger problem can be solved.

For the Doppler-velocity, GMF and FastGMF give quite similar estimates. The top panel of Fig. 4.10 shows that in most cases the difference is less than 50 m s^{-1}, and very rarely exceeds 150 m s^{-1}. This is about the same as our rough estimate for the velocity error, so it seems that as far as the Doppler-velocity is concerned, we do not lose much accuracy if we use solely the FastGMF algorithm. The middle panel in Fig. 4.10 shows that there is more difference between GMF and FastGMF algorithms in the RR than in the directly measured Doppler-velocity. This is consistent with the assumption of inherently larger errors in RR, irrespective of the algorithm.

Figure 4.11 compares the acceleration estimate in the GMF and FastGMF analysis, see also the bottom panel of Fig. 4.6. Irrespective of the GMF algorithm, most of the deduced acceleration values differ significantly from the expected circular-orbit values. This shows that it is problematic to get a
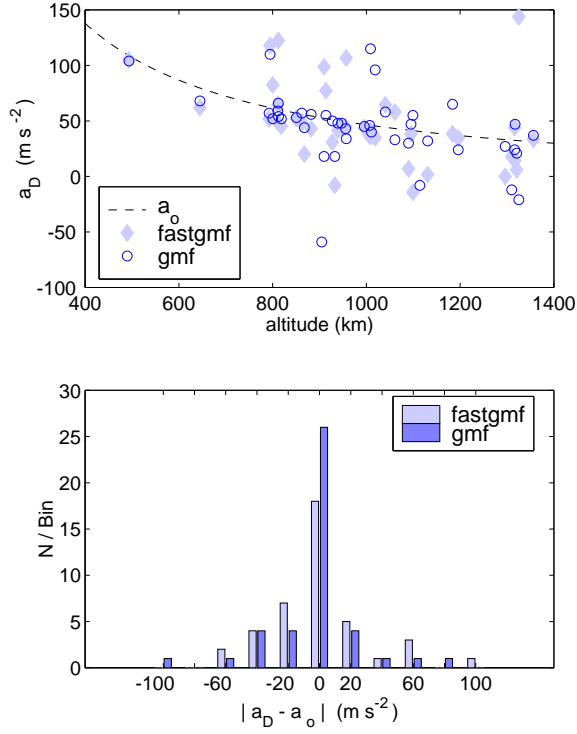
Figure 4.11: Effect of GMF algorithm on acceleration estimate. The top panel shows the acceleration $a_D = dv_D/dt$ for the cp1lt events in the test campaign. The dashed line shows the circular-orbit acceleration a_O. The bottom panel shows the deviation from a_O.

reliable estimate of the acceleration by fitting a line to the Doppler-velocity data. The problem is probably related to the large errors in the velocity estimates. Our strategy at this stage has been to fix the acceleration parameter to the circular-orbit value in the GMF. We have also made a few attempts to estimate the acceleration directly from individual measurements, but it has been difficult to make any clear conclusion based on them.

We do not show any comparison between the GMF algorithms for the remaining basic parameters, the time and range of the maximum signal strength. Both methods give essentially identical results, within the resolution of the spatial and temporal grids used. We take this to mean that the absolute accuracy of range determination is limited by the sampling interval, rather than any estimation errors. In the few cases where the tabulated events indicate a substantial difference in range, there is also clear difference in time. In these cases, the $\mathcal{R}_{\max}(t)$ curve has two or more well separated peaks of almost same height, possible due to rotating targets, and different peaks have by chance been selected by the different algorithms.

4.4 Comparison with a catalogued object

The strongest of the high-altitude tau2 events during the test campaign began at 22:19:06, 20 February 2001. The event was studied by M Landgraf from ESOC [15]. He identified the target as a large catalogued object, COSPAR ID 1994-11G. According to Landgraf, the object has total mass 1390 kg and cylindrical shape, with diameter 2.7 m and height 2.2 m. It has radar cross section 8.3 m^2. The target should have passed the EISCAT beam at the off-axis distance of 1.27°. Figure 4.12 shows our analysis summary plot of the event. The analysis was done with FastGMF, using 0.27 s integration.

The top panel of Fig. 4.12 shows the peak ratio $\mathcal{R}_{max}(t)$ in dB scale. Markers indicate those scans where $\mathcal{R}_{max}$ was larger than the threshold 4.5. The antenna sidelobe pattern is clearly visible. Even the third sidelobe is seen, at about time 4 s, although the corresponding sidelobe on the other side of the main lobe is no more observed. To find the off-centre distance of the transit, we fitted theoretical EISCAT antenna pattern to the $\mathcal{R}_{max}$ data. The theoretical pattern uses the specified diameters and mutual distance of the UHF antenna main reflector and subreflector, but does not account for the effects of the supporting structures. The best fit is shown in Fig. 4.13. It was achieved by assuming that the beam transit took place at 0.52° offset from the optical axis. The actual beam shape differs somewhat from the ideal shape. It is known, for example, that the third sidelobes are actually asymmetric and weaker than the theoretical value. However, even taking this uncertainty into account, it seems difficult to reconcile the predicted offset of 1.27° with the inferred value 0.52°. It is possible that the EISCAT pointing was not what we believed it was, azimuth 183.3° and elevation 77.1°. Normally, EISCAT pointing is known to be accurate to within 0.1°, but the February campaign took place after an exceptionally long maintenance and system upgrade period, after which no pointing calibration had been done. Due to peculiarities of EISCAT pointing hardware, the *azimuth* pointing can fairly easily get an accidental offset during maintenance.

Panels (2) and (3) of Fig. 4.12 show the measured range and Doppler-velocity. The solid dark curves represent quadratic and linear fits to the good ($\mathcal{R}_{max} > 4.5$) points. The large dots represent the predicted values. The measured range is about 7 km larger than predicted. Assuming roughly south-to-north flight direction, 7 km would require an elevation offset about 0.6°. This appears unreasonably large. It is difficult to explain the 7 km from a timing error either. The situation with the velocity data is similar. The slope of the velocity curve is as predicted, but there is a conspicuous and consistent discrepancy of about 0.1 km s^{-1} in the actual velocity values. The circular-orbit, vertical-beam acceleration is 27.5 m s^{-2}, which is consistent with the value 27.1 m s^{-2} from the velocity fit. According to Landgraf, the timing accuracy for the catalogued objects is of the order of 10 s, while we believe our timing to be accurate to within about 0.1 s. However, the
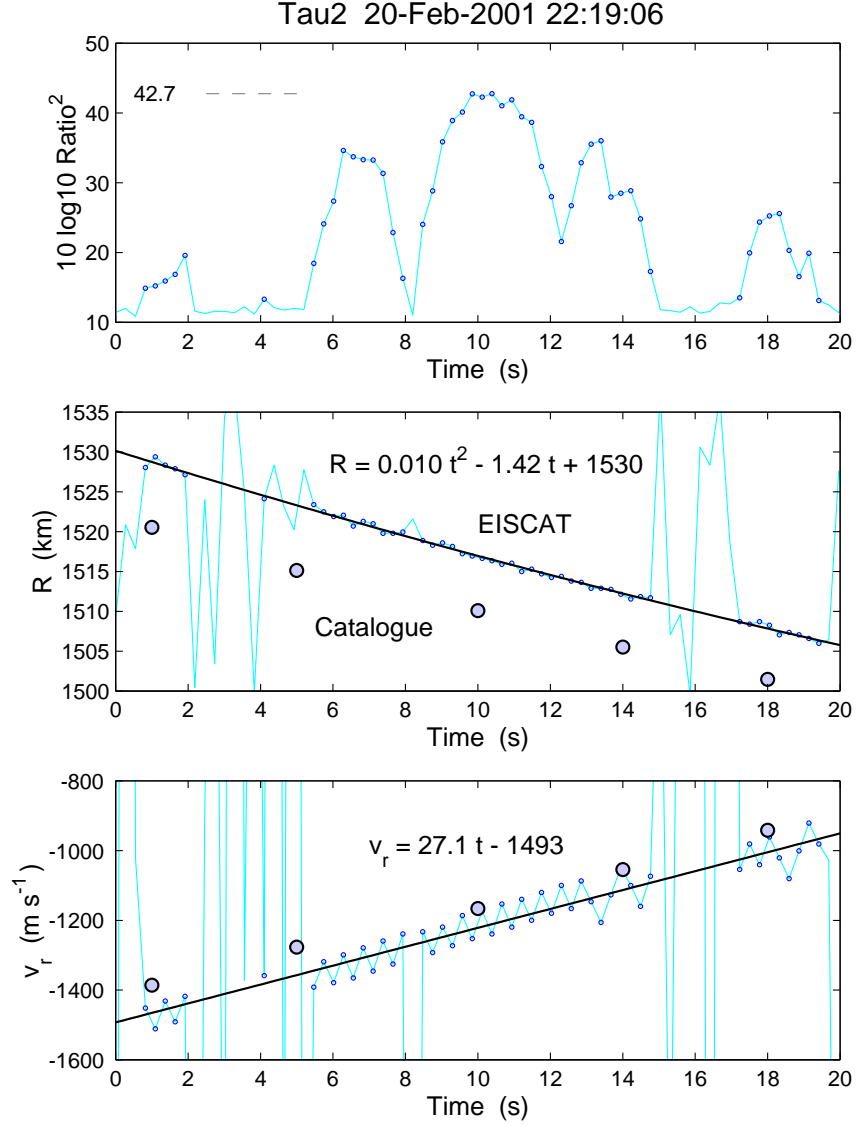
Figure 4.12: Event at 22:19:06, 20 February 2001, catalogue comparison. The top panel shows $\mathcal{R}_{\max}$ in dB. The middle panel shows the measured range (small circles), a parabolic fit, and the catalogue prediction (large circles). The bottom panel shows the measured Doppler-velocity (small circles), a linear fit, and the catalogue prediction for the range rate (large circles). The measurement was analysed with the FastGMF algorithm.
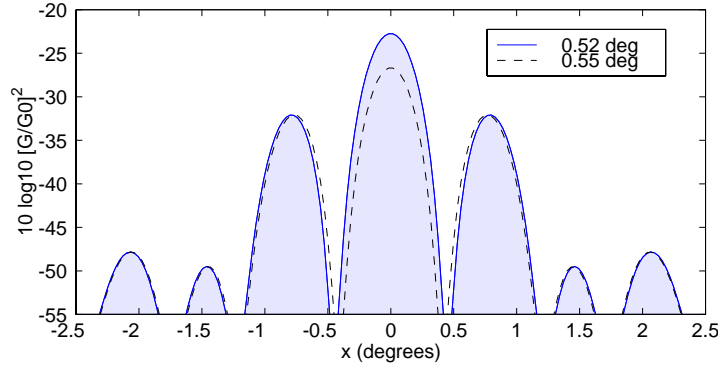
Figure 4.13: A cut of the theoretical EISCAT UHF antenna pattern at off-axis distance $0.52°$. The y-axis is the antenna gain divided by the gain on the optical axis.

range and velocity discrepancies cannot be removed simply by adjusting the relative timing, because the required correction would be about six seconds for the range data, but only about three seconds for the Doppler-data.

As the final comparison, we consider the signal level. We get from Eq. (2.60) an estimate of the SNR, normalized for 1 MHz noise bandwidth. We use the value $\mathcal{R}_{\max} = 137$ from the top panel of Fig. 4.12, and get the number of signal samples $N \approx 24000$ (per 1 μs sampling interval) from the integration time 0.273 s and the duty cycle 8.8%. We find

$$\text{SNR}_{0.52°} = \frac{137^2}{24000} = 0.782.$$

We estimate from Fig. 4.13 that there is about 23 dB loss in received signal power because the target was off-axis, and use that value to correct the SNR estimate. We get

$$\text{SNR}_{\text{axis}} = 0.782 \cdot 10^{2.3} = 156.$$

This is the SNR we would have got if the target would have gone trough the beam centre. We use SNR_{axis}, the measured range 1520 km, and the radar equation to compute the radar cross section. We get cross section 4.2 m^2, which is 2.0 times smaller than the catalogued value. Possible reasons for the discrepancy are our assumptions of 60 K noise temperature and 1.0 MW transmission power and, especially, our estimate of the off-centre loss, 23 dB. The last number alone could be wrong by the required 3 dB.

Chapter 5

Discussion and Conclusions

Measuring SD with EISCAT radars

In this study we have designed, implemented, and tested a system for using the EISCAT incoherent scatter radars to detect and characterize small-size space debris. Our aim has been to develop a system that can observe space debris unobtrusively in the background when the radar is doing its normal ionospheric work. Our solution introduces a new piece of hardware, the SD receiver. The interface between the host radar and the SD receiver is clean and minimal. The radar does not need to know what the SD receiver is doing. It only needs to provide the following minimal and cheap resources:

- ▷ a copy of its received analog signal
- ▷ a copy of its transmission waveform
- ▷ antenna pointing direction and transmission power as a network service
- ▷ optionally, a copy of the site's GPS 1-second pulse.

The SD receiver, in turn, does not need to know the details of the running radar experiment beforehand. Our data processing requires the transmission waveform as input; that we measure directly. The SD receiver needs not to know anything about the host radar's own signal processing.

The SD receiver maps its input analog frequency band, which usually contains several narrowband frequency channels with a total spread of a few MHz, to the band's digital image, represented by a single stream of complex samples. A workstation connected to the receiver transfers the samples to hard disk. Because of the required high sampling rate for routine operations, it appears necessary to perform SD target detection in real time, so that only data containing targets need to be saved. The required computing speed for real-time detection is in excess of one Gflop/s, but is within the reach of a modest number (3–4) of today's workstations, working in parallel. In this preliminary study, we have worked entirely off-line, using the data recorded on disks.

Our SD detection and data analysis is based on statistical inversion, which we have applied via a tool we call the GMF. We have shown that the GMF, with different approximations and resolutions, suffices both for the initial detection and the final parameter estimation. The foundation of the GMF method is firm and simple, and the parameter estimates, at least in special cases, are equivalent to well-known maximum likelihood estimates.

The problem of parameter estimation accuracy

As we have discussed above, the analytic prediction of parameter errors fails for small SNR. The errors can, however, be estimated from the fluctuation of the deduced parameters during the time the target moves through the radar beam. In order to get a more reliable analytic parameter error estimate, we will have to investigate the Gaussianity of the posteriori density and the coherence of the received signal.

The event summary plots of the test campaign show the variation of all target parameters from integration to integration. Jumps in the velocity time series are from a few meters per second to several tens of meters per second, after the linear trend due to acceleration is removed. The linearised-theory curve in Fig. 2.4, on the other hand, gives r.m.s. error 0.16 m s^{-1} for a 2.0 cm target at 1000 km range with 0.1 s integration. According to the results of the test campaign, such a target is just barely detectable even with 0.3 s integration.

Issues of coherent integration

Early on in our two-year study, we proposed to use coherent integration to increase detection sensitivity. The requirement of coherent integration then lead us to adopt the GMF technique. Coherent integration surely is essential for sensitivity. It is essential also for the parameter accuracy; we note from Eq. (4.4) that the posteriori density becomes better localised in proportion to the SNR_N and not merely to the SNR. Therefore, it is unfortunate that we cannot in practice integrate longer than a small fraction of a second, say ≈ 0.1 s, before starting to lack behind the predicted sensitivity gain. We do not know what causes this apparent loss of coherence. It could be that the targets do not scatter coherently, but for a small target deep in the Rayleigh region this would be unexpected in itself. It might be that ionospheric electron density fluctuations cause perceptible irregularities to the phase of the signal. Clutter or interference could disturb the signal. And we can never completely exclude phase drifts due to the hardware, though the GMF method automatically cancels any drift that is common both to the transmission and the reception. This cancellation requires that the transmission samples are of good quality and really represent the transmitted signal.

Real signals are never completely coherent, but we expect our signal model, which is already second order in time, to remove at least most of the long-term phase drifts. On the other hand, the signal model cannot take into account rapid back-and-forth phase variations (phase jitter) during integration. We did check the phase jitter of the SD receiver's clock input signal during the test campaign, and concluded that it would not cause problems for us.

We have also checked, by using perfectly coherent simulated data, that the detection sensitivity increases as expected when integration time is increased. This observation demonstrates that our data analysis software is implemented correctly. In the light of the error estimation problem, it would be desirable to check also velocity and range error behaviour with simulated data. There is a good place for further study here, via simulations, theory model improvements, and more extensive study of actual data.

Test campaign

On the basis of our February 2001 test campaign at EISCAT UHF radar, we can draw several conclusions, which we, in spite of the smallness of the analysed data set, believe to be reasonably robust.

1. We get 10–15 events per hour.
2. Most events are seen in the altitude band 800–1200 km.
3. The detection limit is 2 cm at 1000 km range, with 1 MW transmission peak power and 0.3 s coherent integration.
4. Range can be determined essentially with the accuracy determined by the gate separation, about 0.1 km.
5. Doppler-velocity can be determined with about 0.1 km s^{-1} accuracy or better in most cases.
6. It has not been possible to get good estimate for acceleration, but the circular-orbit value is good enough for the detection purposes as long as almost vertical antenna pointing is used.
7. Our approximate FastGMF algorithm can be used both for detection and final parameter estimation, without significant loss of sensitivity or accuracy.
8. The required sustained computing speed is about 1 Gflop/s for real-time detection in cp1lt and tau2 experiments, on the same level of sensitivity and accuracy that we had in the test campaign.

Next steps

The test campaign showed that meaningful target parameters are obtained at EISCAT with our SD measuring and analysis system, even though the parameter accuracy achieved so far is not as good as we had expected. However, for the practical purpose of routine monitoring of small-size SD, to achieve centimeters per second velocity accuracy is not top priority. It could be that the larger-than-expected errors relate to the apparent lack of coherence, and via it, to detection sensitivity, which definitely has an impact on the measured size distribution. We are confident that the problems with estimation accuracy and coherent integration will be understood, and that more sensitive and accurate measurements can be achieved in the future. This might require more computing power than we have estimated in this study.

One of the main benefits of the statistical inversion approach is that it provides a firm foundation for error analysis. We should extend the measuring model to include several integrations and get a proper error estimate for the event parameters. Another interesting aspect for future work is the use of a priori information. Statistical inversion gives a well-defined way to feed a priori probability information into the detection problem. We have not made use of this option so far, but it is not inconceivable of doing so

later, when we have better ideas about the parameters of the SD popula-
tion. Correct prior information is an extra helping hand, and could prove
useful for large-scale statistical studies of small-size SD populations, where
we might want to push down the detection limit.

The test campaign showed that using our SD receiver, the SD measure-
ments can readily proceed in the background of standard EISCAT iono-
spheric measurements. There are a few things that need to be taken care of
before routine operations could begin.

> ▷ We need real-time detection capability. This requires that key parts
>   of our MATLAB-based software is coded with a compiled language. We
>   also need an environment that allows distributing the detection com-
>   putation among a few networked workstation on scan-by-scan basis.
> ▷ We need to streamline our data analysis, to make it considerably more
>   automated. We need to include non-stationary noise into our mea-
>   surement model to handle noise injections. We need to include valid
>   error estimates into our event parameter listings.
> ▷ We need to analyse the data further, to extract orbital elements based
>   on circular-orbit approximation, and to provide error estimates for
>   them.

In the longer term, we should not forget the unique tristatic property of
EISCAT UHF system, which allows the determination of true orbit and
true scattering cross section from a single measurement, without any addi-
tional assumptions. However, tristatic SD measurements, as well as all other
changes in EISCAT standard experiments like using variable-length IPPs,
require EISCAT's active participation, and are therefore at odds with the
low-profile, no-interference approach that we have adopted in this study.

We have shown that there are no technical nor operational reasons why
EISCAT could not help substantially in space debris monitoring, perhaps
on the level of several thousand analysed events per year. This can hap-
pen without causing any interference with EISCAT's primary ionospheric
mission, and with minimal extra cost.

Chapter 6

Acknowledgements

Appendix A

Multi-frequency GMF

A.1 Generalized matched filtering

When searching the radar echo of a moving point target in received data, we essentially want to find a time-shifted and Doppler-shifted replica of the transmission pattern in the reception time series. The target detection thus amounts to a pattern search (pattern match) operation. A well-known and optimal way of finding a known pattern in noisy data is by matched filtering (MF), where one matches the pattern in time (or, equivalently for a pointlike radar target, in range). When the pattern in addition of the unknown start time also possesses an unknown shift in frequency, the proper tool is the generalized match filter. The generalized MF is capable of matching the pattern both in range and frequency simultaneously, and involves the computation of the generalized match function, GMF.

Denoting the transmission by $T(t)$ and the reception by $S(t)$, the GMF for our detection problem is

$$\text{GMF}(r, \omega) = \left| \int dt \, [S(t) + \Gamma(t)] \cdot \overline{T}(t - 2r/c) e^{-i\omega t} \right|, \qquad (A.1)$$

where $\Gamma(t)$ is noise and the the overbar denotes complex conjugation. In general, a transmission can contain multiple frequency channels, that is, several modulation patterns on different frequencies. The channels can be sent during the same IPP or in the cause of several IPPs. Here we assume that the channels are sent almost consecutively during a single IPP. The transmission on channel n can be written as

$$T_n(t) = \epsilon_n(t) e^{i\omega_n t}, \qquad (A.2)$$

where the complex-valued transmission envelope $\epsilon_n(t)$ describes the code pattern on DC level, a phase code say, and ω_n is the radio-frequency emitted from the antenna. The complete transmission is then

$$T(t) = \sum T_n(t) \qquad (A.3)$$

and the reception is

$$S(t) = \sum S_n(t), \qquad (A.4)$$

where S_n is the reception that would be obtained if only T_n were transmitted.

A.2 GMF with a single-frequency transmission

We first handle the case of a single-channel experiment. The reception S is essentially the transmitted pattern, shifted in time by the pulse propagation delay $2R/c$, where R is the target distance, and shifted in frequency by the Doppler-frequency ω_{D},

$$S(t) \propto \epsilon_1(t - 2R/c) \cdot e^{i(\omega_1 - \omega_{\mathrm{D}})t}. \qquad (A.5)$$

The integrand in the definition of GMF is then

$$
\begin{aligned}
S\overline{T}e^{-i\omega t} &\propto \epsilon_1(t-2R/c)e^{-i(\omega_1-\omega_\mathrm{D})t}\cdot\overline{\epsilon}_1(t-2r/c)e^{i\omega_1 t}\cdot e^{-i\omega t}\\
&= \epsilon_1(t-2R/c)\overline{\epsilon}_1(t-2r/c)\cdot e^{-i(\omega-\omega_\mathrm{D})t}.
\end{aligned} \tag{A.6}
$$

Note that due to the complex conjugation in $\overline{T}$, the rapid time variation $e^{i\omega_1 t}$ cancels from $S\overline{T}$. From (A.6) it is evident that without noise, we maximize the GMF ("get best match") at the point $(r,\omega) = (R,\omega_\mathrm{D})$.

Summarizing the case of a single-frequency transmission, we have seen that

$$
\mathrm{GMF}(r,\omega) \propto \left|\ \int dt\,\epsilon(t-2R/c)\overline{\epsilon}(t-2r/c)e^{-i(\omega-\omega_\mathrm{D})t}\ \right|, \tag{A.7}
$$

which has maximum at the point corresponding to the target range R and the target Doppler-shift ω_D. To get some idea about the behaviour of the GMF, we inspect the cuts of the GMF-surface through the maximum point, in the range and frequency directions.

In the range direction through the maximum the GMF, without noise, reduces to the magnitude of the autocorrelation function of the transmission envelope

$$
\mathrm{GMF}(r,\omega=\omega_\mathrm{D}) \propto \left|\ \int dt\,\epsilon(t-2R/c)\,\overline{\epsilon}(t-2r/c)\ \right|. \tag{A.8}
$$

Typically this is a rather slowly varying function of r around the maximum at $r = R$, with width of the order of twice the length of the constant-phase element of the code (the baud length). The autocorrelation function is identically zero when $2r/c$ differs from $2R/c$ by more than the full duration of the transmission (supposing that the r is not so large that pulse-to-pulse correlation becomes an issue). In the next section we will show that when there are multiple frequencies in the transmission, an additional modulating factor $\rho(r-R)$ appears to the r.h.s. of Eq. (A.7), with faster variation.

In the frequency direction through the maximum, the GMF reduces to

$$
\mathrm{GMF}(r=R,\omega) \propto \left|\int_{\mathcal{T}} dt\, e^{-i(\omega-\omega_D)t}\ \right|. \tag{A.9}
$$

The integral in (A.9) is over the interval (or intervals) of time when the echo is received.

In the simplest case, the GMF-integral is evaluated only over one IPP, of duration P, so that the integration limits in Eq. (A.7) would be from 0 to P. One of the main benefits of the GMF detection method is that, technically, the integral can as easily, though of course not as fast, be evaluated over M IPPs. If the echo arrival time and other properties change only very slowly from IPP to IPP, a factor of M improvement in detection sensitivity in noise-dominated environment can be achieved.

We denote by L the duration of the envelope ϵ_1. Then for a single IPP, Eq. (A.9) becomes

$$
\begin{aligned}
\text{GMF}(r = R, \omega) \quad &\propto \quad |\int_{2R/c}^{L+2R/c} dt\, e^{-i(\omega - \omega_{\mathrm{D}})t} | \\
&= \quad L \cdot \text{sinc}(\frac{L(\omega - \omega_{\mathrm{D}})}{2}).
\end{aligned} \tag{A.10}
$$

The main maximum has width $\Delta\omega = 4\pi/L$ or $\Delta f = 2/L$, measured from zero to zero around the maximum.

If the GMF-integral is taken over M IPPs, each of duration P and all having sufficiently identical transmission of length L, the sinc factor in (A.10) becomes modulated by the Dirichlet-kernel

$$
|\sum_{m=0}^{M-1} e^{im(\omega - \omega_{\mathrm{D}})P} | = \left| \frac{\sin[(\omega - \omega_{\mathrm{D}})MP/2]}{\sin[(\omega - \omega_{\mathrm{D}})P/2]} \right|. \tag{A.11}
$$

The expression (A.11) has maximum equal to M at $\omega = \omega_{\mathrm{D}}$ and zeros at the points $f_m = \frac{m}{MP}$. The qualification "sufficiently identical" especially means the requirement of phase coherence, or constant phase difference, between transmission $\epsilon_T(t)$ and the corresponding echo $\epsilon_R(t)$ from IPP to IPP, so that

$$
\epsilon_R(t) \cdot \overline{\epsilon_T(t)} \approx \epsilon_R(t + mP) \cdot \overline{\epsilon_T(t + mP)} \quad \text{for m} = 1 \dots \text{M}. \tag{A.12}
$$

Equation (A.12) by no means requires the codes themselves to be identical from IPP to IPP, only basically that their magnitudes are equal, so a set of well-behaving phase codes is allowed. Where (A.12) will in practice typically go wrong, even in a fully phase coherent radar system, is when the target itself introduces an extra (other than Doppler), fast varying phase to the echo signal. Another way to destroy the necessary phase coherence, in the flexible EISCAT system at least, is to perform sufficiently "clever" frequency hopping from IPP to IPP.

For a sufficiently coherent system (target+radar), it follows from (A.10) and (A.11) that

$$
\text{GMF}(R, \omega_{\mathrm{D}}) \propto ML. \tag{A.13}
$$

This says that the maximum value of GMF is proportional to the total length of signal, and signifies coherent integration.

In this section, we have assumed that there is only one frequency in the transmission. Most of the treatment and conclusions of this section carry over unmodified to the multi-frequency case. However, we did silently ignore a couple of phase terms in Eq. (A.6). In the case of a single frequency this was OK because those factors can be factored out of the integral, and have magnitude equal to unity. In the case of multiple frequencies, there exists

the possibility of "interference", and we need to proceed more carefully. In particularly, we will show that, while all the rapid time variations listed in Eq. (A.1) will still cancel out, the GMF possesses an interference-type oscillation in the range direction. This interference pattern in range is ultimately due the fact that the phase path (distance in terms of wavelength) to a given range r changes by different amounts for the different frequencies, when r is varied. This interference is purely an artifact of the GMF method. It comes about only when the GMF is computed. There of course cannot be any physical interference, because the transmitted pulses are spatially separated (also) after reflection from the point target.

A.3   GMF with multiple frequencies

We model the reception as an attenuated replica of the transmission, shifted in time by the pulse propagation delay 2R/c,

$$S_n(t) = bT_n(t - 2R/c). \tag{A.14}$$

This is certainly a valid expression for a stationary point target. The multiplier b may be a complex number, and includes all the usual stuff from the radar equation. Target motion, $R = R(t)$, will change the length of the phase path, a change that shows up as the Doppler-velocity. Crucially, in the following we will use (A.14) also for a such a (slowly) moving target to track the phase change.

For the short period of time - a few IPPs at most - that we are interested in, the target range does not change much so the factors coming of the radar equation will not change significantly. Therefore we will take b to be a constant.This might not be a good approximation for real targets due to their internal properties such as rotation. The constant b could also be channel dependent, but as it will anyway assumed to be constant, the dependency can be ignored.

We take time $t = 0$ to correspond to the lower integration limit of the GMF integral, and assume constant radial velocity v_0, thus taking

$$R = R(t) = R_0 + v_0 t. \tag{A.15}$$

Inserting $R(t)$ from (A.15) into (A.14) and using the representation (A.2) of $T(t)$ gives

$$
\begin{aligned}
S_n(t) &= b \cdot \epsilon_n(t - 2R/c) \cdot e^{i\omega_n(t - 2R/c)} \\
&= b \cdot \epsilon_n(t - 2R/c) \cdot e^{i\omega_n t} \cdot e^{-i\omega_n \cdot \frac{2R_0}{c}} \cdot e^{-i\omega_n \cdot \frac{2v_0}{c} t}. \tag{A.16}
\end{aligned}
$$

For a short interval of time, we can replace $R(t)$ inside the envelope $\epsilon_n(t)$ by an average value, or simply by R_0. This approximation suffices here when we just want to inspect the general properties of GMF. Those properties

depend more on the phase factors than the precise shape of the envelope. The exponent of the last phase factor in the expansion (A.16) includes the target's Doppler-frequency

$$\omega_{\mathrm{D},n} = \omega_n \cdot 2v_0/c. \tag{A.17}$$

The Doppler-frequency is proportional to the transmission frequency, and thus is in principle different for the different channels. However, the channel frequency difference in EISCAT is a few MHz at most, and the transmission frequency is on the order of a GHz. Thus, also for the Doppler-frequency, we have $\Delta\omega_{\mathrm{D}}/\omega_{\mathrm{D},n} \sim 10^{-3}$, and we can ignore the difference.

With these notations and cosmetic approximations, we can write the reception (A.16) in the form

$$S_n(t) = b \cdot e^{-i\omega_n \frac{2R_0}{c}} \cdot \epsilon_n(t - 2R_0/c) \cdot e^{i\omega_n t} \cdot e^{-i\omega_{\mathrm{D}}t}. \tag{A.18}$$

The other factor that we need in the GMFintegral, $T_n(t - 2r/c)$, is also expanded based on the definition (A.2), but this time the range parameter is kept fixed. We get an expression that looks like a simplified form of Eq. (A.18),

$$T_n(t - 2r/c) = e^{-i\omega_n \frac{2r}{c}} \cdot \epsilon_n(t - 2r/c) \cdot e^{i\omega_n t}. \tag{A.19}$$

The quantities $T_n(t)$ and $S_n(t)$ above represent the transmission and reception on the RF level. Nevertheless, we have treated them as complex (detected) signals right from the beginning. Detection just removes either the negative-frequency spectral component or the positive-frequency component from the original real signals. As a reminder that the GMF method really operates on the detected signals, on some intermediate frequency (IF) level, we at this point include the frequency-shift from transmission frequencies down to IF by a some local oscillator frequency ω_{LO}, by replacing

$$S_n(t) \quad \rightarrow \quad S_n(t) \cdot e^{-i\omega_{\mathrm{LO}}t} \tag{A.20}$$

$$T_n(t) \quad \rightarrow \quad T_n(t) \cdot e^{-i\omega_{\mathrm{LO}}t}. \tag{A.21}$$

These changes add an extra factor $e^{-i\omega_{\mathrm{LO}}t}$ to the r.h.s. of Eq. (A.18) and an extra factor $e^{-i\omega_{\mathrm{LO}}(t-2r/c)}$ to the r.h.s. of Eq. (A.19). Note that the frequency shift by ω_{LO} is not channel dependent. Taking these added factors into account, we get the following expression for the most general product $S_n\overline{T_m}$ that can appear in the GMF integral in the case of multiple frequencies.

$$\begin{aligned} S_n\overline{T}_m e^{i\omega t} &= b \cdot e^{-i\omega_n \frac{2R_0}{c}} \cdot \epsilon_n(t - 2R_0/c) \cdot e^{+i\omega'_n t} \cdot e^{+i\omega_{\mathrm{D}}t} \cdot e^{i\omega t} \\ &\times e^{i\omega_m \frac{2r}{c}} \cdot \overline{\epsilon_m(t - 2r/c)} \cdot e^{-i\omega'_m t} \cdot e^{-i\omega_{\mathrm{LO}}2r/c}. \end{aligned} \tag{A.22}$$

We denoted by ω'_n the IF frequency $\omega_n - \omega_{\mathrm{LO}}$ of channel n. The time-dependent part of (A.22) is

$$e^{i(\omega'_n - \omega'_m)t} \cdot e^{i(\omega - \omega_{\mathrm{D}})t}. \tag{A.23}$$

In the area of (r, ω)-plain which we are interested in, ω will be of the order of ω_{D}, a few kHz. Instead, the frequency separation of the transmission frequencies is typically of the order a few hundred kHz. Under these condition, the first term in Eq. (A.23) oscillates much more rapidly than the second term, if $n \neq m$, and greatly suppress the value of the integral. This means that, to more or less good approximation, we may ignore those terms where $n \neq m$. With this assumption of approximate orthogonality between the frequency channels, we can compute the GMF from

$$\mathrm{GMF} = \left| \sum_n \int dt S_n \overline{T}_n e^{i\omega t} \right|. \tag{A.24}$$

Inserting $S_n \overline{T_n}$ from (A.22) into (A.24), cancelling the terms $e^{\pm i\omega'_n t}$ and noting that $|e^{-i\omega_{\mathrm{LO}} 2r/c}| = 1$, we get

$$\mathrm{GMF}(r, \omega) = b \left| \sum_n \left\{ e^{-i\omega_n \frac{2(r-R_0)}{c}} \int dt\, E_n(t, r, R_0) \cdot e^{i(\omega - \omega_{\mathrm{D}})t} \right\} \right|. \tag{A.25}$$

We denoted by E_n the autocorrelation product of the transmission envelope on channel n,

$$E_n(t, r, R_0) = \epsilon_n(t - 2R_0/c) \overline{\epsilon_n(t - 2r/c)}. \tag{A.26}$$

Within the approximations leading to (A.25) we note that even in the multi-frequency case, the GMF has the absolute maximum at the point $(R_0, \omega_{\mathrm{D}})$. Clearly, (A.25) is consistent with Eq. (A.7) that we got earlier for the single frequency case. However, now we have also got an extra range-dependent factor in front of the integral. This factor factors out if there is only a single frequency, but gives an interference pattern in range if there are multiple frequencies. In the general case, the interference pattern depends both on the transmission envelopes and the frequencies, but if all envelopes are identical, $E_n = E_1$, the integral can be taken out of the sum in (A.25), and the GMF can be written as the product

$$\mathrm{GMF}(r, \omega) = b \cdot \rho(r - R_0) \cdot \left| \int dt E_1(t, r, R_0) \cdot e^{i(\omega - \omega_{\mathrm{D}})t} \right|, \tag{A.27}$$

where

$$\rho(r - R_0) = \left| \sum_n e^{i\omega_n \frac{2(r-R_0)}{c}} \right|. \tag{A.28}$$

For example, for the ESR tau0 experiment we have two identical transmissions in an IPP. In this case

$$\rho = \left| 2 \cos[(\omega_1 - \omega_2)\frac{r - R_0}{c}] \right|. \tag{A.29}$$

In this two-frequency case, $\rho(r)$ is periodic, with the period Δr determined from

$$|\omega_1 - \omega_2| \cdot \frac{\Delta r}{c} = \pi \qquad (A.30)$$

as

$$\Delta r = \frac{c}{2|f_1 - f_2|}. \qquad (A.31)$$

For example, if the channel frequency separation $|f_1 - f_2|$ is 250 kHz as in our November 1999 data, $\Delta r = 600m$. In the cut $\text{GMF}(\omega = \omega_D, r)$, Eq. (A.27), this rapid oscillation of $\rho(r)$ is superimposed on top of the much slower variation of the magnitude the autocorrelation of the transmission envelope.

Appendix B

Experiment cp1lt

B.1 Modulation and timing

The cp1lt transmission consists of unmodulated 350 μs long Long Pulses (LP) and 16×21 μs (16 bits, 21 μs duration) Alternating Codes (AC). In addition, the transmission contains a few short, unmodulated pulses, which we have ignored in the SD search. Fig. B.1 shows the timing details of the cp1lt transmission. The frequencies F_n given in the figure are the frequencies that we used during measurements in February 19.

In cp1lt, the transmission-reception periods (IPPs) are of two different lengths. The IPPs that start with the transmission of a LP have duration 5201 μs while the IPPs containing the ACs have duration 4663 μs. The transmission frequencies are used in such a way that the shortest period where the transmission repeats identically is 19530 μs for the LPs, and much longer for the ACs where the phase code bit pattern cycle is 2×32 IPPs long. This means that target ranges up to 19530×0.15 km $= 2930$ km can be determined unambiguously with both modulation. Figure 3.2 shows raw data from a few consecutive IPPs.

B.2 Matching windows and blind zones

When the radar is transmitting, no reception can occur. In addition, EISCAT experiment designers use rather long periods around the actual transmission where the receiver is protected, that is, isolated from the antenna. These *transmission windows* are marked in Fig. B.1. There are altitude zones from which we cannot have complete information, because part of the the echo is cut off by a transmission window. We call these zones the *blind zones*, and conversely, refer to the altitude zones from which we can receive a SD echo completely, the *matching windows*. The latter name comes about because in a matching window, the whole transmission can conveniently be used in the GMF. Because in cp1lt both the IPPs and the transmissions have slightly different lengths, the matching windows are located somewhat differently for different modulations. The situation is shown in Fig. B.2.

With our SD receiver, we receive and record continuously, picking input from the transmission sample channel when that channel is detected active, and from the reception channel at other times. With this data it is possible formally to compute the ratio $\mathcal{R}$ for any modulation at any time (range), also outside the matching windows. We present such an overview in Fig. B.3, where we plot the ratio $\mathcal{R}$ for LPs and ACs separately. We compute these code-wise ratio curves by using in the GMF only the appropriate piece of the full transmission.

Matching a transmission against another transmission in the blind zone shows up in Fig. B.3 as an increase of $\mathcal{R}$. It is not clear why this happens, and especially, why the two curves are so different. The two involved frequencies

should be widely different, and no actual match should take place.

When the transmission to be matched starts to overlap a transmission window, it sees the greatly reduced noise level there. This shows up as $\mathcal{R}$ decrease, most clearly in the AC curve towards the end of the lower matching window. It would in principle be possible to increase the altitude coverage a little by making use also of this "partially blind" region, but we have not done so in this work.

There is a hard target echo visible in Fig. B.3, towards the end of the upper matching window. The scan in Fig. B.3 is near the maximum of the second event in Table B.1.

B.3   Handling the noise injection

When we match a transmission against a noise injection burst we find an increase in $\mathcal{R}$, shown in Fig. B.3. In cp1lt, EISCAT injects a 330 μs long burst of wide band noise of known temperature into the receiver front end, 2.63 ms after the start of each AC transmission, as can be seen Fig. 3.2. EISCAT uses the noise injection for receiver gain calibration. Also we should ultimately use it, but in this work we have instead relied on the presumably known system noise temperature as the means of finding signal power in physical units. Conversion of signal power to physical units anyway is needed only when estimating the effective target diameter.

On the face of our definition of the ratio $\mathcal{R}$, Eq. (2.38), it may look surprising that an increase of the noise power in the denominator should lead to an increase of the ratio. Indeed, this only happens because we have oversimplified our measurement model. In our derivation of the GMF formalism, we explicitly assumed the noise to be stationary. Therefore, our software estimates the noise power as the average of the squared magnitudes of the first N samples of the matching window,

$$P_\gamma = \frac{1}{N} \sum_{n=j_1}^{j_1+N-1} |m_n|^2, \tag{B.1}$$

where j_1 refers the first gate in the matching window and N is the number of signal samples. (It is then also hoped that those N points do not contain much signal contribution, but this is a separate issue, see below.) Clearly, this is not the proper way to go if the noise, in actual fact, is not stationary.

In Fig. B.3 the noise injection bump shows up only in the first matching window in the alternating code curve, and only in the second matching window in the LP data, as it should. Also note that the LP forms a narrower bandwidth filter than the AC, and so the ratio curve is more smooth for the by LP filtered data.

The noise injection bumps present a problem for $\mathcal{R}_{\mathrm{max}}$-based detection and target size estimation. Suppose the maximum of the noise injection

bump has $\mathcal{R} = 4$, while the background ratio level is about 2. This means that the detection threshold must be set at least to a value around 5, in order to avoid a lot of false alarms near the top of the bump. Suppose that halfway up the slope of the bump there is a promising spike that rises 2 units from the edge of the bump, to a maximum $\mathcal{R} = 4.5$. What should we conclude about the spike? If we conclude it is a target, what strength do we assign to it? How do we force an automatic threshold detector to notice the small spike in the shadow of the large bump?

Our off-line detection software circumvents the noise injection problem by estimating the shape of the bump from a large number of scans, subtracting the mean bump from all scans, and applying the detector only to this cleaned data. The procedure probably is fine from the detection point of view, but did we distort the size estimate? Whether this solution is right or wrong surely is no big issue as long as we are only accepting very clear events as the analysable events. But if we someday would want to work nearer the noise floor, say to count really weak targets, these matters must be addressed seriously.

B.4   Detection setup

We have scanned about 10000 seconds of continuous recording of the cp1lt data from February 19, 2001. The sampling interval was 0.6 μs in this case, so this set contains 16.7×10^9 complex samples, taking 62 Gbytes on disk.

The range coverage of the SD search is shown in Fig. B.3. Between altitudes 396 km and 1377 km, only the interval 682–785 km could not be searched. To get this coverage, we scanned the whole data set three times.

1. Both the LPs and the ACs were used to scan the range intervals 396–666 km (386–649 km altitude) and 855–1413 km (833–1377 km).
2. The higher window was extended downwards by matching only with ACs on interval 806–900 km (785–877 km).
3. The lower window was extended upwards by matching only with the LPs, on range interval 600–700 km (585–682 km).

Other scan parameters were set in the following way.

▷ Range step was 15 gates, 9 μs, which seems to be small enough for the 21 μs bit length of the alternating codes (see Fig. 2.11).
▷ Length of coherent integration was 64 IPPs, 0.31 s.
▷ Searched velocity interval was $-5 \ldots + 5$ km s^{-1}.
▷ Acceleration: For each range j, only a single value $\alpha(j)$ was used in GMF. The acceleration was computed from Eq. (2.63). The antenna elevation was 77.1°, so we simply used the range in place of the altitude in the formula. Even though the acceleration was thus not measured

in individual scans, a (mostly bad) estimate for it was still found sub-
sequently in most cases, from the change of the radial velocity from
scan to scan.

▷ Scan period was 0.5 s, meaning that after each scan, 0.2 s was skipped
before starting next scan. All the detected targets stayed more than
a second about the detection threshold. Therefore, we do not think
that the 0.2 s gaps caused many events to be missed.

▷ At the time of our campaign, the UHF transmitters were not operating
perfectly. The transmission was rather stable most of the time, but
the maximum power was estimated to be slightly less than 1 MW only.
We have used this single value in all estimates of the effective target
diameter.

B.5   Detection

The detector software (programs SPDSCAN and SPDVIEW), with threshold
5, found 52 events. Of these, we deemed seven to be not analysable. In
most cases the rejection was because the target obviously was outside the
matching window and triggered the detector only because of the rather large
width of the detection peak. Raw data for the remaining 45 events were
first moved to a separate place, using the archiver subsystem of SPDVIEW.
A suitable length of raw data, from about 5 to about 30 seconds per event,
was transferred. This reduced the data size needed to be kept online from
62 Gbytes to about 3 Gbytes.

B.6   Parameter estimation

After detection, the analyser program (SPDVIEW) was used to re-scan the
event data around the known range with full resolution and without gaps in
time. We used both the FastGMF and the standard GMF algorithms. The
standard GMF algorithm involved a complex 520800 points FFT for each
range gate. The FastGMF algorithm squeezed this number down to 3712.
Analysis with the full GMF was desperately slow. The analyser produced
2+2 summary plots per event, which we saved to the same directory as the
event's raw data. The event summary plots for all events are available via
our SD web page at the URL `www.sgo.fi/~jussi/spade/`.

Data quality plot

Example of the first type of event summary plot is in Fig. B.4. It shows
the results of the scan which corresponds to the maximum ratio $\mathcal{R}_{\max}$ of
the event. The time of the scan, which we more precisely define to be the
start time of the 0.31 s data segment used for the scan, is taken as the time

label for the event in event listing (Table B.1), and as the time origin in the time series plots (Fig. 3.5). The absolute accuracy of the time label is of the order of 10 ms.

The top data panel of Fig. B.4 shows the first few IPPs of the raw data of the scan. The next panel shows the ratio $\mathcal{R}(\text{r})$ near its maximum at $r = R_0$. The bottom panel shows the velocity slice of the GMF through its maximum position. That is, the panel displays the function

$$\omega \mapsto \frac{\text{GMF}(R_0, \omega, a_{\text{O}})}{\mathcal{P}_{noise}}, \tag{B.2}$$

with the Doppler-frequency ω converted to velocity units. The $a_{\text{O}} = a_{\text{O}}(R_0)$ is the circular-orbit acceleration value.

The noise power estimate (B.1) uses N points starting from the first point in the search window. In the high resolution re-scans, the search window is narrow and is more or less centred to the target position, as shown by the vertical lines in the top panel of Fig. B.4. The estimate $\mathcal{P}_{noise}$ will contain signal, and the ratio $\mathcal{R}$ and the velocity slice (B.2) will become systematically underestimated. For high signal levels, this suppression gets serious and must be removed, or else it affects the target size estimate. For the time series plot in Fig. 3.5, there is a simple remedy: the analyser uses the minimum of the individual noise estimates. For properly chosen data set, there will be a few scans which contain no signal, and they will yield an uncontaminated estimate of the noise.

Parameter estimation plot

An example of the second type of analyser plots is in Fig. 3.5, see Section 3.2.5.

The FastGMF algorithm introduces errors to the parameter estimates, very prominent in $v_{\text{D}}(t)$ in Fig. 3.5, but also affecting $r(t)$. Moreover, FastGMF systematically underestimates the ratio. Making a fit to several scans hopefully improves the estimates of r and v_{D}.

B.7 Figures and tables

Listing of cp1lt analysis results is in Tables B.1, B.2 and B.3.
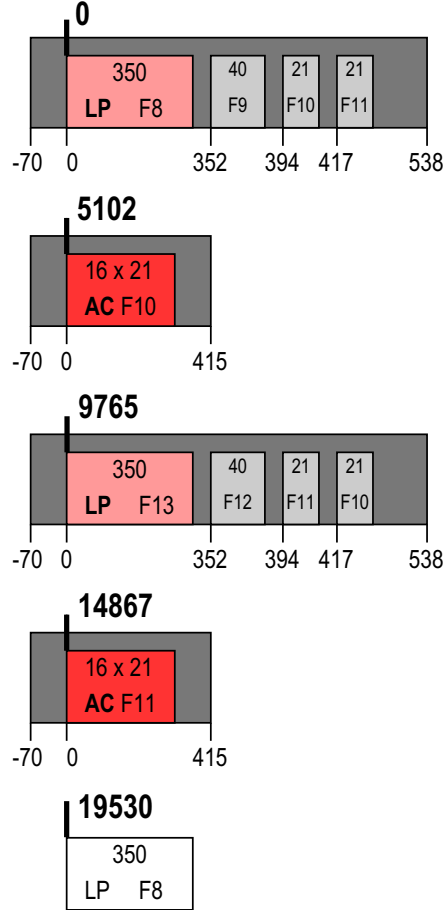
Figure B.1: Transmission in experiment cp1lt, 19 February 2001. The draw-ing is not to scale. The top four panels show transmission during the first four interpulse periods of the transmission loop (times are in μs). The IPPs have lengths 5102 μs, 4663 μs, 5102 μs, and 4663 μs. After this 19.530 ms cycle, the transmission repeats identically, expect that the alternating code phase-flip pattern keeps changing over a set of 32 different codes. The full loop is 32×19.530 ms $= 0.625$ s. The coherent integration is done in integer multiples of the 19.530 ms cycle. The dark gray area around the transmission pulses indicates receiver protection, during which receiver is isolated from the antenna. In addition to the 350 μs long Long Pulses (LP) and 336 μs long Alternating Codes (AC), the transmission contains short unmodulated pulses, but we have not used them in SD search.
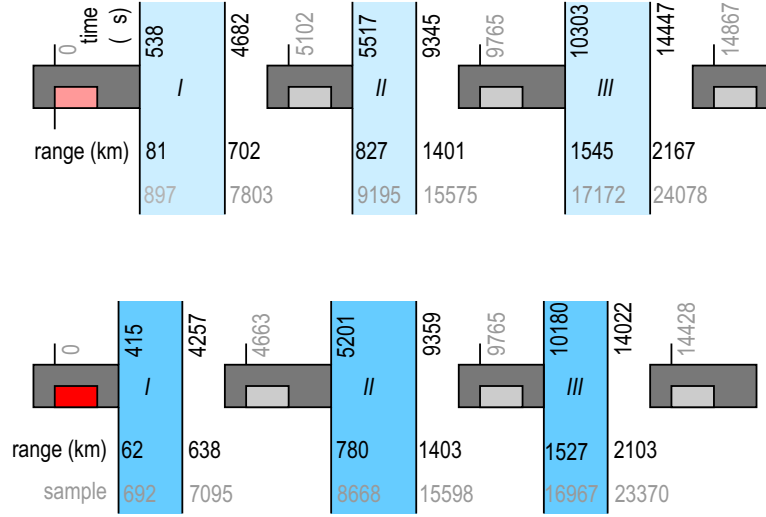
Figure B.2: Space debris matching windows (tall blocks, blue) in experiment cp1lt. Note that the drawing is not to scale. The top panel shows the matching windows when only Long Pulses are matched, the bottom panel when only Alternating Codes are matched. When both are matched, the usable window is the intersection of the LP and AC windows: lowest window is from 81 km to 638 km in range, next window from 827 km to 1401 km in range. If only alternating codes are matched, the seconds window starts already at 780 km range. Note that the high altitude boundary of each window is placed a pulse length before the receiver protection kicks in. The boundary could be moved slightly forward by matching only part of the pulse, but then one would need to compensate for the reduced weight. Also, some 200 μs from the low altitude boundary of each window is contaminated by strong clutter, and has been skipped in actual SD search. The last row of numbers in each panel give the sample numbers or "gate" numbers of the window boundaries when sampling interval is 0.6 μs, to top row gives the corresponding times in μs, counted from the beginning of the respective transmission.
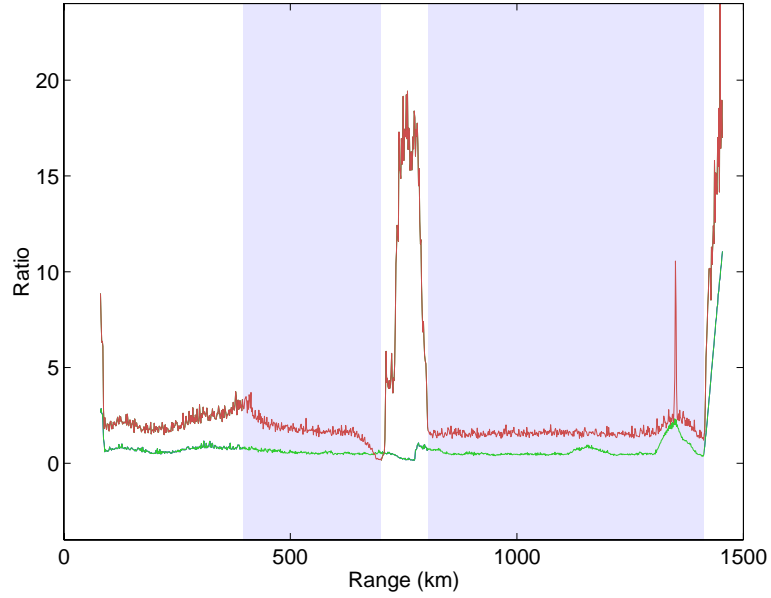
Figure B.3: Actual SD search windows in cp1lt (blue strips). The lower window covers ranges 396–666(–700) km and the upper window covers (806–)855–1413 km. The range shown in parentheses was matched only with the long pulses in the first window and only with the alternating codes in the second window. Other parts where matched with both set of codes. The top curve (red) shows matching with only the alternating codes, the bottom curve (green) matching only with the long pulses. The lower matching window is rather heavily 'corrupted' due to alternating codes seeing the noise injection — there is noise injection only in the IPPs which transmit alternating codes. The FastGMF algorithm was used in this example. The data comes from near the maximum of the SD event at 22:23:48.3, 19 February. There is a debris echo visible at about range 1350 km.

19-Feb-2001 22:23:48.4 cp1lt_A25_fi@uhf

SD20010219_222344_00006.927317<520800>

FastGMF(10) Tc=0.31 s dR=0.09 km

R=1350.3 km

$(a=31.4 \text{ m s}^{-2})$

$v=1.25 \text{ km s}^{-1}$

p=26.5 d=4.5 cm

Figure B.4: The centre scan of an event in experiment cp1lt, 19 February 2001

Table B.1: SD events, cp1lt, 19 February 2001, 22:23–24:00.

UT is time of maximum signal strength, R is range at the given UT, v_D is the Doppler-velocity, positive away from the radar, d_{eff} is the effective target diameter. RR is range rate (dR/dt), a_D is Doppler-drift dv_D/dt, a_O is circular-orbit acceleration at R. The final column gives the event's event number in the compilation of summary plots, 'ac' means matching only with alternating codes. There are two rows for each event: the first row shows the event analysed with FastGMF, the second row shows the event analysed with GMF.

The list continues in Table B.2.

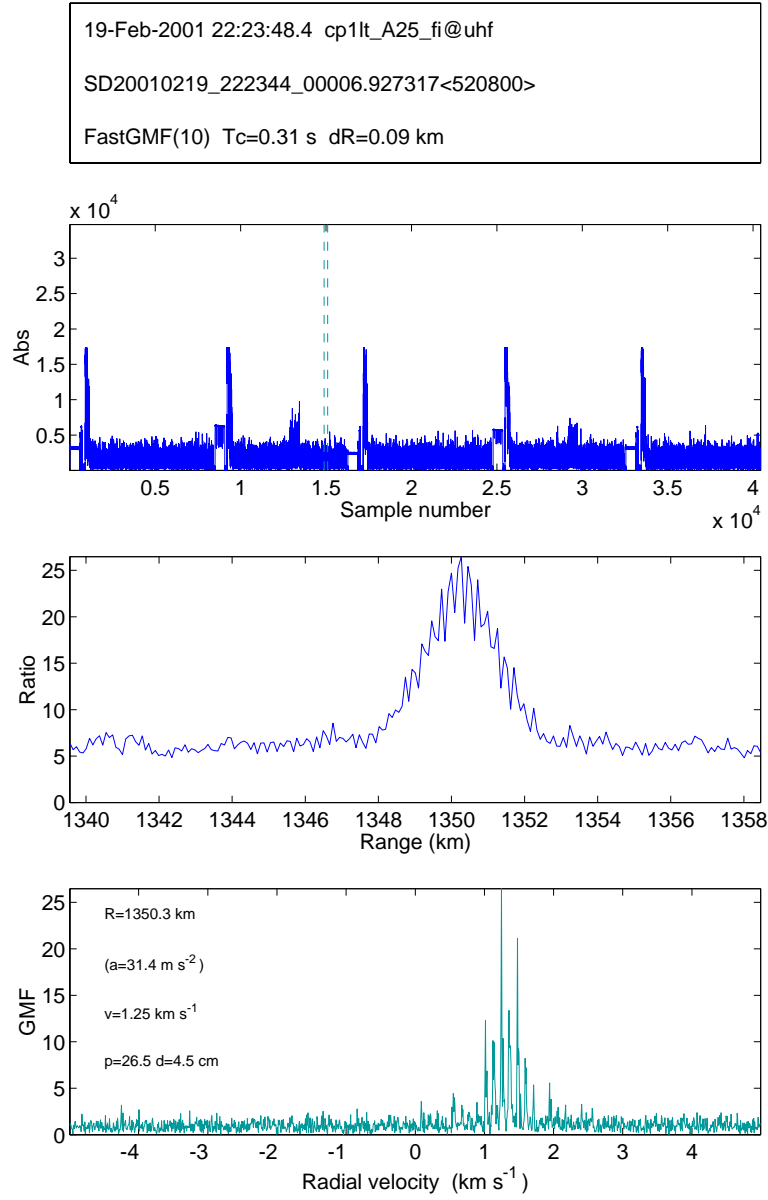| UT		R	v_D	d_{eff}	RR	a_D	a_O	Ref	
		km	km/s	cm	km/s	m/s^2	m/s^2		
2001 02 19	22:23:26.2	821.0	+0.86	3.9	+0.73	82	59.6	1	ac
		820.9	+0.85	3.9	+0.81	52			
2001 02 19	22:23:48.4	1350.2	+1.29	4.6	+1.28	44	31.4	2	
		1350.3	+1.31	4.6	+1.28	24			
2001 02 19	22:25:38.1	1127.6	-1.22	5.4	-1.12	-14	39.9	3	
		1127.6	-1.18	5.3	-1.19	55			
2001 02 19	22:26:03.2	981.8	-1.26	3.4	-1.21	107	47.6	4	
		981.7	-1.26	3.4	-1.21	34			
2001 02 19	22:31:08.6	1142.7	+0.71	3.0	+0.71	123	39.2	5	
		1143.0	+0.68	3.1	+1.14	-8			
2001 02 19	22:31:43.3	951.3	-1.34	3.2	-1.55	31	49.6	6	
		951.4	-1.33	3.2	-1.52	50			
2001 02 19	22:33:24.5	1045.0	+1.50	3.5	+1.50	35	44.0	7	
		1045.1	+1.40	3.6	+1.32	96			
2001 02 19	22:35:54.3	1123.7	+0.89	3.0	+1.06	38	40.1	8	
		1123.6	+0.84	3.0	+0.89	47			
2001 02 19	22:39:45.7	833.4	-0.59	2.9	—	—	58.5	9	ac
		833.5	-0.57	3.0	—	—			
2001 02 19	22:43:22.1	1160.0	-1.22	5.1	-1.54	2	38.4	10	
		1160.0	-1.22	5.2	-1.25	32			
2001 02 19	22:44:13.0	1214.5	+1.53	3.6	+1.57	38	36.2	11	
		1214.3	+1.54	3.7	+1.40	65			
2001 02 19	22:51:01.9	670.8	+1.37	2.4	+3.06	62	75.9	12	lp
	:50:58.7	661.5	+1.20	2.4	+3.07	68			
2001 02 19	22:54:18.7	506.3	+1.20	3.5	+1.13	105	105.6	14	
		506.3	+1.21	3.5	+1.22	104			

Table B.2: SD events cp1lt, 19 February 2001, 23:00–24:00.

The list continues in Table B.3.

UT			R	v_D	d_{eff}	RR	a_D	a_O	*Ref*	
			km	km/s	cm	km/s	m/s^2	m/s^2		
2001 02 19	23:02:54.7		1088.2	-1.22	13.5	-1.14	58	41.8	16	
			1088.2	-1.22	14.5	-1.22	33			
2001 02 19	23:06:01.1		933.3	-1.35	3.5	-1.27	99	50.8	17	
			933.3	-1.38	3.5	-1.22	18			
2001 02 19	23:18:29.7		812.8	-1.03	46.5	-0.99	53	60.3	18	ac
			813.0	-1.04	51.1	-1.00	57			
2001 02 19	23:22:03.9		1037.2	-1.10	2.7	-1.32	45	44.4	19	
			1036.9	-1.18	2.8	-1.36	40			
2001 02 19	23:32:27.6		963.1	+0.49	13.0	+0.54	42	48.8	20	
	:27.9		963.4	+0.54	14.8	+0.56	48			
2001 02 19	23:33:07.0		1226.9	-0.31	3.6	-0.35	35	35.7	21	
			1226.9	-0.34	3.6	-0.29	24			
2001 02 19	23:33:43.2		839.8	+1.04	42.3	+0.96	45	57.9	22	ac
			839.8	+0.96	45.2	+0.96	52			
2001 02 19	23:40:34.3		956.8	+1.37	3.7	+1.46	-8	49.2	24	
			956.8	+1.33	3.7	+1.26	18			
2001 02 19	23:43:51.4		872.8	+0.78	4.2	+0.80	52	55.2	25	
			872.8	+0.79	4.2	+0.81	53			
2001 02 19	23:47:49.6		1359.2	-0.97	3.4	-1.04	144	31.1	26	
			1359.2	-1.00	3.5	-1.12	-21			
2001 02 19	23:51:49.8		971.0	-1.30	3.4	-1.29	44	48.3	27	
	:49.5		971.5	-1.35	3.5	-1.15	48			
2001 02 19	23:56:49.0		928.2	-1.47	2.3	—	—	51.1	28	
	:49.3		928.2	-1.47	2.3	-1.30	—			
2001 02 19	23:56:59.3		1391.4	+0.87	4.7	+0.95	33	30.2	29	
			1391.4	+0.91	4.7	+1.00	37			
2001 02 19	23:59:50.2		1033.1	+1.26	14.4	+1.26	36	44.6	30	
			1033.1	+1.22	14.1	+1.27	46			

Table B.3: SD events cp1lt, 20 February 2001, 00:00–01:11.

UT		R	v_D	d_{eff}	RR	a_D	a_O	Ref	
		km	km/s	cm	km/s	m/s^2	m/s^2		
2001 02 20	00:00:51.8	1353.2	-1.33	6.0	-1.35	6	31.3	31	
	:50.6	1354.8	-1.33	5.9	-1.28	21			
2001 02 20	00:01:18.4	884.4	-0.86	4.3	-0.84	54	54.3	32	
		884.4	-0.87	4.4	-0.84	57			
2001 02 20	00:01:51.3	1343.9	-1.33	3.3	-1.31	18	31.6	34	
		1343.9	-1.32	3.3	-1.29	-12			
2001 02 20	00:07:02.6	1329.4	-1.33	3.6	-1.66	0	32.1	35	
		1329.4	-1.34	3.7	-1.44	27			
2001 02 20	00:11:58.1	1351.3	-1.35	4.0	-1.18	14	31.4	36	
		1351.3	-1.34	4.0	-1.20	47			
2001 02 20	00:12:59.4	1021.4	-1.36	46.6	-1.35	44	45.3	37	
		1021.4	-1.36	49.4	-1.36	45			
2001 02 20	00:22:45.1	832.2	+0.97	5.1	+0.99	65	58.6	38	ac
		832.2	+0.97	5.0	+0.98	59			
2001 02 20	00:24:41.8	980.8	-1.36	4.1	-1.34	45	47.7	40	
	:42.1	980.4	-1.35	4.1	-1.32	43			
2001 02 20	00:26:31.9	833.3	+0.69	6.5	+0.63	123	58.5	41	ac
		833.3	+0.64	6.5	+0.63	66			
2001 02 20	00:33:16.2	905.1	-1.34	2.5	-1.41	43	52.8	44	
	:16.8	904.9	-1.32	2.6	-1.42	56			
2001 02 20	00:34:17.5	1118.8	+1.21	4.4	+1.21	7	40.3	45	
	:16.8	1118.0	+1.19	4.4	+1.06	30			
2001 02 20	00:40:29.6	937.4	+0.66	2.8	+0.77	77	50.5	46	
		937.4	+0.69	2.8	+0.68	55			
2001 02 20	00:48:48.2	1067.3	+1.06	4.9	+1.08	65	42.8	47	
		1067.3	+1.07	5.0	+1.10	58			
2001 02 20	00:49:23.4	815.1	-1.17	2.5	-0.84	118	60.1	48	ac
		815.1	-1.17	2.5	-0.81	110			
2001 02 20	00:59:26.8	833.7	+0.94	4.8	+0.96	48	58.5	49	ac
		833.8	+0.93	4.8	+0.95	54			
2001 02 20	00:59:40.0	872.6	-0.83	2.4	—	—	55.2	50	
		872.6	-0.73	2.4	—	—			
2001 02 20	01:01:58.6	890.1	+0.70	3.3	+0.63	20	53.9	51	
		890.1	+0.71	3.2	+0.48	44			
2001 02 20	01:04:13.2	1035.0	+1.28	2.9	+1.54	41	44.5	52	
		1035.0	+1.28	3.0	+1.86	115			

Appendix C

Experiment tau2

C.1   Modulation and timing

The EISCAT UHF tau2 transmission consists of 16×36 μs alternating code, on two frequencies that are sent on alternating interpulse periods, 6510 μs apart. The duty cycle is 8.8%. Timing diagram of the experiment is in Fig. C.1. In this experiment all IPPs are of same length. The blind zones are larger than in cp1lt, the lowest one is between about 880 km and 1080 km. With more careful programming of the scanner, the lower part of this gap could in fact be accessed, but with reduced accuracy.

C.2   Detection and analysis

From the recorded 1.7 hours of tau2 data we analysed the contiguous 1.3 h set 20 February 2001, 21:26–22:40. We searched SD in the two lowest matching windows, starting from 485 km range. In the lower window we used 0.9 s coherent integration, scan repetition period 1.0 s and range step 1.5 km. In the higher window we used coherent integration 0.27 s, scanning period 0.5 s and range coverage 1095–1800 km with 1.5 km resolution (471 gates). We used only a single acceleration value per range gate.

We did the final analysis both with GMF and, for comparison, also with FastGMF, using 0.3 s integration and range steps 0.15 km or 0.075 km, the latter corresponding to the 0.5 μs sampling interval.

We found five clear events in the lower window. We found seven clear events in the higher window, but one of these was actually a "ghost" of the very strong event 5 in Table C.1 from the low-altitude window. The strength of the echo in the lower window ($\mathcal{R}_{\mathrm{max}}= 640$) was more than 30 dB stronger than the ghost echo ($\mathcal{R}_{\mathrm{max}}= 10$). In tau2 experiment, the same phase code pattern K is sent on two consecutive IPPs, although in different frequencies. The frequency difference is only 900 kHz, so it is not inconceivable that K could match the 900 kHz shifted copy of itself on the -30 dB level.

Listing of the tau2 events is in Table C.1. We display graphically various properties of the events in figures 2.2, 3.4, 4.1, 4.2, 4.3, 4.4, 4.5 and 4.12.

Timing diagram
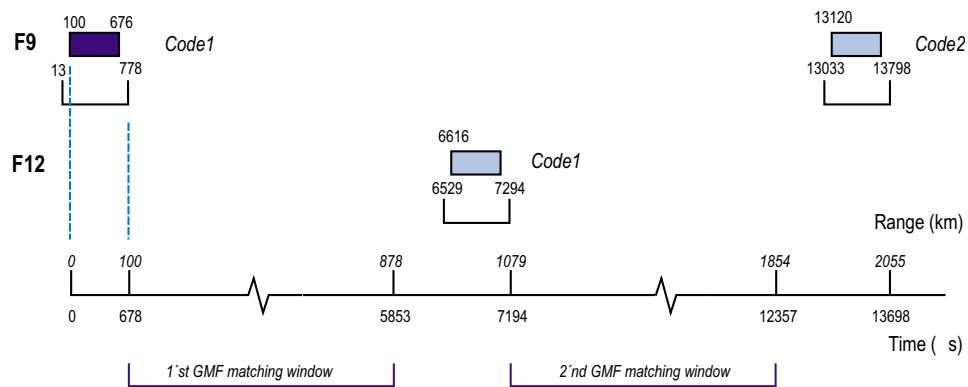
Structure of transmission window
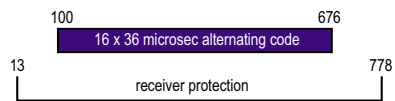
Transmission order

Code1	Code1	Code2	Code2	... Code32	Code32
F9	F12	F9	F12	F9	F12

F9 = 928.7 MHz

F12 = 929.6 MHz

Figure C.1: Tau2 transmission timing diagram.

Table C.1: SD events, tau2, 20 Feb 2001, 21:26–22:40.

UT is time of maximum signal strength, R is range at the given UT, v_D is the Doppler-velocity, positive away from the radar, d_{eff} is the effective target diameter. RR is range rate (dR/dt), a_D is Doppler-drift dv_D/dt, a_O is circular-orbit acceleration at R. The final column gives the event's event number in the compilation of summary plots. There are two rows for each event: the first row shows the event analysed with FastGMF, the second row shows the event analysed with GMF.

UT		R km	v_D km/s	d_{eff} cm	RR km/s	a_D m/s^2	a_O m/s^2	Ref
2001 02 20	21 30 09.4	609.1	+1.47	1.8	—	—	85.2	1
		609.1	+1.47	1.9	—	—	85.2	
2001 02 20	21 31 34.4	656.2	+0.57	2.5	+0.58	144	78.0	2
		656.2	+0.55	2.6	+0.43	92	78.0	
2001 02 20	21 38 02.9	1140.1	-1.17	4.7	-1.12	33	39.3	3
		1140.1	-1.15	4.7	-1.24	38	39.3	
2001 02 20	21 46 50.4	1124.0	-1.15	2.6	—	—	40.0	4
		1123.7	-1.11	2.6	—	—	40.1	
2001 02 20	21 52 39.2	575.5	-1.42	12.9	-1.40	70	91.1	5
		575.5	-1.42	13.3	-1.41	87	91.1	
2001 02 20	22 07 05.0	692.3	-0.81	2.3	—	—	73.2	6
		692.3	-0.82	2.4	—	—	73.2	
2001 02 20	22 16 21.2	1209.4	-1.20	3.4	-1.21	40	36.4	7
		1209.5	-1.21	3.5	-1.26	50	36.4	
2001 02 20	22 19 16.2	1516.5	-1.22	52.0	-1.22	25	26.8	8
		1516.5	-1.22	55.4	-1.21	25	26.8	
2001 02 20	22 23 15.4	1353.1	+1.29	3.1	—	—	31.3	9
		1353.4	+1.32	3.1	—	—	31.3	
2001 02 20	22 25 09.2	836.0	-0.57	2.1	-0.62	6	58.2	10
		836.0	-0.57	2.2	-0.62	32	58.2	
2001 02 20	22 26 45.1	1402.2	-1.05	15.6	-1.04	29	29.8	11
		1402.3	-1.04	15.6	-1.07	28	29.8	

Bibliography

[1] M. Baron *The EISCAT facility.* J. atmos. terr. Phys. 46, 469, 1984.

[2] M. Baron *EISCAT progress 1983–1985.* J. atmos. terr. Phys. 48, 767, 1986.

[3] G. Wannberg, I. Wolf, L.-G.Vanhainen, K. Koskenniemi, J. Röttger, M. Postila, J. Markkanen, R. Jacobsen, A. Stenberg, R. Larssen, S. Eliassen, S. Heck and A. Huuskonen *The EISCAT Svalbard radar: A case study in modern incoherent scatter radar system design.* Radio sci., 32, 2283, 1997.

[4] A. Pellinen-Wannberg and G. Wannberg *Meteor observations with the European incoherent scatter UHF radar.* J. Geophys. Res., 99, 11379, 1994.

[5] A. Pellinen-Wannberg *The high power large aperture radar method for meteor observations.* Proceedings of the Meteoroids 2001 Conference, Swedish Institute of Space Physics, ESA SP-495, 443-450, 2001.

[6] M. Lehtinen *Statistical theory of incoherent scatter radar measurements.* EISCAT Techn. Note 86/45, Eur. Incoherent Scatter Sci. Assoc., Kiruna, Sweden, 1986.

[7] A. Carlson *Communications systems, second edition.* McGraw-Hill, Singapore, 1975.

[8] M.I. Skolnik *Introduction to radar systems, second edition.* McGraw-Hill, Singapore, 1981.

[9] 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.

[10] Morchin *Radar engineer's sourcebook.* Artech House, 1993.

[11] A. Tarantola *Inverse problem theory.* Elsevier Science Publishers, Netherlands, 1987.

[12] A. Westman, Private communication, 2002.

[13] *ESA space debris handbook.*

[14] M. Lehtinen *Method and system for measuring radar reflectivity and Doppler-shift my means of a pulse radar.* US Patent No. US 6,232,913 B1, May 15, 2001.

[15] M. Landgraf, Private communication, 2001.