PARAMETER ESTIMATION FOR LINEAR AM/FM SINUSOIDS USING FREQUENCY DOMAIN DEMODULATION

Axel Röbel
IRCAM-CNRS-STMS, Analysis-Synthesis Team, Paris, France
email : roebel(at)ircam.fr

ABSTRACT

This article deals with the estimation of sinusoidal parameters for non stationary sinusoids. It will be shown that for linear amplitude and frequency modulation only the frequency modulation creates additional estimation bias for the standard sinusoidal parameter estimator. Then a new algorithm for frequency domain demodulation of spectral peaks is proposed that can be used to obtain an approximate maximum likelihood estimate of the frequency slope, and an estimate of the amplitude, phase and frequency parameter with significantly reduced bias. An experimental evaluation compares the new estimation scheme with some previously existing methods. It shows that significant bias reduction is achieved for a large range of slopes and zero padding factors. A real world example demonstrates the benefits of the new method.

KEY WORDS

sinusoidal parameter estimation, estimation bias, FM demodulation, AM/FM sinusoid.

1 Introduction

The sinusoidal signal model is widely used for signal analysis and/or signal transformation of speech and music signals [1, 2]. An important step for the creation of a sinusoidal model is the estimation of the sinusoidal parameters (amplitude, frequency and phase) from the spectral peaks of the Fourier transform. This estimation is straightforward as long as the sinusoids are stationary. The log amplitude and unwrapped phase of the maximum bin of the spectral peak and its two neighbors are interpolated by means of a quadratic polynomial and the height and frequency position of the maximum of the polynomial provide amplitude and frequency estimates. The phase is obtained from the phase polynomial evaluated at the frequency position determined in the previous step. This estimator is denoted as quadratically interpolated FFT (QIFFT) estimator [3].

For non stationary sinusoids the parameter estimation becomes more difficult. The QIFFT algorithm is severely biased whenever the frequency is not constant. The term bias refers to the systematic estimation error. It describes the offset of the estimator that exists even if no measurement noise is present. For the high order partials in natural vibrato signals the estimation bias of the QIFFT estimator accounts for a significant amount of residual energy. It is the major reason for the perceived voiced energy in the residual of vibrato signals. A number of algorithms with low estimation bias for non stationary sinusoids have been proposed. Some of these algorithms implement a maximum likelihood estimate (MLE) by means of a grid search for the demodulator of the frequency rate that maximizes the amplitude of the demodulated signal [4]. These algorithms use time domain demodulation, and therefore they cannot be applied if the signal contains more than a single sinusoid. Algorithms for multi-component signals generally rely on the fact that the analysis window is approximately Gaussian such that a mathematical investigation becomes tractable [5, 6, 3]. The method presented in [3] is special in that it tries to extend its range to other analysis windows by means of a set of linear bias correction functions. The resulting estimator is computationally very efficient and achieves small bias for standard windows as long as the zero padding factor is sufficiently large (≥ 3) and the frequency chirp rate is relatively small.

In the following we present a bias correction scheme for sinusoidal parameter estimation of non stationary sinusoids. First it will be shown that demodulation of the frequency variation is the key to bias reduction when the sinusoid is non-stationary. We then propose to use frequency domain demodulation to achieve an approximate MLE of the frequency slope and the other partial parameters. The main advantage of the frequency domain demodulation is the fact that the method can be applied to multi-component signals. After demodulation a standard QIFFT algorithm can be used to find the fundamental parameters. Due to the fact that the QIFFT estimator has small bias for constant frequency sinusoids the resulting estimate is significantly improved. A simple version of the algorithm has been presented in [7]. It has been shown that the demodulation can be achieved by means of spectral deconvolution using only the peak to be analyzed and a properly selected and scaled demodulation kernel. In the original version the frequency slope estimate was entirely handled by the frequency slope estimator in [3].

The version to be presented here is a refined version of the original demodulation algorithm. The enhancements include a new procedure to improve the initial estimate of the frequency slope reducing the remaining bias for large frequency slopes. Furthermore, the constraint to use the same analysis window for the signal spectrum and the demodulation kernel has been removed. Accordingly, it becomes possible to trade-off bias against noise sensitivity. A computationally efficient version of the algorithm using precomputed and linearly interpolated demodulation kernels is presented.

The organization of the article is as follows. In section **2** we will show how the bias of the standard estimators is related to the frequency slope. In section **3** we will describe the demodulation scheme and the approximate ML frequency slope estimator. In section **4** we present experimental results and a comparison of the frequency demodulator estimator with existing algorithms. In section **5** we conclude with an outlook on further improvements.

2 Estimation bias

The signal model that will be used here assumes a linear evolution for amplitude and frequency trajectories such that a complex discrete time sinusoid can be represented as

$$s(n) = (A + an)\exp(i(\phi + 2\pi\omega n + \pi D n^2)). \qquad (1)$$

Here A is the mean amplitude of the signal and a is the amplitude slope. ϕ is the phase of the sinusoid at time $n = 0$, ω is its mean frequency and D is the frequency slope. Note, that all frequency values are normalized with respect to the samplerate. The center of the analysis window is is located at time 0 such that an ideal estimator should provide (A, ω, ϕ) as estimates for amplitude, frequency, and phase.

To start the investigation we summarize the sources of bias that are known to exist for the standard QIFFT estimator and discuss there implications in the current context.

First, there is the use of a second order model for interpolating the spectral bins. While this is systematically wrong for the present sinusoidal model, it does not have any direct relation to the fact that the sinusoidal parameters are varying. Because the QIFFT algorithm will be used extensively, it is nevertheless important to reduce this type of bias as far as possible. This can be achieved by means of zero padding the analysis window or, as demonstrated recently, by means of simple bias correction functions [8].

Second, there is the bias due to other sinusoidal components. The technique that is generally used to reduce this bias is windowing. The analysis window reduces the sidelobes of the sinusoidal components such that the cross component bias of distant sinusoidal components can be effectively reduced. Note however, that the reduction of the sidelobe amplitudes always is accompanied by an increased mainlobe width. Accordingly windowing increases the cross component bias for nearby components. In the following we will assume that the sinusoidal components are resolved such that the frequency distance between two sinusoids is always larger than the width of the mainlobe of both components. In this case the cross component bias will stay nearly the same for stationary and non-stationary components.

Third, there is the bias due to the non-stationary parameters. For the sinusoidal model in eq. (1) and a Gaussian analysis window the bias has been analyzed mathematically in [6]. The result shows, that the QIFFT algorithm suffers from additional bias due to parameter variation only if the frequency slope $D \neq 0$. In this case, the estimation of all three basic parameters are biased and the bias increases with the absolute value of D.

To study the bias for arbitrary analysis windows we split the sinusoidal model in eq. (1) into two parts, a sinusoid with constant amplitude A and sinusoid with mean amplitude 0 and amplitude slope a. Then we investigate the properties of the spectra of the individual parts and use the linearity of the Fourier transform to draw conclusions for the sinusoidal spectrum. We assume the coordinate system of the amplitude and phase spectra to be shifted such that its frequency origin is located at the sinusoidal frequency ω. For $D = 0$ the amplitude of the spectra of both parts will be an even function with the spectrum of the second part being 0 at the origin. The phase of the first part is even (constant equal to ϕ within the mainlobe) and the phase of the second part is odd (constant equal to ϕ with a phase jump of π at the origin. Combining these facts we can conclude that the sum of these spectra has even amplitude and as long as the slope is sufficiently small such that the final spectrum contains a mainlobe the QIFFT estimator provides results that are biased only by the first two sources of bias mentioned above. We conclude that the time varying amplitude by itself does not add any additional bias.

When $D \neq 0$ the phase of both parts are modified. It is easy to verify that the constant amplitude component will no longer have constant phase and that its phase will be different from ϕ at the origin. This creates a bias in the phase estimator. Moreover, when $a \neq 0$ the amplitude spectra keep their symmetry properties. The phase spectra of both parts have a non constant but even symmetric shift superimposed. By consequence the amplitude spectrum of the complete signal will have its amplitude maximum been shifted away from the origin. Accordingly, the QIFFT estimator suffers from additional bias quite similar as has been shown for the Gaussian window in [6].

3 Reducing the bias

In the previous section we saw that the source of the bias of the QIFFT estimator is the frequency slope of the sinusoid. A simple approach to estimate the parameters (A, ϕ, ω) of a sinusoid related to a spectral peak requires two steps:

1. estimate the frequency slope,

2. demodulate the sinusoid and use the QIFFT estimator to find the sinusoidal parameters.

Note, that this approach is in principle equivalent to the MLE described in [4]. Because the demodulation is used during the refinement of the frequency slope estimate we will first discuss the frequency domain demodulation algorithm. In a second step the frequency slope estimation is described.

3.1 Demodulation

The main objective of the present algorithm is to provide a means to demodulate the sinusoid using only the observed part of the spectral peak. Initially, we assume we are given a frequency slope estimate $\hat{D} = D$ for a peak that is part of a signal spectrum.

In time domain the demodulation can be achieved simply by multiplication with a demodulator signal

$$y(n) = \exp(-i\pi\hat{D}n^2). \tag{2}$$

Multiplication of the signal in eq. (2) with the signal eq. (1) will remove the frequency slope and keep all other parameters unchanged such that the QIFFT algorithm can be applied. However, the signal we are interested in is observable only via the part of its mainlobe that constitutes the observed spectral peak.

The demodulation algorithm that uses the observed peak to demodulate the sinusoid will be described in the frequency domain using as sources the spectral peak to be analyzed and the spectrum of the deconvolution signal. Assume $S(k)$ is the N-point DFT of the sinusoid to be analyzed and $Y(k)$ the DFT of the demodulator signal. All DFT spectra are calculated such that the origin of the DFT basis functions is in the center of the analysis window. The signal analysis window is $w_s(n)$ and the demodulator signal is windowed using $w_y(n)$. To obtain the demodulated sinusoid spectrum $X(k)$ we would need to compute the circular convolution

$$X(k) = C\frac{S(k) \circledast Y(k)}{N}, \tag{3}$$

where C is a normalization factor taking into account windowing effects. The demodulator window w_d will be multiplied with the signal window such that the resulting spectrum contains as effective window the product window $w_y(n)w_s(n)$. Therefore, proper normalization would be achieved by means of setting $C = 1/\sum_n(w_y(n)w_s(n))$.

Due to the fact that only part of the sinusoid spectrum is available the normalization factor needs to be adapted. Assume the peak under investigation is denoted by $P(k)$. $P(k)$ is part of the spectrum $S(k)$ and covers B bins. To estimate the impact of the missing part we create a spectral model of the observed sinusoid assuming the initial slope estimate is correct

$$P_m(k) = \sum_n w_s(n)\exp(i\pi\hat{D}n^2)\exp(-\frac{2\pi j}{N}kn), \tag{4}$$

and select a subset $\bar{P}_m(k)$ of B bins around the center frequency $k = 0$. [1] The required normalization factor can now be approximately estimated as

$$C = \frac{1}{\max_k(|\bar{P}_m(k) \circledast Y(k))}. \tag{5}$$

Now we can replace $S(k)$ in eq. (3) by $P(k)$ and demodulate using the corrected normalization factor C. Some remarks are in order:

- The correction factor will be more precise (=lower bias) for demodulator windows that concentrate more energy in the B-bin wide band around frequency 0 of the spectrum. This calls for low side lobes. The demodulator window, however, will as well be applied to the signal. As a result the estimator sensitivity to noise will increase. Accordingly the demodulator window allows to trade-off noise sensitivity and bias. The experimental investigation suggests that the use of the Hanning window as demodulator window w_d is a favorable choice for all analysis windows w_s.

- It is preferable if the spectral peak $P(k)$ that is used for demodulation is delimited by amplitude values of the same magnitude. Otherwise the additional asymmetry that is due only due to cutting the peak out of its complete spectrum will add to the estimation bias that is due to the fact that only part of the peak is used for demodulation. In our experiments we found significant improvements for low and mid SNR if we enforced to cut the spectral peak such that left and right amplitude border are at approximately the same level.

- For parameter estimation from demodulated peaks with the QIFFT estimator it is essential to use the bias correction functions proposed in [8] with correction factors adapted to the effective window $w_y(n)w_s(n)$.

Our investigation shows, that the demodulation kernels $Y(k)$ can be pre-calculated for a fixed grid of frequency slopes and then linearly interpolated to obtain an approximate demodulation kernel for any given slope. If the length of the analysis windows is M a frequency slope grid with step size $0.025/M^2$ is sufficient to produce estimates that are nearly indistinguishable from the results produced with the non interpolated kernels. To use the complete information that is available in the observed peak we use deconvolution kernels of length $2B+1$ centered around the maximum of the deconvolution spectrum.

3.2 Frequency slope estimation

As mentioned above the maximum likelihood (ML) frequency slope estimator maximizes the amplitude of a demodulated peak [4]. Accordingly the maximization of the amplitude of the demodulated peak using the demodulation algorithm described above is a means to achieve an approximate MLE.

To avoid the search of a large grid of frequency slopes we propose to use an approximate initial estimate of the frequency slope $\hat{D}$, and to use this slope estimate and two slopes with $\hat{D} \pm D_o$ to create three different demodulations of the observed peak. From the amplitudes of these demodulated peaks a 2nd order polynomial model of the relation between frequency slope and demodulated amplitude can be derived. The maximum of this polynomial is expected to provide a refined estimate of the frequency slope.

The open question we need to address is: how do we get an approximate estimate of the frequency slope? Given

[1] If B is even the resulting model is not symmetric!

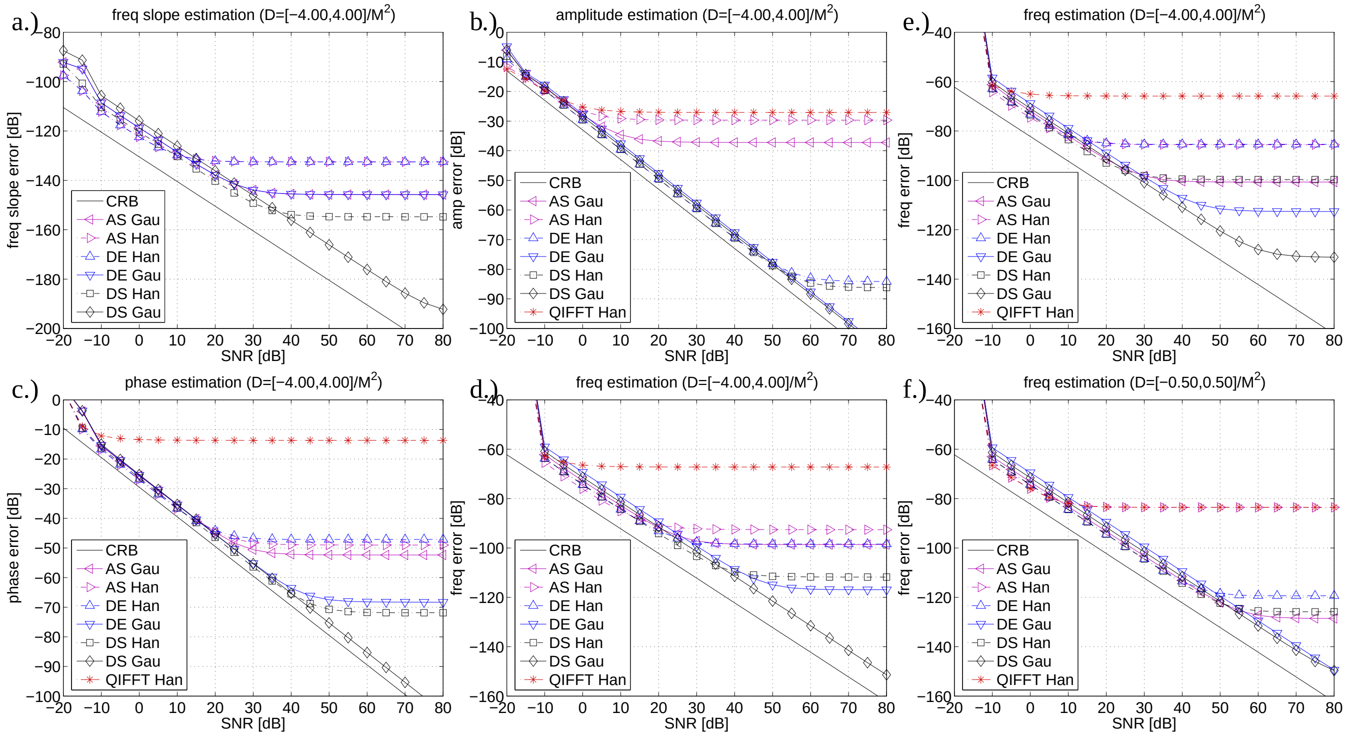

Figure 1. Comparison of the estimation errors for the different parameter estimators using window size $M = 1001$ and sinusoids with strong (a-e) and weak (f) amplitude and frequency slope parameters. DFT size is $N = 4096$ (a-d), and $N = 1024$ (e-f), The CRB for constant amplitude polynomial phase signals is displayed as lower limit. Algorithms using a Gaussian/Hanning window are distinguished by means of solid/dashed lines. See text form more details.

the highest order coefficients α_ϕ and α_A of the QIFFT polynomial for amplitude (A) and phase (ϕ) of the peak under investigation the frequency slope estimate for a Gaussian analysis window is [3, 9]

$$\hat{D} = \frac{\alpha_\phi}{\alpha_\phi^2 + \alpha_A^2}. \tag{6}$$

Note, the remarkable fact that the same estimator has been obtained for exponential amplitude evolution in [3] and for a first order approximation of the spectrum of a sinusoid with linear amplitude evolution in [9]. The fact that the amplitude evolution does have only a very small impact on the frequency slope estimator leads us to suppose that that eq. (6) will provide useful estimates for other windows than the Gaussian window as well. The argument here is that the signal that is obtained after the analysis window has been applied can always be considered to be equivalently generated by means of a Gaussian analysis window and a sinusoid with appropriately modified amplitude evolution. Because the desired frequency estimate does not change with the amplitude evolution of the sinusoid and because the estimator eq. (6) appears not to be sensitive to small changes of the amplitude evolution of the sinusoid it will be considered as approximate estimator for the frequency slope for arbitrary analysis windows (see section 4).

The free parameter to select is the frequency slope offset D_o. In general a polynomial approximation improves when the approximation range is decreased. This would call for a small D_o. In the present case, however, the relation between demodulation slope and amplitude of the demodulated peak is covered by measurement noise (due to estimation errors of the amplitude of the demodulated peak, due to the partially observed sinusoidal spectrum, and due to the sampling of the Fourier spectrum by the DFT). Selecting the slope offset from the range $D_o = [0.1/M^2, 1/M^2]$ provides a good compromise between validity of the 2nd order model and reduction of the impact of the measurement noise. For the examples in the experimental section we used $D_o = 0.3/M^2$. The refined frequency slope estimate is then again used to demodulate the observed peak. This final demodulation yields the complete set of estimates for amplitude, frequency and phase.

The precision of the frequency slope estimate that is obtained from the maximum of the polynomial is slightly, but consistently improved if the polynomial model is not constructed for the demodulated amplitudes $\hat{A}_i$ but for $\hat{A}_i\sqrt{C_i}$ where C_i is the normalization factor from eq. (5). Up to now a theoretical explanation of this experimental finding has not yet been found. Using $\sqrt{C}$ to calculate the demodulated amplitudes will obviously create biased amplitude estimates. For the problem of slope estimation it appears to improve the fit of the polynomial model and therefore, it will be preferred. After the slope has been determined from the maximum of the polynomial a renormalization can be performed if the amplitude of the supporting points is required.

4 Experimental evaluation

The proposed parameter estimation procedure will be evaluated by means of comparing it to the bias correction estimate proposed in [3] (denoted as *AS*) for which Gaussian and Hanning analysis windows are used. Furthermore we use the original version of the demodulation estimator according to [7]. The results of this algorithm are denoted as *DE*. The new version that includes the slope enhancement and uses the Hanning window for all demodulation kernels is denoted as *DS*.

The window type that is used will be indicated by adding the letter *G* for Gaussian or *H* for Hanning or *X* for both to the estimator shortcut. The Gaussian analysis window is cut such that it has a length of 8σ with σ being the standard deviation of the Gaussian. The results of the QIFFT estimator are shown as reference as well as the Cramer Rao bounds for second order polynomial phase estimation that have been presented in [10]. Note, however, that these bounds have been found for constant amplitude polynomial phase signals, such that they can only be used to provide an approximate idea of the estimator efficiency.

In the experiments we use synthetic test signals with a single sinusoid according to eq. (1) with $A = 1$, ω randomly sampled from a uniform distribution over the frequency range $[0.2, 0.3]$, ϕ randomly chosen from a uniform distribution between $[-\pi, \pi]$, and varying slopes a and D. The analysis window covers $M = 1001$ samples in all cases.

In the first experiment displayed in fig. 1 (a.-d.) the frequency slope D is selected from a uniform distribution over the interval $[-4/M^2, 4/M^2]$. Similarly the amplitude slope a is sampled from a uniform distribution over the range $[-1/M, 1/M]$. The slope ranges are considered realistic for real world signals. Note, that in harmonic signals the frequency slope scales with the partial number such that for high partials extreme slopes may arise. The FFT size that has been used for this experiment is $N = 4096$.

Due to the equivalence between the frequency slope estimators used in ASH and DEH as well as in ASG and DEG these algorithms show no difference with respect to frequency slope estimation in fig. 1 (a). As expected the Hanning window estimators have larger estimation bias which is due to the fact that the slope estimation from eq. (6) is used for a non Gaussian windows. The bias is increased only by 15dB. For this setup the DSX estimators have by far the lowest bias, however, are slightly more sensitive to noise by about 2-3db.

As expected the amplitude estimate (b) of ASX is strongly biased due to the fact that the amplitude trajectory model does not match the signal. DEX and DSX are both similar and better then ASX. Note, that the improved frequency slope estimate of DSX hardly improves the amplitude estimate compared to DEX and that the increase of the noise sensitivity of DEX and DSX is negligible. For frequency d.) and phase estimation c.) DSX has by far the smallest bias (compared to the other estimators using the same analysis window). DEH and ASH perform approximately similar for both for frequency and phase estimation . Given that DEX and ASX estimators both use the same frequency slope estimate this shows that the bias of these two estimators is due to the error in the frequency slope estimate which is improved by the refined slope estimate of DSX.

The increase of the noise sensitivity for the demodulation algorithms is negligible for phase estimation. For the frequency estimator the use of the Hanning window instead of the analysis window is clearly diminishing the noise sensitivity when the analysis window is Gaussian.

There are 2 more results provided to demonstrate the robustness of the method with respect to small zero padding. In both examples FFT size is reduced to an zero padding factor of approximately 1 ($N = 1024$). Due to space constraints only the frequency estimation will be discussed as an example. Besides the reduced FFT size the experiment in plot e.) is identical to the one in d.). While the estimation bias is increased for all algorithms, DSX still has the lowest bias. Noise sensitivity is hardly changed compared to experiment d.). For the last experiment f.) the slope ranges are reduced to $a \in [-0.15/M, 0.15/M]$, and $D \in [-0.5/M^2, 0.5/M^2]$. These settings approximately reflect the range of values that has been used to select the correction parameters for the ASH estimator [3]. The small zero padding factor, however, is a problem for this estimator and the experiment confirms that for small zero padding factors the ASH estimator cannot outperform the standard QIFFT estimator. The DSX estimator again has the smallest bias and is significantly better than the DEX estimator.

4.1 A real world example

To demonstrate that the advantages of the proposed estimator are effective in real world situations we have implemented the bias reduction methods in a complete additive modeling system. The theoretical investigation has been restricted to cover the case of resolved sinusoids, only. For real world applications, however, the algorithm has to prove that it will act gracefully when the underlying model no longer holds (transients, unresolved sinusoids due to reverberation, ...). The major problem we encountered is the enhanced frequency slope estimation described in 3.2. When the underlying signal model is no longer valid the relation between frequency slope and amplitude of the demodulated peak becomes arbitrary and quite often the method degenerates and tries to model transients or nearby sinusoids by means of extreme slopes.

To prevent these situations we have adopted a simple strategy. First we test whether the polynomial model of the relation between amplitude and slope indicates a maximum or a minimum amplitude for the optimal slope. Second we limit the slope offset that is allowed to happen to achieve the maximum amplitude. If either the amplitude is a minimum

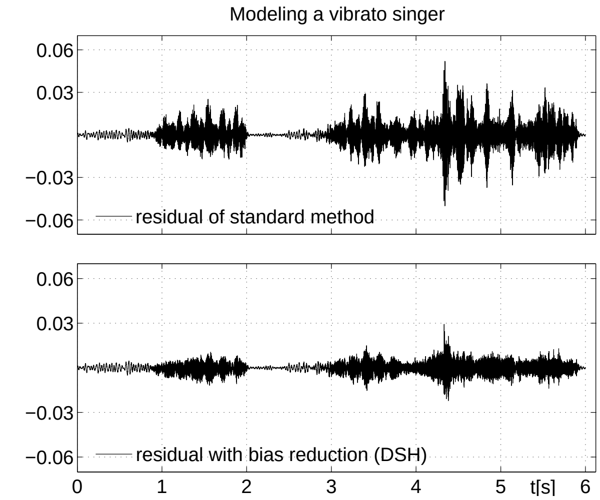

Figure 2. Residual signal of a vibrato tenor singer using QIFFT estimator (top) and the enhanced demodulation method DSH (bottom).

or the required slope offset is outside the user supplied threshold (here $\pm 2D_o$) the DSH method can no longer be applied and we use the DEH method as a fallback. Note, that the implementation of the algorithm uses linear demodulation kernel interpolation as discussed in section 3.1.

We compare the estimators by means of the residual energy of an harmonic model of a tenor signal. The signal contains strong vibrato, and therefore, the bias due to the non-stationary parameters is significant. The harmonic models contain 30 sinusoids and are estimated using a window size of 800 samples and FFT size 4096. We calculated the variance of the residual signal for the QIFFT, DEH, DSH and ASH methods. Due to the fact that is different in the different frequency bands the reduction of the residual error changes with the frequency band. Compared to the QIFFT estimator the reduction of the residual energy is rather significant. For the complete signal we get a reduction of $4.19/4.72/5.04$dB for ASH/DEH/DSH respectively. For the band 2-4kHz the reductions are $7.32/8.4/9.33$dB for ASH/DEH/DSH respectively. The residual signals for the QIFFT and DSH estimator are shown in figure fig. 2. The reduction of the residual is clearly visible.

5 Conclusions

In the present paper we have shown that an efficient bias reduction strategy for estimation of sinusoidal parameters consists in a frequency slope estimation and demodulation prior to application of the standard QIFFT estimator. The procedure significantly reduces the bias of the standard estimator. It can work with arbitrary windows and does show much less dependency on the zero padding factor than a recently proposed algorithm. The computational costs are significantly higher then those for the standard estimator ($\approx$ factor 10). However, they are sufficiently low such that real time estimation of some tenth of sinusoids from audio signals can be achieved. A real world example with a vibrato signal has shown that compared to the standard estimator the enhanced slope demodulation estimator can in some frequency bands achieve a residual error reduction of up to 9dB.

References

[1] X. Amatriain, J. Bonada, A. Loscos, and X. Serra. Spectral processing. In U. Zölzer, editor, *Digital Audiuo Effects*, chapter 10, pages 373–438. John Wiley & Sons, 2002.

[2] T. F. Quatieri and R. J. McAulay. Speech transformation based on a sinusoidal representation. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 34(6):1449–1464, 1986.

[3] M. Abe and J. O. Smith. AM/FM rate estimation for time-varying sinusoidal modeling. In *Proc. Int. Conf. on Acoustics, Speech and Signal Processing*, pages 201–204 (Vol. III), 2005.

[4] T. Abatzoglou. Fast maximum likelihood joint estimation of frequency and frequency rate. In *ICASSP*, pages 1409–1412 VOL. II, 1986.

[5] J. S. Marques and L. B. Almeida. A background for sinusoid based representation of voiced speech. In *Proc. Int. Conf. on Acoustics, Speech and Signal Processing*, pages 1233–1236, 1986.

[6] G. Peeters and X. Rodet. SINOLA: A new analysis/synthesis method using spectrum peak shape distortion, phase and reassigned spectrum. In *Proc. Int. Computer Music Conference*, pages 153–156, 1999.

[7] A. Röbel. Estimation of partial parameters for non stationary sinusoids. In *Proc. Int. Computer Music Conference (ICMC)*, 2006.

[8] M. Abe and J. O. Smith. Design criteria for the quadratically interpolated FFT method (I): Bias due to interpolation. Technical Report STAN-M-117, Stanford University, Department of Music, 2004. available at `http://ccrma.stanford.edu/STANM/stanms/stanm114/index.html`.

[9] G. Peeters. *Modèles et modification du signal sonore adapté à ses charactéristiques locales*. PhD thesis, Univertsité Paris 6, 2001. available at `http://recherche.ircam.fr/equipes/analyse-synthese/peeters/ARTICLES/Peeters_2001_PhDThesisv1.1.pdf`, french only.

[10] B. Ristic and B. Boashash. Comments on “The Cramer-Rao lower bounds for signals with constant amplitude and polynomial phase”. *IEEE Transactions on Signal Processing*, 46(6):1708–1709, 1998.