bima memo 46.

DC Bias in the One-Bit BIMA Samplers

By Jerome A. Hudson, UC Berkeley Radio Astronomy Laboratory

Abstract

This is a quick look at how a slight voltage bias affects the BIMA samplers. For brevity, only the one-bit case has been considered, but it is probably reasonable to assume the same general conclusions hold for the two-bit case. The effect of moderate DC bias on wide-band noise is a slight offset of the spectrum, with mild distortion. The DC bias can be approximately removed using a simple algorithm, with residuals in the spectrum for Nyquist-sampled lags of about 0.5%.

Introduction

The one- and two-bit samplers used in the BIMA correlator necessarily make some assumptions about the voltage of the input signal: (1) that the voltage fluctuations have zero mean, (2) that the fluctuations obey a Gaussian distribution, and (3) that the fluctuations are stationary; that is, that the statistical parameters are constant for a reasonable length of time.

In normal crosscorrelation mode, there is no problem with (1). Any DC is taken out of the signal by switching the polarity of the signal in the Walsh pattern that also reduces the chance for cross-talk between signals, and additionally encodes the quadrature components of the signal, permitting sideband separation. In autocorrelation mode, the Walshing no longer serves either purpose. Consequently, we sometimes observe some DC offsets in the autocorrelations, and frequently also see harmonics of 50 MHz oscillations. Neither effect is very obnoxious, but recently we are having some severe requirements placed on the system as observers attempt to subtract pairs of on/off source spectra.

To illustrate the problem, Fig. 1 shows autocorrelations in which one of the data sets (sixth window) exhibits a DC offset. In the same scan, the last window shows evidence of pickup of some oscillation, as does the first, to a lesser degree. The windows represent 100 MHz bandwidths, with frequency offsets of 100 MHz, beginning at 100 MHz in the I.F. Fig. 2 shows minor DC offsets and oscillations changing in time (over a few minutes). The plots are differences between scans (first minus second, first minus third). The vertical scales in both cases are dimensionless; the uncorrected data have been normalized (divided by the gate counts).

The main question I would like to clear up in this note is whether the drifting DC offsets frequently observed in autocorrelations are distorting the spectral passbands, and whether (without some herculean hardware fixups) a simple correction scheme can be used to correct the problems observed.

Model

Mathematically, correlations of clipped Gaussian noise can be modelled via the joint probability distribution:

$$f_\rho(u,v) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left[\frac{-(u^2+v^2-2\rho uv)}{2(1-\rho^2)}\right]. \tag{1}$$

where ρ is the correlation coefficient, and u and v are random variates with unit R.M.S. The autocorrelation case does not differ from this, since u and v in that case are samples from the same time stream which are partially correlated.

It is helpful to rotate the coordinate system by 45° by way of a change of variables:

$$\begin{aligned} x &= \frac{\sqrt{2}}{2}u + \frac{\sqrt{2}}{2}v, \\ y &= -\frac{\sqrt{2}}{2}u + \frac{\sqrt{2}}{2}v. \end{aligned} \tag{2}$$

Let

$$a = \frac{1}{\sqrt{2(1+\rho)}},$$
$$b = \frac{1}{\sqrt{2(1-\rho)}}.$$

Then (1) becomes

$$f_\rho(x,y) = \frac{ab}{\pi}\exp[-(a^2x^2+b^2y^2)]. \tag{3}$$

In the case of one-bit clipping, each variate is compared to ground, and encoded as either a +1 or −1. Then the clipped correlation between u and v as recorded by our correlator is the expectation value for the various events:

$$\rho_1 = (+1)P_{1,1} + (+1)P_{-1,-1} + (-1)P_{-1,1} + (-1)P_{1,-1}, \tag{4}$$

where, for instance, $P_{1,1}$ is the probability for having event $u \geq 0$ and $v \geq 0$. $P_{1,1}$ is f integrated over the area between the u and v axes, which are lines running at 45° angles in the rotated coordinates:

$$P_{1,1} = \frac{2ab}{\pi}\int_0^\infty dx\exp(-a^2x^2)\int_0^x dy\exp(-b^2y^2). \tag{5}$$

This nicely symmetrical case is readily integrated. We make the change of coordinates to polar, wherein the y axis is stretched by the amount b/a; let:

$$r^2 = x^2 + y^2,$$
$$\frac{b}{a}dxdy = rdrd\theta,$$

so that (5) becomes

$$P_{1,1} = \frac{2a^2}{\pi}\int_0^\theta d\theta\int_0^\infty rdr\exp(-a^2r^2).$$

The second integral evaluates simply to $1/2a^2$, giving us

$$P_{1,1} = \frac{\theta}{\pi}.$$

θ is simply the angle between the transformed v axis and the x axis, whence

$$P_{1,1} = \mathrm{atan}(b/a)$$
$$= \sqrt{\frac{1+\rho}{1-\rho}}.$$

From the complementary angle,

$$P_{-1,1} = \frac{1}{\pi}[\frac{\pi}{2} - \theta],$$

and the other probabilities are equal, respectively, to these. Thus,

$$\rho_1 = 2P_{1,1} - 2P_{-1,1}$$
$$= \frac{4}{\pi}\theta - 1$$
$$\tan[\frac{\pi}{4}(1+\rho_1)] = \sqrt{\frac{1+\rho}{1-\rho}}.$$

The result simplifies nicely if we make use of the trigonometric identity

$$\tan\frac{\phi}{2} = \sqrt{\frac{1-cos\phi}{1+cos\phi}}.$$

The result is the Van Vleck correction for 1-bit clipping:

$$\rho = \sin \frac{\pi}{2} \rho_1. \tag{6}$$

The Van Vleck correction is used in all one-bit computations in the BIMA software. The presence of DC complicates the situation, however. Instead of the symmetric case above, the u and v axes no longer pass through the origin, but are displaced by an amount $c = (\sqrt{2}/2)V_{DC}$. Fig. 3 shows the domain of integration for computing $P_{1,1}$. Contours of the probability distribution are sketched for a case of $rho \approx 0.6$.

Recomputing $P_{1,1}$ in the presence of DC:

$$\begin{aligned} P_{1,1} &= \frac{2ab}{\pi} \int_c^\infty dx \exp(-a^2x^2) \int_0^{x-c} dy \exp(-b^2y^2), \\ &= \frac{a}{\sqrt{\pi}} \int_c^\infty dx \exp(-a^2x^2) \mathrm{erf}(bx - bc). \end{aligned} \tag{7a}$$

Similarly,

$$P_{-1,-1} = \frac{a}{\sqrt{\pi}} \int_{-\infty}^{c} dx \exp(-a^2x^2) \mathrm{erf}(bc - bx), \tag{7b}$$

$$\begin{aligned} P_{-1,1} = \frac{a}{\sqrt{\pi}} \{ &\int_{-\infty}^{c} dx \exp(-a^2x^2) \mathrm{erfc}(bc - bx) \\ &+ \int_c^\infty dx \exp(-a^2x^2) \mathrm{erfc}(bx - bc) \}, \end{aligned} \tag{7c}$$

where

$$\mathrm{erfc}(x) = 1 - \mathrm{erf}(x).$$

These were evaluated numerically using Gauss-Legendre quadrature. The results are illustrated in Fig. 4, where the differences between the computed ρ_1 curves and the Van Vleck case (DC=0) are plotted. The worst case shown has an assumed DC value of 0.15 times the RMS voltage, matching fairly well the worst window in Fig. 1.

There is a slight glitch at ρ=0.99 is due to numerical difficulties evaluating the above integrals for an extremely flattened domain. A 16th-order Legendre polynomial was used in the Gauss integrations, and this borders on the limit of acceptable precision. The curve for 0 DC is a useful numerical check on the computation.

Sample Data

While the curves in Fig. 4 seem to indicate a real problem exists, let us actually investigate what would happen if the worst case model were fitted to the data in window 6, Fig. 1. Data that were taken in the absence of DC (unfortunately collected on a different date, Dec. 1994, from the date the biased data were taken, June 1995) were taken as the error-free input model. The lag data for window 6 were first Van Vleck-corrected, then fed as inputs (true values of ρ) to the 0.15 Vrms model. The output lags did not differ significantly from the uncorrected starting data.

The resulting spectra, both made with Van Vleck corrections only, are plotted in Fig. 5. We see that the primary effect of the DC bias was to shift the spectrum uniformly downward. What saves us is that both curves agree at ρ=1; where they disagree, the spread in ρ_1 is sufficiently small to avoid distorting the lag curve, hence only the shift at ρ=0 is noticed. An effect which can be cured by Hanning smoothing is the cyclic ripple resulting from the DC offset (and failing to sample the $N + 1$st lag point).

What would happen with another common case: oversampling in the time domain? The example in Fig. 5 was taken at the Nyquist rate. But what would happen if we were to take, say, 8 times as many time samples? This would force the lag curve to sample a more or less full gamut of values from about −0.2 to 1.0 and assure the loss of linearity. It turns out that even this case does fairly well at DC = 0.15. The lag curves and spectra are compared in Fig. 6 a and b.

A Cure?

An obvious thing to attempt is to estimate the DC offset by averaging the upper half of the lag domain, and then to subtract $DC(1 - \rho 1)$ from every lag sample. This approximates the curves of Fig. 4 as being slanted arcsines. Trying out the scheme on the above two cases (Nyquist and over-sampled noise spectra) produces the results plotted in Fig. 7a and 7b. In 7a, the unbiased and approximately corrected curves are indistinguishable; a slight difference is discernable in Fig. 7b.

The residuals are plotted as Fig. 8, and from this it is clear that the approximation is only that; we can still have 0.5% differences in the spectra, and worse. This doesn't necessarily mean 0.5% differences in short, successive integrations, but the error is a systematic one that depends upon the severity of the DC offsets.

Figure Captions

Fig. 1 Auto-correlation lag data, Antenna 3, showing DC offset in window 6. (June 1995)

Fig. 2 Auto-correlation lag data, Antenna 3, differences of 3 scans over several minutes. (June 1995)

Fig. 3 Domain of integration for computing $P_{1,1}$ for a DC offset $\sqrt{2}c$.

Fig. 4 One-bit correlator response curves for varying amounts of DC

Fig. 5 Comparison of two spectra, one with simulated DC offset

Fig. 6 Comparison of (a) autocorrelations and (b) power spectra of an oversampled, filtered noise signal.

Fig. 7 Comparison of spectra made with (a) Nyquist-sampled and (b) Over-sampled lags, plotting the "true" spectrum as a solid line, and the approximation to it with short dashes. DC offset was 0.15 Vrms.

Fig. 8 Residuals for the approximate removal of DC bias from the lag domain, after computation of spectra.

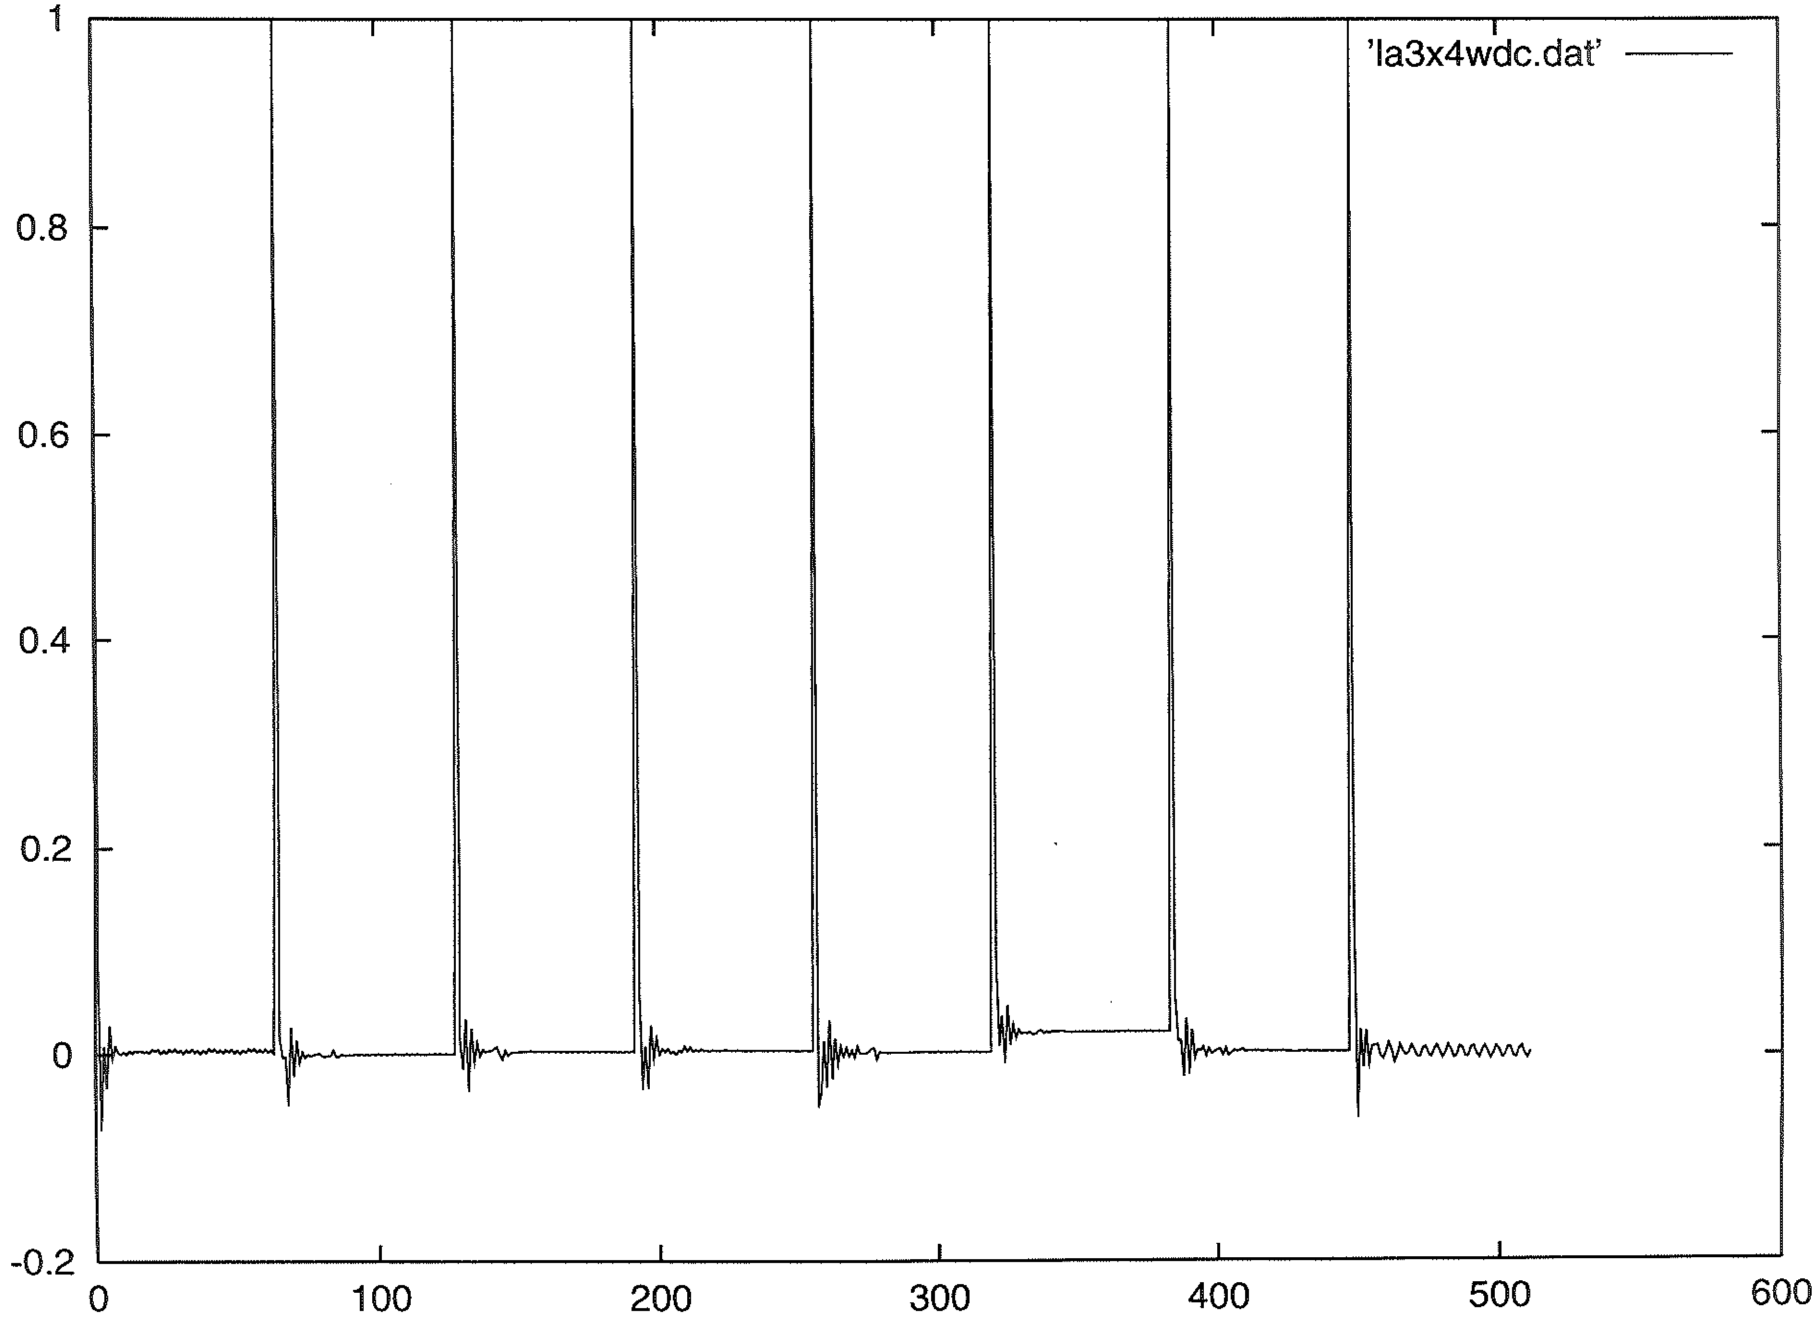

Fig 1

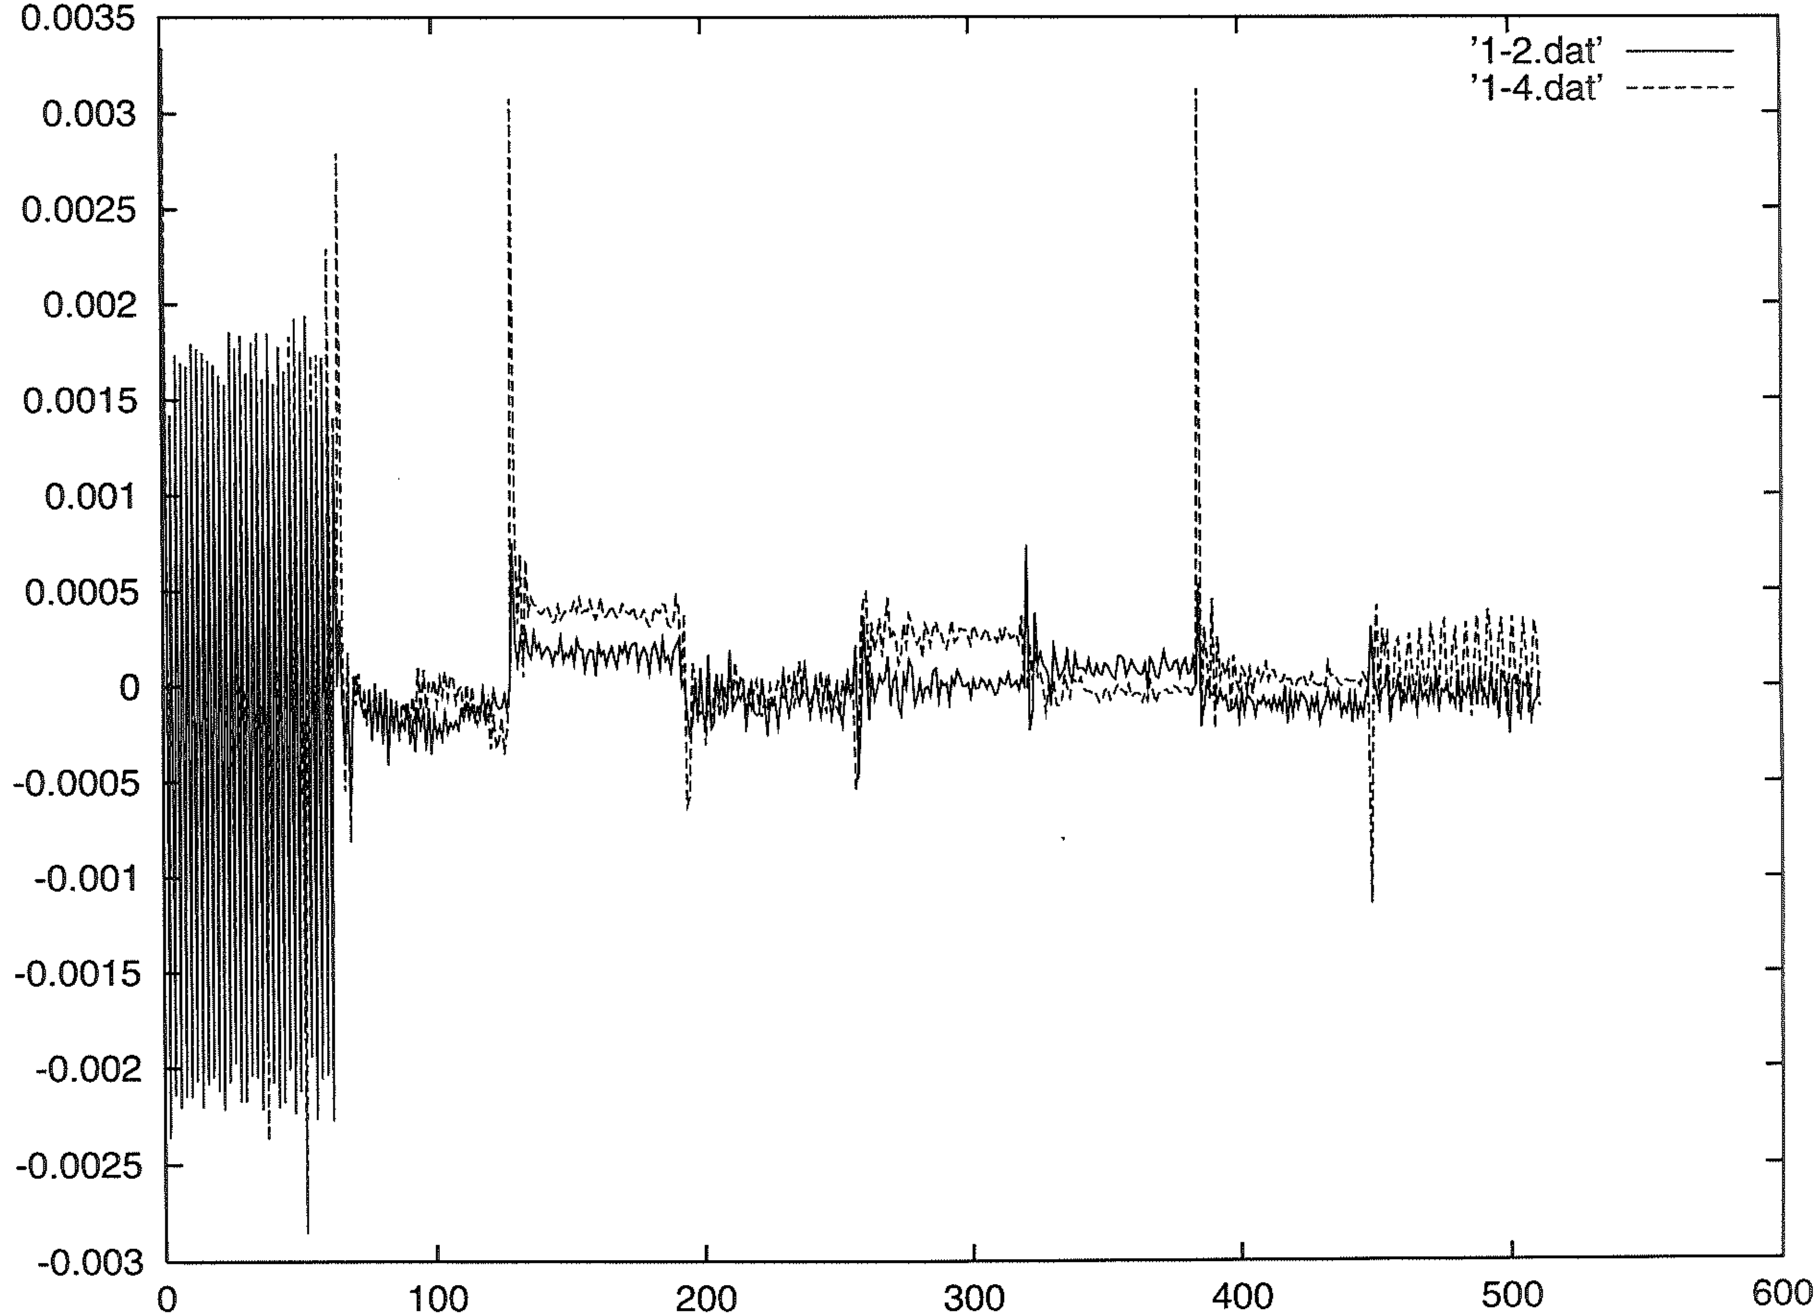
'1-2.dat'
'1-4.dat'
0.0035
0.003
0.0025
0.002
0.0015
0.001
0.0005
0
-0.0005
-0.001
-0.0015
-0.002
-0.0025
-0.003
0
100
200
300
400
500
600

Fig 2

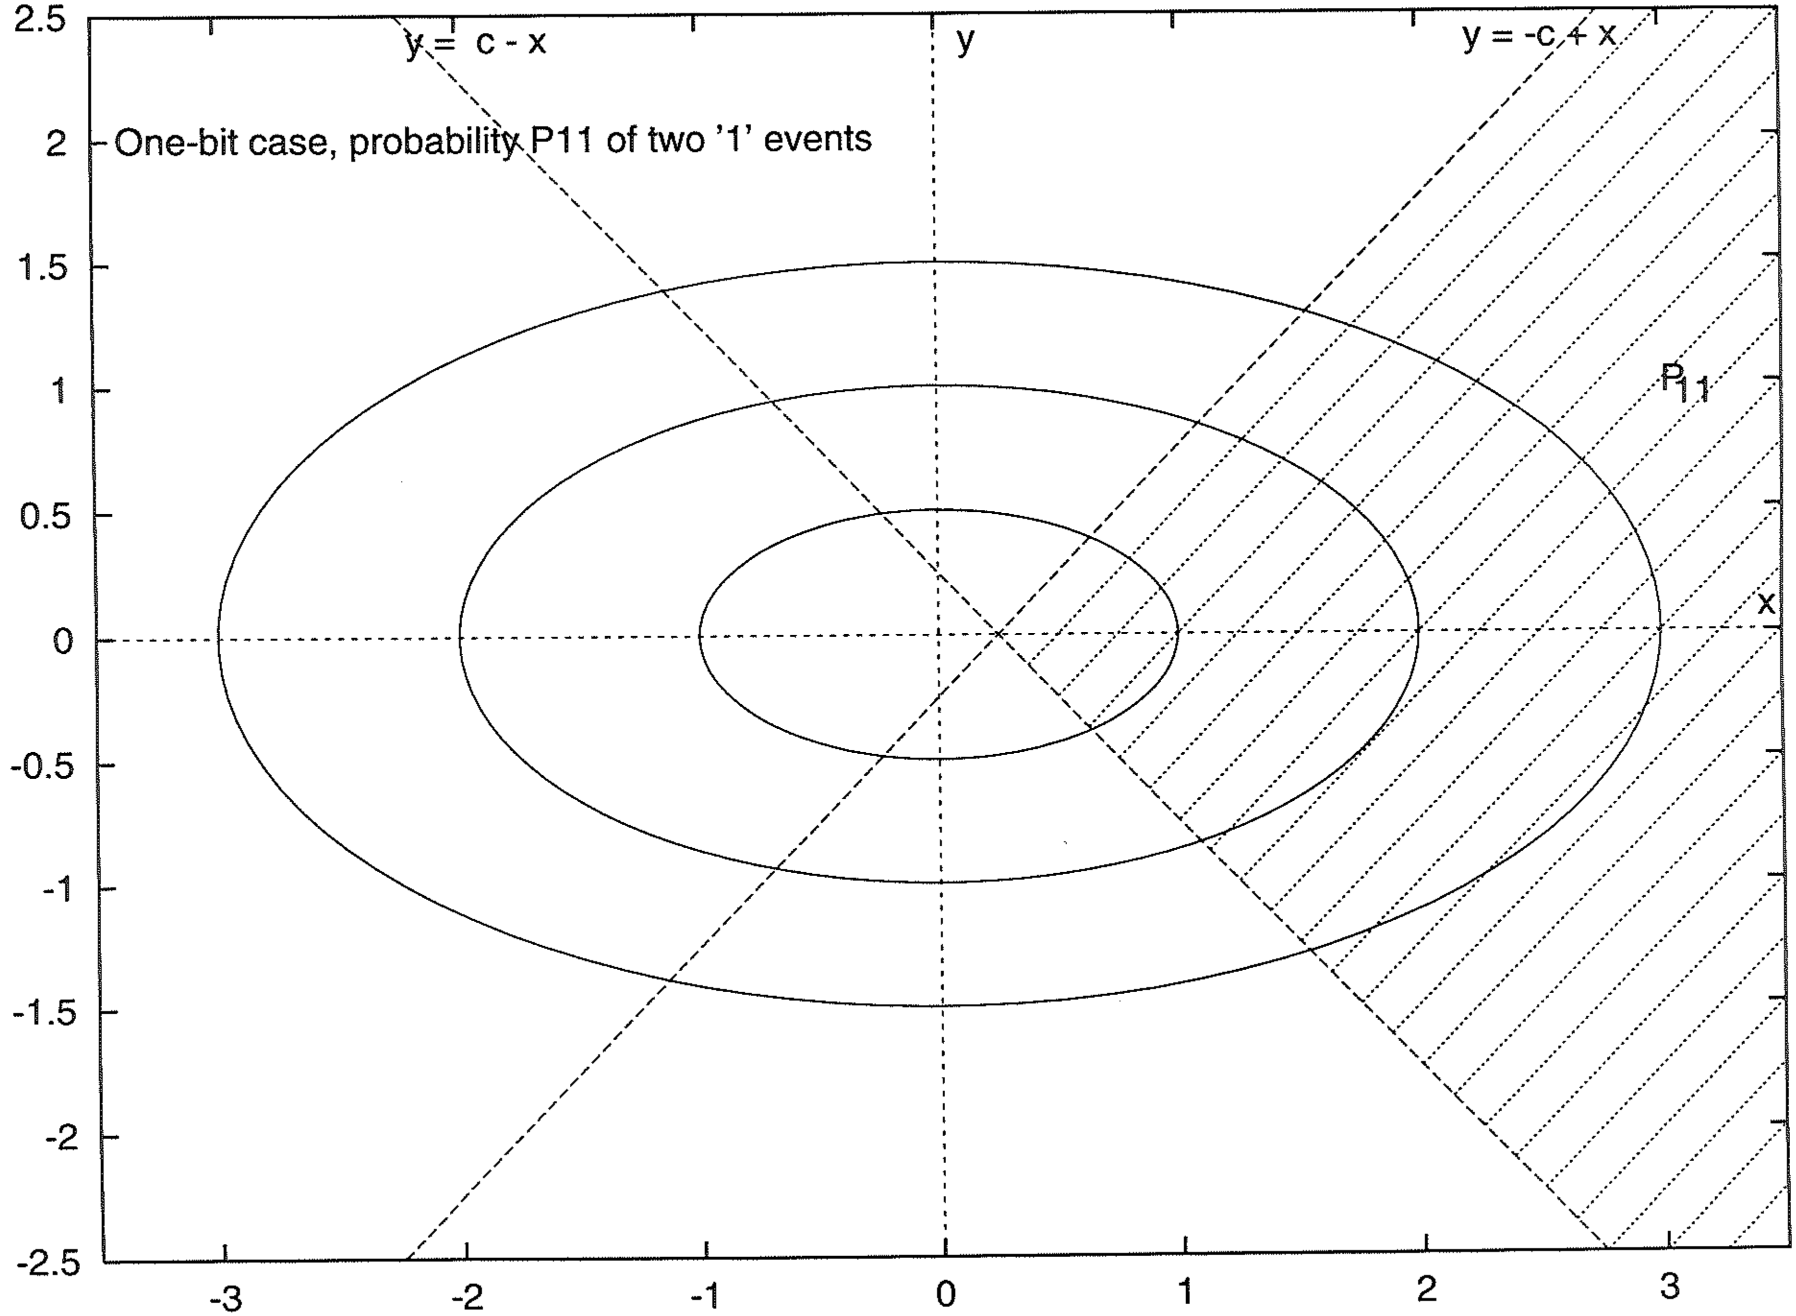
y = c - x
y
y = -c + x
One-bit case, probability P11 of two '1' events
P_{11}
x
2.5
2
1.5
1
0.5
0
-0.5
-1
-1.5
-2
-2.5
-3
-2
-1
0
1
2
3

Fig. 3

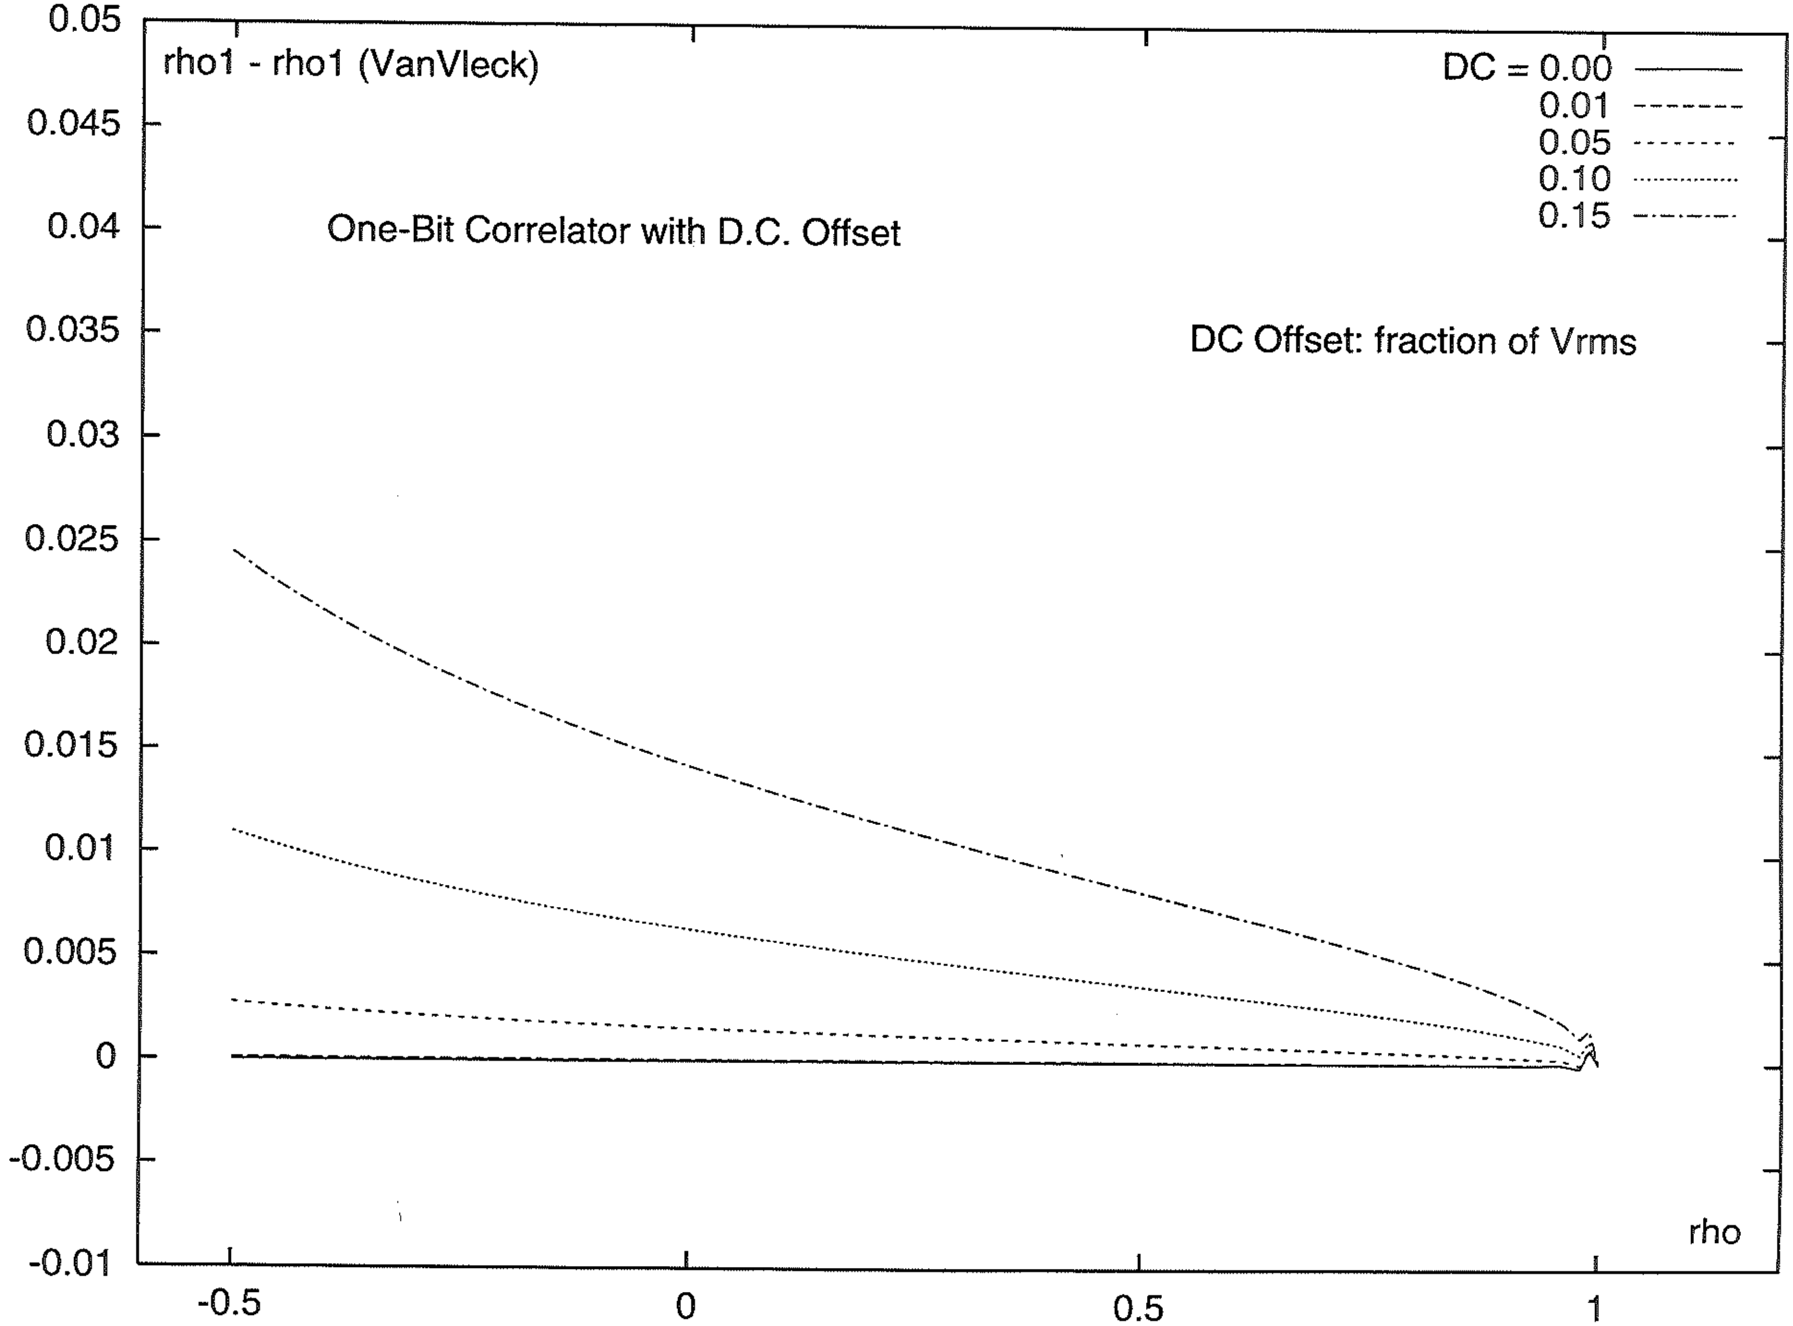
rho1 - rho1 (VanVleck)
One-Bit Correlator with D.C. Offset
DC Offset: fraction of Vrms
DC = 0.00
0.01
0.05
0.10
0.15
rho
0.05
0.045
0.04
0.035
0.03
0.025
0.02
0.015
0.01
0.005
0
-0.005
-0.01
-0.5
0
0.5
1
Fig. 4

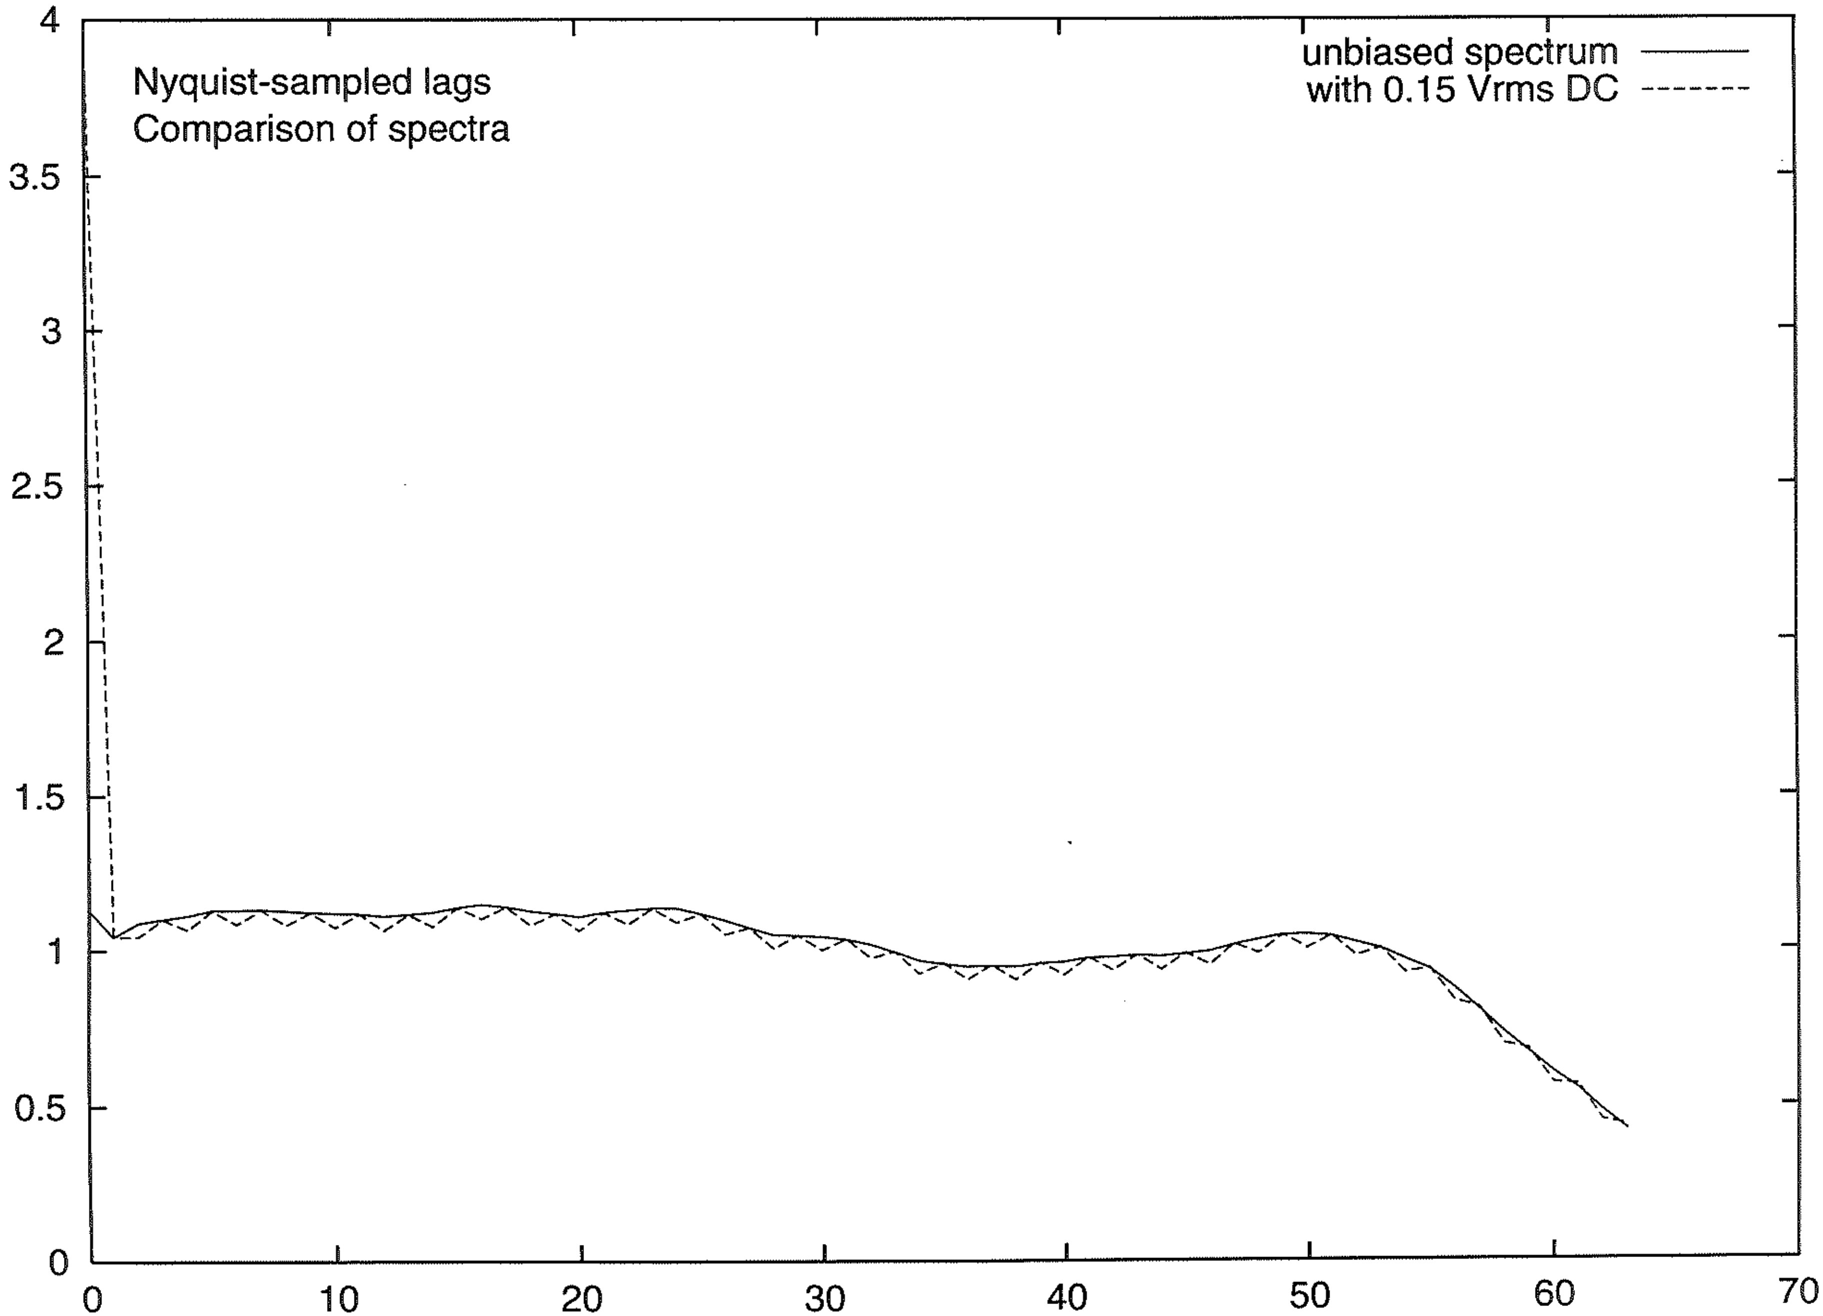
Nyquist-sampled lags
Comparison of spectra
unbiased spectrum
with 0.15 Vrms DC
0
0.5
1
1.5
2
2.5
3
3.5
4
0
10
20
30
40
50
60
70

Fig. 5

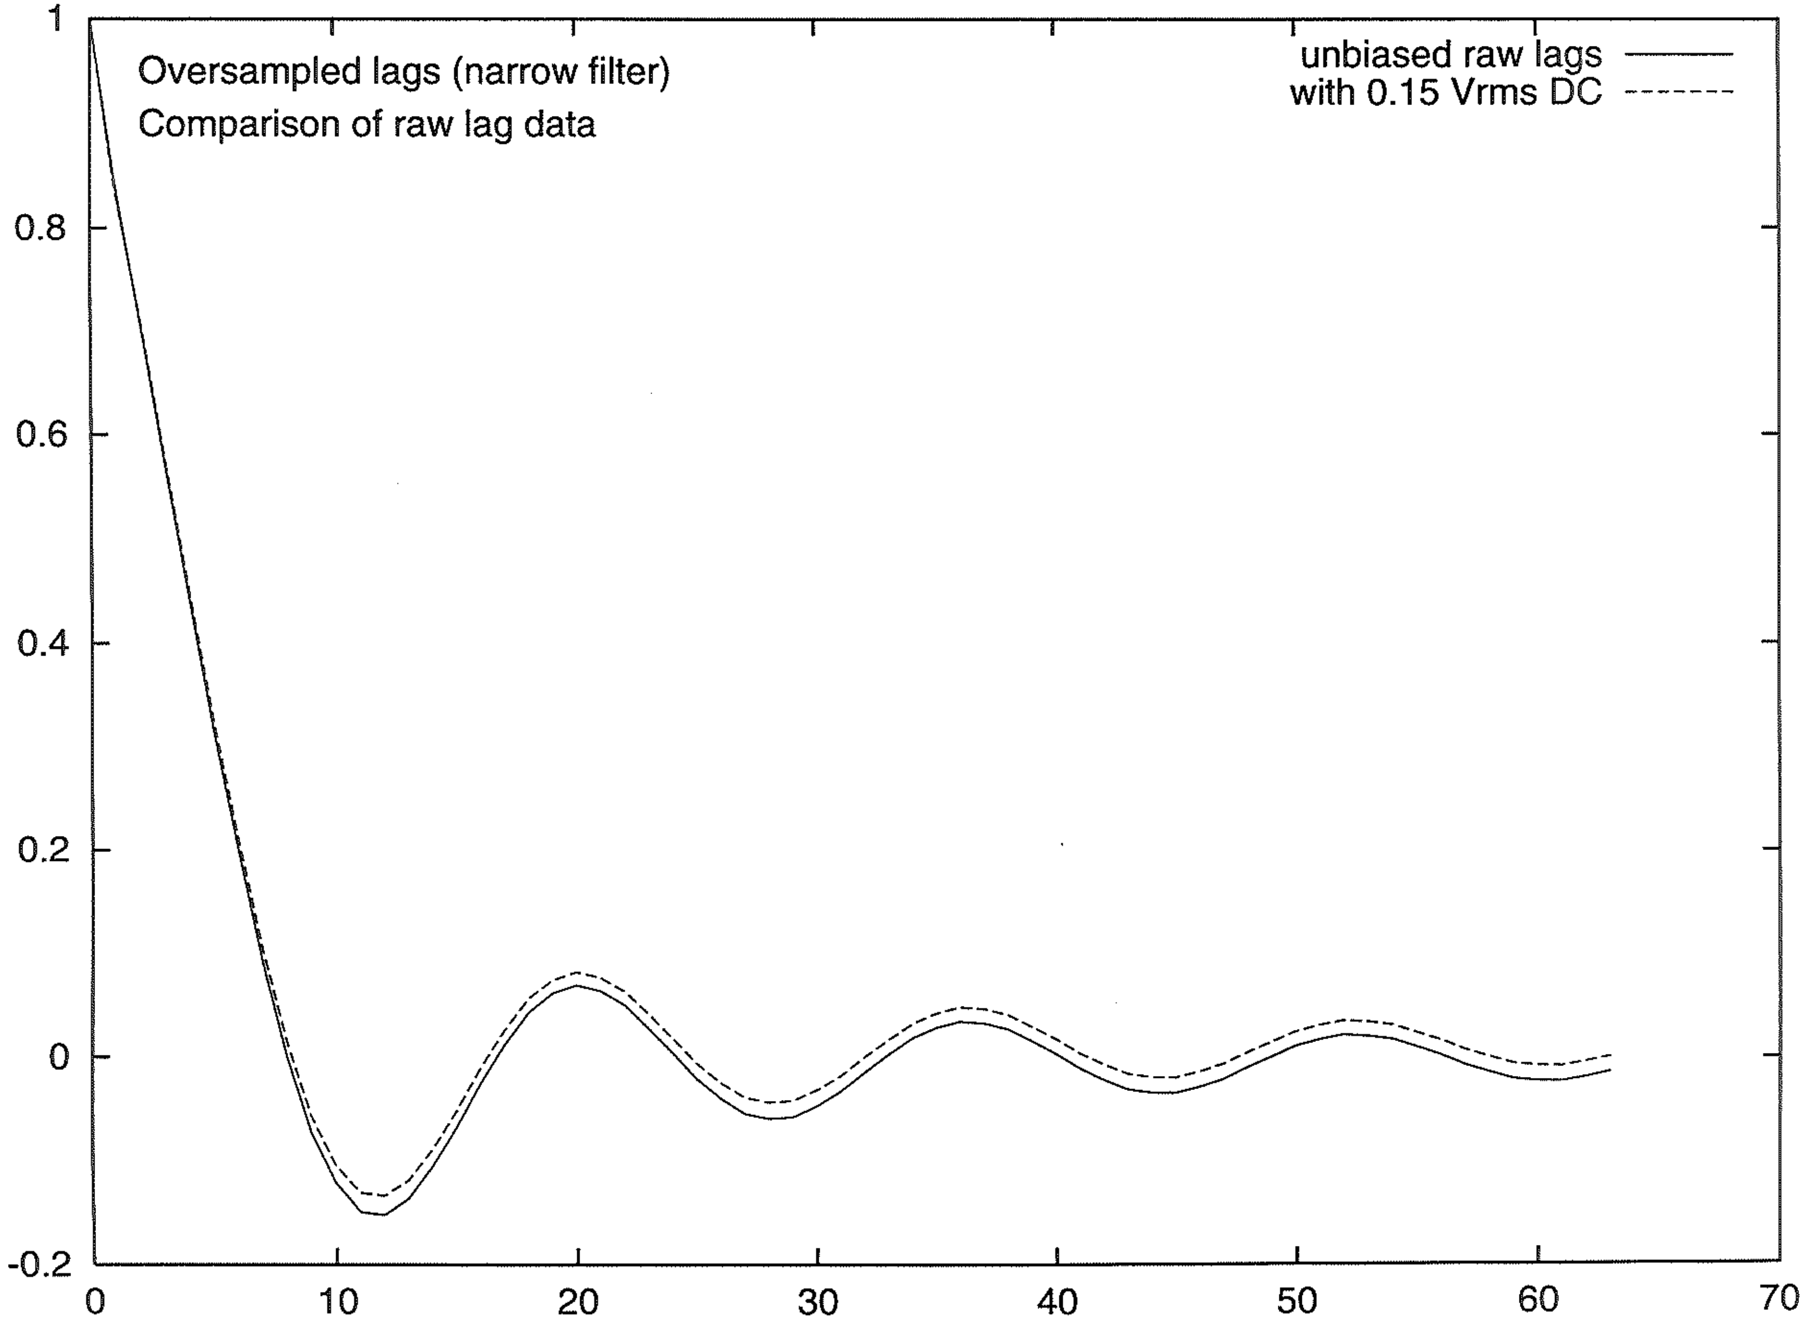
Oversampled lags (narrow filter)
Comparison of raw lag data
unbiased raw lags
with 0.15 Vrms DC
1
0.8
0.6
0.4
0.2
0
-0.2
0
10
20
30
40
50
60
70

Fig. 6a

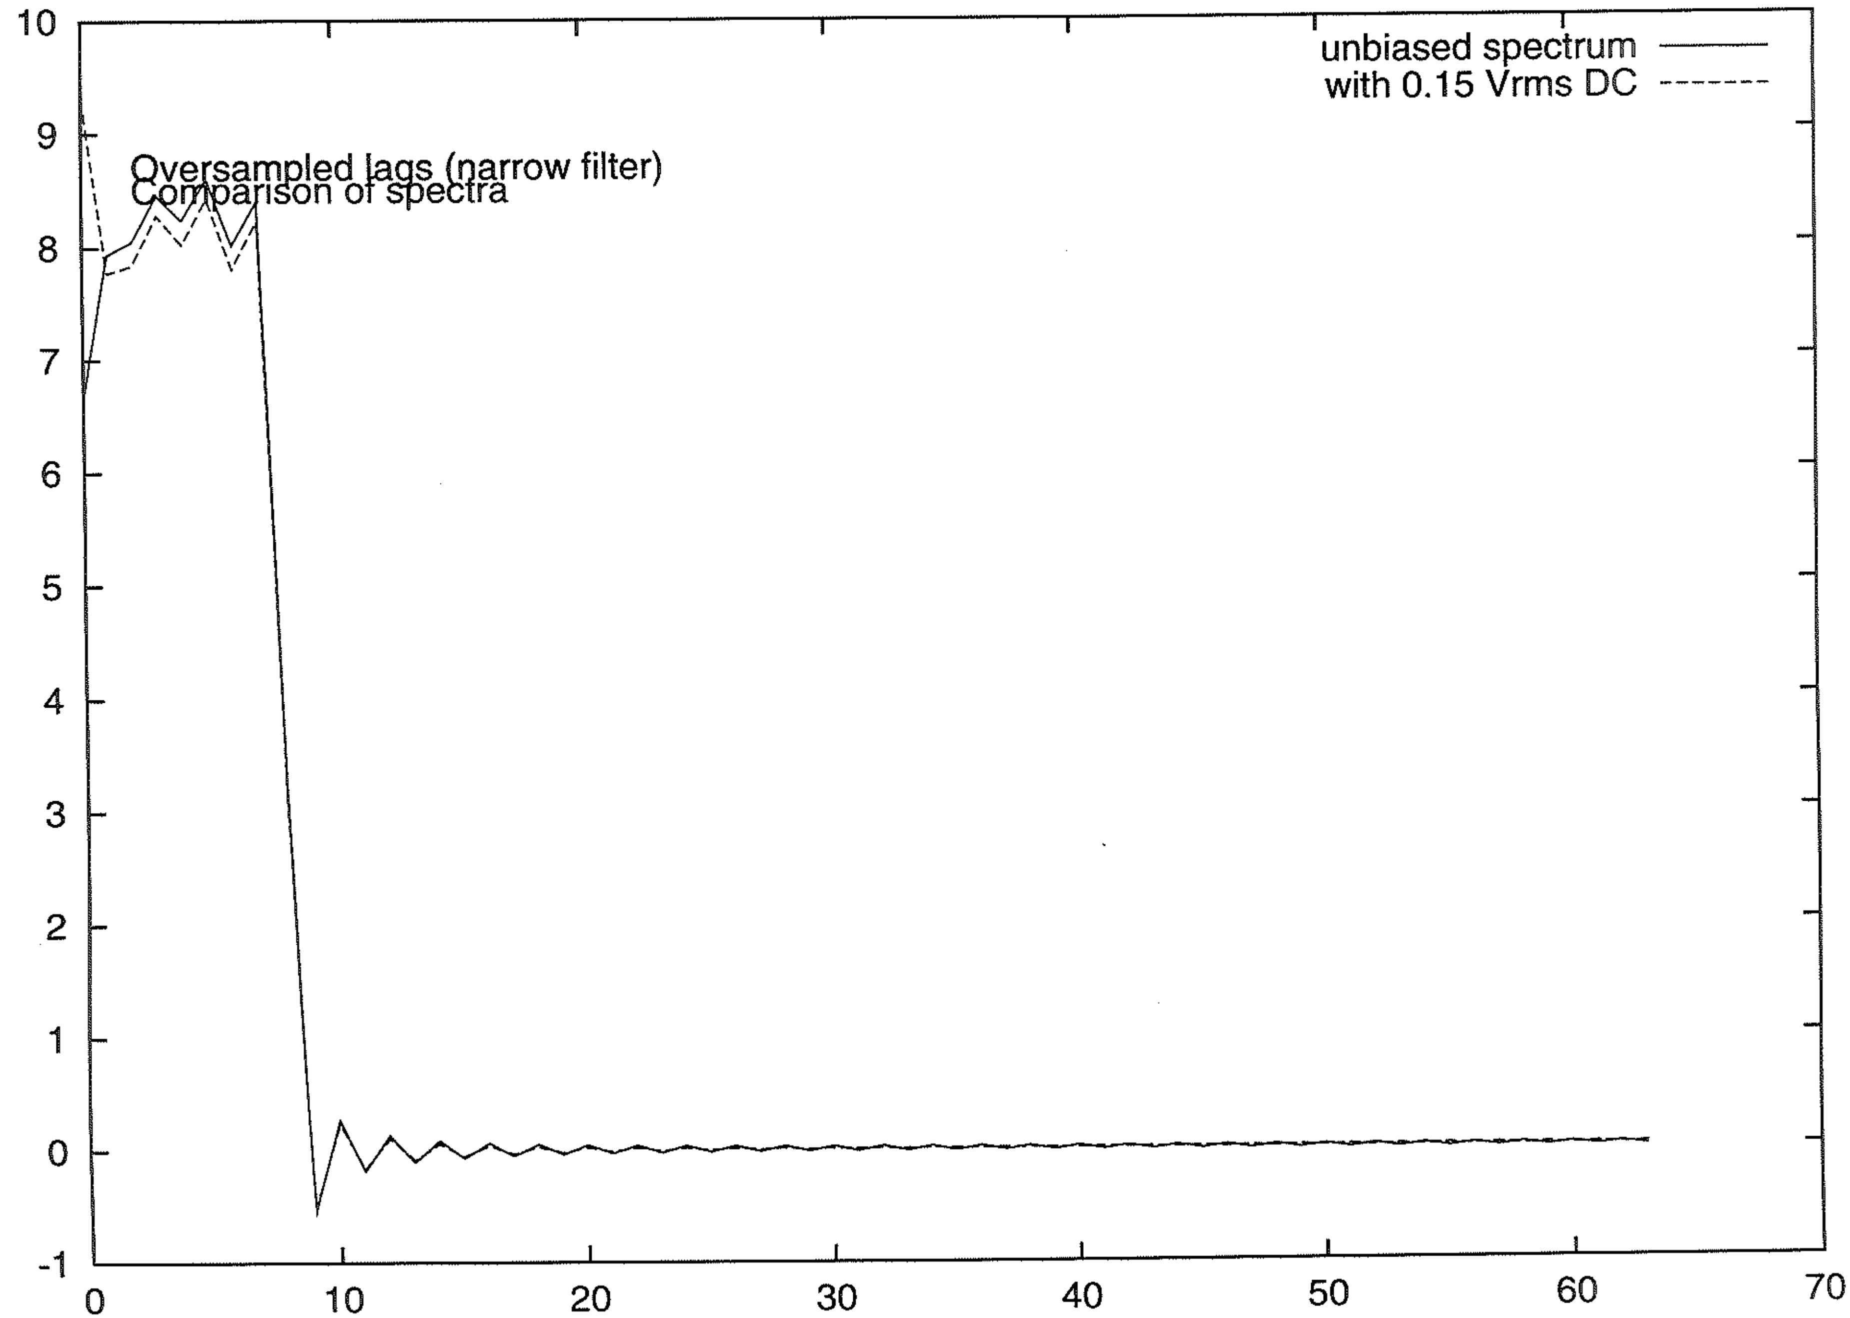
Oversampled lags (narrow filter)
Comparison of spectra
unbiased spectrum
with 0.15 Vrms DC
10
9
8
7
6
5
4
3
2
1
0
-1
0
10
20
30
40
50
60
70

Fig 6b

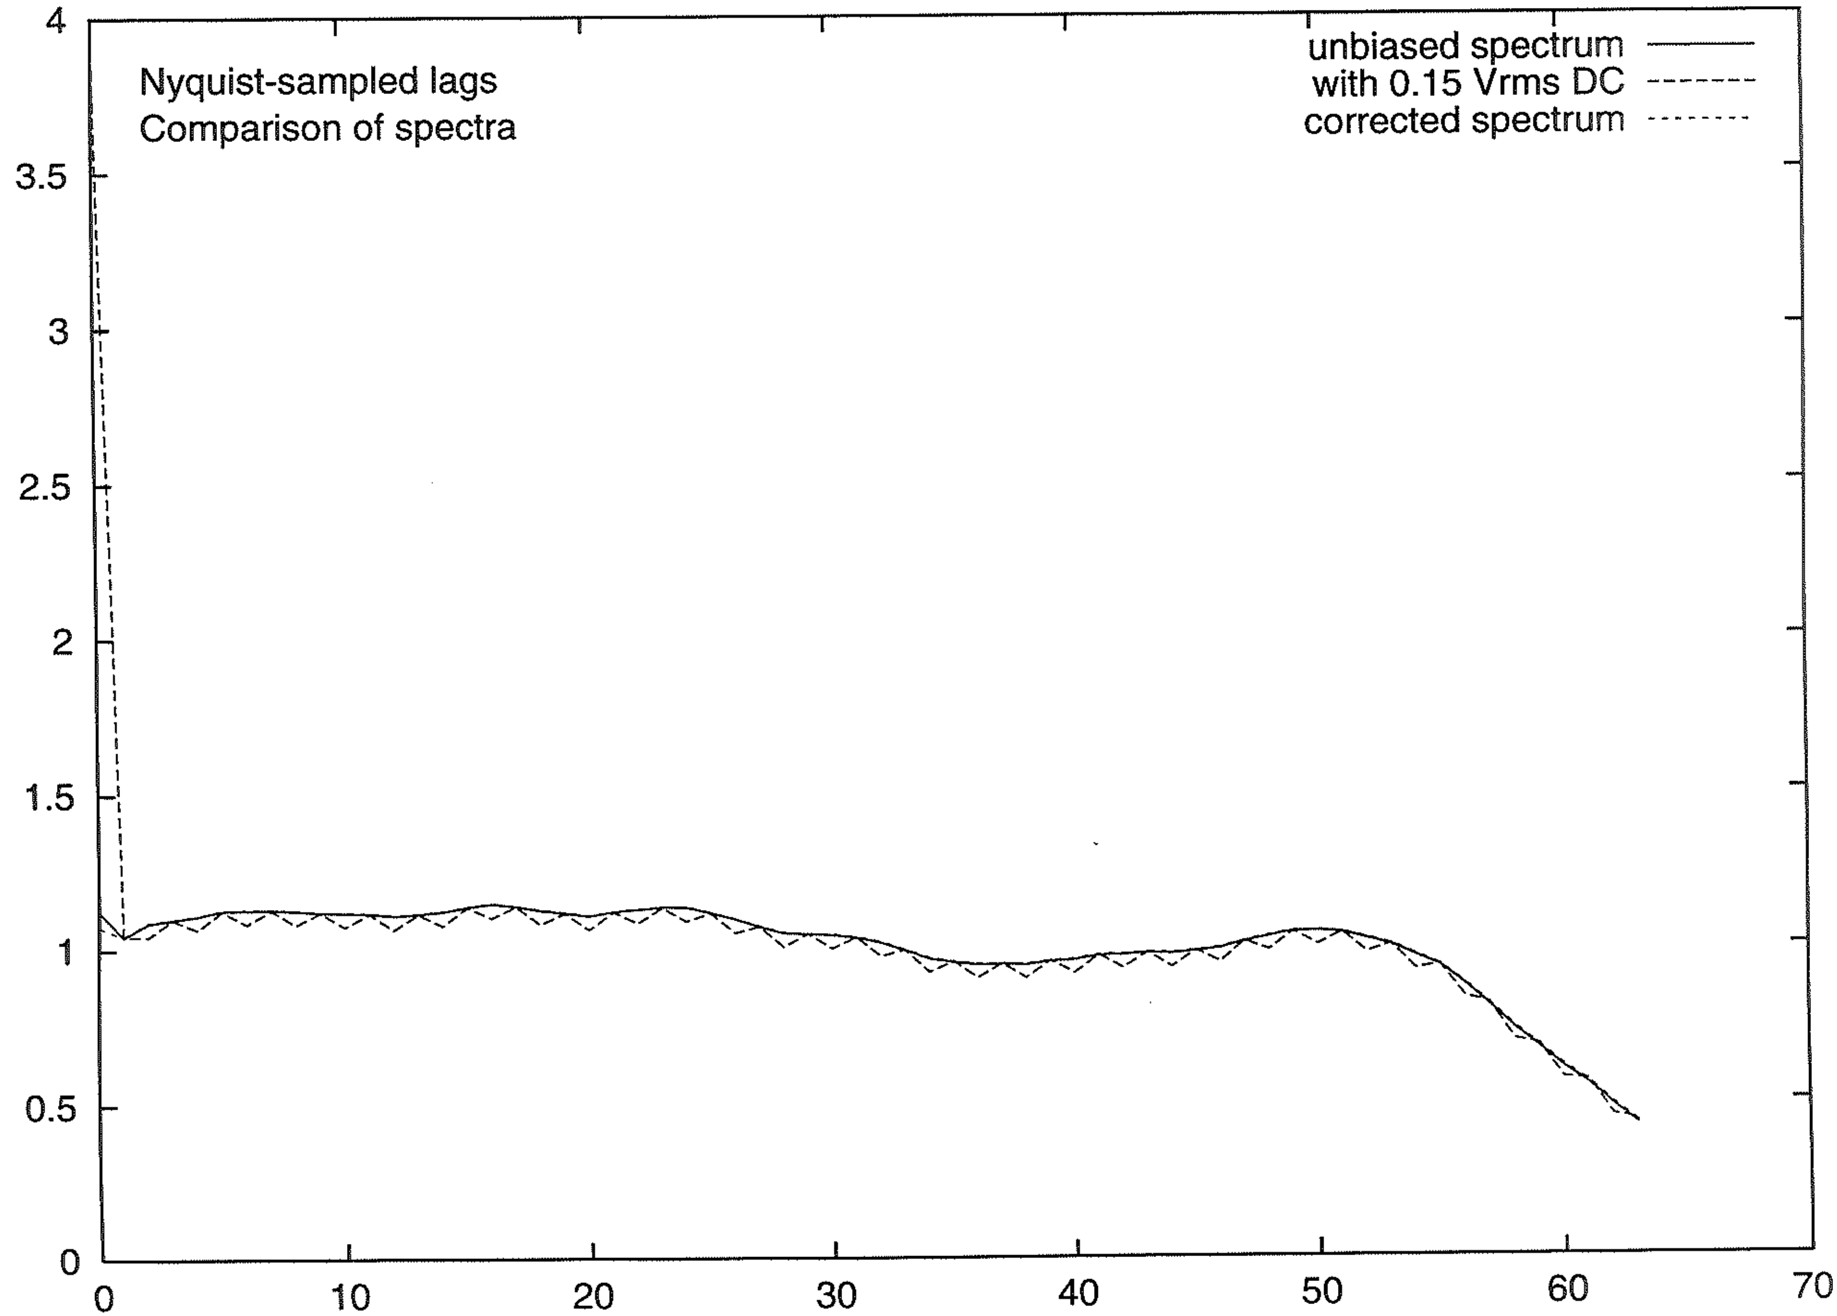
Nyquist-sampled lags
Comparison of spectra
unbiased spectrum
with 0.15 Vrms DC
corrected spectrum
4
3.5
3
2.5
2
1.5
1
0.5
0
0
10
20
30
40
50
60
70

Fig. 7a

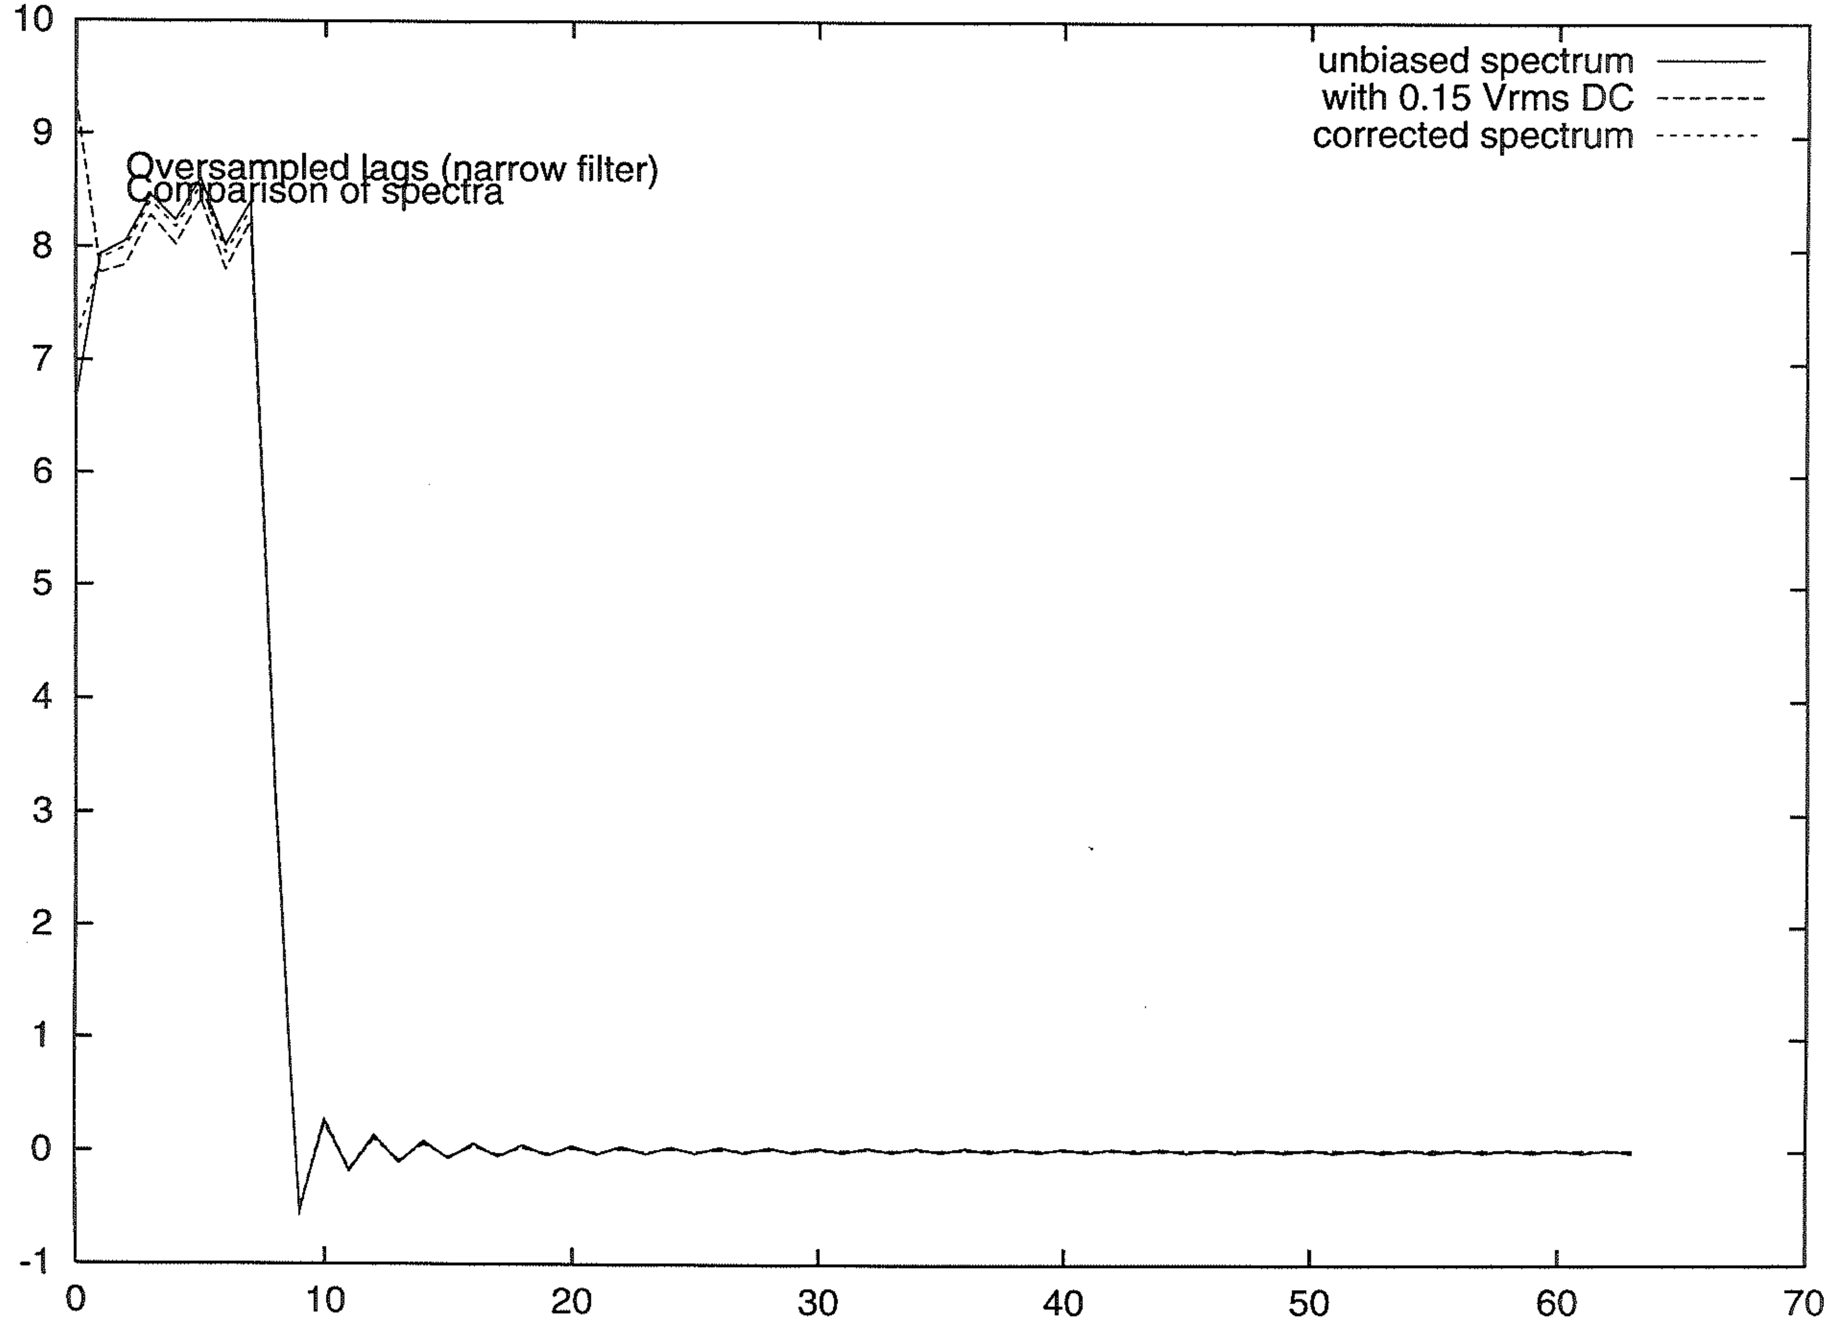
Oversampled lags (narrow filter)
Comparison of spectra
unbiased spectrum
with 0.15 Vrms DC
corrected spectrum
10
9
8
7
6
5
4
3
2
1
0
-1
0
10
20
30
40
50
60
70

Fig. 7b

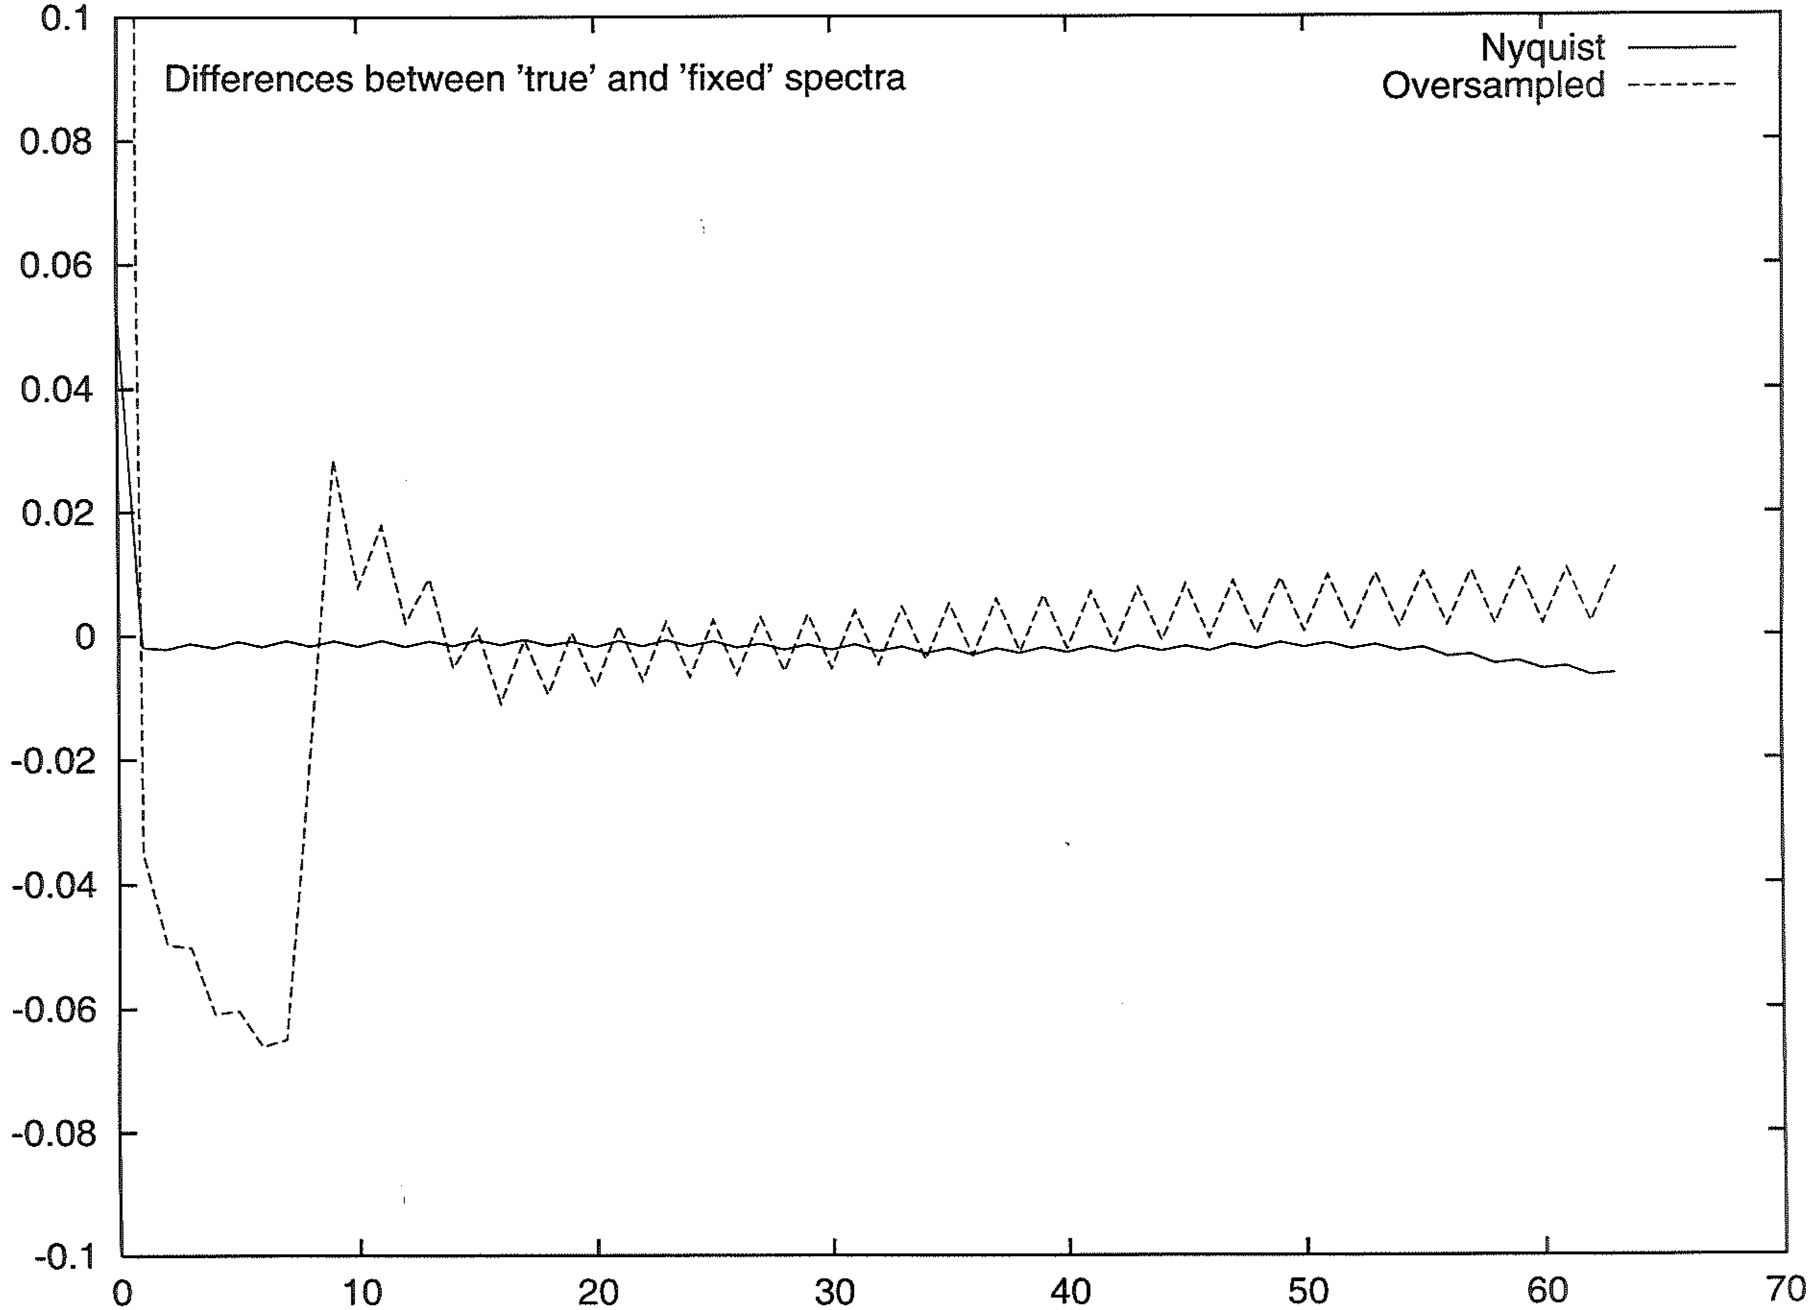
Differences between 'true' and 'fixed' spectra
Nyquist
Oversampled
0.1
0.08
0.06
0.04
0.02
0
-0.02
-0.04
-0.06
-0.08
-0.1
0
10
20
30
40
50
60
70

Fig. 8