Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Computational Physics 202 (2005) 134–149
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

Sine-fit procedure for unevenly sampled, multiply clocked signals

Edoardo Milotti *

Dipartimento di Fisica, Università di Udine and I.N.F.N., Sezione di Trieste, Via delle Scienze, 208 – I-33100 Udine, Italy

Received 28 December 2003; accepted 6 July 2004
Available online 25 August 2004

Abstract

When signals are unevenly sampled and homodyne or heterodyne setups are used that involve multiple clocks, it is not possible to use the standard algorithms to produce accurate spectral density estimates. The PVLAS experiment is an example of such an experimental scheme and here it is shown how to implement an algorithm which generalizes the standard sine-fit and Lomb–Scargle procedures to obtain both amplitude and phase of important Fourier components of the physical signal, with an application to PVLAS data.

1. Introduction

In today's practice, signals are often sampled by analog-to-digital converters (ADC) driven by very precise clocks, and afterwards (or also concurrently if enough computing power is available) the samples are analyzed by Fast Fourier Transform (FFT) or other similar standard techniques (for an overview, see [1]). However, this is not always the case, and when signals are unevenly sampled and homodyne or heterodyne setups are used that involve multiple clocks, it is not possible to use the standard algorithms to produce accurate spectral density estimates. The PVLAS experiment, designed to perform a very delicate test of quantum electrodynamics (QED) [2,3] and to search for light pseudoscalar particles [4], is an example of a multiclock heterodyne scheme, and here it is shown how to implement an algorithm which generalizes the standard sine-fit [5] and Lomb–Scargle [6] procedures to obtain both amplitude and phase of important Fourier components of the physical signal, with an application to PVLAS data. The next section summarizes a few important facts, Sections 3 and 4 outline the theory, Section 5 reports the results obtained with test data, and Section 6 shows how the procedure has been implemented in the case of PVLAS.

* Fax: +39 432 558222.
E-mail address: milotti@ts.infn.it.

0021-9991/$ - see front matter
doi:10.1016/j.jcp.2004.07.005

2. Spectral power density and phase estimates

It is easy to show [1] that the discrete Fourier transform of a set of N samples $\{s_k\}$ is equivalent to a least square fit procedure to the data with the model function

$$s_k^{\text{model}} = \frac{1}{\sqrt{N}} \sum_n F_n \exp\left(\frac{2\pi kn}{N}\right), \tag{1}$$

i.e., with a set of fixed-frequency trigonometric functions, where $\{F_n\}$ is a set of complex parameters. If all the samples have the same statistical variance, one finds that the least square fit yields

$$F_n = \frac{1}{\sqrt{N}} \sum_k s_k \exp\left(-\frac{2\pi kn}{N}\right), \tag{2}$$

which is just the standard definition of the discrete Fourier transform.

This basic idea can and has been extended, e.g., in the Prony method one tries to fit a signal with a small number of complex exponentials

$$s^{\text{model}}(t_k) = \sum_n A_n \exp(\mathrm{i}\omega_n t_k + \mathrm{i}\varphi_n) \exp(-t_k/\tau_n), \tag{3}$$

where the amplitudes A_n, phases ϕ_n, frequencies ω_n, and decay times τ_n are taken as fit parameters; unfortunately the Prony method requires a nonlinear fit procedure, and it is known to be computationally expensive and very sensitive to background noise [7,8].

In general, fitting methods may be more computationally demanding than the standard FFT, but have two important advantages, there is no constraint on the number of samples (the FFT requires the number of samples to be a power of 2, otherwise one falls back on the computationally inefficient Discrete Fourier Transform (DFT)) and they are not limited to evenly spaced samples. One such example is the method of Lomb and Scargle (LS) [6] which has been developed in the first place to obtain spectral estimates in the case of astronomical data: this method is equivalent to fitting a single sine to data while sweeping its frequency over the range of interest. Since the LS sine has – at each step – a fixed frequency, the fit procedure leads to simple linear equations that are easily solved. In addition, it is easy to estimate the variance of the results [9], data may be unevenly spaced and in this case the spectral estimate acquires additional immunity from aliasing effects [10]. Sine-fit methods are often used to measure the performance of ADC circuits, but this is usually done with equally spaced samples and with nonlinear fitting algorithms to estimate frequency as well as amplitude and phase [5].

In addition to the spectral density estimate, the data analysis of the PVLAS experiment must return the phase of some selected spectral lines, and the FFT is not well suited for this purpose (see Appendix A), while the sine-fit method returns the phase just as well as the amplitude.

3. Generalized sine-fit procedure

We consider the case in which a physical effect may be modulated by one or more clocks which run at a variable rate, i.e., the phase function of each term is known up to an additive constant while the frequency appears to change. We also assume damping to be negligible, so that the model signal (3) becomes

$$\begin{aligned} s^{\text{model}}(t_k) &= \sum_n A_n \cos\left[\varphi_n(t_k) + \varphi_{0,n}\right] = \sum_n \left[A_n \cos\varphi_n(t_k) \cos\varphi_{0,n} - A_n \sin\varphi_n(t_k) \sin\varphi_{0,n}\right] \\ &= \sum_n \left[x_n \cos\varphi_n(t_k) + y_n \sin\varphi_n(t_k)\right], \end{aligned} \tag{4}$$

where the sample times t_k may be unevenly spaced, the phase functions $\varphi_n(t_k)$ are known, and $x_n = A_n \cos\varphi_{0,n}$, $y_n = -A_n \sin\varphi_{0,n}$. In order to find the parameters x_n, y_n, and then the amplitudes $A_n = \sqrt{x_n^2 + y_n^2}$ and phase constants $\varphi_0 = -\arctan(y_n/x_n)$, we minimize the χ^2 function

$$\chi^2 = \sum_n \frac{\left[s_k - s^{\text{model}}(t_k)\right]^2}{\sigma_n^2} \tag{5}$$

and obtain the equations

$$\begin{aligned}
&\sum_k \left[x_k \sum_n \cos\varphi_l(t_k)\cos\varphi_n(t_k) + y_k \sum_n \cos\varphi_l(t_k)\sin\varphi_n(t_k) \right] = \sum_k s_k \cos\varphi_l(t_k), \\
&\sum_k \left[x_k \sum_n \sin\varphi_l(t_k)\cos\varphi_n(t_k) + y_k \sum_n \sin\varphi_l(t_k)\sin\varphi_n(t_k) \right] = \sum_k s_k \sin\varphi_l(t_k).
\end{aligned} \tag{6}$$

If we define the matrix elements

$$\begin{aligned}
S_{l,n}^{(\text{cs})} &= \sum_k \cos\varphi_l(t_k)\sin\varphi_n(t_k), \\
S_{l,n}^{(\text{sc})} &= \sum_k \sin\varphi_l(t_k)\cos\varphi_n(t_k), \\
S_{l,n}^{(\text{cc})} &= \sum_k \cos\varphi_l(t_k)\cos\varphi_n(t_k), \\
S_{l,n}^{(\text{ss})} &= \sum_k \sin\varphi_l(t_k)\sin\varphi_n(t_k)
\end{aligned}$$

and the vector elements $b_n^{(\text{c})} = \sum_n s_k \cos\varphi_l(t_k)$, and $b_n^{(\text{s})} = \sum_n s_k \sin\varphi_l(t_k)$ then we can write the matrix equation

$$\begin{pmatrix} \mathbf{S}^{(\text{cc})} & \mathbf{S}^{(\text{cs})} \\ \mathbf{S}^{(\text{sc})} & \mathbf{S}^{(\text{ss})} \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} = \begin{pmatrix} \mathbf{b}^{\text{c}} \\ \mathbf{b}^{\text{s}} \end{pmatrix}. \tag{7}$$

It is easy to see that the $2N \times 2N$ matrix $\mathbf{S}$ is symmetric and moreover if $\mathbf{T} = \mathbf{S}^{-1}$, then $\mathbf{T}$ is symmetric and the solution of the linear system (7) is $\mathbf{Tb}$. Thus the solution vector is a linear combination of the samples $\{s_k\}$, and if we assume that the sampling error is the same for all samples, and also that it has a Gaussian distribution with variance σ^2 and is uncorrelated [11], then the solution vector has a Gaussian distribution as well, and the variances and the covariances of the individual x_ns and y_ns are [12]:

$$\text{var}(x_n) = \sigma^2 \sum_k \left(\frac{\partial x_n}{\partial s_k} \right)^2, \tag{8}$$

$$\text{var}(y_n) = \sigma^2 \sum_k \left(\frac{\partial y_n}{\partial s_k} \right)^2, \tag{9}$$

$$\text{cov}(x_n, y_n) = \sigma^2 \sum_k \left(\frac{\partial x_n}{\partial s_k} \right) \left(\frac{\partial y_n}{\partial s_k} \right). \tag{10}$$

The partial derivatives can be obtained directly from the solution of the linear system (6):

$$\frac{\partial x_n}{\partial s_k} = \sum_{l=1}^{N} \left[T_{n,l} \cos\varphi_l(t_k) + T_{n,l+N} \sin\varphi_l(t_k) \right], \tag{11}$$

$$\frac{\partial x_n}{\partial s_k} = \sum_{l=1}^{N} [T_{n+N,l} \cos \varphi_l(t_k) + T_{n+N,l+N} \sin \varphi_l(t_k)] \tag{12}$$

and finally, using the matrix identity $\mathbf{ST} = \mathbf{I}$, and after a few straightforward but lengthy passages, we find the following formulas for the variance of each x_n, y_n:

$$\sigma_{n,x}^2 = \mathrm{var}(x_n) = \sigma^2 T_{n,n}, \tag{13}$$

$$\sigma_{n,y}^2 = \mathrm{var}(y_n) = \sigma^2 T_{n+N,n+N}, \tag{14}$$

$$\rho_n \sigma_{n,x} \sigma_{n,y} = \mathrm{cov}(x_n, y_n) = \sigma^2 T_{n+N,n}. \tag{15}$$

4. Variance of estimated amplitude and phase

Using the variances in Eqs. (13)–(15), we can calculate the variances of the estimated amplitudes and phases. When the ratios $x_n/\sigma_{n,x}$ and $y_n/\sigma_{n,y}$ are well above 1, we can assume that the amplitudes and phases have a Gaussian distribution as well, and use the usual approximate formulas [12] to evaluate the variances; then we find

$$\mathrm{var}\, A_n = \sigma_{n,x}^2 \frac{x_n^2}{A_n^2} + \sigma_{n,y}^2 \frac{y_n^2}{A_n^2} + 2\rho_n \sigma_{n,x} \sigma_{n,y} \frac{x_n y_n}{A_n^2}, \tag{16}$$

$$\mathrm{var}\, \varphi_{0,n} = \sigma_{n,x}^2 \frac{x_n^2}{A_n^4} + \sigma_{n,y}^2 \frac{y_n^2}{A_n^4} - 2\rho_n \sigma_{n,x} \sigma_{n,y} \frac{x_n y_n}{A_n^4}, \tag{17}$$

$$\mathrm{cov}(A_n, \varphi_{0,n}) = \sigma_{n,x}^2 \frac{x_n y_n}{A_n^3} - \sigma_{n,y}^2 \frac{x_n y_n}{A_n^3} + \rho_n \sigma_{n,x} \sigma_{n,y} \left(-\frac{x_n^2}{A_n^3} + \frac{y_n^2}{A_n^3} \right) \tag{18}$$

and we see that $\mathrm{var}\, \varphi_{0,n} \approx \mathrm{var}\, A_n / A_n^2$, as expected on intuitive grounds.

However, when the ratios $x_n/\sigma_{n,x}$ and $y_n/\sigma_{n,y}$ are close to or smaller than 1, we can no longer use Eqs. (16)–(18) to evaluate the variances, since the probability distribution of amplitudes and phases is no longer Gaussian. The joint probability density function (PDF) of each pair of variates x, y, is the usual Gaussian bivariate distribution

$$G(x, y; x_0, y_0, \sigma_x, \sigma_y) = \frac{1}{2\pi \sigma_x \sigma_y \sqrt{1-\rho^2}} \exp\left[-\frac{(x-x_0)^2}{2(1-\rho^2)\sigma_x^2} + \rho \frac{(x-x_0)(y-y_0)}{(1-\rho^2)\sigma_x \sigma_y} - \frac{(y-y_0)^2}{2(1-\rho^2)\sigma_y^2} \right], \tag{19}$$

where we have introduced the correlation coefficient $\rho = \mathrm{cov}(x, y)/\sigma_x \sigma_y$, and x_0, y_0 are the values obtained from the fit. The joint amplitude-phase PDF is

$$P(A, \varphi_0) = A G(A \cos \varphi, A \sin \varphi; x_0, y_0, \sigma_x, \sigma_y) \tag{20}$$

and the marginal PDFs of amplitude and phase are

$$P_A(A) = A \int_0^{2\pi} G(A \cos \varphi, A \sin \varphi; x_0, y_0, \sigma_x, \sigma_y)\, \mathrm{d}\varphi, \tag{21}$$

$$P_\varphi(\varphi) = \int_0^{\infty} A G(A \cos \varphi, A \sin \varphi; x_0, y_0, \sigma_x, \sigma_y)\, \mathrm{d}A. \tag{22}$$

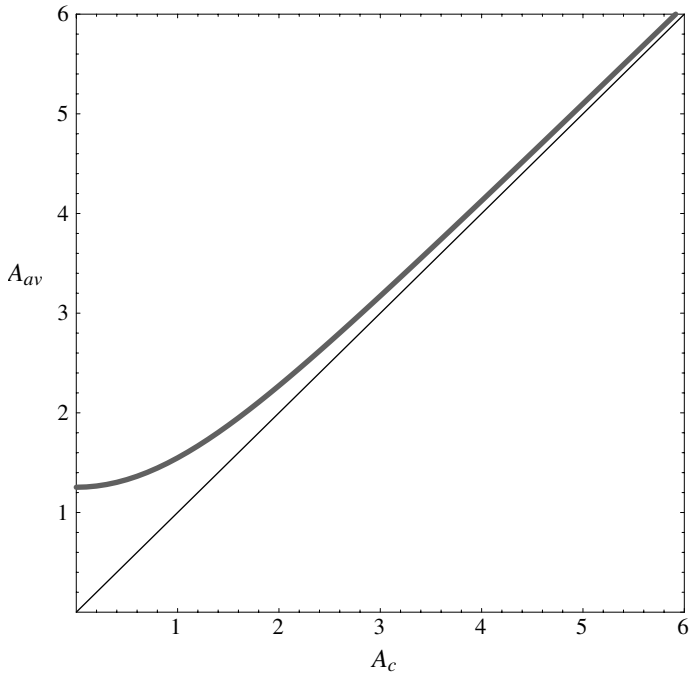

Fig. 1. Plot of the average amplitude A_{av} vs. central amplitude A_{c}.

When x_0/σ_x, y_0/σ_y are small, then the amplitude estimate $A_0 = \sqrt{x_0^2 + y_0^2}$ does not correspond to the most probable value nor to the actual average of A: Fig. 1 is a plot of the average $A_{av} = \int_0^\infty A P_A(A)\,\mathrm{d}A$ vs. the central value A_{c} which corresponds to the centroid of the Gaussian distribution (and which would be the "true" estimate), and it shows the deviation when the center of the PDF is close to the origin. The conclusion is that the amplitude-phase pair is not well-suited to small signal-to-noise ratios and in this case one should rather rely on the x_0s and y_0s alone.

Before moving on to the numerical tests of the algorithm, it is important to remark that in the case of equally spaced sampling times the matrix $\mathbf{S}$ is approximately diagonal (it becomes exactly diagonal when the phase difference over the total acquisition time is an integer multiple of π), and that in this case the diagonal elements are of the order of $N/2$, while the off-diagonal elements are at most of the order of 1. This means that also for samples that are approximately equally spaced the matrix $\mathbf{S}$ is approximately diagonal, and that the matrix $\mathbf{T}$ is another almost diagonal matrix with diagonal elements of the order of $2/N$ and with off-diagonal elements at most of the order of $1/N^2$; then the xs and ys are almost uncorrelated, and their variance is nearly the same, so that the joint PDF has a nearly circular symmetry.

5. Numerical tests

The fit procedure has been implemented and has been tested on simulated model data: all tests have been carried out with $2^{20} = 1{,}048{,}576$ samples, and the timescale has been set so that a frequency of 1 (in AFU = arbitrary frequency units) corresponds to 32 samples/period (these values have been chosen so that the sample size and the time unit are similar to those in PVLAS, see below); this means that in an ordinary FFT analysis the one-sided spectral density estimate ranges from 0 to 16 AFU and that the frequency

resolution is $1/2^{15} \approx 3.05 \times 10^{-5}$ AFU (this also means that any phase fluctuation smaller than $\approx 0.0055°$ is completely contained in a single FFT bin). The simulated data include white noise with standard deviation σ: this is the standard deviation of the noise signal *integrated* over the useful frequency band, so that the average mean squared fluctuation in each bin of the folded noise spectrum in an ordinary FFT analysis is $\sigma^2_{\mathrm{FFT}} = \sigma^2/2^{19} \approx 1.91 \times 10^{-6}\sigma^2$.

Figs. 2 and 3 show the results for reconstructed phase and amplitude vs. input amplitude in the case of a single fixed-frequency sine plus Gaussian white noise. The figures show that the fit procedure behaves quite well, but it is natural to ask how does it compare to a standard DFT analysis. Obviously, in this case the DFT would be nearly as good in pinpointing the amplitude, but unless the sine frequency corresponds *exactly* to one of the DFT frequencies, the DFT is unable to determine the phase of the input sine (see Appendix A). Moreover, the sine-fit procedure works also when the frequency is unstable (but the phase is known) and with unequally spaced samples: the DFT cannot handle these cases. Finally, the sine-fit procedure works with any number of samples, while one can achieve FFT acceleration only for sample sizes that are powers of 2.

Several other tests have been carried out and one always finds that the algorithm performs just as well with a single sine plus noise with the addition of a relevant phase jitter (the case with average sample separation 11°.25 and phase jitter ±5° as been tested without any appreciable degradation of the output values).

Figs. 4 and 5 show the results of fitting a simple model with two sines; the sines have fixed frequency in the model of Fig. 4, while the sines in the model of Fig. 5 have an additional, individual, phase jitter that brings about a frequency jitter approximately 9 times as large as the frequency resolution, and this produces a power leakage that cannot be reduced by windowing. As a result the weaker peak in the spectrum of Fig. 4 all but disappears in Fig. 5, while the stronger peak shows a clear broadening. While the FFT analysis cannot detect the weaker peak, the sine-fit procedure can, and the fit parameters of the

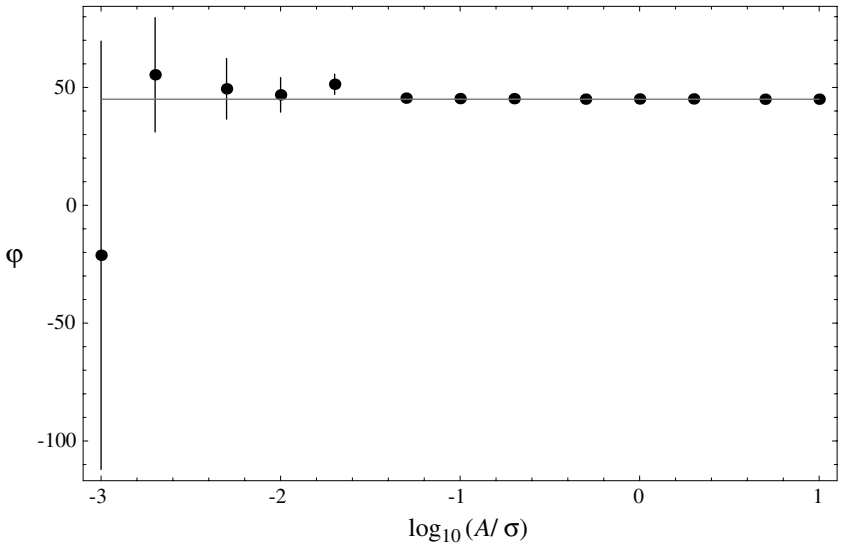

Fig. 2. The figure shows the result of fitting several sets of simulated data, each with $2^{20} \approx 10^6$ simulated samples; the model function is a simple sine with variable amplitude A plus noise; the input phase angle is 45°. Here φ is the phase angle obtained from the fit procedure while A is the input sine amplitude, and σ is the noise standard deviation; in this case the noise variance in a single FFT bin is related to the noise standard deviation by the formula $\log_{10}\sigma_{\mathrm{FFT}} \approx \log_{10}\sigma - 2.85$, and we see that the phase estimate is good only as long as $A/\sigma_{\mathrm{FFT}} \gtrsim 1$.

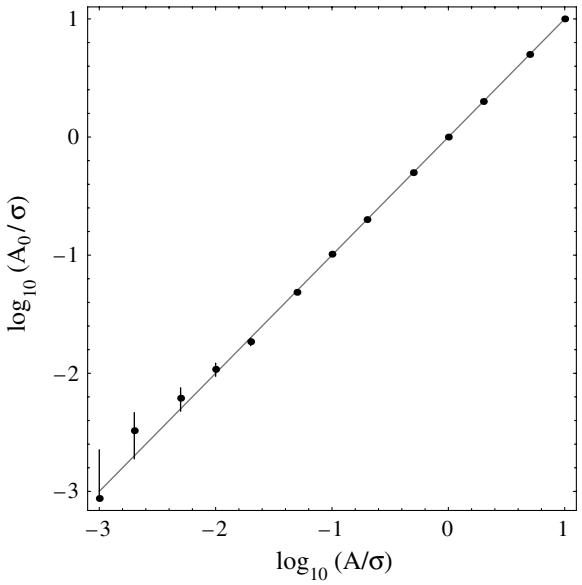

Fig. 3. The figure shows the result of fitting $2^{20} \approx 10^6$ simulated samples as in Fig. 2; the model function is a simple sine plus noise as in Fig. 2. Here A_0 is the amplitude obtained from the fit procedure while A is the input sine amplitude, and σ is the noise standard deviation; just as in the previous figure we see that the amplitude estimate is good only as long as $A/\sigma_{\mathrm{FFT}} \gtrsim 1$.

weaker sine do not seem to be affected by the nearby stronger peak. This parameter independence is fairly well preserved even in the case of a fit which does not include the stronger sine: Fig. 6 shows that when the fit model does not include a nearby peak the parameters of the fitted line are affected only when the peaks are very close to each other. In particular, Figs. 6 and 7 show contamination in the case of a unfitted sine with equal amplitude and 10-fold higher amplitude, respectively. Similarly, Fig. 8 shows the response to a wrong hypothesis (one sine in the wrong place). While the model sines in Figs. 6–8 have fixed frequency, Fig. 9 shows what happens when the frequency is not fixed but jitters (as described in the caption of Fig. 5): contamination extends further, but is very much reduced with respect to the fixed-frequency case. This is easy to understand: because of the leakage produced by the phase and frequency jitter, a ''wrong'' frequency Fourier component is always present in the nearby bins, but, at the same time, the nonlinear phase function sets narrow bounds on acceptable Fourier components.

The lesson learned from these simulations is that one should try to fit the data including all prior information on the spectral density, but if this is not possible the algorithm still returns reasonable results. However, if one neglects strong peaks close to fitted sines, then some contamination is unavoidable, and in the worst case it may lead to false peaks. This contamination is much reduced when the spectral components have different phase functions.

6. Analysis of PVLAS data

Here I sketch very briefly how the methods described above shall be used to analyze the signals recorded in the PVLAS experiment [3]; the actual physics results shall be described in forthcoming PVLAS collaboration papers.

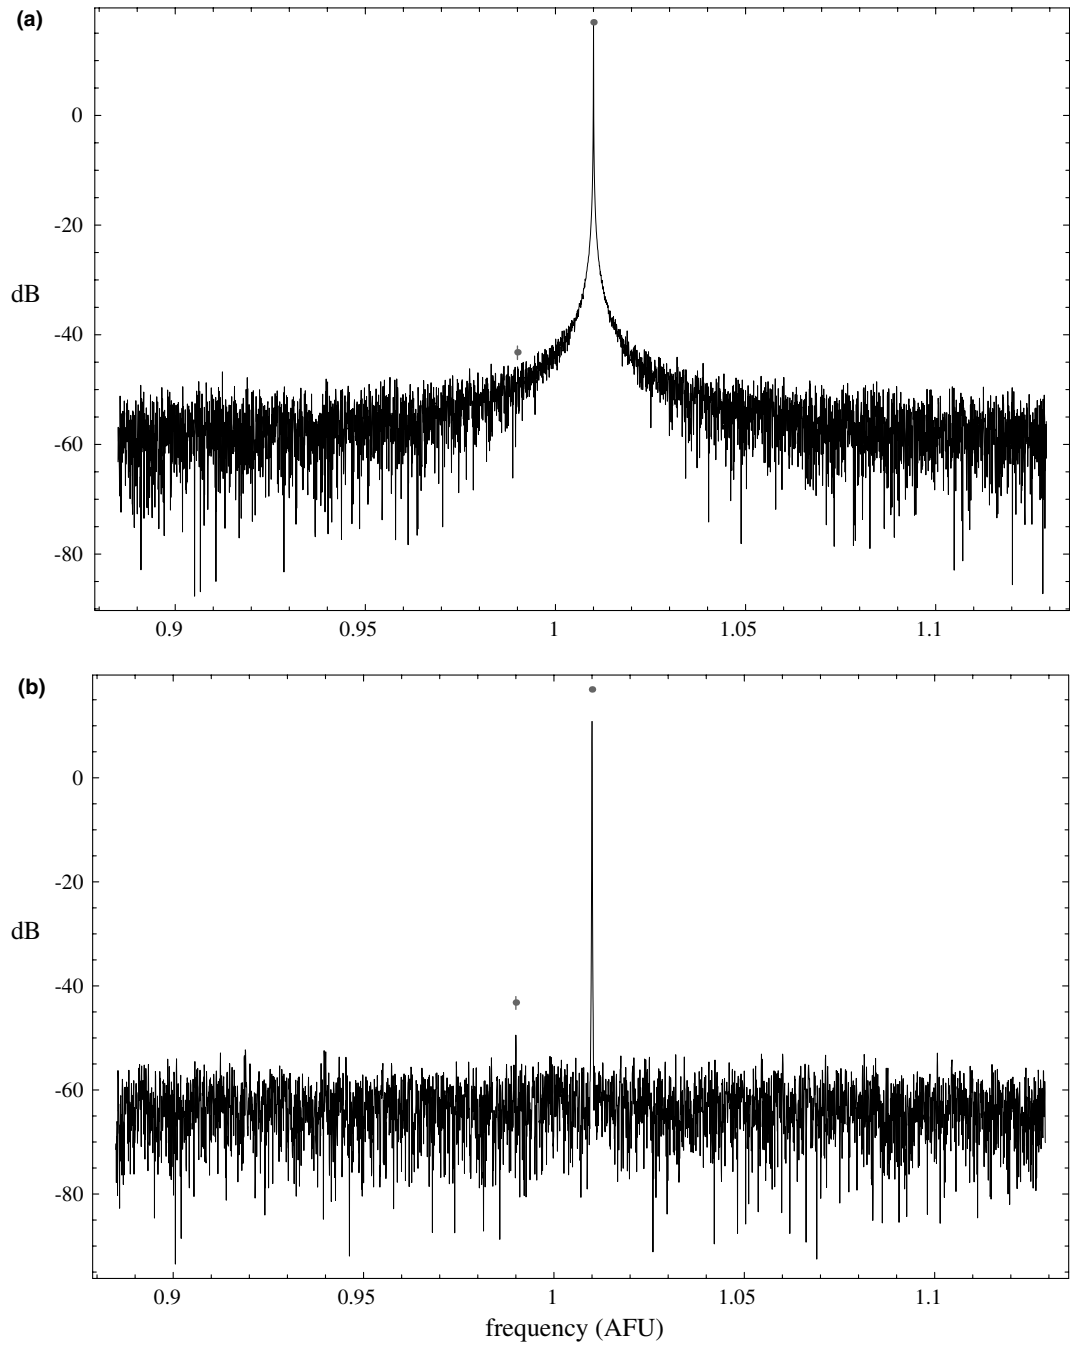

Fig. 4. Fit of a simple model with two fixed-frequency sines (amplitudes 0.01σ, 10σ, frequencies 0.99, 1.01 AFU, phase constants 0°, 90°, respectively) plus Gaussian white noise with amplitude standard deviation σ. In this case, both cosines have been fit, and the fitted quantities are: amplitudes $(0.0098 \pm 0.0014)\sigma$, $(10.0012 \pm 0.0014)\sigma$, phases $2.35 \pm 8.07°$, $90.0156 \pm 0.0079°$. The figure shows: (a) fit results (dots) superimposed on the unwindowed FFT estimate of the spectral density; (b) as the previous figure but with a Hamming window. The fit results coincide with the FFT estimate in the unwindowed case (in the windowed case the FFT shows the usual amplitude decrease).

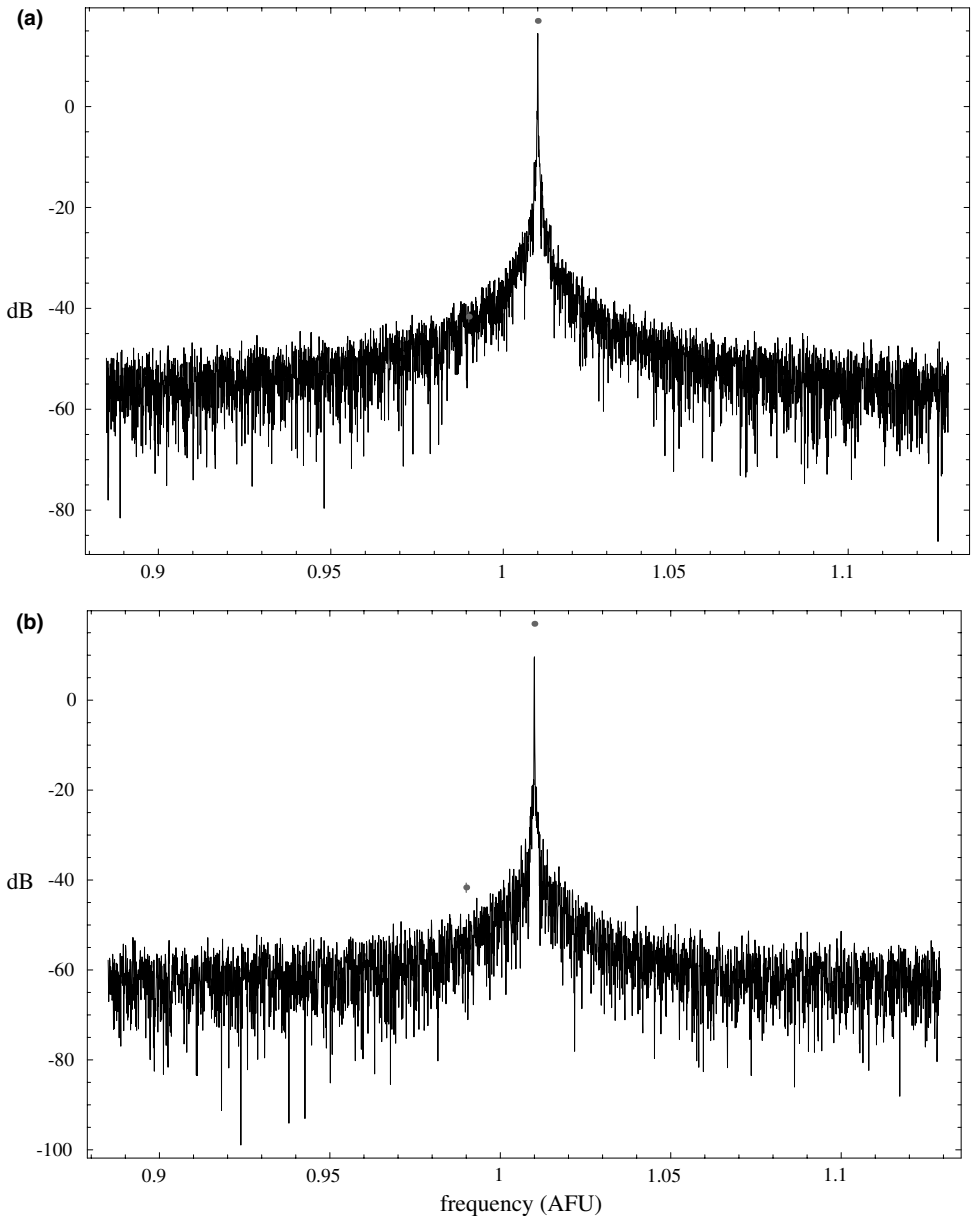

Fig. 5. Fit of a simple model with two sines plus Gaussian white noise with amplitude standard deviation σ, as in Fig. 4, but with the addition of a small phase jitter, 0.1° for both sines (phase does not increase linearly with time, but performs an additional random walk, so that the actual phase change at each time step is $\Delta\varphi = \omega\Delta t + r \cdot 0.1°$, where r is a uniform pseudorandom number distributed between -0.5 and 0.5). Just as in Fig. 4 both sines have been fit, and the fitted quantities are: amplitudes $(0.0117 \pm 0.0014)\sigma$, $(9.9991 \pm 0.0014)\sigma$, phases $-4.30 \pm 6.76°$, $89.9930 \pm 0.0079°$. The figure shows: (a) fit results (dots) superimposed on the unwindowed FFT estimate of the spectral density and (b) as the previous figure but with a Hamming window. Now the FFT estimates are affected by a systematic error due to power leakage to nearby bins, and this is a kind of power leakage that cannot be compensated by windowing. Here the standard deviation of phase jitter is approximately 9 times as large as the frequency resolution, and as a result the bigger peak is considerably wider in the FFT spectrum, while the smaller peak is completely wiped out. On the other hand the sine-fit is still quite good, with a phase uncertainty which is dictated by the background noise only.

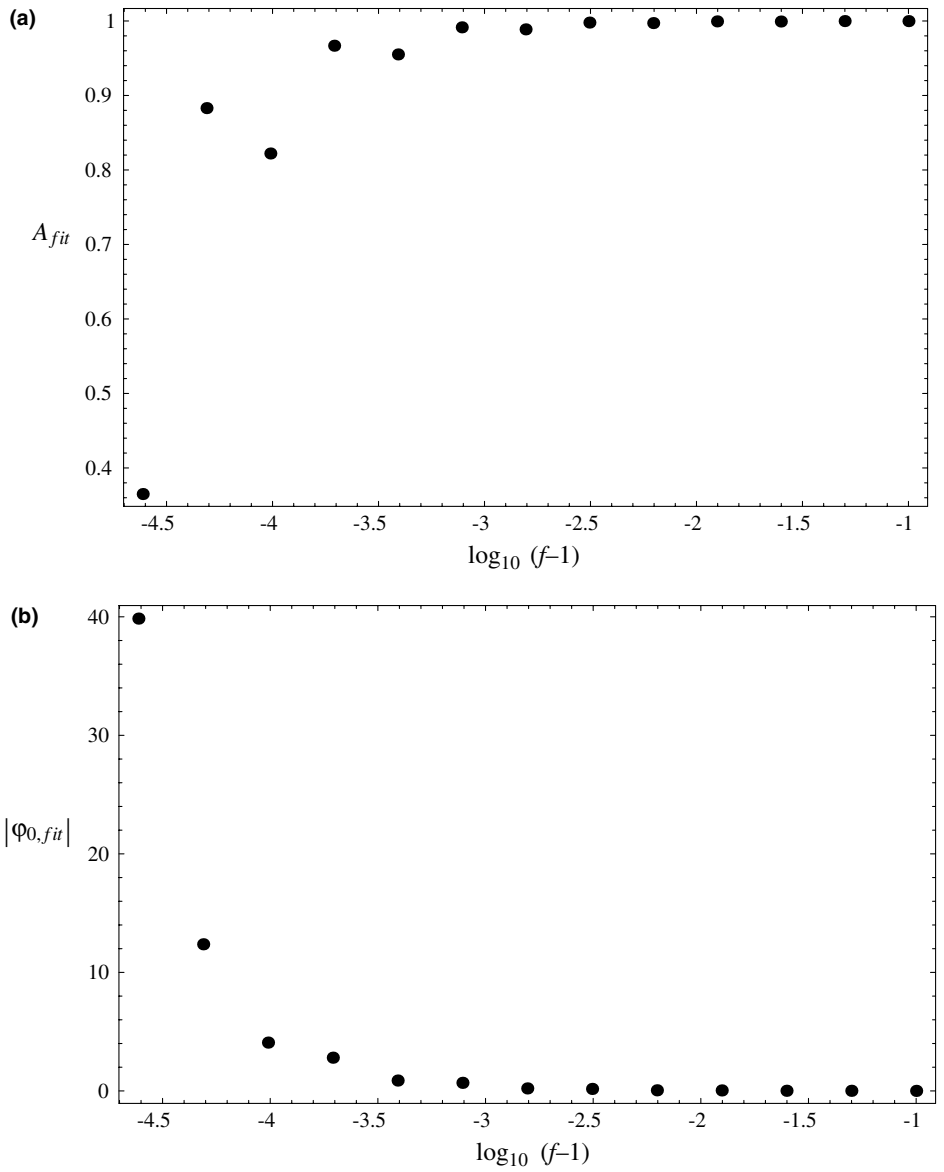

Fig. 6. Test of contamination from a nearby unfitted peak: the input model has two sines (amplitudes 1 a.u., phase constants 0° and 90°, respectively) and no noise, while the fit model has just one sine; the lower lying peak (fitted) has fixed frequency 1 AFU, while the higher frequency peak has variable frequency f AFU. (a) Plot of the reconstructed amplitude vs. the logarithm of the frequency difference $\log(f-1)$ and (b) plot of the reconstructed phase constant vs. the logarithm of the frequency difference $\log(f-1)$. In this case, the bin size of an ordinary FFT would be about 3×10^{-5} AFU, which corresponds to about -4.52 on the logarithmic scale: thus we see that the single sine model shows no significant contamination as long as the unfitted sine is at least 10 FFT bins away from the fitted sine.

PVLAS has been designed to detect light-field scattering, a subtle QED effect predicted long ago by Euler and Heisenberg, and Weisskopf [13]: it seeks to achieve this difficult goal by repeatedly bouncing a polarized beam of light back and forth through a strong magnetic field, and measuring the ensuing polarization

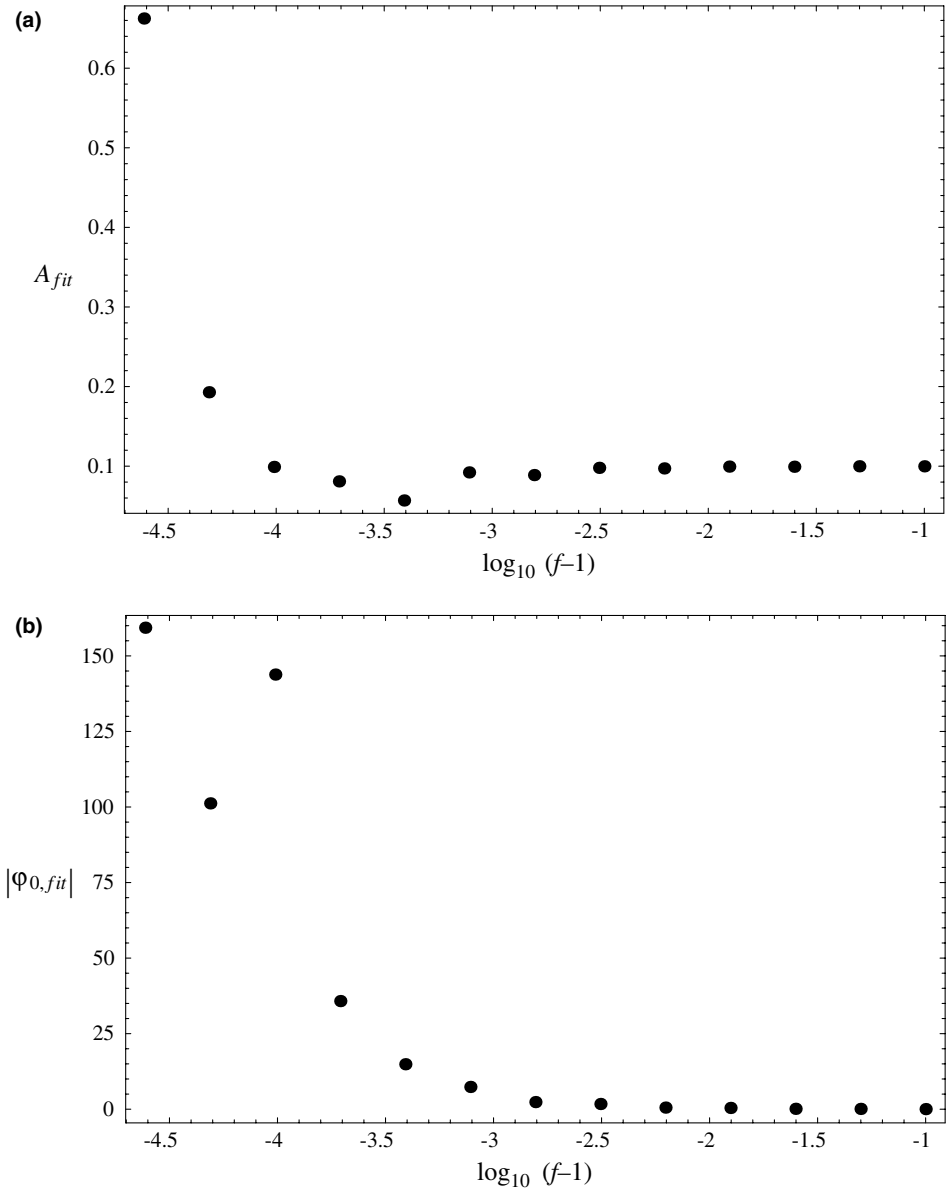

Fig. 7. Test of contamination from a nearby strong unfitted peak: the input model has two sines (amplitudes 0.1 and 1 a.u., phase constants 0° and 90°, respectively) and no noise, while the fit model has just one sine; the weaker peak (fitted) has fixed frequency 1 AFU, while the stronger peak has variable frequency f AFU. (a) Plot of the reconstructed amplitude vs. the logarithm of the frequency difference $\log(f-1)$ and (b) plot of the reconstructed phase constant vs. the logarithm of the frequency difference $\log(f-1)$. In this case, contamination is larger and is felt as far as 20–30 FFT bins.

change. The QED vacuum loaded by the magnetic field should act like a uniaxial birefringent medium. The apparatus is schematically shown in Fig. 10: a stabilized laser emits a beam that is linearly polarized by a polarizer, then it enters a Fabry–Perot cavity, and as it exits the cavity it goes through a polarization

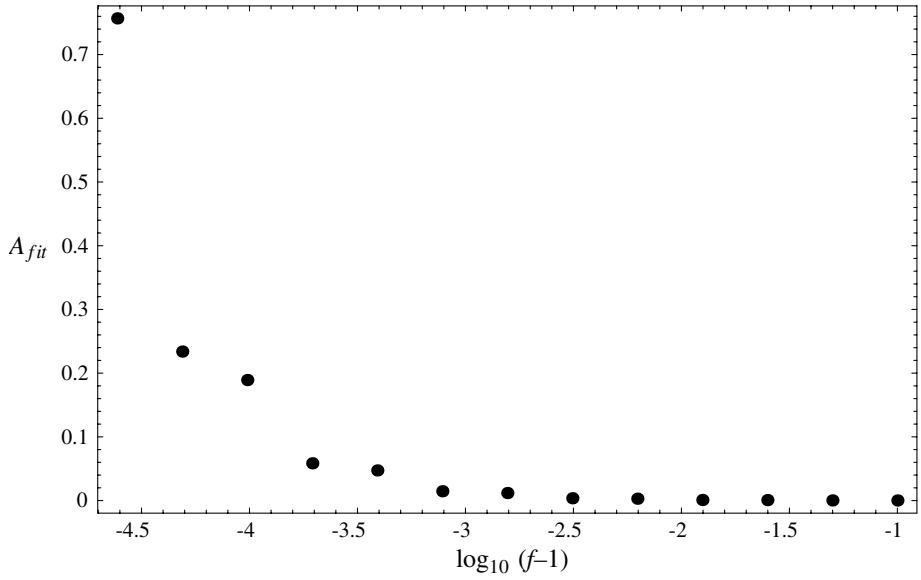

Fig. 8. Another test of contamination from a nearby unfitted peak: in this case the input model has just one sine (amplitude 1 a.u., phase constant 0°), while the fit model has one sine at the wrong frequency (fixed at 1 AFU). Here we see that the presence of the close unfitted sine leads to a wrong algorithmic response just as in Figs. 6 and 7; however here contamination produces a response to a wrong hypothesis, while in Figs. 6 and 7 it leads to wrong fit parameters.

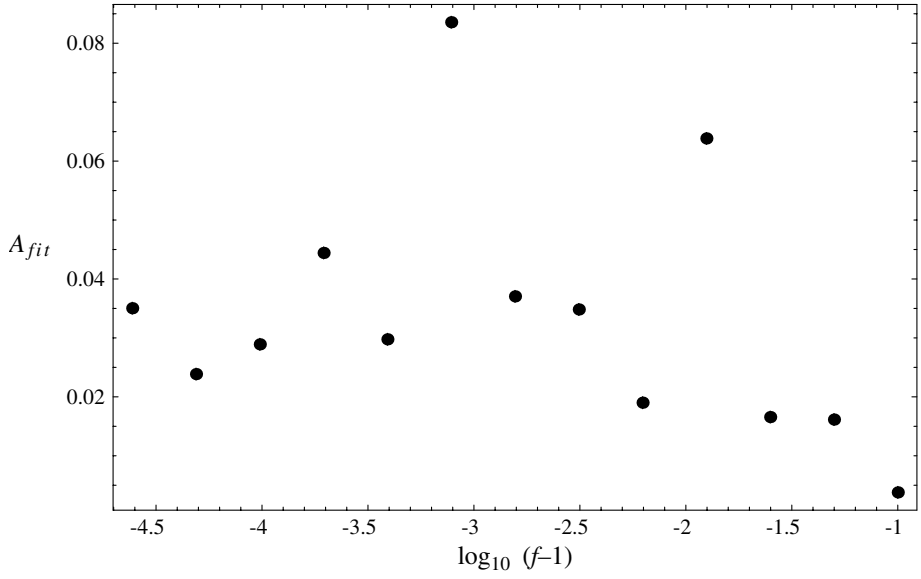

Fig. 9. Yet another test of contamination from a nearby unfitted peak as in Fig. 8: in this case the input model has just one sine (amplitude 1 a.u., phase constant 0°), while the fit model has one sine at the wrong frequency (fixed at 1 AFU); in addition the fit hypothesis includes a large phase jitter (±5°). Here the large phase fluctuation leads to a wider contamination window because phase fluctuations increase leakage, however contamination is reduced overall because the disturbing sine has a different phase function that can never quite adapt to the assumed phase function.

Fig. 10. Simplified layout of the PVLAS ellipsometer setup (see main text for an explanation).

modulator and finally through an analyzer: at the end the beam is detected by a low-noise photodiode. The birefringence of vacuum induced by the field is modulated at low frequency by mechanically rotating the large dipole magnet, and a higher frequency birefringence modulation is added by the polarization modulator. The photodiode current is proportional to the light intensity [2,3]

$$I(t) \approx I_0\left[\sigma_P^2 + (\Psi(t) + \eta(t) + \Gamma(t))^2\right], \tag{23}$$

where σ_P is the polarizer extinction, $\Psi(t)$ is the tiny ellipticity due to the QED effect, $\eta(t)$ is the ellipticity due to the modulator, and $\Gamma(t)$ is the additional and slowly varying ellipticity introduced by the optical elements in the beam path. If we represent modulation as follows:

$$\Psi(t) = \Psi_0 \cos\left(2\varphi_L(t) + 2\varphi_{L,0}\right),$$
$$\eta(t) = \eta_0 \cos\left(\varphi_H(t) + \varphi_{H,0}\right),$$

where $\varphi_L(t) + \varphi_{L,0}$ is the phase angle of the rotating table that supports the dipole magnet, then, expanding (23), one finds [2,3] that the photodiode current is

$$\begin{aligned}\frac{I}{I_0} = {} & \sigma_P^2 + \tfrac{1}{2}\Psi_0^2 + \tfrac{1}{2}\eta_0^2 + \Gamma^2(t) + 2\Gamma(t)\Psi_0 \cos\left(2\varphi_L(t) + 2\varphi_{L,0}\right) + \tfrac{1}{2}\Psi_0^2 \cos\left(4\varphi_L(t) + 4\varphi_{L,0}\right) \\ & + \eta_0\Psi_0\left[\cos\left(\varphi_H(t) + 2\varphi_L(t) + \varphi_{H,0} + 2\varphi_{L,0}\right) + \cos\left(\varphi_H(t) - 2\varphi_L(t) + \varphi_{H,0} - 2\varphi_{L,0}\right)\right] \\ & + 2\eta_0\Gamma(t)\cos\left(\varphi_H(t) + \varphi_{H,0}\right) + \tfrac{1}{2}\eta_0^2 \cos\left(2\varphi_H(t) + 2\varphi_{H,0}\right).\end{aligned} \tag{24}$$

Eq. (24) is the physical model of the photodiode signal, and is of the form discussed in the previous sections: in order to fit it we must know first the phase functions $\varphi_L(t)$ and $\varphi_H(t)$. The rotating table that supports the magnet is driven by an oleodynamic drive and the mechanical rotation rate is fairly stable, and is characterized by small but significant random irregularities that depend both on the previous rotation history and on the environmental parameters: Fig. 11 shows a record of the rotation frequency in one data acquisition run (run 503, with about 1,300,000 samples taken with an ADC clock frequency of 8200 Hz). The position of the rotating table is monitored by means of a LED-photodiode pair that detects a series of 32 (almost)

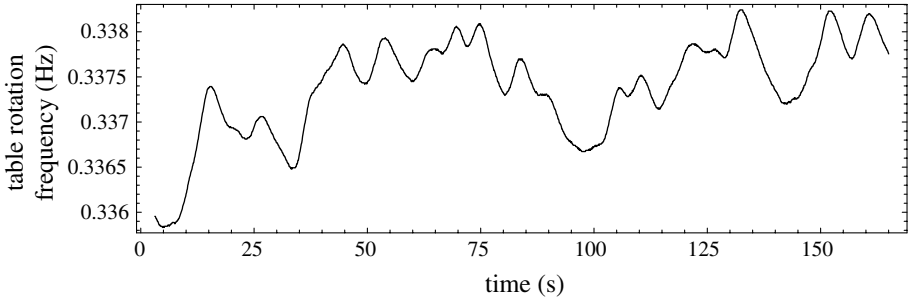

Fig. 11. Plot of the table rotation frequency during one short data acquisition run (run 503).

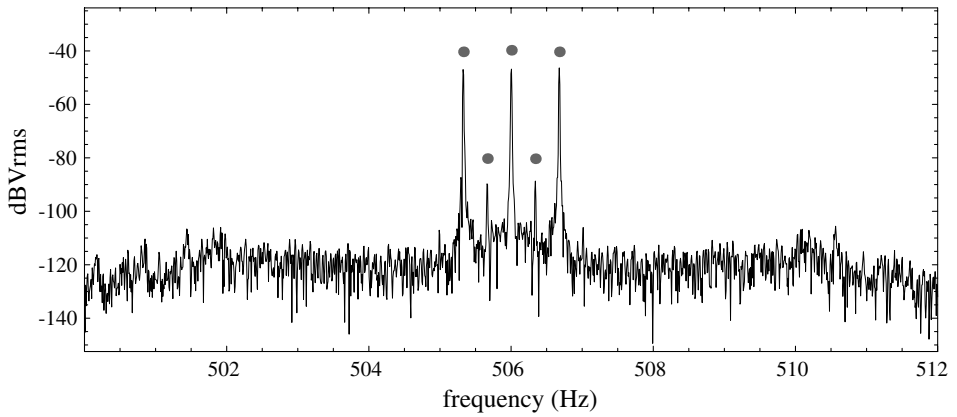

Fig. 12. Detail of the spectral density of the photodiode signal in one run with gas (run 503 with N_2 at low pressure); the central peak corresponds to the $2\eta_0\Gamma(t)\cos(\varphi_H(t)+\varphi_{H,0})$ term in Eq. (24), while the sidebands correspond to $\cos(\varphi_H(t)\pm 2\varphi_L(t)+\varphi_{H,0}\pm 2\varphi_{L,0})$ (as in Eq. (24)) and to $\cos(\varphi_H(t)\pm\varphi_L(t)+\varphi_{H,0}\pm\varphi_{L,0})$ (spurious signals). The dots mark the amplitude values found by the sine-fit procedure: part of the increase with respect to the FFT peaks is attributable to windowing (a Hamming window was used with the FFT), but not all of it, as it appears from the different gain obtained for the two pairs of sidebands.

evenly spaced ticks on the table edge. The tick signal is used to reconstruct the table phase angle $\varphi_L(t)$. The polarization modulator is driven by a precise sine-wave generator (HP3325B): the signal is very clean and it is fairly easy to precisely reconstruct the phase function $\varphi_H(t)$. Fig. 12 shows both a detail of the FFT estimate of the spectral density of the windowed signal, and the fit results. In this short run the table rotation frequency is very stable and the peaks are very strong: here the physical effect is mimicked by the Cotton-Mouton effect in nitrogen [14]. Both the physical peaks and the smaller peaks from spurious signals (due to mechanical vibrations in the apparatus) have been fit, and in this case there is an excellent agreement between the FFT estimate and the sine-fit procedure (most of the FFT peak reduction is due to windowing). This is not always the case, especially when the physical signal is very weak and when the table rotation frequency is not as smooth as here.

7. Conclusions

The algorithm studied in this paper is a variant of the sine-fit method and it provides an excellent amplitude and phase determination of selected spectral components. It has all the advantages of similar methods: samples do not have to be evenly spaced, samples may be any number and not a power of 2, the method provides an error estimate as well as an accurate determination of amplitude and phase. The numerical tests show that the algorithm performs reasonably well even when the theoretical model is not well defined. All the tests in the paper have been performed with (almost) evenly spaced points and have been compared to the ordinary FFT because the FFT representation is so familiar, but it is clear that in extreme cases neither the FFT nor other methods that require evenly spaced samples, like AR models [1], may compete with this sine-fit algorithm. The generalized sine-fit procedure successfully returns both amplitude and phase of all the components present in the model when the time-dependent part of the phase functions is given. It is at fault when the model does not include some component which is both strong and very close to the sines included in the model. However, the greater the difference between the different phase functions, the better becomes the rejection of unwanted components. As an alternative, one might try to use Bayesian methods [15], but these would not bring any improvement: in fact the sine-fit procedure described here is a model containing all of the available prior information and a Bayesian method could at best be equivalent to it.

Acknowledgement

I thank all the members of the PVLAS collaboration and especially Emilio Zavattini, for many useful discussions.

Appendix A. Phase determination with the FFT

Let

$$f_k = A_0 \cos(\omega_0 t_k + \varphi_0) \tag{A.1}$$

be a fixed-frequency signal, and assume a sampling interval Δt, so that $t_k = k\Delta t$ and the sampling frequency is $\nu_S = 1/\Delta t$. Then, if we take N evenly spaced samples, the frequency resolution of an FFT analysis is $\Delta\nu = \nu_S/N$, and using $\Delta\nu$ as frequency unit we see that the angular frequency ω_0 corresponds to $n_0 = \omega_0/\Delta\omega = \nu_0/\Delta nu = N\nu_0/\nu_S$. In general n_0 is not an integer, and

$$f_k = A_0 \cos(\omega_0 t_k + \varphi_0) = A_0 \cos\left(\frac{\omega_0}{\Delta\omega}\Delta\omega k\Delta t + \varphi_0\right) = A_0 \cos\left(\frac{2\pi k n_0}{N} + \varphi_0\right). \tag{A.2}$$

The DFT of the signal (A.2) is easily shown to be

$$\begin{aligned} F_n &= \frac{1}{\sqrt{N}} \sum_k A_0 \cos\left(\frac{2\pi k n_0}{N} + \varphi_0\right) \exp\left(-\frac{2\pi k n}{N}\right) \\ &= \frac{A}{2\sqrt{N}} \left\{ e^{i[\varphi_0 + \pi(n-n_0) + \pi(n-n_0)/N]} \frac{\sin(\pi(n-n_0))}{\sin(\pi(n-n_0)/N)} + e^{i[-\varphi_0 - \pi(n+n_0) + \pi(n+n_0)/N]} \frac{\sin(\pi(n+n_0))}{\sin(\pi(n+n_0)/N)} \right\} \end{aligned} \tag{A.3}$$

so that the phase associated to each FFT bin in the proximity of the two peaks is a linear function:

$$\varphi_+(k) = \varphi_0 + \pi(n-n_0) + \pi\frac{n-n_0}{N}, \tag{A.4}$$

$$\varphi_-(k) = -\varphi_0 - \pi(n+n_0) + \pi\frac{n+n_0}{N}. \tag{A.5}$$

From Eqs. (A.4) and (A.5), it would seem possible to extract the phase information of signal (A.1), but in actual practice this is quite difficult or nearly impossible because of the presence of many other peaks and background noise, and because the actual determination of phase can only be made mod 2π. In fact the determination of phase requires prior determination of the *exact peak position* n_0, otherwise even a very small error leads a very large phase error: for instance, in the case of the PVLAS data, n_0 is of the order of 10^6, thus, even a relative error of 2×10^{-6} causes a phase error of the order of 2π and total ignorance about the actual (folded) phase φ_0. In practice, the background noise and the presence of other close peaks makes this method unfeasible, and therefore the FFT is not capable of providing any useful phase information in an experimental situation like that of PVLAS.

References

[1] S.M. Kay, S.L. Marple, Spectrum analysis – a modern perspective, Proc. IEEE 69 (1981) 1380.
[2] E. Iacopini, B. Smith, G. Stefanini, E. Zavattini, On a sensitive ellipsometer to detect the vacuum polarization induced by a magnetic field, Nuovo Cimento 61B (1981) 21.
[3] D. Bakalov, et al., The measurement of vacuum polarization: The PVLAS experiment, Hyperfine Int. 114 (1998) 103.
[4] R. Cameron, et al., Search for nearly massless, weakly coupled particles by optical techniques, Phys. Rev. D 47 (1993) 3707.

[5] G.N. Stenbakken, J.P. Deyst, Time-base nonlinearity determination using iterated sine-fit analysis, IEEE Trans. Instrum. Meas. IM-47 (1998) 1056.
[6] N.R. Lomb, Further development and properties of the spectral analysis by least-squares, Astrophys. Space Sci. 39 (1976) 447; J.D. Scargle, Studies in astronomical time series analysis. ii statistical aspects of spectral analysis of unevenly spaced data, Astrophys. J. 263 (1982) 835.
[7] R.N. McDonough, W.H. Huggins, Best least-squares representation of signals by exponentials, IEEE Trans. Aut. Cont. AC-13 (1968) 408.
[8] M.R. Osborne, G.K. Smyth, A modified prony algorithm for exponential function fitting, SIAM J. Sci. Comp. 16 (1995) 119.
[9] J.H. Horne, S.L. Baliunas, A prescription for period analysis of unevenly sampled time series, Astrophys. J. 302 (1986) 757.
[10] F.J. Beutler, Alias-free randomly timed sampling of stochastic processes, IEEE Trans. Inf. Theory IT-16 (1970) 147.
[11] In practice we assume a white noise background. We also assume that any additional noise, like the ADC quantization noise is also uncorrelated, and that its statistical distribution is well approximated by a Gaussian distribution.
[12] W.T. Eadie, D. Drijard, F.E. James, M. Roos, B. Sadoulet, Statistical Methods in Experimental Physics, North-Holland, Amsterdam, 1982.
[13] W. Heisenberg, H. Euler, Folgerungen aus der Diracschen Theorie des Positrons, Z. Phys. 98 (1936) 714;
V. Weisskopf, Über die elektrodynamik des vakuums auf grund der quantentheorie des elektrons, Kgl. Danske Videnskab. Selskabs. Mat.-fys. Medd. 14 (1936) 6.
[14] C. Rizzo, A. Rizzo, D.M. Bishop, The Cotton-Mouton effect in gases: experiment and theory, Int. Rev. Phys. Chem. 16 (1997) 81.
[15] G.L. Bretthorst, Bayesian Spectrum Analysis and Parameter Estimation, Springer-Verlag, New York, 1988.