Journal of Magnetic Resonance **143,** 1–16 (2000)
doi:10.1006/jmre.1999.1960, available online at http://www.idealibrary.com on IDEAL®

Frequency-Selective Quantification of Biomedical Magnetic Resonance Spectroscopy Data

Leentje Vanhamme,* Tomas Sundin,† Paul Van Hecke,‡ Sabine Van Huffel,* and Rik Pintelon§

**Department of Electrical Engineering (ESAT), Katholieke Universiteit Leuven, Kard. Mercierlaan 94, 3001 Leuven, Belgium; †Systems and Control Group Uppsala University, P.O. Box 27, SE-75103 Uppsala, Sweden; ‡Biomedical NMR Unit, Katholieke Universiteit Leuven, Gasthuisberg, 3000 Leuven, Belgium; and §Department of ELEC, Vrije Universiteit Brussels, Pleinlaan 2, 1050 Brussels, Belgium*

E-mail: leentje.vanhamme@esat.kuleuven.ac.be

Received December 30, 1998; revised July 20, 1999

In this paper the possibility of obtaining accurate estimates of parameters of selected peaks in the presence of unknown or uninteresting spectral features in biomedical magnetic resonance spectroscopy (MRS) signals is investigated. This problem is denoted by frequency-selective parameter estimation. A new time-domain technique based on maximum-phase finite impulse response (FIR) filters is presented. The proposed method is compared to a number of existing approaches: the application of a weighting function in the time domain, frequency domain fitting using a polynomial baseline, and the time-domain HSVD filter method. The ease of use and low computational complexity of the FIR filter method make it an attractive approach for frequency-selective parameter estimation. The methods are validated using simulations of relevant ^{13}C and ^{31}P MRS examples.

Key Words: **MRS; frequency-selective quantification; FIR filter; AMARES; NLLS.**

INTRODUCTION

Accurate and efficient quantification of magnetic resonance spectroscopy (MRS) signals is the essential step prior to the conversion of the estimated signal parameters into biochemical quantities (e.g., concentration, pH). The quantification of the signal can be done directly in the measurement domain or alternatively in the frequency domain after transformation by the discrete Fourier transform (DFT).

The time-domain function often used to model the N measured data points is the following sum of exponentially damped complex sinusoids (Lorentzians):

$$y(n) = \hat{y}(n) + w(n) = \sum_{k=1}^{K} a_k e^{j\varphi_k} e^{(-\alpha_k + j2\pi f_k)n\Delta t} + w(n)$$

$$n = 0, \ldots, N-1, \qquad [1]$$

where $j = \sqrt{-1}$, a_k is the amplitude, φ_k the phase, α_k the damping, and f_k the frequency of the kth sinusoid ($k = 1, \ldots, K$); Δt the sampling interval, and w is circular complex white Gaussian noise. The caret on the y indicates that this quantity represents the model function rather than the actual measurements. The time-domain estimation methods can be divided into two classes. On the one hand, there are the so-called black-box methods (examples can be found in (*1–4*)). Minimal user interaction and limited incorporation of prior knowledge are inherent to this type of methods. Recent variants have been proposed in which some forms of prior knowledge can be imposed (*5–7*). On the other hand, interactive methods exist that are iterative, require user involvement, and allow inclusion of prior knowledge. These algorithms minimize the difference between the data and the nonlinear model function:

$$\min_{a_k,\varphi_k,\alpha_k,f_k} \sum_{n=0}^{N-1} |y(n) - \hat{y}(n)|^2. \qquad [2]$$

This nonlinear least-squares (NLLS) approach leads to maximum likelihood parameter estimates if the underlying assumptions concerning the model function and noise distribution are satisfied. VARPRO (*8*) and the more recent AMARES (*9*) are examples of this type of methods.

The frequency-domain methods can also be divided into two classes. The *nonparametric* frequency-domain methods are based on integration of the peak area of the frequency-domain signal (*10*). The advantage of these methods is that no assumptions have to be made concerning the signal, but the accuracy of the integration-based methods requires appropriate phasing (which is far from trivial) and is dependent on the definition of the integration area, especially in those cases where the peaks are not well separated, when missing time-domain data or acquisition artifacts, or when unidentified broad resonances distort the frequency-domain baseline. More advanced *parametric* frequency-domain methods have been presented that rely on a model function for the metabolite peaks (see, e.g., (*11–13*)). The methods are often based on the time-domain model given in Eq. [1]. The model is transformed into a frequency-domain model by the DFT resulting in the expression

1090-7807/00 $35.00

$$\mathcal{F}\{\hat{y}(n)\} = \hat{Y}(\nu_l) = \sum_{k=1}^{K} a_k e^{j\varphi_k} \frac{1 - e^{(-\alpha_k + j2\pi(f_k - \nu_l))N\Delta t}}{1 - e^{(-\alpha_k + j2\pi(f_k - \nu_l))\Delta t}}$$

$$\nu_l = l/(N\Delta t), \quad l = 0, \ldots, N-1. \qquad [3]$$

To obtain the parameters, the difference between the DFT of the data and the frequency-domain model function is minimized:

$$\min_{a_k, \varphi_k, \alpha_k, f_k} \sum_{l=0}^{N-1} |Y(\nu_l) - \hat{Y}(\nu_l)|^2. \qquad [4]$$

The solution of this NLLS problem is the same as the one obtained by minimizing Eq. [2] directly in the time domain. If a model other than the Lorentzian one is used (e.g., Gauss or Voigt), a simple exact analytical expression for the DFT of the a model function is not available (*14*). It is, however, always possible to numerically compute the model function in the frequency domain by taking the DFT of the model function in the time domain (*15*).

If the spectrum contains clusters of peaks of which little is known and of which no model function is available, parameter estimation becomes difficult. These nuisance signals hamper the accurate quantification of the peaks of interest. Broad underlying components are usually dealt with by deleting the first data points. Peaks having approximately the same damping as the peaks of interest have to be treated differently. In this paper we focus on the quantification of a few selected peaks of the spectrum in the presence of other peaks, whose influence cannot be reduced by deleting a few data points. The issue is denoted by frequency-selective (FS) parameter estimation in the following paragraphs.

The problem has been studied previously and some methods that can be used in combination with black-box time-domain methods have been proposed in the literature (*16, 17*). However, since *in vivo* MRS signals are characterized by a low signal-to-noise ratio (SNR), prior knowledge often needs to be imposed to obtain relevant parameter estimates. Therefore, time-domain (or frequency-domain) methods based on model function fitting will be considered here. A limited number of results obtained in this field can be found in the literature. One example is (*18*) where a time-domain weighting in combination with VARPRO was proposed to minimize the influence of nuisance peaks on parameter estimates of peaks that are relatively well separated from the nuisance peaks. Recently, a simple preprocessing method, ER filter (*19*), has been proposed. Although this technique inherently distorts the signal it can be used in cases where the spectral region of interest is small compared to the full spectral width and the number of data points is large. Other examples can be found in methods that have been applied to a special case of FS estimation, i.e., solvent suppression in ^{1}H MRS. In this well-studied field a number of preprocessing techniques have been presented that can be applied to remove the influence of nuisance peaks prior to a model-fitting procedure. The HSVD filter method (*20*) is an example of such a technique that can be used to estimate, reconstruct, and finally subtract the signal part corresponding to the nuisance peaks from the original signal. Other examples based on different filtering (convolution) techniques are given in Refs. (*21–23*). In Ref. (*24*) a powerful technique based on the use of a maximum-phase finite impulse response (FIR) filter followed by fitting an adapted time-domain model function to the metabolites of interest was proposed to suppress the water peak. FS parameter estimation is usually considered to be very easy in the frequency domain. Indeed, fitting of a frequency-domain model function in a well-defined frequency region is straightforward to implement. The situation is, however, complicated by the fact that the "tails" of the nuisance peaks overlap with the frequency region of interest. This is usually handled by modeling the nuisance peaks in the frequency region of interest as an additional baseline by some choice of basis functions (e.g., polynomials, sinusoids, damped sinusoids) (*12, 15, 25*).

In this paper the FIR filter method presented in (*24*) is adapted for general use in FS parameter estimation. The method is compared to time-domain weighting, HSVD preprocessing, and frequency domain fitting using a polynomial baseline. The paper is organized as follows: In the next section the FS quantification methods under investigation are presented. The following section contains an exhaustive numerical evaluation of the presented methods. This is partly done by a simple simulation example but also by applying the methods to two relevant biomedical MRS scenarios. Finally, in the last section, the main conclusions are formulated.

FREQUENCY-SELECTIVE QUANTIFICATION METHODS

In this section four methods that can be used for FS parameter estimation are described. All methods are based on model fitting by the NLLS solver used in AMARES (*9*). The Jacobian and function values are calculated analytically for all methods taking the imposed prior knowledge and preprocessing steps into account. Three methods are used in combination with time-domain parameter estimation: FIR filtering (convolution), time-domain weighting, and HSVD filtering. A detailed discussion covers the adaptation of the FIR filter method presented in (*24*) for FS parameter estimation. Finally, a frequency-domain model-fitting method based on minimizing in certain frequency intervals the difference between the DFT transformed data and the model function of Eq. [3] (with a polynomial baseline added to the model) is described.

Note that AMARES$_w$ and AMARES$_H$, which are described below, are implemented within the software package MRUI (*26, 27*), a graphical user interface for facilitating the use of sophisticated analysis routines for MRS data quantification in biomedical/biochemical laboratories and the clinical environ-

ment. The newly developed method AMARES$_f$ will be incorporated in a future release of the MRUI package.

FIR Filter Method: AMARES$_f$

Basic theory of the method. An FIR filter is defined by the convolution sum

$$y_f(n) = \sum_{m=0}^{M-1} h_m y(n-m) \quad n = 0, \ldots, N-1, \qquad [5]$$

where $\{h_m\}_{m=0,\ldots,M-1}$, are the constant (possibly complex) filter coefficients (see, e.g., (*29*)). To obtain the first $(M-1)$ samples of $y_f(n)$ it is usually assumed that the signal is zero outside the time window or cyclic. Either of these assumptions leads to distortions in the first $(M-1)$ samples of $y_f(n)$, which therefore have to be discarded resulting in the signal

$$y_f(n) = \sum_{m=0}^{M-1} h_m y(n-m+M-1), \quad n = 0, \ldots, N-M. \qquad [6]$$

The samples which are discarded should contain as little signal information as possible. For a decaying signal this can be achieved as follows. One of the important effects of applying the FIR filter defined in Eq. [5] is to delay the time-domain signal. This delay, which is a function of the frequency, is specified by the group delay of the filter. Therefore, by using a filter with a group delay equal to the filter length, practically no information contained in the first high-energy samples of the original signal is lost by deleting the first $(M-1)$ samples of y_f. Assuming that the data have been collected long enough such that the $(M-1)$ last signal samples contain mainly noise, there is practically no loss of signal energy due to the fact that these samples are shifted out of the time window. A filter with a group delay equal to the filter length is in general not exactly realizable, but for a given magnitude response of a filter a maximum-phase filter has the largest possible group delay (*28*).

Such an FIR filter can be used as follows to eliminate peaks in certain frequency regions. The influence of applying a filter on the exponentially damped sinusoids can be examined by applying Eq. [6] to the signal model of Eq. [1]:

$$\begin{aligned}\hat{y}_f(n) &= \sum_{m=0}^{M-1} h_m (x_1(n-m+M-1) \\ &\quad + x_2(n-m+M-1) \\ &\quad + \cdots + x_K(n-m+M-1)) \\ &= \bar{h}b_1 x_1(n) + \bar{h}b_2 x_2(n) + \cdots + \bar{h}b_K x_K(n) \\ &= \sum_{k=1}^{K} \bar{h}b_k x_k(n) \quad n = 0, \ldots, N-M, \end{aligned} \qquad [7]$$

where

$$x_k(n) = a_k e^{j\varphi_k} e^{(-\alpha_k + j2\pi f_k)n\Delta t} \quad k = 1, \ldots, K$$

denote the individual, exponentially damped complex sinusoids,

$$\bar{h} \triangleq (h_{M-1} \cdots h_0)$$

and

$$b_k \triangleq (1 \; e^{(-\alpha_k + j2\pi f_k)\Delta t} \cdots e^{(-\alpha_k + j2\pi f_k)(M-1)\Delta t})^{\mathrm{T}}.$$

The superscript T denotes the transpose. From Eq. [7] it is seen that the filtered NMR signal consists of the original damped sinusoids (same frequency and damping) altered by a complex scalar $\bar{h}b_k$. The filter can thereby be designed to suppress the nuisance peaks (i.e., make $|\bar{h}b_k| \approx 0$) including their frequency-domain tails (see, e.g., (*24, 30*)) while leaving the peaks of interest practically undistorted (i.e., make $|\bar{h}b_k| \approx 1$). Signals not exactly obeying the Lorentzian model are also efficiently removed if they can be approximated by a sum of (damped) complex sinusoids with a frequency located in the stop band of the filter.

If the model function of Eq. [1] is fitted to the filtered signal, the estimated amplitude $\tilde{a}_k$ and phase $\tilde{\varphi}_k$ must be corrected as follows:

$$\begin{aligned} a_k &= \frac{\tilde{a}_k}{|\bar{h}b_k|} \\ \varphi_k &= \tilde{\varphi}_k - \tan^{-1}\left(\frac{\mathrm{imag}(\bar{h}b_k)}{\mathrm{real}(\bar{h}b_k)}\right), \end{aligned} \qquad [8]$$

where real(.) and imag(.) denote the real and imaginary part of (.), respectively. If, however, prior knowledge concerning the amplitudes and/or phases is to be imposed, the filter influence must be taken into account directly in the estimation process and Eq. [2] must be modified to

$$\min_{a_k,\varphi_k,\alpha_k,f_k} \sum_{n=0}^{N-M} \left| y_f(n) - \sum_{k=1}^{K} \bar{h}b_k x_k(n) \right|^2. \qquad [9]$$

The method that solves the above-stated minimization problem is referred to as AMARES$_f$ in the following.

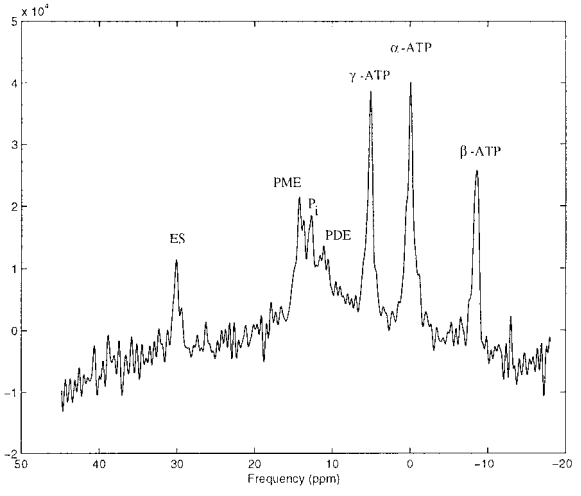

FIG. 1. *In vivo* ^{31}P signal from perfused rat liver, acquired at 4.7 T (81 MHz). The real, zero-, and first-order phase-corrected parts of the spectrum are shown. Peaks of the external standard (ES), phosphomonoesters (PME), inorganic phosphate (P_i), phosphodiesters (PDE), and the α, β, and γ phosphorus of adenosine triphosphate (ATP) can be observed. α-ATP is taken as the reference (0 ppm).

The phase changes introduced by applying the FIR filter also influence the visualization of the filtered signals. One solution is to display the magnitude of the spectra. To display the phased, real spectrum, the DFT transformed signal has to be corrected not only for the zero- and first-order phase but also for the frequency-dependent phase response of the filter. This is illustrated for an *in vivo* ^{31}P signal of the perfused rat liver, acquired at 4.7 T (81 MHz). In Fig. 1 the real part of the original signal is displayed after zero- and first-order phase correction. In the left-hand side of Fig. 2 the β-ATP signal has been extracted from the signal using the filter and phase corrected to zero and first order. In the right-hand side the filtered signal is shown with an additional phase correction to compensate for the influence of the filter.

In summary, by using a FIR filter it is possible to eliminate peaks from the spectrum, including their frequency-domain tails. The distortions introduced by the filter are known and can be taken into account in the model function. The use of a maximum-phase FIR filter ensures that the loss of information-rich samples due to filtering is minimized. The following paragraph provides additional information about the design and the proper use of the maximum-phase FIR filters.

Design and use of the maximum-phase FIR filter. The standard implementation of a maximum-phase FIR filter is based on the technique of (*31*) and proceeds as follows. A linear-phase FIR filter is designed and transformed to a maximum-phase filter by rooting the filter polynomial and reflecting the zeros outside the unit circle followed by reconstruction of the impulse response of the filter. To design the linear-phase filter, the constrained least-squares algorithm proposed in (*32*) was used. The parameters specifying the filter design are the filter order M, cutoff frequency f_c, stop-band suppression sup, and passband ripple r. In (*24*) this FIR filter method was used for water signal suppression in ^{1}H MRS. In the latter paper, an automatic filter design scheme was presented in which the only parameter to be specified by the user is the approximate frequency of the peak of interest that lies closest to the water peak. The problem at hand here is more difficult since the user may want to suppress one or more regions in the spectrum containing nuisance peaks. Therefore the design proposed in (*24*) was adapted in the following way:

1. The user defines the filter cutoff frequencies fl_r and fh_r by selecting the frequency region(s) containing the nuisance peaks, $[fl_r, fh_r]$, $r = 1, \cdots, R$, where R is the number of frequency regions.
2. An estimate $\tilde{\sigma}$ of the noise standard deviation is calcu-

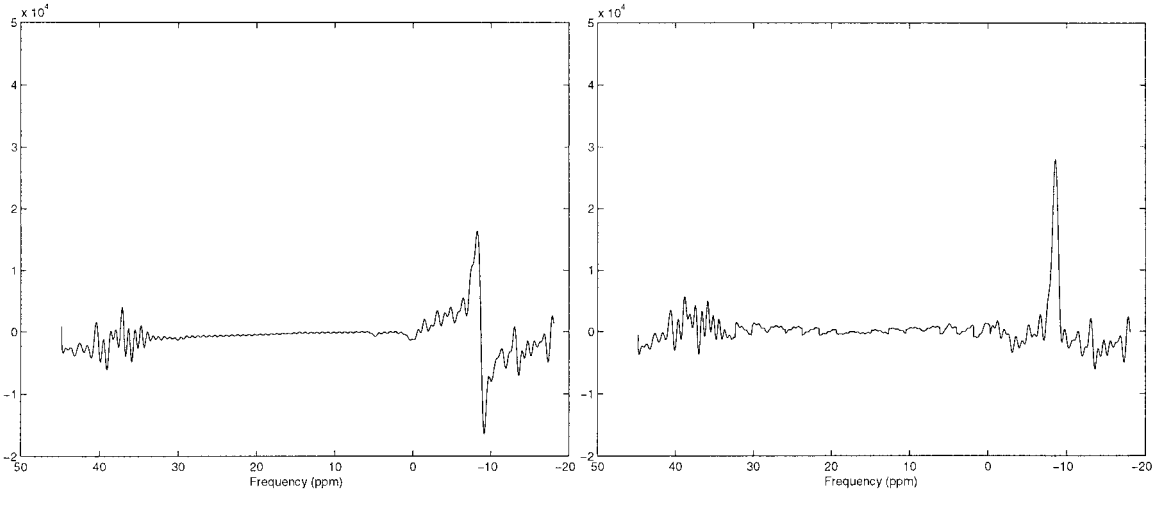

FIG. 2. Real part of the maximum phase FIR-filtered ^{31}P spectrum. Left: zero- and first-order phase corrected as in Fig. 1. Right: corrected for zero- and first-order phase as in Fig. 1 and for the frequency-dependent phase of the filter.

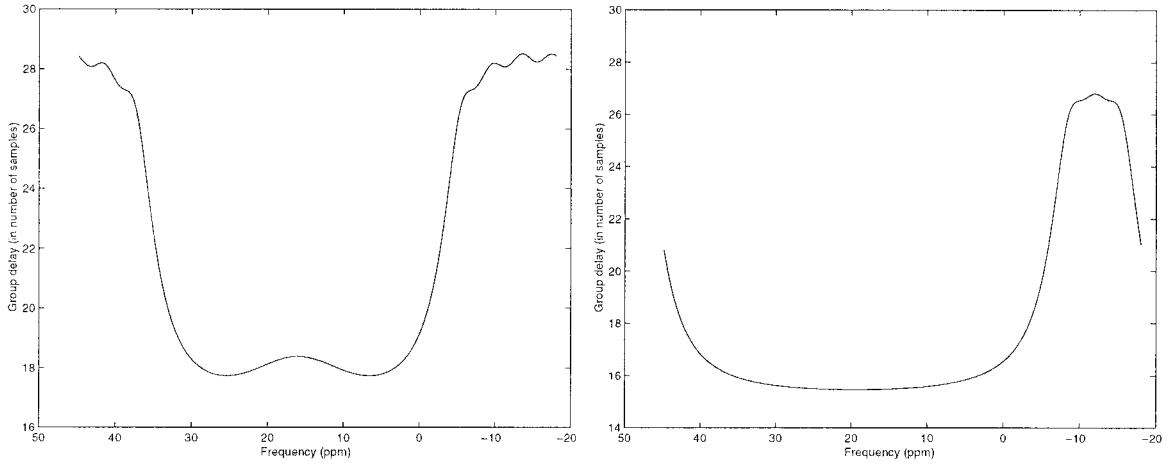

FIG. 3. Group delay of the maximum-phase FIR filter designed to extract the β-ATP peaks of Fig. 1. Left: lower cutoff frequency = −1.43 ppm, upper cutoff frequency = 31.9 ppm. Right: lower cutoff frequency = −5.14 ppm, upper cutoff frequency = 44.2 ppm.

lated as the standard deviation of the last samples of the signal $y(n)$. The peak value, s_0, is computed as the value of the strongest nuisance peak,

$$s_0 = \max\left|\frac{Y(\nu_l)}{\sqrt{N}}\right|,$$

where $\nu_l \in [fl_r, fh_r] r = 1, \cdots, R$.

3. The initial filter suppression sup is taken equal to

$$\text{sup} = \frac{2\tilde{\sigma}}{s_0}$$

and starting values for the filter order (e.g., $M = 30$) and passband ripple (e.g., $r = 5\%$) are chosen.

4. The linear-phase filter is designed and transformed into a maximum-phase filter.

5. The signal is filtered and the maximum value s_{m} in the frequency regions of the nuisance peaks is calculated,

$$s_{\text{m}} = \max\left|\frac{Y_f(\nu_l)}{\sqrt{N}}\right|,$$

where $\nu_l \in [fl_r, fh_r]$, $r = 1, \cdots, R$. If

$$s_{\text{m}} > 2\tilde{\sigma},$$

the filter suppression (e.g., sup = sup/5) is increased and the process is restarted from 4. If s_{m} does not decrease in the next iteration, the filter order is increased (e.g., $M = M + 10$).

If the nuisance peaks are relatively narrow (slowly decaying), the above scheme results in a suitable filter in the first iteration. A higher filter suppression is required if some of the nuisance peaks have a fast decay. In this case the iterative procedure will find a filter with a sufficient suppression within a few iterations (<10). The computational complexity is thereby increased compared to the original filter design scheme. The advantage is, however, that a suitable filter is found independent of the shape and signal energy of the nuisance peaks. See (*24*) for a more detailed discussion on the computational complexity associated with the design of the maximum-phase FIR filter.

In (*24*), a maximum-phase filter is constructed that suppresses the water signal down to the noise level while keeping the stop band of the filter as narrow as possible based only on the specification of the first peak of interest closest to the water peak. A narrow stop band is advisable since in general it leads to a higher group delay in the passband than a filter with the same order, stop-band suppression, and passband ripple. This principle is illustrated using the same ^{31}P example described above (Fig. 1). First the automatic design scheme is applied with −1.43 ppm as a lower bound and 31.9 ppm as an upper bound for the region to be suppressed (Fig. 2). The group delay of the resulting maximum phase filter of length 31 is shown in the left-hand side of Fig. 3. The group delay of the filter with −5.14 ppm and 44.2 ppm as lower and upper bounds, respectively, is shown in the right-hand side of the figure. Two effects can be observed. First, the filter with the larger stop band has an overall smaller group delay in the passband and, second, the peak closest to the filter lower bound is located in the transition band of the filter, leading to an even smaller group delay for this specific peak. The smaller group delay for the frequencies of interest will lead to an overall loss in SNR and a decrease in resulting parameter accuracy. This shows that care must be taken on how to specify the lower and upper bounds of the region(s) to be suppressed and it is therefore advisable to inspect the group delay of the filter in the passband in difficult scenarios (large stop band and/or bound close to frequency of interest). When choosing the bound closest to the region to be suppressed, it is important to make sure that no signal parts of the nuisance region have a frequency within the transition band

of the filter. Otherwise, the nuisance peaks would be insufficiently suppressed, introducing bias in the parameters of interest.

If the upper and lower bounds of the region(s) to be suppressed are chosen with some care, the nuisance peaks are efficiently removed and the bias of the estimates is eliminated. This is accomplished with practically no loss of signal-to-noise ratio and with a low computational complexity. The influence of the choice of the upper and lower bounds in the filter design scheme is further examined under Numerical Examples.

Time-Domain Weighting: AMARES$_w$

The influence of nuisance peaks in NLLS parameter estimation techniques such as VARPRO and AMARES was studied in (*18*). The bias term of the amplitude estimates, assuming correct estimates of frequency and damping, was derived and it was seen that this term could be reduced by introducing an appropriate weighting vector $v(n)$, in the NLLS fit:

$$\min_{a_k,\varphi_k,\alpha_k,f_k} \sum_{n=0}^{N-1} |v(n)[y(n) - \sum_{k=1}^{K} x_k(n)]|^2. \qquad [10]$$

The choice of the weighting vector is a trade-off between reducing the nuisance peaks influence and the loss of SNR. The weighting vector should be matched with the frequency distance, amplitude, and damping of the nuisance peaks to find the optimal trade-off. This is, however, not feasible in practice and instead a generic weighting vector is applied that is expected to work reasonably well in most scenarios. In (*18*) a weighting consisting of a quarter-wave sinusoid for the first (and last) 20 samples was recommended:

$$v(n) = \begin{cases} \sin\left(\dfrac{n\pi}{40}\right) & n \in [0,\ 19] \\ 1 & n \in [20,\ N-21] \\ \sin\left(\dfrac{(N-n-1)\pi}{40}\right) & n \in [N-20,\ N-1] \end{cases}. \qquad [11]$$

Inserting the weighting function of Eq. [11] into Eq. [10] and solving it is referred to as AMARES$_w$ in the following. The method is expected to give good results for relatively well-separated peaks. However, the technique always leads to a loss of SNR resulting in an increased variance of the parameter estimates. These properties are illustrated under Numerical Examples.

HSVD Filter Method: AMARES$_H$

Another approach to solve the problem of FS estimation is to model the nuisance peaks by a black-box method and subtract the reconstructed time-domain signal from the original signal prior to parameter estimation. This technique has already been applied for solvent suppression using various techniques to estimate the water signal. A black-box method often used in this context is the HSVD method (*20*). HSVD is a subspace-based parameter estimation method in which the noisy signal space is subdivided into a "signal" and a "noise" subspace using a SVD of a Hankel data matrix (*33*). The signal subspace is found by truncating the SVD of this matrix to rank M, the number of exponentials that models the underlying signal. In general HSVD provides a mathematical fit of the data by a sum of exponentially damped complex-valued sinusoids. HSVD can therefore be used to approximate complicated features of the nuisance peaks. The signal corresponding to the fitted regions is subsequently subtracted from the original signal.

In this paper the following scheme is used to process spectra with HSVD for FS estimation:

1. The user specifies the model order M and frequency regions where the nuisance peaks have a frequency located $[fl_r, fh_r]$, $r = 1, \cdots, R$, where R is the number of frequency regions.
2. HSVD is used to model the original signal by a sum of M exponentially damped complex-valued sinusoids.
3. The modeled peaks lying within the user-defined regions are used to reconstruct the nuisance peaks. Afterward the reconstructed nuisance signal is subtracted from the original signal.
4. The residual signal is quantified with AMARES.

The complete procedure is denoted as AMARES$_H$ in the following.

A drawback of this method is the large computational complexity since it requires the computation of the SVD of the Hankel data matrix. Fast versions of this HSVD algorithm have been developed in which the computation of a full SVD is circumvented; see (*34*, *35*). The gain in efficiency of these fast methods, however, decreases when the number of data points decreases and/or the model order increases (*35*). The choice of model order M in HSVD is not trivial in case the nearby nuisance peaks have an unknown or other than Lorentzian model function. The choice of the correct model order is important (and difficult) since undermodeling or overmodeling can lead to a deterioration in parameter accuracy (see (*24*)). Methods exist that make an automatic choice of model order based on various information criteria and on the dominant singular values (see, e.g., (*37*) and reference therein). In this paper the minimum description length (MDL) criterion used in (*37*) is examined in the context of FS parameter estimation. The MDL criterion, however, requires the knowledge of all of the singular values. Consequently, the entire SVD must be calculated when the MDL criterion is used to estimate the model order and no fast methods such as those presented in (*34*, *35*) can be used.

Frequency-Domain Fitting: AMARES$_{FREQ}$

FS estimation using frequency-domain model fitting is done by minimizing the difference between the DFT transformed data and a model function for the frequency region(s) of interest. The influence from the other peaks in the spectrum is not zero in the frequency region(s) of interest and must be taken into account. This is normally done by introducing extra parameters to model the tails of the nuisance peaks. Here we choose to use a polynomial basis function. The complete model function used to model the spectrum in the frequency region(s) of interest is given by

$$\hat{Y}(\nu_l) = \sum_{k=1}^{K} a_k e^{j\varphi_k} \frac{1 - e^{(-\alpha_k + j2\pi(f_k - \nu_l))N\Delta t}}{1 - e^{(-\alpha_k + j2\pi(f_k - \nu_l))\Delta t}} + \sum_{p=0}^{P} d_p \nu_l^p$$
$$\nu_l \in [fl_r, fh_r], \quad r = 1, \cdots, R, \qquad [12]$$

where P is the model order of the polynomial, d_p are the complex polynomial coefficients, and $[fl_r, fh_r]$, $r = 1, \ldots, R$ is (are) the frequency region(s) of interest.

To obtain the parameters of interest, the minimization is now only carried out over the frequency region(s) of interest instead of over the entire frequency range as done in Eq. [4]:

$$\min_{a_k,\varphi_k,\alpha_k,f_k,d_p} \sum_{\substack{\nu_l \in [fl_r, fh_r] \\ r=1,\ldots,R}} |Y(\nu_l) - \hat{Y}(\nu_l)|^2,$$
$$k = 1, \ldots, K, \quad p = 0, \cdots, P. \qquad [13]$$

The minimization procedure of Eq. [13] is denoted by AMARES$_{\text{FREQ}}$ in the following. The user must specify the frequency range of the regions of interest and as a function thereof the polynomial order must be chosen. In difficult cases where the peaks of interest are not lying on a flat baseline the choice of the polynomial order is not trivial. If the polynomial model order is chosen appropriately, i.e., if the polynomial is a good model for the underlying baseline in the frequency region included, the estimation bias is expected to be low. The introduction of these extra parameters to be estimated inevitably increases the variance of the parameter estimates. If the polynomial is not a good model for the baseline, bias is introduced in the estimated parameters. The influence of the choice of the frequency region and the corresponding polynomial order are further investigated under Numerical Examples. Note that other choices of basis functions (e.g., (damped) sinusoids) might in some cases lead to a better modeling of the baseline.

NUMERICAL EXAMPLES

The above methods are compared for three examples. First an artificial two-peak example is constructed to examine the estimation properties of the different methods as a function of the amplitude of the nuisance peaks and the frequency distance between the nuisance peaks and the peaks of interest. The performance of AMARES$_{\text{f}}$ in case of a non-Lorentzian lineshape is also examined. Thereafter the methods are evaluated for two relevant biomedical MRS examples. Unless otherwise specified the model function of Eq. [1] is used and the added complex noise is circular, white, and Gaussian distributed. The SNR for each peak is expressed in decibels (dB) and defined as

$$\text{SNR peak } k \triangleq 20 \log\left(\frac{a_k}{\sigma}\right),$$

where σ is the noise standard deviation. The quality of the amplitude estimates is measured as the relative root mean squared error (RRMSE) in percentages:

$$\text{RRMSE peak } k \triangleq 100 \sqrt{\frac{1}{L} \sum_{l=1}^{L} \frac{(a_k - \tilde{a}_k^l)^2}{a_k^2}},$$

where L is the number of simulation runs and $\tilde{a}_k^l$ denotes the estimate of a_k obtained in simulation run l. In the following examples, the number of simulation runs is taken equal to 300. The relative bias (RBias) in percentages

$$\text{RBias peak } k \triangleq 100 \frac{a_k - \frac{1}{L}\sum_{l=1}^{L} \tilde{a}_k^l}{a_k}$$

and the relative standard deviation (RSTD) in percentages

$$\text{RSTD peak } k \triangleq 100 \sqrt{\frac{1}{L-1} \sum_{l=1}^{L} \frac{\left(\tilde{a}_k^l - \frac{1}{L}\sum_{l=1}^{L} \tilde{a}_k^l\right)^2}{a_k^2}}$$

are also used to express the theoretical expectations about the performance of the methods. The RRMSE is compared to the relative Cramér-Rao lower bound (CRB). The CRB is calculated from a model consisting of the peaks of interest only. The CRB indicates the best possible accuracy of an estimate for *any* unbiased estimator (see, e.g., (*36*)).

Two-Peak Example

In this section we consider an example consisting of two exponentially damped sinusoids in which only one is of interest. The peak of interest (peak1) has the following parameter values:

$$f_1 = 20 \text{ Hz}$$

$$\alpha_1 = 10 \text{ Hz}$$

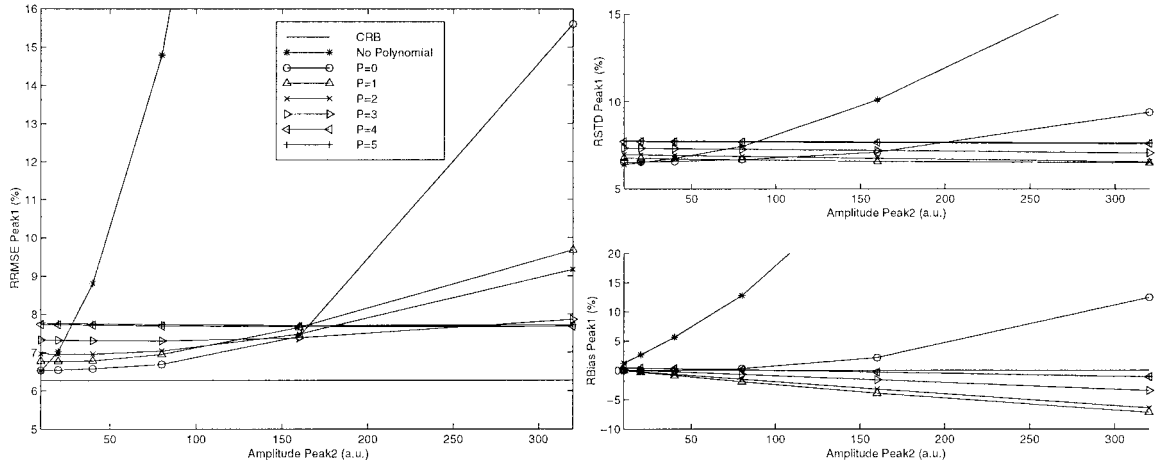

FIG. 4. Two-peak example: RRMSE, RSTD, and RBias of amplitude estimates of peak1 ($a_1 = 20$, $f_1 = 0.02$ kHz) as a function of the amplitude of the nuisance peak (peak2) for the frequency-domain fitting procedure $\text{AMARES}_{\text{FREQ}}$ using a polynomial baseline fit with different polynomial orders P, SNR of peak1 = 10 dB. Left: RRMSE. Top right: RSTD. Bottom right: RBias. CRB denotes the theoretical lower bound.

$$\varphi_1 = 0 \text{ degrees}$$

$$a_1 = 20 \text{ au}$$

The damping and phase of the nuisance peak (peak2) are fixed to

$$\alpha_2 = 10 \text{ Hz}$$

$$\varphi_2 = 0 \text{ degrees},$$

whereas the amplitude and frequency of peak2 are varied. The sampling frequency is 1 kHz and the number of data points is 512. Unless otherwise specified the SNR used in the simulations for peak1 is equal to 10 dB.

Part A: Frequency-Domain Baseline Fitting Using $\text{AMARES}_{\text{FREQ}}$

Procedure. The purpose is to assess the influence of the polynomial baseline fit on the accuracy of the parameters of peak1 estimated with $\text{AMARES}_{\text{FREQ}}$. Simulations were run with and without inclusion of a polynomial baseline. The amplitude of peak2 is varied between 10 and 320 au while its frequency is fixed to 0.1 kHz. Polynomial orders varied from $P = 0$–5 and the frequency region included in Eq. [13] was -0.2 to 0.04 kHz.

The RRMSE, RSTD, and RBias of the amplitude estimates of peak1 are shown as a function of the amplitude of peak2 for different polynomial orders in Fig. 4.

To get an idea of the influence of the size of the frequency region included in Eq. [13], the amplitude of peak2 is fixed and the lower and upper bounds of the frequency interval are varied between -0.5 and 0.01 kHz and between 0.03 and 0.09 kHz, respectively. The corresponding parameter accuracy of peak1 (RRMSE, RSTD, RBias) with the amplitude of peak2 fixed to 320 au is shown in Fig. 5 (for $P = 2$) and Fig. 6 (for $P = 4$).

Results and discussion. A higher polynomial order is required to achieve good estimates as the amplitude of peak2 increases (Fig. 4). It is seen that the bias almost vanishes if the polynomial order is chosen to be sufficiently high. The price for including additional polynomial parameters is an increased standard deviation on the estimates.

The standard deviation is overall lower for $P = 2$ (Fig. 6), but the bias is in general much larger than the bias for $P = 4$ (Fig. 6), thereby confirming the result of Fig. 4. It can also be observed that if the region of interest is small, the standard deviation is higher since the number of points included in the optimization problem is smaller. The polynomial of order $P = 2$ is seen to be a bad model for the baseline for almost any choice of the frequency region. The polynomial model of order $P = 4$ is a good approximation for most choices of the frequency region but also breaks down when the frequency region becomes too large.

As the amplitude of peak2 decreases, a lower polynomial order suffices to model the baseline and the results becomes less sensitive to the frequency range included in Eq. [13].

Part B: Time-Domain FIR Filtering Using AMARES_f

Procedure. The purpose is twofold: first to assess the influence of the value of the upper and lower bound of the region to be suppressed in the AMARES_{f} filter design scheme and second to assess the performance of the method in case the nuisance peak does not have a Lorentzian lineshape.

The amplitude of peak2 is fixed to 320 au and the upper and lower bounds are varied between 0.11 and 0.5 kHz and 0.03 and 0.09 kHz, respectively. The RMSE, RSTD, and RBias of

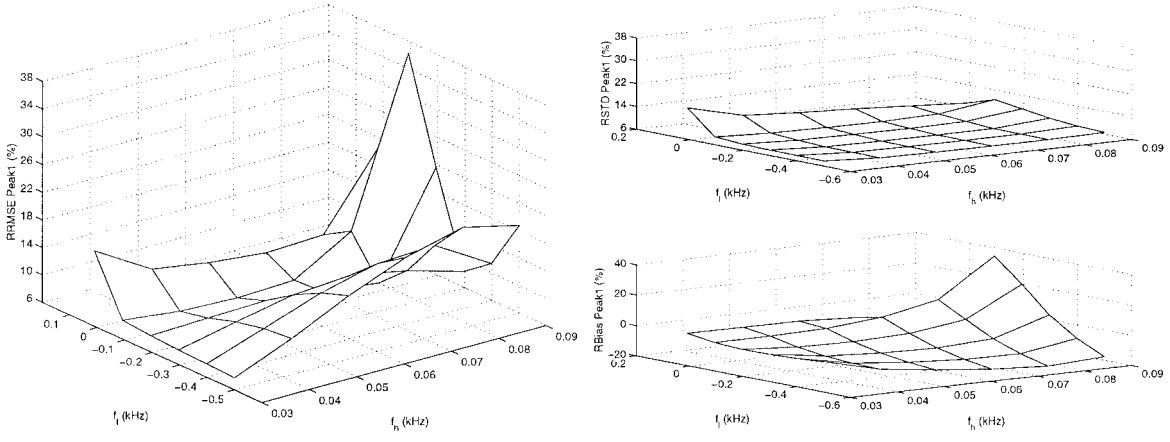

FIG. 5. Two-peak example: RMSE, RSTD, and RBias of the AMARES$_{\mathrm{FREQ}}$ amplitude estimates of peak1 ($a_1 = 20$, $f_1 = 0.02$ kHz) as a function of the width of the frequency region included in Eq. [13]. The lowest frequency varies between -0.5 and 0.01 kHz and the highest frequency varies between 0.03 and 0.09 kHz. The polynomial order is $P = 2$. The parameters of peak2 are $a_2 = 320$, $f_2 = 0.1$ kHz. SNR of peak1 = 10 dB. Left: RRMSE. Top right: RSTD. Bottom right: RBias.

the amplitude estimates of peak1 are shown in Fig. 7 as a function of the bounds.

To test the performance of AMARES$_{\mathrm{f}}$ in case the lineshape is non-Lorentzian, an example is used consisting of two Gaussian peaks: $\sum_{k=1}^{2} a_k e^{-g_k(n\Delta t)^2} e^{j(2\pi f_k n\Delta t+\varphi_k)}$, with

$$f_1 = 20 \text{ Hz}$$

$$a_1 = 20/(\sqrt{\pi \ln 2})\text{au}$$

$$a_2 = 230/(\sqrt{\pi \ln 2})\text{au}$$

$$g_1 = g_2 = 10^2/(4 \ln 2)\text{Hz}^2$$

$$\varphi_1 = \varphi_2 = 0 \text{ degrees}$$

and f_2 varying between 0.06 and 0.22 kHz. Note that the Gaussian peaks have the same full width at half height and the same maximum value of the real part of the DFT spectrum as Lorentzian peaks with amplitudes of 20 and 320 au and dampings of 10 Hz. The SNR of peak1 is equal to 6.6 dB. The filter design scheme is used with upper and lower bounds of the region to be suppressed equal to the frequency of peak2 plus and minus 0.02 kHz, thereby fulfilling the basic principles as explained under Design and Use of the Maximum-Phase FIR Filter. Equation [9] was adapted to compensate for the filter influence in case of a Gaussian lineshape. The results are displayed in Fig. 8.

Results and discussion. The bias is effectively removed for almost all combinations of upper and lower bounds used in the

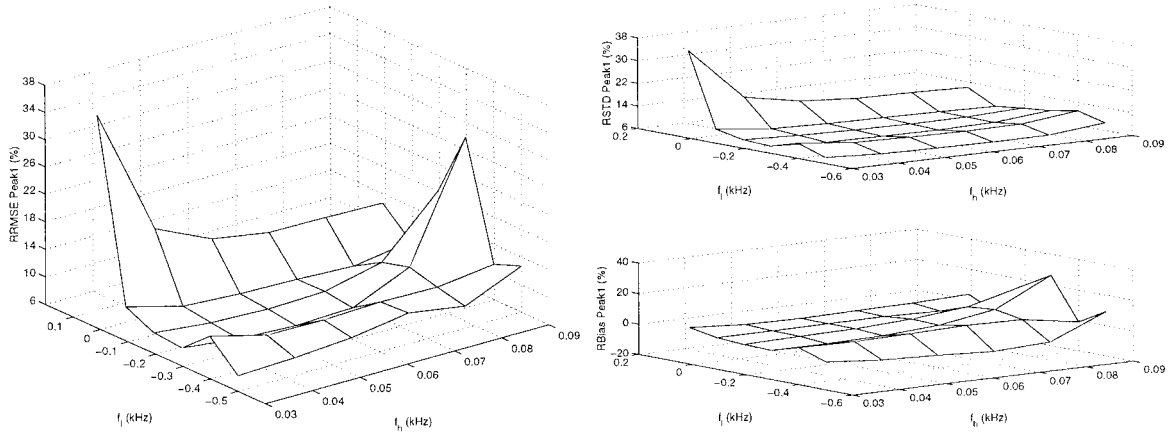

FIG. 6. Two-peak example: RMSE, RSTD, and RBias of the AMARES$_{\mathrm{FREQ}}$ amplitude estimates of peak1 ($a_1 = 20$, $f_1 = 0.02$ kHz) as a function of the width of the frequency region included in Eq. [13]. The lowest frequency varies between -0.5 and 0.01 kHz and the highest frequency varies between 0.03 and 0.09 kHz. The polynomial order is $P = 4$. The parameters of peak2 are $a_2 = 320$, $f_2 = 0.1$ kHz. SNR of peak1 = 10 dB. Left: RRMSE. Top right: RSTD. Bottom right: RBias.

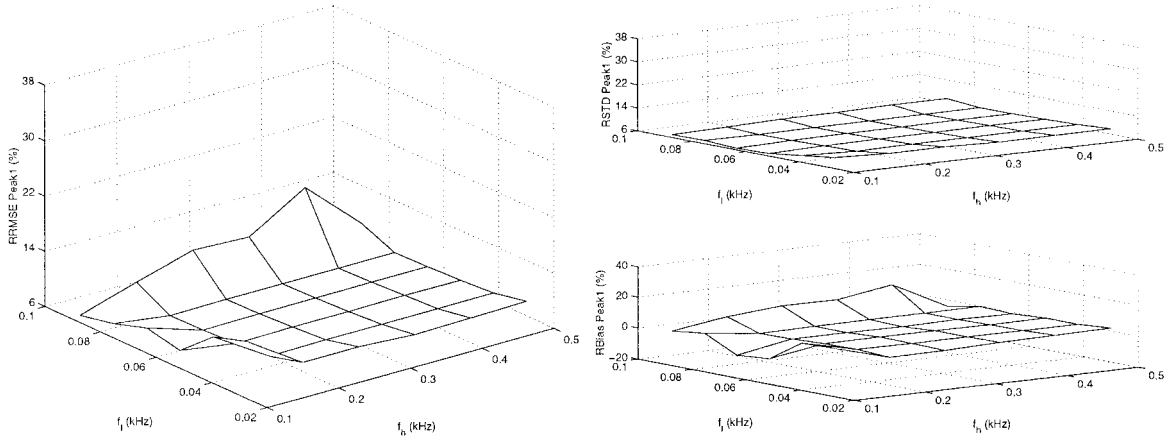

FIG. 7. Two-peak example: RMSE, RSTD, and RBias of the AMARES$_f$ amplitude estimates of peak1 ($a_1 = 20$, $f_1 = 0.02$ kHz) as a function of the lower and upper bounds in the filter design scheme. The parameters of peak2 are $a_2 = 320$, $f_2 = 0.1$ kHz. SNR of peak1 = 10 dB. Left: RRMSE. Top right: RSTD. Bottom right: RBias.

filter design scheme (Fig. 7). The method breaks down when the upper and/or lower bound is chosen too close to the frequency of peak2, resulting in insufficient suppression of the nuisance peak. A slight increase in standard deviation can be observed when the lower cutoff frequency is chosen closer to the peak of interest. This is explained by the resulting lower group delay for peak1 leading to a loss of SNR and subsequently a higher standard deviation.

AMARES$_f$ has no problems in dealing with the Gaussian nuisance peak (Fig. 8). The bias is removed and the RRMSE follows the CR bound very closely.

Part C: Method Comparison

Procedure. The accuracy of the four frequency-selective estimation methods is compared.

The frequency of peak2 is varied between 0.06 and 0.22 kHz while its amplitude is kept fixed at 320 au. The results are displayed in Fig. 9. For AMARES$_{FREQ}$, the polynomial order that eliminates the bias and has the lowest standard deviation is shown for every frequency. The frequency region from -0.15 to 0.04 kHz is included in Eq. [13] for all frequencies of peak2. It was verified that for the polynomial orders shown, the width of this frequency interval was not critical. For AMARES$_f$, the upper and lower bounds in the filter design scheme are chosen as the frequency of peak2 plus and minus 0.02 kHz, thereby fulfilling the basic principles explained under Design and Use of the Maximum-Phase FIR Filter. In AMARES$_H$ the model order is taken equal to 2, the theoretical one.

Results and discussion. For AMARES$_w$ a "breakdown" frequency is found below which the assumptions of a "large"

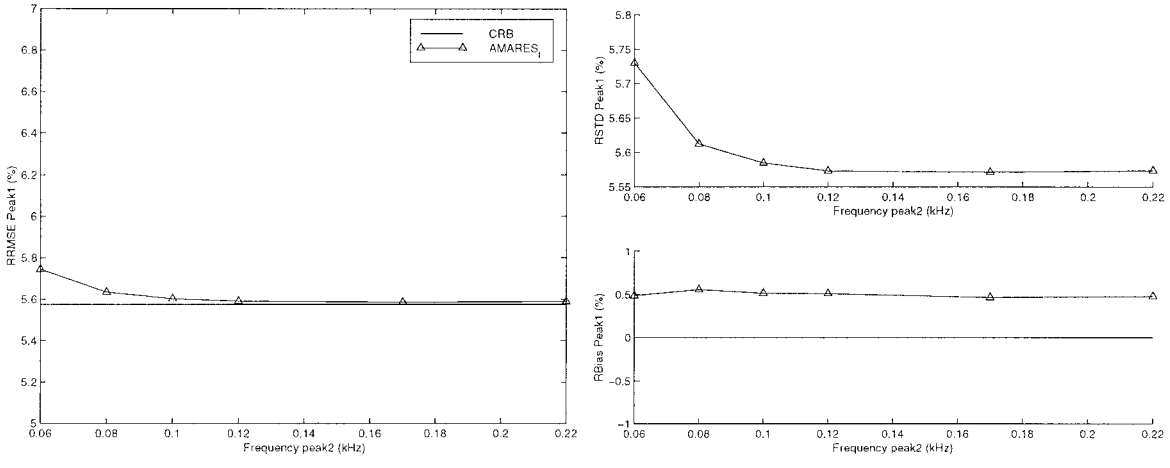

FIG. 8. Two-peak example: RRMSE, RSTD, and RBias of the AMARES$_f$ amplitude estimates of peak1 ($a_1 = 13.6$, $f_1 = 0.02$ kHz) as a function of the frequency of the nuisance peak (peak2, $a_2 = 216.9$). The two peaks have a Gaussian lineshape. SNR of peak1 = 6.6 dB. Left: RRMSE. Top right: RSTD. Bottom right: RBias. CRB denotes the theoretical lower bound.

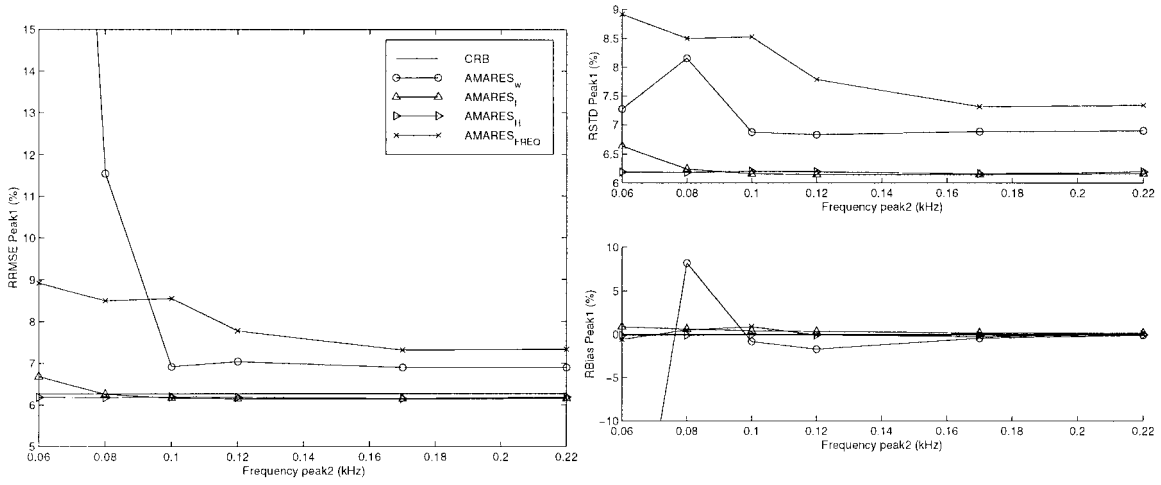

FIG. 9. Two-peak example: RRMSE, RSTD, and RBias of the amplitude estimates of peak1 ($a_1 = 20$, $f_1 = 0.02$ kHz) as a function of the frequency of the nuisance peak (peak2, $a_2 = 320$) for different FS estimation procedures. SNR of peak1 = 10 dB. Left: RRMSE. Top right: RSTD. Bottom right: RBias. CRB denotes the theoretical lower bound.

frequency separation no longer hold. The method is successful in eliminating the estimator bias outside this region, but the loss of SNR associated with this approach increases the standard deviation of the estimates. AMARES$_{FREQ}$ is capable of handling close nuisance peaks if the polynomial order is chosen high enough. This is confirmed by the low estimator bias. The standard deviation is, however, increasing as a function of the polynomial order thereby degrading the estimates. The time-domain methods based on FIR filtering and HSVD filtering are clearly outperforming the other methods and their RRMSEs are very close to the theoretical lower bounds even for a small frequency separation.

Frequency-Selective Quantification of *in vivo* ^{13}C Spectrum

Procedure. A simulation study is performed using signals derived from an *in vivo* proton decoupled ^{13}C spectrum of subcutaneous adipose tissue in the human forearm. The spectrum was acquired at 4.7 T (50.3 MHz) using a 5-cm-diameter double-tuned surface coil. In some applications one is only interested in obtaining estimates of the two peaks around 128 and 130 ppm, respectively. The peak at 130 ppm corresponds to carbons in double bonds of mono- and polyunsaturated fatty acid chains while the peak at 128 ppm corresponds exclusively to carbons in double bonds of polyunsaturated chains. The amplitudes of these peaks can be used to study noninvasively the relative amounts of poly- and monosaturation of fatty acids in human tissue. The real, phased spectrum of the original *in vivo* signal is displayed in the left-hand side of Fig. 10 while the real, phased spectrum of the FIR-filtered signal is shown on the right-hand side.

The simulation signal was derived as follows. The *in vivo* signal was quantified using AMARES. The model fitted to the signal consisted of 15 Lorentzian peaks and no prior knowledge was imposed. The estimated parameters were then used as parameters for the noiseless simulation signal. The sampling frequency is 10 kHz and the number of data points is 512. It is important to note that the region around 30 ppm was difficult to model correctly. In that region some signal features are still present in the residual, while outside that frequency region the residual looks like white noise. In this case it is certainly advantageous to use FS estimation, since no satisfactory model is present for the entire signal.

The accuracy of the four FS estimation methods is compared for the peak at 130 ppm for different noise levels. For AMARES$_{FREQ}$, different polynomial orders and widths of frequency regions were tested. The results obtained for different model orders in AMARES$_H$ were compared and the performance of the automatic order estimation scheme was examined. In AMARES$_f$ the influence of the choice of the upper and lower bounds in the filter design scheme was studied. Figure 11 shows results for AMARES$_{FREQ}$ with $P = 1$, lower bound equal to 117 ppm, upper bound equal to 147 ppm. The results displayed for AMARES$_f$ were generated with suppression of the region below 73 ppm and above 167 ppm. For AMARES$_H$ the results were obtained for a model order of 15.

Results and discussion. In AMARES$_{FREQ}$ quantification of the ^{13}C spectrum turned out to be impossible without inclusion of a polynomial. A polynomial order of $P = 0$ is only valid in a small region around the two peaks of interest. A polynomial order of $P = 1$ provides a bias close to zero in a large frequency region: the lower bound can vary between 87 and 117 ppm and the upper bound between 137 and 161 ppm. There is no significant change in standard deviation and RMSE in this region. There is only a slight increase in RRMSE by using a higher polynomial order. For AMARES$_H$, there are no

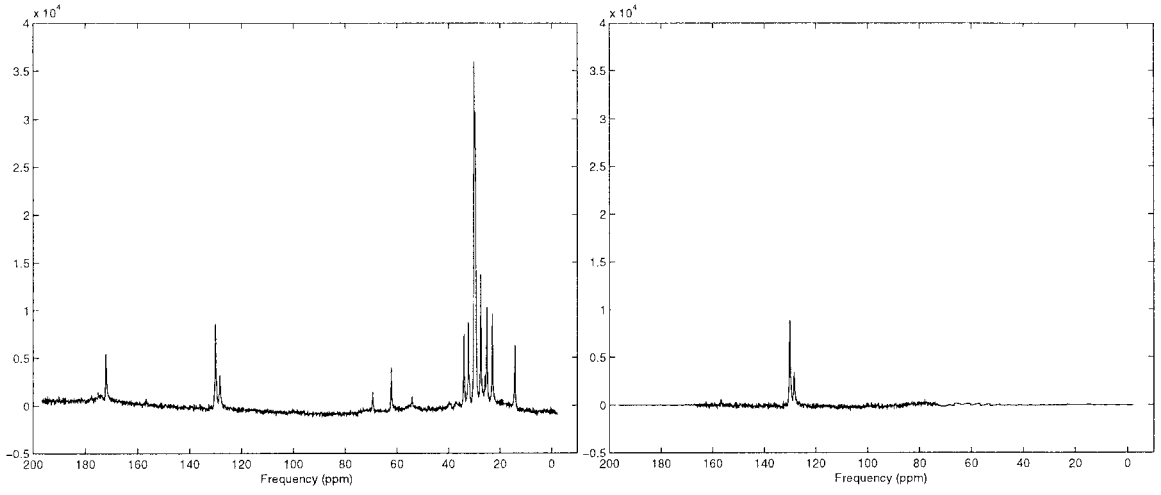

FIG. 10. *In vivo* proton-decoupled ^{13}C spectrum of subcutaneous adipose tissue in the human forearm (acquired at 4.7 T (50.3 MHz)). The peaks to be quantified are at 128 and 130 ppm. Left: real, phased part of ^{13}C spectrum. Right: real, phased part of the same signal after maximum phase FIR filtering.

significant differences between the results obtained for the theoretical model order and those determined by the automatic model order selection criterion. The results are very insensitive to the choice of the lower and upper bounds in the filter design scheme. Only if the bounds are chosen close to the peaks of interest is there a degeneration in the results. As seen in Fig. 11, the results for both AMARES$_f$ and AMARES$_H$ are very close to the theoretical CR bound. Since the peaks of interest are well separated from the other peaks in the spectrum, the other methods perform only slightly worse, as expected. Note also that all methods are able to remove the bias and that the RSTD contributes most to the RRMSE.

FS estimation in this case also leads to a reduction in calculation time. Quantifying all 15 peaks of the spectrum with AMARES takes about 7.6 s, while it takes 0.6 s to analyze the two peaks of interest with AMARES$_f$. The filter design itself takes less than 1 s in this case. Overall, FS estimation in this example reduces the computation time with a factor of 5.

Frequency-Selective Quantification of *in vivo* ^{31}P Spectrum

Procedure. The simulation study in this section is based on a signal derived from an *in vivo* ^{31}P signal of the perfused rat liver, acquired at 4.7 T (81 MHz), displayed in the left-hand side of Fig. 1. The β-ATP signal is used for the quantification of the ATP concentration, since, unlike the other ATP signals, it is essentially free from underlying contributions from other low-concentration molecules. Therefore, in some applications, it is interesting to extract only the parameters of the β-ATP

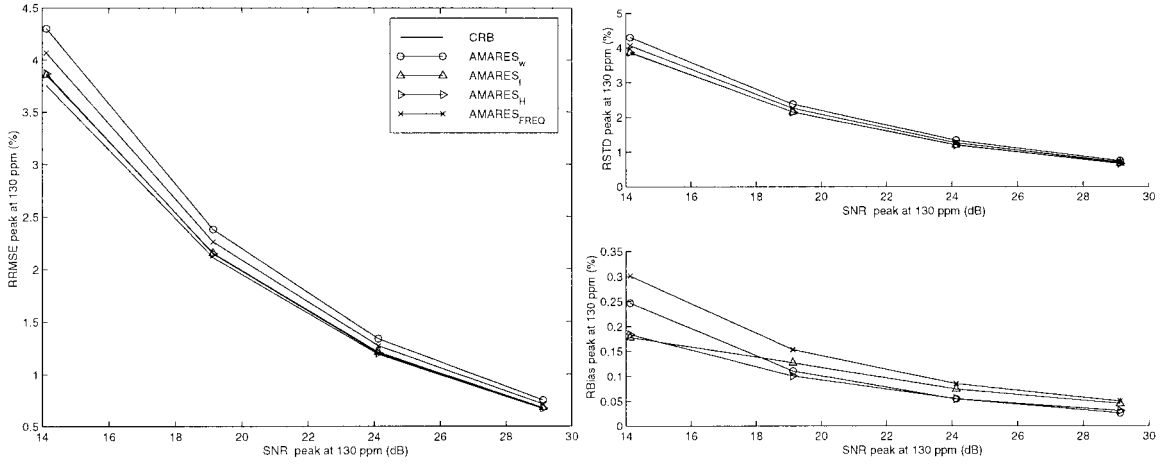

FIG. 11. RRMSE, RSTD, and RBias of amplitude estimates of the peak located at 130 ppm in the simulated ^{13}C spectrum derived from Fig. 10 for different SNRs using different FS estimation methods. The polynomial order P used in AMARES$_{FREQ}$ is 1 for all SNR levels. Left: RRMSE. Top right: RSTD. Bottom right: RBias. CRB denotes the theoretical lower bound.

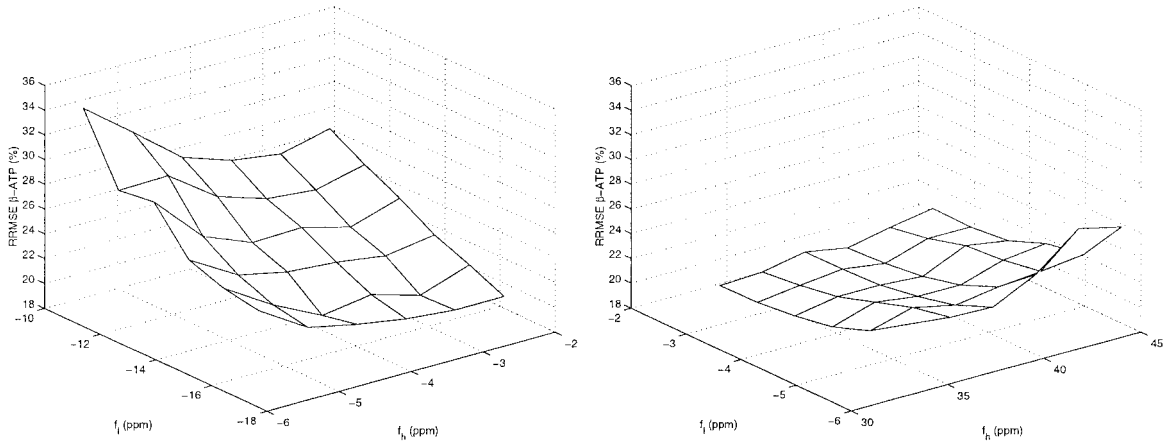

FIG. 12. RMSE of the amplitude estimate of β-ATP in case their concentration has dropped to 20% of their original value. Left: a function of the frequency region included in Eq. [13] ($AMARES_{FREQ}$). The results are displayed for $P = 1$. Right: a function of the choice of the lower and upper bounds in the filter design scheme ($AMARES_f$).

triplet. In Fig. 2, how the triplet can be extracted using the FIR-filter method is illustrated.

To derive the parameters of the simulation signal, the *in vivo* signal was quantified using AMARES. A Lorentzian lineshape was used and the first data points were excluded from the fit to reduce the influence of the broad contribution. The following prior knowledge was imposed:

1. amplitude ratios: 1/1 in doublets of α- and γ-ATP, 1/2/1 in β-ATP triplet;
2. frequency splittings of 16 Hz within the multiplets;
3. dampings of all ATP peaks equal;
4. phases of all peaks equal.

The noise level in the simulations is similar to the one in the *in vivo* signal. In all four compared methods the dampings and phases of the three β-ATP peaks were constrained to be equal and an amplitude ratio of 1/2/1 and frequency splittings of 16 Hz were imposed. The sampling frequency is 5 kHz and the number of data points is 128. The amplitude of the ATP peaks was varied from 20 to 100% of the originally derived value. These variations correspond to changes typically encountered in *in vivo* ^{31}P signals. The accuracy of the FS methods is compared. Different polynomial orders ($P = 0$–5) and widths of frequency regions were tested for $AMARES_{FREQ}$. The upper bound of the frequency region included in Eq. [13] was varied between -5.88 and -2.5 ppm and the lower bound between -17.5 and -11.3 ppm. The influence of the specified upper and lower bounds in the filter design scheme was tested. The lower bound of the frequency region to be suppressed was varied between -5.88 ppm and -2.5 ppm and the upper bound between 33.2 ppm and 44.2 ppm. The sensitivity of the RRMSE of the β-ATP peaks w.r.t. the choice of these parameters in $AMARES_{FREQ}$ and $AMARES_f$ is shown in Fig. 12. The results are displayed for the lowest ATP value and in $AMARES_{FREQ}$ the polynomial order is $P = 1$. The influence of the choice of the model order in $AMARES_H$ and the automatic order selection method were investigated and the results of this analysis are shown in Fig. 13 for the β-ATP peaks. In Fig. 14 the 4 FS parameter estimation methods are compared. The results of $AMARES_{FREQ}$ corresponding to $P = 1$ are shown. A lower bound of -2.5 ppm and an upper bound of 33.2 ppm were used in $AMARES_f$. The model order in $AMARES_H$ is equal to 11, the theoretical model order.

Results and discussion. A polynomial order of $P = 1$ is needed to get reasonable results—low bias and standard deviation—for $AMARES_{FREQ}$ (Fig. 12). The choice of the region included in Eq. [13] is in this case important since the model is only valid in a rather small frequency region. Higher polynomial orders, however, lead to a large increase in standard deviation, thereby worsening the results even more. The choice of lower and upper bound in the $AMARES_f$ filter design scheme is not critical. Only if the lower bound is chosen close to the peaks of interest is there a degradation of the amplitude estimates.

For $AMARES_H$ the choice of the model order has a strong influence on the accuracy of amplitude estimates of the β-ATP peaks (Fig. 13). Slight undermodeling ($M = 8$) gives the same results as using the correct model order. Overmodeling, on the other hand, leads to much worse results. A similar phenomenon occurs when HSVD is used to model the water peak (*24*) in proton spectra. A wrong choice of model order in that application also leads to a deterioration of parameter accuracy. The MDL criterion for automatic order estimation works very well for ATP values between 40 and 100% of the original value, but breaks down for values below 20%, i.e., in low SNR cases.

The results for $AMARES_H$ and $AMARES_f$ are very similar if the right model order is chosen in $AMARES_H$ (Fig. 14). $AMARES_w$ and $AMARES_{FREQ}$ perform worse for all ATP amplitudes. Again the bias is effectively removed for all

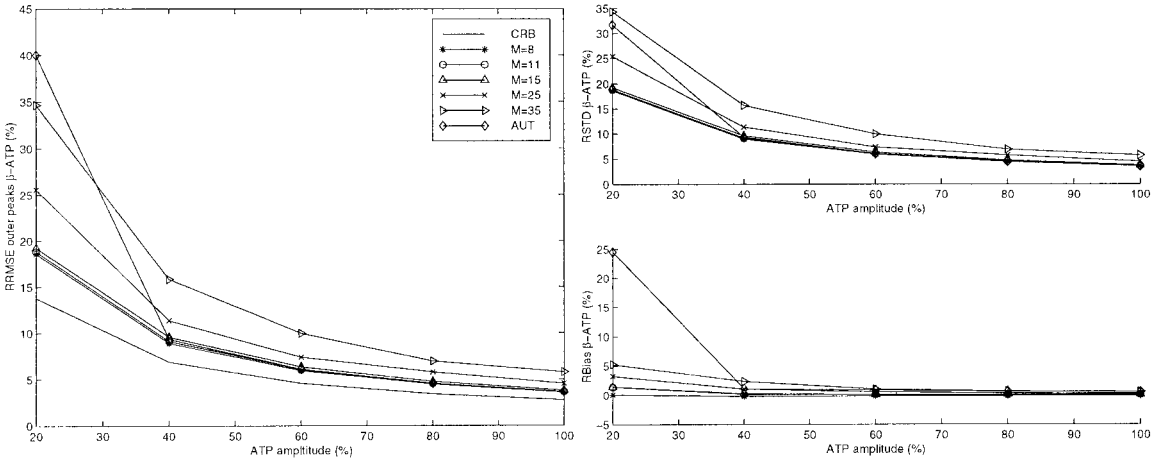

FIG. 13. RRMSE, RSTD, and RBias of the estimates of the amplitudes of the β-ATP peaks for different choices of model order (M) in $AMARES_H$. AUT denotes the results obtained by the automatic model order selection criterion. The results are obtained for ATP amplitudes ranging between 20 and 100% of the maximum value. The corresponding SNR of the two outer peaks in the triplet varies linearly between −4 and 10 dB. CRB denotes the theoretical lower bound.

methods but the standard deviation of $AMARES_w$ and $AMARES_{FREQ}$ is higher.

The gain in calculation time in this particular example is less pronounced than in the ^{13}C example. In the ^{13}C example 2 of 15 peaks are quantified. Since no prior knowledge is used in the ^{13}C example, $AMARES_f$ minimizes a cost function with 8 variables while AMARES used on the entire signal minimizes a cost function with 60 variables. In this particular example, since we impose a lot of prior knowledge, the number of variables to be fitted by AMARES used on the entire signal is 20, compared to 8 variables for $AMARES_f$. Quantifying the entire spectrum using AMARES takes about 1.5 s, while analyzing the β-ATP region only with $AMARES_f$ takes 0.3 s. The filter design in this case also took less than 1 s. Overall, the FIR filter approach still leads to a slight reduction in calculation time.

CONCLUSIONS

In this paper frequency-selective quantification of biomedical MRS data is studied. The influence of nuisance peaks is in most cases not negligible and must be taken care of prior to parameter estimation. A number of methods that can be used in combination with time-domain or frequency-domain model-fitting procedures are revisited: time-domain weighting, HSVD filtering, fitting in the frequency domain using a polynomial baseline. A new method based on maximum-phase FIR filter-

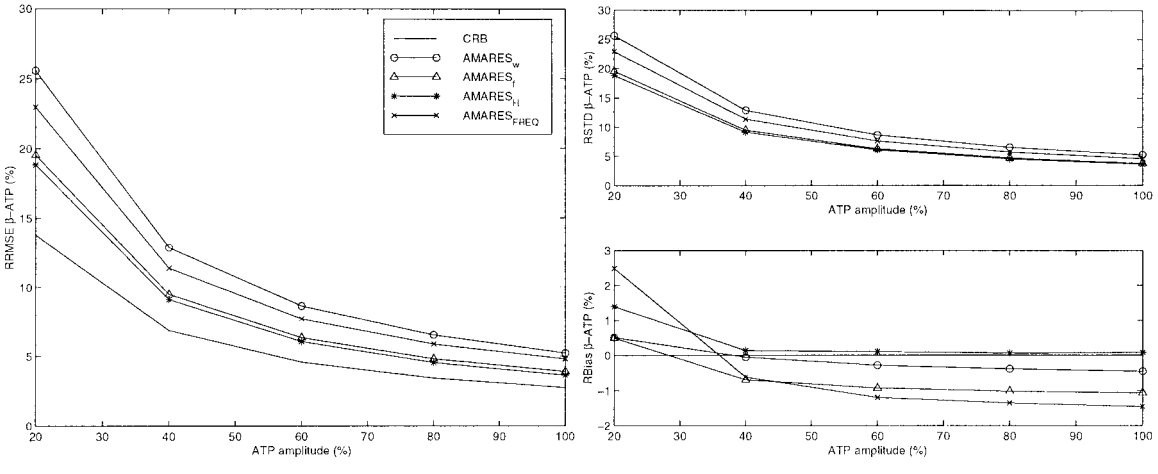

FIG. 14. RRMSE, RSTD, and RBias of the estimates of the amplitudes of the β-ATP peaks for four different FS estimation methods. The results are obtained for ATP amplitudes ranging between 20 and 100% of the maximum value. The corresponding SNR of the two outer peaks in the triplet varies linearly between −4 and 10 dB. The polynomial order used in $AMARES_{FREQ}$ is 1 for all β-ATP amplitudes. Left: RRMSE. Top right: RSTD. Bottom right: RBias. CRB denotes the theoretical lower bound.

ing is presented in this context and compared with the other FS methods.

The time-domain weighting procedure, $AMARES_w$, can be used to reduce the estimator bias if the peaks of interest are well separated from the nuisance peaks. The price is, however, a loss of SNR thereby increasing the estimator variance. Moreover, the method breaks down when the peak of interest is close to the nuisance peaks.

If the appropriate model order is chosen, the HSVD filtering method $AMARES_H$ is seen to be a very accurate technique for FS estimation. A drawback of the method is the sensitivity to the chosen model order in more difficult cases. The use of the minimum description length criterion to automatically determine the model order was examined. MDL works well for a reasonable SNR but breaks down when the SNR decreases. A drawback of the use of MDL is that an entire singular value decomposition of a matrix needs to be computed, which is time consuming.

Analyzing only a part of the spectrum is straightforward to implement for frequency-domain fitting methods (e.g., $AMARES_{FREQ}$) suggesting that these methods are suitable for FS estimation. However, the situation is complicated by the influence of the nuisance peaks which have to be modeled in some way. Using a polynomial to model the tails of the nuisance peaks is one way of reducing this influence. If the polynomial badly models the baseline in the region of interest, bias is introduced in the estimates and the choice of polynomial order in combination with the size of the frequency region is seen to be critical. These extra baseline model parameters also increase the variance of the estimates.

If the upper and lower bounds of the frequency region(s) to be suppressed are chosen with some care, $AMARES_f$ is able to keep the bias and the standard deviation of the parameters of interest low. The automatic filter design scheme makes the method easy to use and the computational complexity of $AMARES_f$ is low.

Two important MRS applications were examined. The first one is the extraction of two close peaks, remotely located from the other peaks in a ^{13}C spectrum. Since this is a rather easy example, all FS methods perform very similarly. The use of the FIR filter approach leads in this case to a reduction in computational time when compared to fitting the entire spectrum with AMARES. In the ^{31}P example, the parameters of the β-ATP peaks were extracted. The FIR filter method leads to the best results. HSVD filtering performs well only for an appropriate choice of model parameters and the accuracy of the results is seen to be very sensitive to the model order. For $AMARES_{FREQ}$ the choice of the polynomial order and frequency region included in the fit were rather critical for this example.

In summary, the low computational complexity, ease of use, and high parameter accuracy make $AMARES_f$ an attractive tool for FS parameter estimation.

ACKNOWLEDGMENTS

This work is supported by the Belgian Programme on Interuniversity Poles of Attraction (IUAP-4/2 and 24), initiated by the Belgian State, Prime Minister's Office for Science, Technology, and Culture, by the EU Programme "Training and Mobility of Researchers" project (Contract ERBFMRXCT970160) entitled "Advanced Signal Processing for Medical Magnetic Resonance Imaging and Spectroscopy," by a Concerted Research Action (GOA) project of the Flemish Community, entitled "Model-Based Information Processing Systems," by F.W.O. Fund for Scientific Research–Flanders Grant G.0360.98, and by the Swedish Foundation for Strategic Research. LV is a Ph.D. student funded by the IWT (Flemish Institute for Support of Scientific–Technological Research in Industry). SVH is a Research Director with the F.W.O. This work was carried out at ESAT-SISTA/COSIC, Katholieke Universiteit Leuven, Heverlee, Belgium.

REFERENCES

1. R. Kumaresan and D. Tufts, Estimating the parameters of exponentially damped sinusoids and pole-zero modeling in noise, *IEEE Trans. Accoust. Speech Signal Process.* **ASSP 30,** 833–840 (1982).
2. H. Barkhuysen, R. de Beer, W. Bovée, and D. van Ormondt, Retrieval of frequencies, amplitudes, damping factors, and phases from time-domain signals using a linear least-squares procedure, *J. Magn. Reson.* **61,** 465–481 (1985).
3. S. Kung, K. Arun, and D. B. Rao, State–space and singular-value decomposition-based approximation methods for the harmonic retrieval problem, *J. Opt. Soc. Am.* **73,** 1799–1811 (1983).
4. H. Barkhuysen, R. de Beer, and D. van Ormondt, Improved algorithm for noniterative time-domain model fitting to exponentially damped magnetic resonance signals, *J. Magn. Reson.* **73,** 553–557 (1987).
5. H. Chen, S. Van Huffel, D. van Ormondt, and R. de Beer, Parameter estimation with prior knowledge of known signal poles for the quantification of NMR spectroscopy data in the time domain, *J. Magn. Reson. A* **119,** 225–234 (1996).
6. H. Chen, S. Van Huffel, A. van den Boom, and P. van den Bosch, Subspace-based parameter estimation of exponentially damped sinusoids using prior knowledge of frequency and phase, *Signal Process.* **59,** 129–136 (1997).
7. S. Van Huffel, Subspace-based exponential data modeling using prior knowledge, *in* "Proc. of the IEEE Benelux Chapter Signal Processing Symposium (IEEEBSPS)," pp. 211–214, Leuven, Belgium (1998).
8. J. van der Veen, R. de Beer, P. Luyten, and D. van Ormondt, Accurate quantification of *in vivo* ^{31}P NMR signals using the variable projection method and prior knowledge, *Magn. Reson. Med.* **6,** 92–98 (1988).
9. L. Vanhamme, A. van den Boogaart, and S. Van Huffel, Improved method for accurate and efficient quantification of MRS data with use of prior knowledge, *J. Magn. Reson.* **129,** 35–43 (1997).
10. R. Meyer, M. Fisher, S. Nelson, and T. Brown, Evaluation of manual methods for integration of in vivo phosphorus NMR spectra, *NMR Biomed.* **1,** 131–135 (1988).
11. F. Abildgaard, H. Gesmar, and J. Led, Quantitative analysis of complicated nonideal Fourier transform NMR spectra, *J. Magn. Reson.* **79,** 78–89 (1988).
12. Y. Hiltunen, M. Ala-Korpela, J. Jokisaari, S. Eskelinen, K. Kiviniitty, M. Savolainen, and Y. Kesäniemi, A lineshape fitting model for ^{1}H NMR spectra of human blood plasma, *Magn. Reson. Med.* **21,** 222–232 (1991).

13. Y.-L. Martin, A global approach to accurate and automatic quantitative analysis of NMR spectra by complex least-squares curve fitting, *J. Magn. Reson. A* **111,** 1–10, (1994).
14. J.-P. Grivet, Accurate numerical approximation to the gauss-lorentz lineshape, *J. Magn. Reson.* **125,** 102–106 (1997).
15. J. Slotboom, C. Boesch, and R. Kreis, Versatile frequency domain fitting using time domain models and prior knowledge, *Magn. Reson. Med.* **39,** 899–911 (1998).
16. I. Dologlou, S. Van Huffel, and D. Van Ormondt, Frequency-selective MRS data quantification with frequency prior knowledge, *J. Magn. Reson.* **130,** 238–243 (1998).
17. J. Tang and J. Norris, LP-ZOOM, a linear prediction method for local spectral analysis of NMR signals, *J. Magn. Reson.* **79,** 190–196 (1988).
18. A. Knijn, R. D. Beer, and D. Van Ormondt, Frequency-selective quantification in the time domain, *J. Magn. Reson.* **97,** 444–450 (1992).
19. S. Cavassila, B. Fenet, A. van den Boogaart, C. Remy, A. Briguet, and D. Graveron-Demilly, ER-filter: A preprocessing technique for frequency-selective time-domain analysis, *J. Magn. Reson. Anal.* **3,** 87–92 (1997).
20. A. van den Boogaart, D. van Ormondt, W. Pijnappel, R. de Beer, and M. Ala-Korpela, Removal of the water resonance from ^{1}H magnetic resonance spectra, *in* "Mathemetics in Signal Procenings III" (J. G. McWhirter, Ed.), p. 175–195. Uarendon Pren, Oxford, 1994.
21. K. Cross, Improved digital filtering technique for solvent suppression, *J. Magn. Reson. A* **101,** 220–224 (1993).
22. G. Sobering, M. Kienlin, C. Moonen, P. van Zijl, and A. Bizzi, Post-acquisition reduction of water signals in proton spectroscopic imaging of the brain, *in* "Proc. SMRM," p. 771 (1991).
23. M. Deriche and X. Hu, Elimination of water signal by postprocessing, *J. Magn. Reson. A* **101,** 229–232 (1992).
24. T. Sundin, L. Vanhamme, P. Van Hecke, I. Dologlou, and S. Van Huffel, Accurate quantification of ^{1}H spectra: From FIR filter design for solvent suppression to parameter estimation, *J. Magn. Reson.* **139,** 189–204 (1999).
25. G. Della Lunga, R. Pogni, and R. Basosi, A simple method for baseline correction in EPR spectroscopy, *J. Magn. Reson. A* **108,** 65–70 (1994).
26. http://azur.univ-lyon1.fr/TMR/tmr.html.
27. http://www.mrui.uab.es/mrui/mruiHomePage.html.
28. J. G. Proakis and D. G. Manolakis, "Digital Signal Processing: Principles, Algorithms and Applications," Prentice Hall, Upper Saddle River, NJ (1996).
29. A. V. Oppenheim and R. W. Schafer, "Discrete-Time Signal Processing," Prentice Hall, Eaglewood Cliffs, NJ (1989).
30. C. Craven and J. Waltho, The action of time-domain convolution filters for solvent suppression, *J. Magn. Reson. B* **106,** 40–46 (1995).
31. O. Herrmann and H. Schüßler, Design of nonrecursive digital filters with minimum phase, *Electron. Lett.* **6,** 329–330 (1970).
32. I. Selesnick, M. Lang, and C. Burrus, Constrained least square design of FIR filters without specified transition bands, *IEEE Trans. Signal Process.* **44**(8), 1879–1892 (1996).
33. H. Barkhuysen, R. de Beer, and D. van Ormondt, Improved algorithm for noniterative time-domain model fitting to exponentially damped magnetic resonance signals, *J. Magn. Reson.* **73,** 553–557 (1987).
34. W. Pijnappel, A. van den Boogaart, R. de Beer, and D. van Ormondt, SVD-based quantification of magnetic resonance signals, *J. Magn. Reson.* **97,** 122–134 (1992).
35. L. Vanhamme, R. Fierro, S. Van Huffel, and R. de Beer, Fast removal of residual water in proton spectra, *J. Magn. Reson.* **132,** 197–203 (1998).
36. T. Söderström and P. Stoica, "System Identification," Prentice Hall, London (1989).
37. Yuang-Ya Lin, P. Hodgkinson, M. Ernst, and A. Pines, A novel detection–estimation scheme for noisy NMR signals: Applications to delayed acquisition data, *J. Magn. Reson.* **128,** 30–41 (1997).