Journal of Magnetic Resonance 203 (2010) 167–176

Contents lists available at ScienceDirect

Journal of Magnetic Resonance

journal homepage: www.elsevier.com/locate/jmr

Spectral estimation of irregularly sampled exponentially decaying signals with applications to RF spectroscopy ☆

Erik Gudmundson [a,*], Petre Stoica [a], Jian Li [b], Andreas Jakobsson [c], Michael D. Rowe [d] John A.S. Smith [d], Jun Ling [b]

[a] Department of Information Technology, Uppsala University, P.O. Box 337, SE-751 05 Uppsala, Sweden
[b] Department of Electrical and Computer Engineering, University of Florida, Gainesville, FL 32611-6130, USA
[c] Department of Mathematical Statistics, Lund University, P.O. Box 118, SE-221 00 Lund, Sweden
[d] King's College London, Division of Engineering, Strand, London WC2R 2LS, UK

ARTICLE INFO

Article history:
Received 11 October 2009
Revised 8 December 2009
Available online 24 December 2009

Keywords:
Missing data
Spectral estimation
Nonparametric spectral estimation
Irregularly sampled data
Spectroscopic techniques

ABSTRACT

The problem of estimating the spectral content of exponentially decaying signals from a set of irregularly sampled data is of considerable interest in several applications, for example in various forms of radio frequency spectroscopy. In this paper, we propose a new nonparametric iterative adaptive approach that provides a solution to this estimation problem. As opposed to commonly used methods in the field, the damping coefficient, or linewidth, is explicitly modeled, which allows for an improved estimation performance. Numerical examples using both simulated data and data from NQR experiments illustrate the benefits of the proposed estimator as compared to currently available nonparametric methods.

1. Introduction

Spectral estimation is a classical problem with applications in a wide variety of fields, such as astronomy, communications, economics, medical imaging, radar, spectroscopic techniques, e.g., nuclear magnetic resonance (NMR) and nuclear quadrupole resonance (NQR), and much more; consequently, the literature pertaining to the problem is very rich, see, e.g., [1] and the references therein. However, the majority of the works focuses on estimating the spectrum from a finite sequence with evenly sampled data, whereas in many applications it is often not possible or suitable to sample in such a way. This leads to sets of data that we categorize as either "gapped" or having "missing" samples, where by the former we mean regularly sampled data in which blocks of samples are missing, and by the latter we mean that there is no particular structure in the sampling. Applying traditional spectral estimation algorithms to such sets of data may lead to a large loss in performance compared to having a complete set of data [2,3].

The problem of estimating the spectral parameters of exponentially damped signals has attracted significant attention during recent decades as such signals arise naturally in several applications, for example in NMR and NQR. Traditionally, the frequencies and damping coefficients, or linewidths, of such signals are estimated using the periodogram, where the dampings are given as the width of the peaks. When only a few samples are available, this estimate is highly biased as the peak width in the periodogram depends on the number of samples (see, e.g., [1]). To improve the spectrum estimation, both parametric methods, e.g., the Matrix Pencil method (MP) [4], and nonparametric methods, e.g., the damped Capon algorithm [5], have been proposed. Commonly, the free induction decay (FID) of such a spectroscopic signal is modeled as a sum of exponentially damped sinusoidals. However, the use of more refined signal models such as using a Voigt lineshape for the spectral lines have also been shown to be beneficial [6–8]. Model-based methods making strong use of the assumed signal model will be sensitive to deviations to the assumed model, whereas nonparametric methods, not making use of such explicit models, will generally offer more robust spectral estimates. As an example, we will here examine stochasticly excited spectroscopic signals, for which the measurements can be well modeled as FID's.

Stochastic or noise excitation has found application in NMR [9,10], electron paramagnetic resonance (EPR) [11] and NQR

☆ This work was supported in parts by the Swedish Research Council (VR), Carl Trygger's Foundation, the Swedish Foundation for International Cooperation in Research and Higher Education (STINT), the US National Science Foundation under Grant No. ECCS-0729727, the UK Home Office, the UK Defence Science and Technology Laboratory, and the European Research Council (ERC).

* Corresponding author. Fax: +46 18 51 19 25.
E-mail addresses: erik.gudmundson@it.uu.se (E. Gudmundson), ps@it.uu.se (P. Stoica), li@dsp.ufl.edu (J. Li), andreas.jakobsson@matstat.lu.se (A. Jakobsson), michaeldavidrowe@hotmail.com (M.D. Rowe), john.smith@kcl.ac.uk (J.A.S. Smith), lingjun@ufl.edu (J. Ling).

1090-7807/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmr.2009.12.015

[12,13]. Often, there are advantages in acquiring more than one complex data point after each radio frequency (RF) pulse of a stochastic excitation sequence but the resultant FID signal will then have gaps where data is missing between the blocks of acquired data [9,10]. The gaps in the data correspond to the times when the RF pulses are applied plus the following dead-times before which the signal can be acquired. In this situation, there are two dwell, or sampling, times. One is the actual sampling time between the acquired complex data points in each block of data, while the other is the time between equivalent points of each block of data. Here, we refer to the latter as the stochastic dwell or sampling time and define it as the gap (missing data) time plus the acquisition time of each data block. It is thus the time between the first data point of one block and the first point of the next. In general, the stochastic dwell time is not an integer multiple of the actual data dwell time. In such a case, the noise-free signal can be well modeled as a damped sinusoidal signal, where the data is given as blocks of evenly sampled data but where the gaps between the blocks are not necessarily an integer multiple of the sampling time within the blocks. Therefore, the data can sometimes be considered *gapped* but is often highly irregularly sampled.

The problem of estimating spectra when data is missing has been considered, for example, in the astronomical literature, where several parametric and nonparametric methods have been proposed, see, e.g., [2,14,15]. Irregular sampling schemes have also drawn attention in NMR, as they allow for faster data acquisition in multidimensional NMR. Numerous methods for spectral estimation have been proposed, for example linear prediction (LP) (see, e.g., [16]), the filter diagonalization method [17], PARAFAC, and other multidimensional decomposition methods (see [18] for a good tutorial). These methods will first reconstruct the missing data to yield a full data sequence without any missing samples. Then, in a second step, the spectral content is estimated from the reconstructed data set. If the data has been sampled irregularly, this might necessitate the reconstruction of the data on a very fine sampling grid, possibly making the so-obtained full data set very large and cumbersome to process.

Methods that directly estimate the spectrum from the given set of data, independent of the sampling grid, have also been proposed. The common approach is to fit the spectrum to the set of available data through either iterative maximization of the spectral entropy, such as the maximum entropy reconstruction method [19] (see also [20] for a comparison with LP) and the forward maximum reconstruction method [21], or by iterative minimization of a norm, for example the l1-norm [22]. These methods have the shortcoming that they require the selection of user parameters, such as the noise level. This can often be a difficult task and it might potentially degrade the performance of the spectrum estimation. Furthermore, such methods do not allow for accurate estimation of the signal decay, which is often of interest in applications such as NMR and NQR. In effect, when only a few samples are available, the quality of the signal decay estimate will typically be quite poor as the width of the spectral mainlobe will be limited by the resolution of the spectral estimator. Algorithms that exploit a Lorentzian signal model, such as MP, often yield satisfactory spectral estimates if the data is well described by the model, but do not work for irregular sampling schemes. Along the same lines, we have earlier proposed an algorithm for accurately reconstructing the missing samples [23], but this will also only work for the gapped data case.

Recently, a new approach, termed the iterative adaptive approach, or IAA, was proposed in [24], and its usefulness has been shown in several areas: MIMO radar imaging [25], source detection [24], spectral analysis of real-valued data [26], and coherence spectrum estimation [27]. This nonparametric algorithm does not require the specification of any user parameters. Furthermore, no assumptions have to be made on the sampling scheme and the spectrum can be directly estimated from the possibly irregularly sampled data. In this paper, we extend the IAA algorithm to also explicitly model the signal decay. Different from the previous derivations of IAA, where a weighted least-squares approach was used, we here derive it using a linear estimation approach. We also note that IAA can, if needed, be extended to reconstruct missing data, leading to the missing-data IAA, or MIAA, algorithm [28]. This step is also possible for the algorithm proposed here, but is beyond the scope of this paper.

To illustrate the power and potential of the proposed algorithm, we examine the results on both simulated data and on measurements from stochastic NQR (sNQR) experiments examining some common explosives.

The paper is organized as follows. In the next section (Section 2) we derive the dIAA algorithm and in Section 3 we propose a subspace algorithm that provides an initial estimate. The performance of the proposed algorithms is studied in Section 4, using both simulated and measured data. Finally, Section 5 contains our conclusions.

2. The dIAA algorithm

In this section, we propose the nonparametric decay IAA (dIAA) algorithm that allows for a two-dimensional (2D) spectral representation of arbitrarily sampled 1D measurements, without needing the selection of any user parameter. The 2D representation of the spectral content of the signal allows for an accurate representation of both the line frequencies and their spectral widths.

Consider a signal described by the sum of the contributions from each frequency grid point $\{\omega_k\}_{k=1}^K$ and decay grid point $\{\beta_d\}_{d=1}^D$:

$$y_{t_n} = \sum_{k=1}^{K}\sum_{d=1}^{D} \alpha_{d,k} e^{(-\beta_d + i\omega_k)t_n}, \quad n = 1,\ldots,N, \tag{1}$$

where $\alpha_{d,k}$ represents the complex-valued amplitude for a possible sinusoidal component with frequency ω_k and decay β_d, whereas N denotes the number of available samples. Note that no assumption is made on the sampling scheme, thereby allowing for arbitrary sampling times, i.e., $\{t_n\}_{n=1}^N$. Furthermore, note that no signal model has been assumed; rather the signal is made from the contribution corresponding to each of the frequency and decay grid points. There is also no corrupting noise term as is typical in model-based methods describing the data as a signal part and a noise part. The contribution of any noise component, or any other interference, is instead implicitly described via its contribution to $\alpha_{d,k}$. Let $\boldsymbol{s}_{d,k}$ denote the (damped) Fourier vector for frequency ω_k and decay β_d, so that

$$\boldsymbol{s}_{d,k} = \left[e^{(-\beta_d + i\omega_k)t_1} \;\cdots\; e^{(-\beta_d + i\omega_k)t_N}\right]^T, \tag{2}$$

and denote the vector of available measurements as

$$\boldsymbol{y} = \left[\, y_{t_1} \;\cdots\; y_{t_N} \,\right]^T. \tag{3}$$

Then, by using (2) and (3), (1) can be rewritten in a more compact form, given by:

$$\boldsymbol{y} = \sum_{k=1}^{K}\sum_{d=1}^{D} \alpha_{d,k}\boldsymbol{s}_{d,k}. \tag{4}$$

Denote the signal energy at the grid point associated with frequency ω_k and damping β_d as

$$P_{d,k} = |\alpha_{d,k}|^2. \tag{5}$$

Also, let

$$\boldsymbol{Q}_{d,k} = \boldsymbol{R} - P_{d,k}\boldsymbol{s}_{d,k}\boldsymbol{s}_{d,k}^*, \tag{6}$$

denote the contribution from all other grid points except the grid point (d,k), where $\boldsymbol{R}$ is the covariance matrix of the available data, given by

$$\boldsymbol{R} = \sum_{k=1}^{K}\sum_{d=1}^{D} P_{d,k}\boldsymbol{s}_{d,k}\boldsymbol{s}_{d,k}^*. \tag{7}$$

Thus, $\boldsymbol{Q}_{d,k}$ can be seen as a form of interference covariance matrix. Consider a general linear estimator of $\alpha_{d,k}$:

$$\hat{\alpha}_{d,k} = \boldsymbol{h}_{d,k}^*\boldsymbol{y}. \tag{8}$$

The weight vector $\boldsymbol{h}_{d,k}$ can be found through

$$\min_{\boldsymbol{h}_{d,k}} \boldsymbol{h}_{d,k}^*\boldsymbol{Q}_{d,k}\boldsymbol{h}_{d,k} \quad \text{s.t. } \boldsymbol{h}_{d,k}^*\boldsymbol{s}_{d,k} = 1, \tag{9}$$

i.e., we design a linear combiner that minimizes the output from all grid points other than (β_d,ω_k), but passes the component with the frequency and damping of interest undistorted. Since $P_{d,k}$ does not depend on $\boldsymbol{h}_{d,k}$ and $P_{d,k}(\boldsymbol{h}_{d,k}^*\boldsymbol{s}_{d,k})(\boldsymbol{h}_{d,k}^*\boldsymbol{s}_{d,k})^* = P_{d,k}$, (9) is equivalent to

$$\min_{\boldsymbol{h}_{d,k}} \boldsymbol{h}_{d,k}^* \underbrace{\left(\boldsymbol{Q}_{d,k} + P_{d,k}\boldsymbol{s}_{d,k}\boldsymbol{s}_{d,k}^*\right)}_{\boldsymbol{R}} \boldsymbol{h}_{d,k} \quad \text{s.t. } \boldsymbol{h}_{d,k}^*\boldsymbol{s}_{d,k} = 1. \tag{10}$$

The minimizer is readily found as (see, e.g., [1])

$$\hat{\boldsymbol{h}}_{d,k} = \frac{\boldsymbol{R}^{-1}\boldsymbol{s}_{d,k}}{\boldsymbol{s}_{d,k}^*\boldsymbol{R}^{-1}\boldsymbol{s}_{d,k}}. \tag{11}$$

An estimate of the amplitude at gridpoint (β_d,ω_k) can thus be found as (see (8))

$$\hat{\alpha}_{d,k} = \frac{\boldsymbol{s}_{d,k}^*\boldsymbol{R}^{-1}\boldsymbol{y}}{\boldsymbol{s}_{d,k}^*\boldsymbol{R}^{-1}\boldsymbol{s}_{d,k}}. \tag{12}$$

As $\boldsymbol{R}$ depends on the amplitudes at each grid point, (12) must be implemented as an iterative algorithm. Here, we use the least-squares (LS) 2D periodogram [5] as an initial estimator, i.e., we initially set $\boldsymbol{R} = \boldsymbol{I}$, the identity matrix. The dIAA spectral representation is thus found by iterating the estimation of $\boldsymbol{R}$ in (7), and the grid point amplitudes in (12), until a suitable stopping criterion is met. Table 1 provides a summary of the required steps. Herein, we iterate until the estimates have practically converged, i.e., the difference between the amplitudes $\hat{\alpha}_{d,k}$ between each step is smaller than some preset threshold ε, which generally requires no more than 10–15 iterations to provide a good solution. There is as of yet no proof of convergence of the IAA, MIAA, or dIAA algorithms; however, in [25], the IAA algorithm has been shown to converge locally. All indications suggest that a similar result would hold also for dIAA. It is worth noting that the difference between dIAA and IAA is that the former establishes a 2D grid mesh, both in frequency and damping domain, whereas the latter only considers a 1D frequency grid. Hence, the resulting dIAA algorithm will be roughly D times more computational demanding than IAA. As can be expected, this also implies that if we set $D = 1$ and $\beta_1 = 0$, omitting the detailing of the signal decays, then dIAA reduces to IAA. Typically for spectroscopic signals, most of the $\alpha_{d,k}$ values are nearly zero, except for a small number of significant elements, and analogously to the IAA case, this suits the dIAA algorithm best [24]. Finally, we note that in order to decrease the computational complexity, one may implement the algorithm by making use of a procedure similar to the one suggested in [29].

Table 1
Outline of the dIAA algorithm.

Step 0:	Initialize using an LS-estimate: $\hat{\alpha}_{d,k} = \frac{\boldsymbol{s}_{d,k}^*\boldsymbol{y}}{\boldsymbol{s}_{d,k}^*\boldsymbol{s}_{d,k}} = \frac{\boldsymbol{s}_{d,k}^*\boldsymbol{y}}{\sum_{t=1}^{N_I} e^{-2\beta_d t}}$.
Step 1:	Compute the covariance matrix $\boldsymbol{R}$ using (7) together with the most recent estimate of $\{\hat{\alpha}_{d,k}\}$.
Step 2:	Compute the amplitude estimate $\{\hat{\alpha}_{d,k}\}$ using (12) with the most recent estimate of $\boldsymbol{R}$.
Step 3:	Repeat steps 1 and 2 until practical convergence.

3. A subspace-based algorithm

As we are herein primarily interested in spectroscopic data exhibiting Lorentzian lineshapes, measured in uniformly sampled subblocks, we now proceed to proposing a fast parametric, i.e., model-based, subspace-based algorithm for obtaining an initial estimate of the range of frequencies and dampings of interested. The algorithm is termed damped block-ESPRIT (dBESP), as it is based on the principles of the ESPRIT algorithm (see, e.g., [1]). It should be stressed that this algorithm makes much stronger assumptions on the assumed signal, making it applicable in only a subset of the cases where dIAA can be used. Nevertheless, for spectroscopic data, the method allows for a way to initialize the region of interest for dIAA. Assume that the noise-free data can be written as

$$y_t = \sum_{p=1}^{P} \alpha_p e^{(-\beta_p + i\omega_p)t}, \tag{13}$$

where P is the number of damped sinusoids in the signal, which is assumed to be known. Furthermore, assume that the data is generated in blocks of N_{bl} regularly sampled points, where the length of the gap between two consecutive blocks is N_{gap} samples, not necessarily an integer, and that S is the number of blocks. Let $\{\boldsymbol{x}_k\}_{k=1}^{S} \in \mathbb{C}^{N_{bl}\times 1}$ denote the vector of data in block k. Then, $\boldsymbol{x}_k$ can be written as

$$\boldsymbol{x}_k = \boldsymbol{A}\boldsymbol{B}\boldsymbol{c}_k, \tag{14}$$

where

$$\boldsymbol{A} = [\boldsymbol{a}_1 \ \cdots \ \boldsymbol{a}_P], \tag{15}$$

$$\boldsymbol{a}_n = \left[e^{(-\beta_n + i\omega_n)} \ \ldots \ e^{(-\beta_n + i\omega_n)N_{bl}}\right]^T, \tag{16}$$

$$\boldsymbol{B} = \mathrm{diag}\left([\alpha_1 \ \cdots \ \alpha_P]^T\right), \tag{17}$$

$$\boldsymbol{c}_k = \left[e^{(-\beta_1 + i\omega_1)(k-1)(N_{bl}+N_{gap})} \ \ldots \ e^{(-\beta_P + i\omega_P)(k-1)(N_{bl}+N_{gap})}\right]^T, \tag{18}$$

and $\mathrm{diag}(\boldsymbol{x})$ is a diagonal matrix with the vector $\boldsymbol{x}$ on its diagonal. The matrix containing each block of data as a column can therefore be written as

$$\boldsymbol{X} = [\boldsymbol{x}_1 \ \cdots \ \boldsymbol{x}_S] = \boldsymbol{A}\boldsymbol{B}\boldsymbol{C}^*, \tag{19}$$

where $\boldsymbol{C}^* = [\boldsymbol{c}_1 \ \cdots \ \boldsymbol{c}_S]$. Let $\boldsymbol{U}$ denote the matrix made from the P dominant left singular vectors of $\boldsymbol{X}$. Then, reminiscent of the ESPRIT algorithm (see, e.g., [1]), we can use $\boldsymbol{U}$ to find an estimate of $\{\beta_k\}_{k=1}^{P}$ and $\{\omega_k\}_{k=1}^{P}$, since range (A) is spanned by $\boldsymbol{U}$. The estimates are found from the magnitude and angle, respectively, of the eigenvalues of

$$\boldsymbol{\Phi}_L = \left(\boldsymbol{U}_1^*\boldsymbol{U}_1\right)^{-1}\boldsymbol{U}_1^*\boldsymbol{U}_2, \tag{20}$$

where $\boldsymbol{U}_1 = [\boldsymbol{I}_{N_{bl}-1} \ \boldsymbol{0}]\boldsymbol{U}$ and $\boldsymbol{U}_2 = [\boldsymbol{0} \ \boldsymbol{I}_{N_{bl}-1}]\boldsymbol{U}$, $\boldsymbol{I}_{N_{bl}-1}$ denotes the identity matrix of size $(N_{bl}-1)\times(N_{bl}-1)$, and $\boldsymbol{0}$ the $(N_{bl}-1)\times 1$ column vector of zeros. When $S < N_{bl}$, using the matrix $\boldsymbol{V}$ made from the right singular vectors of $\boldsymbol{X}$, which span range $(\boldsymbol{C})$, is preferable. The estimates of the dampings and frequencies are then retrieved from the eigenvalues of

$$\boldsymbol{\Phi}_R = \left(\boldsymbol{V}_1^*\boldsymbol{V}_1\right)^{-1}\boldsymbol{V}_1^*\boldsymbol{V}_2, \tag{21}$$

where $\boldsymbol{V}_1$ and $\boldsymbol{V}_2$ are formed similarly to $\boldsymbol{U}_1$ and $\boldsymbol{U}_2$, but with $\boldsymbol{U}$ replaced by the conjugate of $\boldsymbol{V}$. However, aliasing in the frequency

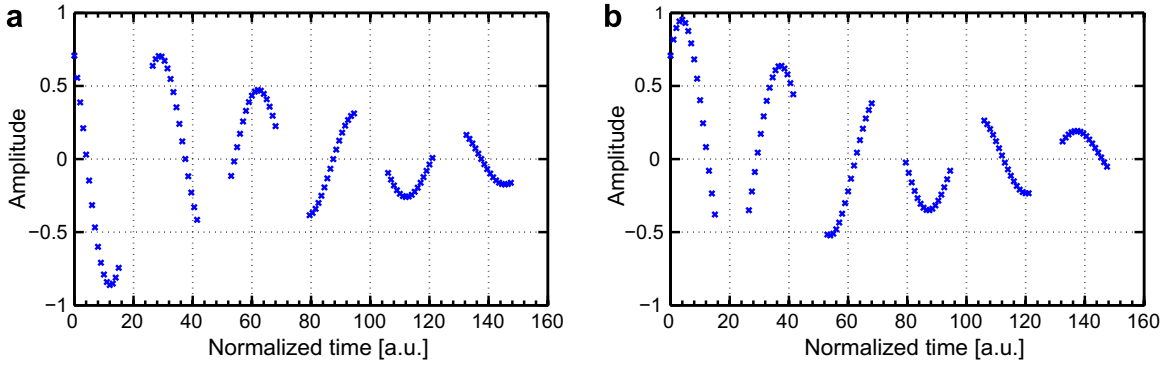

Fig. 1. (a) Real part and (b) imaginary part of a simulated FID without noise.

estimate will likely occur, due to the factor $(N_{bl}+N_{gap})$ in (18). A remedy for this is to use the left singular vectors to get an initial estimate $\{\hat{\omega}_k^L\}_{k=1}^P$, and then use that information to unwrap the phase. Denote the (possibly) wrapped frequency estimates obtained from the right singular vectors by $\{\hat{\psi}_k\}_{k=1}^P$, so that

$$\hat{\psi}_k = \omega_k(N_{bl}+N_{gap}) - 2\pi q_k. \tag{22}$$

An estimate of q_k can be found as

$$\hat{q}_k = \frac{\hat{\omega}_k^L(N_{bl}+N_{gap}) - \hat{\psi}_k}{2\pi}, \tag{23}$$

where $\hat{q}_k$ is rounded off to the closest integer. The estimates of the frequencies, obtained from the right singular vectors, are thus formed as

$$\hat{\omega}_k^R = \frac{\hat{\psi}_k + 2\pi\hat{q}_k}{N_{bl}+N_{gap}}, \quad k=1,\ldots,P, \tag{24}$$

and estimates of the dampings as the natural logarithm of the magnitude of the eigenvalues of $\boldsymbol{\Phi}_R$. We note that the number of sinusoids that can be dealt with by dBESP is limited by the available number of data blocks.

4. Results

In this section, we illustrate the performance of proposed algorithms using both simulated data and experimental data.

4.1. Simulations

We first evaluate the algorithms on simulated data, examining the performance for different signal-to-noise ratios (SNR's), defined as $\text{SNR} = \sigma_w^{-2}\sigma_s^2$, where σ_w^2 and σ_s^2 denote the energy of the noise and of the noise-free signal, respectively. The performance was evaluated using the normalized root mean squared error (nRMSE):

$$\text{nRMSE}(\hat{x}) = \sqrt{E\left\{\frac{(x-\hat{x})^2}{x^2}\right\}}, \tag{25}$$

where E is the expectation operator, x denotes the true parameter and $\hat{x}$ an estimate thereof. Initially, we examine a simplistic signal model, containing only a single damped sinusoid corrupted by a zero mean white Gaussian circularly symmetric noise. The exponentially damped sinusoid has frequency $\omega = 2\pi f, f = 0.03$ Hz, damping (or spectral width) $\beta = 0.012$ Hz, and complex amplitude $\alpha = 1+i$, where the initial phase is uniformly distributed on $[0, 2\pi)$. The data is sampled in blocks of 16 samples with a sampling rate of 1 Hz. The first block starts at 0 and the kth block at 26.5 k. In total, 6 blocks of data were used, giving a total of $N = 96$ samples. Fig. 1 displays the real and imaginary part of the noise-free signal.[1] We compare the proposed algorithms with the periodogram spectral estimator and the maximum likelihood (ML) estimator (see, e.g., [1]). The former method is commonly used for estimation of damping and frequency, where the damping can be estimated from half-width of the peak at half its height. As the data is irregularly sampled, we use the least-squares periodogram (see, e.g., [31]), in the following abbreviated as PerLS. Note that extending the ML algorithm to multiple peak scenarios is not straightforward and, as this paper focuses on nonparametric spectral estimation, this algorithm will not be considered for those scenarios. The spectra were evaluated over $K = 1000$ frequency grid points, and for dIAA and ML the damping grid ranged from 0 to 0.02 in $D = 101$ steps. Furthermore, dIAA was iterated ten times, which empirically has been shown to be enough to provide a good solution.

We note that one could also compare with the periodogram estimated from averaging the sub-periodograms from each block. This implies that each periodogram is estimated from only 16 samples and, as the periodogram is inconsistent and its performance highly depends on the number of samples used, a large bias in the damping estimate is expected. Therefore, this approach is not recommended and will not be further considered.

Fig. 2 shows the nRMSE of the frequency, damping, and amplitude estimates, where the expectation in (25) was empirically evaluated over 100 Monte-Carlo simulations. As a comparison, the Cramér-Rao lower bound (CRB) is also displayed, showing the theoretical lower limit for the variance of any unbiased estimator (see [32] for a reference on the CRB for exponentially damped sinusoids).

As is seen in Fig. 2a, dIAA, PerLS and ML show similar performance for the frequency estimate, and the CRB is attained for $\text{SNR} \geqslant -5$ dB. For dBESP, a higher SNR is needed to achieve the same performance and the CRB is not attained until the signal is significantly stronger ($\text{SNR} \geqslant 5$ dB). For that value of SNR, the nRMSE's of dIAA, PerLS and ML become 0 due the limitation of the frequency grid.

For the estimation of damping, as is shown in Fig. 2b, dIAA shows a large improvement over PerLS, whereas ML outperforms all the other methods. Moreover, dBESP outperforms dIAA for $\text{SNR} \geqslant 5$ dB. The poor performance of PerLS is mainly due to the fact that a small value of N prevents PerLS from obtaining an accurate estimate, despite an increased SNR. Note that the limited range of dampings we search over will give a lower limit of the nRMSE, yielding biased dIAA estimates that appear lower than the CRB at $\text{SNR} = -10$ dB.

[1] The simulation data is chosen to mimic the signal obtained by examining the explosive RDX, excited at 5.192 MHz, using stochastic NQR spectroscopy [13,30].

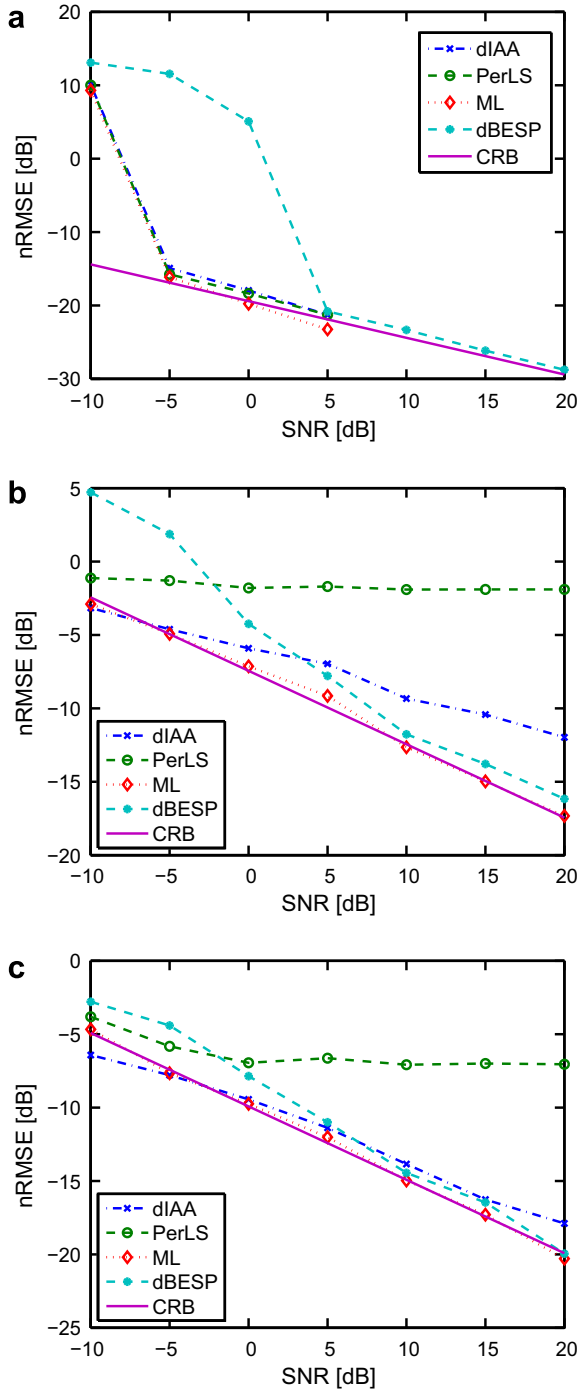

Fig. 2. nRMSE of the estimates of (a) frequency, (b) damping, and (c) amplitude for different SNR's.

In Fig. 2c, we compare the performance of the amplitude estimates. As expected, ML is optimal and attains the CRB for all SNR's considered. For PerLS, the amplitude is estimated using least-squares, using the frequency and damping estimates previously obtained. Again, PerLS shows poor performance, mainly due to the poor damping estimate. For large SNR, dBESP outperforms dIAA, even though the difference is quite small. Intriguingly, for small SNR we find that the nRMSE obtained by dIAA is smaller than the CRB. This is due to the fact that the estimation of the damping is limited, as discussed above, which will yield biased amplitude estimates.

We remark that we expect the difference in performance between dIAA and PerLS to be even larger in a multi-peak scenario, especially for closely spaced peaks where leakage and sidelobe effects may significantly affect the periodogram.

Next, we examine the performance for different dampings, at SNR = 10 dB. The setting is identical to the previous one except for the damping grid that now ranged from 0 to 0.060 in $D = 61$ steps, and that we now used 5 blocks of data, giving a total of $N = 80$ samples.

Fig. 3 displays the result of the nRMSE for the frequency, damping, and amplitude estimates, evaluated over 100 Monte-Carlo simulations. As can be seen, both ML and dBESP are approximately attaining the CRB for all dampings considered at this SNR, except

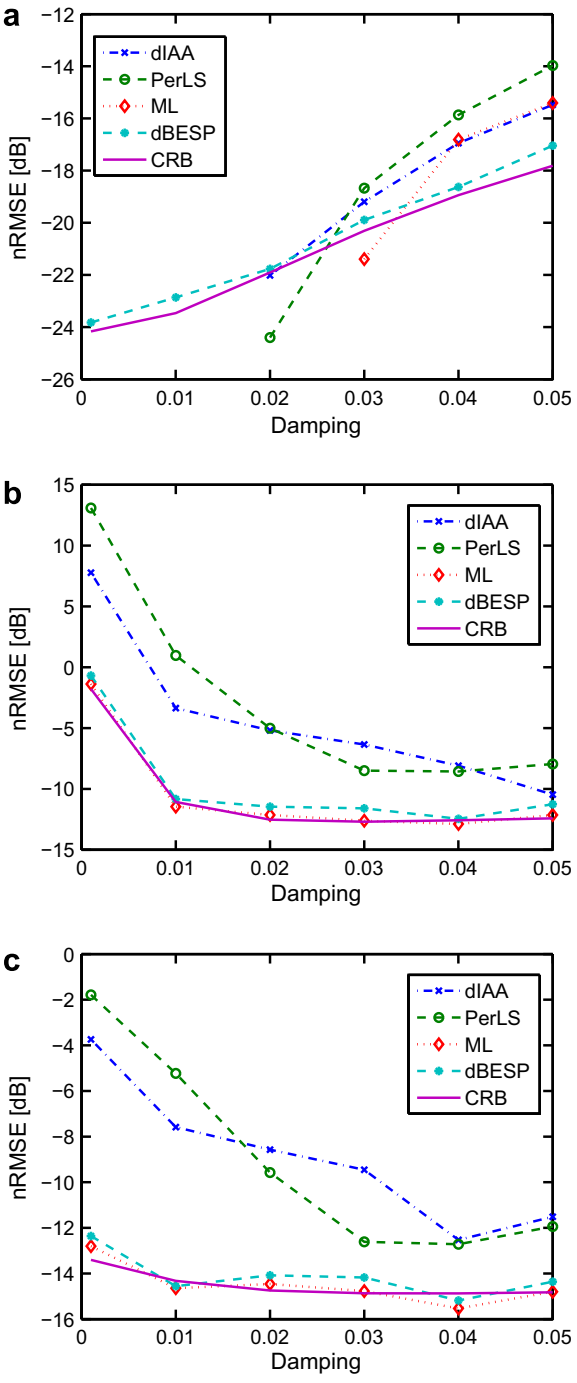

Fig. 3. nRMSE of the estimates of (a) frequency, (b) damping, and (c) amplitude for different dampings at SNR = 10 dB.

for dampings larger than 0.03 Hz, where ML shows a bias in the frequency estimate. The nonparametric methods can only compete with dBESP and ML in the frequency estimation for dampings smaller than 0.03 Hz. Furthermore, for low dampings, the frequency estimates of dIAA, PerLS and ML have nRMSE = 0 for reasons discussed above. For larger damping, there is hardly any information left in the last block of data, causing problems for PerLS, and dIAA yields better estimates than PerLS.

For the estimation of damping, shown in Fig. 3b, we see that the performance is poor for $\beta = 0.001$, which is not that surprising considering that the true parameter is very small; indeed when estimating a parameter with the true value of zero, one would expect the relative error to be infinitely high as long as the variance is larger than zero. PerLS decreases its nRMSE for increasing damping; however, this increase in performance is artificial. The damping estimates from PerLS are somewhat constant, no matter the true damping. As the true damping is increasing, the estimation error decreases and PerLS shows an improvement in nRMSE performance. For large dampings, the lack of information in the last block of data causes the performance to degrade. Again, ML and dBESP are close to attaining the CRB, showing the best performance.

In Fig. 3c, the amplitude estimates are compared, and dBESP shows a good performance attaining the CRB for all dampings considered. The nonparametric methods show similar performance; however, they cannot compete with dBESP. For dampings equal to 0.03 and 0.04 Hz, PerLS shows a better result than dIAA, whereas it is the other way around for the other dampings considered.

As should be expected as the simulation data well match the assumed signal model, the results from the simulations show that dBESP and ML outperform the nonparametric methods; ML shows

Table 2
Table of estimates for the AN data.

Method	$\hat{f}$ (kHz)	$\hat{\beta}$ (kHz)	$\hat{\alpha}$ ($\times 10^5$)
dIAA	496.7	0.32	$3.9 - 4.5i$
PerLS	496.7	0.40	$3.8 - 4.3i$
dBESP	496.8	0.19	$3.6 - 0.8i$
ML	496.7	0.17	$3.0 - 3.5i$

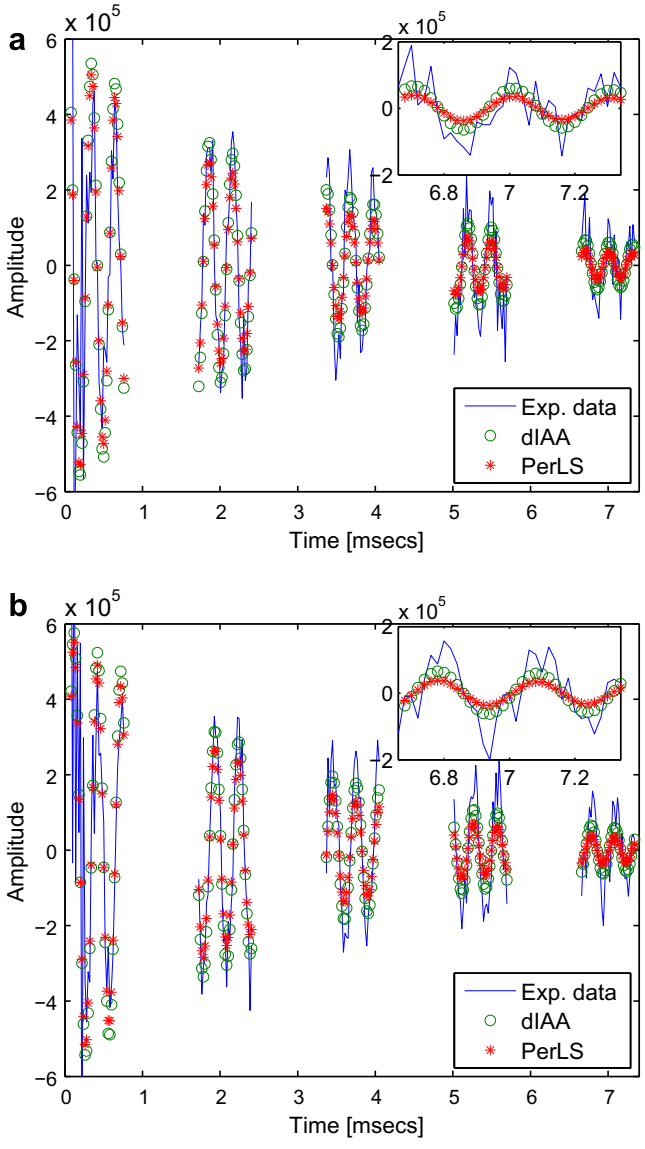

Fig. 4. (a) Real part and (b) imaginary part of the FID from the AN data together with data generated from the estimated parameters using dIAA and PerLS. The insets show an enlargement of the last block of data.

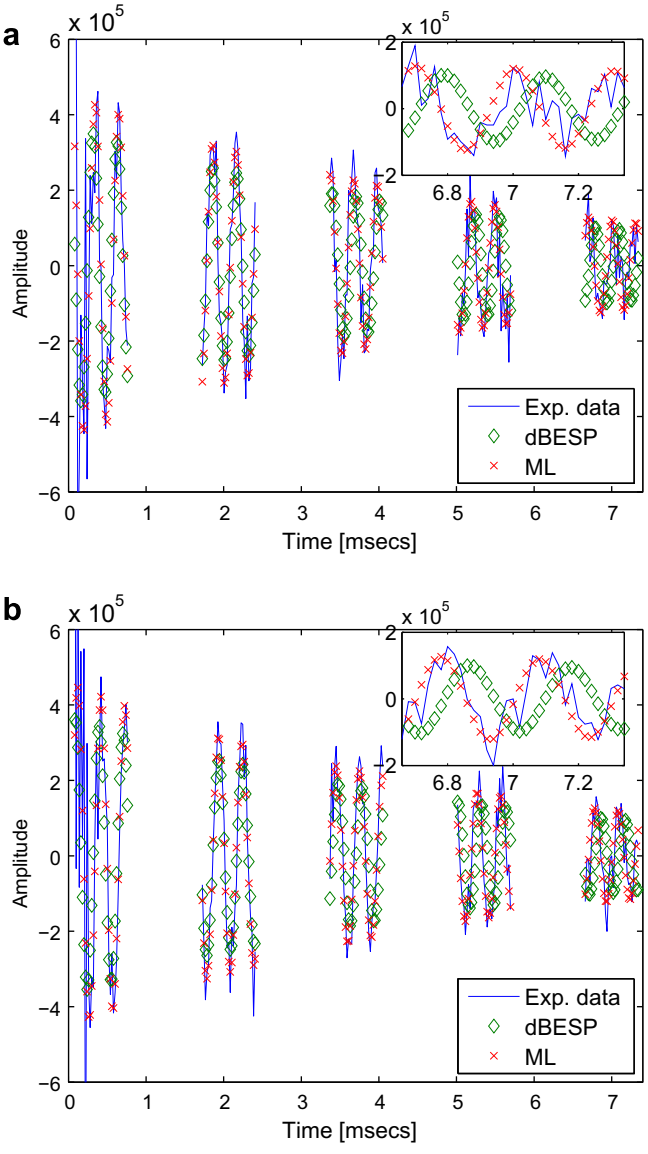

Fig. 5. (a) Real part and (b) imaginary part of the FID from the AN data together with data generated from the estimated parameters using dBESP and ML. An enlarged version of the last block is shown in the insets.

excellent performance for all SNR's, and for SNR larger than about 5 dB, dBESP also works well. We note that the amplitude estimate's relative errors using PerLS and dIAA decrease for increasing damping, up to $\beta = 0.04$ Hz. After that the lack of information in the signal causes a degradation in performance. Furthermore, dIAA outperforms PerLS for smaller dampings, whereas for larger dampings the situation appears to be the opposite. This is due to the fact that for larger dampings, the damping estimate from PerLS is improved for reasons discussed above; therefore, due to the improved damping estimate, the amplitude estimate is also improved. However, for larger dampings, the frequency estimate degrades, which then also affects the amplitude estimate.

To summarize this part: We have seen that the model-based methods have shown excellent performance in the estimation of the parameters, whereas dIAA has shown to be a robust nonparametric method, performing well also for lower SNR. PerLS has shown excellent frequency estimation performance, but greatly suffering from the few number of available samples when estimating the linewidths.

4.2. Experimental data

We proceed to evaluate the performance of the algorithms using measured sNQR data. Two samples were used for the tests. One was 98 g of ammonium nitrate (AN) fertilizer prills and the other was 38 g of PE4, a plastic explosive which contains about 88% RDX as the active component. Signals were obtained from the AN sample alone and from the AN and PE4 samples together. Stochastic excitation was performed with a 511 element MLBS RF pulse train and after each pulse 64 data points were acquired with a 20 μs dwell time. However, to better illustrate the performance of the discussed estimators, we will here only examine 35 out of the collected AN samples (the first four and last 25 samples were removed) and 48 samples of the combination of AN and PE4 (the last 16 samples were removed). The reason for using the so-obtained truncated data sets is to better emphasize the performance difference between the methods; clearly, the performance difference will depend on the effective data lengths, with larger data sets providing more accurate estimates, and thus also reducing the performance difference between the estimates. For the AN sample the stochastic dwell time was 1645 μs, the RF pulse widths were 20 μs with excitation at 500 kHz and 160 mW peak-to-peak RF power. For the AN and PE4 sample together, the stochastic dwell time was 1580 μs, the RF pulse widths were 80 μs with excitation at 499 kHz and 240 mW peak-to-peak RF power. In both cases, the Q of the NQR detection probe was 80–90. The spectra were generated using a frequency grid of $K = 1000$ points, and for dIAA, the damping grid ranged from 0 to 2 kHz in 81 steps. For the AN data, the estimates of the frequency, damping and complex amplitude for the different methods are displayed in Table 2. The frequency estimates are very similar, but the damping and amplitudes differ. As seen from the table, dBESP shows a smaller amplitude compared to the optimal ML method, whereas dIAA and PerLS show a slightly larger damping and amplitude estimate. Figs. 4 and 5 show the cross-correlation domain AN data together with data modeled using the estimated parameters from Table 2. The data modeled with parameters estimated from dBESP appears to become out of phase due to a biased amplitude estimate, whereas

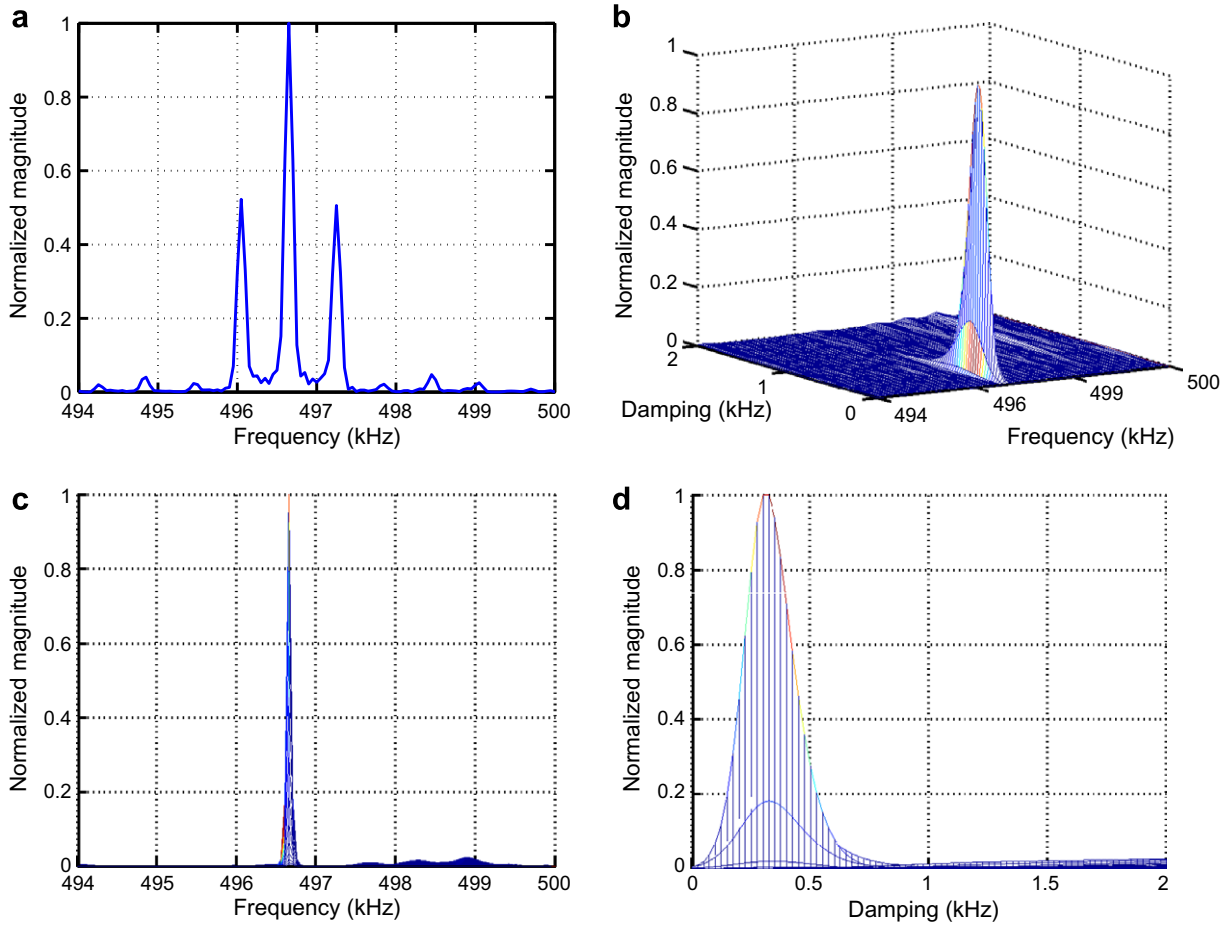

Fig. 6. (a) PerLS and (b) dIAA spectrum of AN. (c) and (d) show the magnitude-frequency and magnitude-damping views, respectively, of the dIAA spectrum in (b). Note the artifacts in (a) introduced by the sampling scheme.

the data from dIAA and PerLS dampens somewhat faster than the experimental data. The ML data models the data well, as expected. From Fig. 6a, we see that some artifacts are introduced in the PerLS spectrum due to the sampling scheme. These artifacts are removed in the dIAA spectrum seen in Fig. 6b and the 497 kHz peak from AN is clearly seen. Fig. 6c and d display the rotated dIAA spectrum in Fig. 6b, rotated such that the frequency and the damping become clearer. We remark that the height of the sampling scheme artifacts depends largely on the size of the block of missing data; the larger the gap, the larger the artifacts.

For the AN and PE4 combined sample data, the estimates of the peaks are displayed in Table 3. The ML estimate is omitted for the reasons discussed above. The AN peak is expected to be at 497 kHz and the RDX peaks at 501 and 503 kHz, although the exact frequencies will depend on the temperature of the sample. Fig. 7 displays the cross-correlation domain data together with data modeled using the estimated parameters from Table 3. As is seen, the data is quite well modeled by the PerLS and dIAA estimates in the first two blocks, but after that the SNR becomes too low to allow for easy interpretation. Fig. 8 shows the spectral estimates obtained using dIAA and PerLS. It is worth noting that the estimates contain several spurious peaks and artifacts resulting from interference signals in the data. Furthermore, the non-uniform sampling also causes artifacts in the PerLS spectrum, as is clearly seen in Fig. 8a. As seen in the table, both dIAA and PerLS offer accurate frequency estimates, well matching the expected frequencies of the RDX and AN peaks. The dBESP estimates, however, seem to suffer from a significant frequency offset for the third peak. Moreover, the estimates of the damping constants are seen to differ significantly between the methods. In an ideal experiment, the two RDX peaks should appear symmetrical. In this experiment, in order to excite all three peaks, a comprise has been made to the excitation frequency, removing the symmetry. Given the earlier simulated signals, we are inclined to trust the dIAA estimates more than the PerLS estimates, indicating that the latter method will underestimate the true signal dampings. However, it is worth noting that simulations with multi-peak data shows that for dIAA, leakage in the damping estimates may also occur, thus possibly causing a bias in the estimate. We also note that for longer data sequences, the performance of the estimates will improve, reducing the influence of the artifacts resulting from the non-uniform sampling. Therefore, it is reasonable to expect that the seen performance gain of dIAA will be more pronounced for shorter measurement sets, with the differences between the estimates diminishing as the data sets grow.

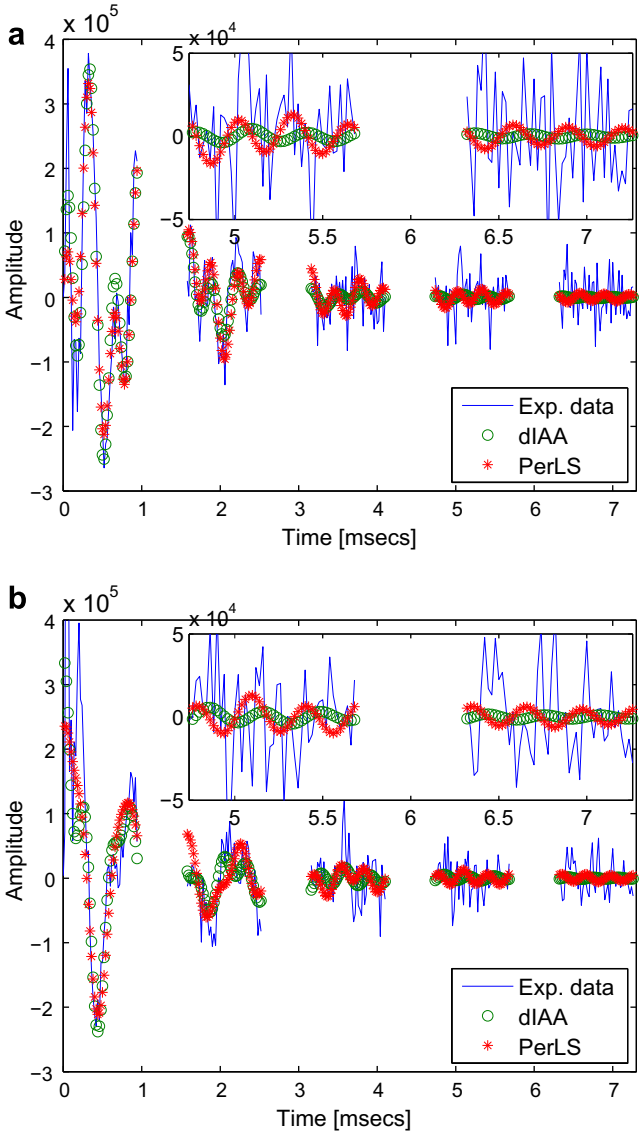

Fig. 7. (a) Real part and (b) imaginary part of the FID from the sample data containing a combination of AN and PE4, together with data generated from the estimates in Table 3. The insets show an enlargement of the two last blocks of data.

5. Conclusions

In this paper, we have derived a nonparametric iterative spectral estimation technique, dIAA, that yields an accurate 2D spectral representation versus frequency and damping, for arbitrarily sampled data. Furthermore, we have proposed a parametric subspace-based method, dBESP, which gives reliable estimates of the parameters of a sum of exponentially decaying sinusoidal components corrupted by additive white noise. The methods have been compared on both simulated and experimental data. In the single-peak case, the studied parametric methods ML and

Table 3
Table of estimates for the AN and PE4 combined sample data.

	dIAA	PerLS	dBESP
$\hat{f}_1$ (kHz)	495.8	495.8	499.0
$\hat{f}_2$ (kHz)	500.4	500.4	507.8
$\hat{f}_3$ (kHz)	502.3	502.2	522.3
$\hat{\beta}_1$ (kHz)	0.60	0.37	0.36
$\hat{\beta}_2$ (kHz)	1.08	0.86	1.43
$\hat{\beta}_3$ (kHz)	1.28	0.96	0.86
$\hat{\alpha}_1$ (a.u.)	$(7.73-2.08i)\times 10^4$	$(6.57-2.73i)\times 10^4$	$(-6.09-5.70i)\times 10^4$
$\hat{\alpha}_2$ (a.u.)	$(-2.12+1.52i)\times 10^5$	$(-2.02+1.71i)\times 10^5$	$(-1.12+1.37i)\times 10^4$
$\hat{\alpha}_3$ (a.u.)	$(1.02+2.01i)\times 10^5$	$(1.02+0.82i)\times 10^5$	$(-1.31+0.27i)\times 10^4$

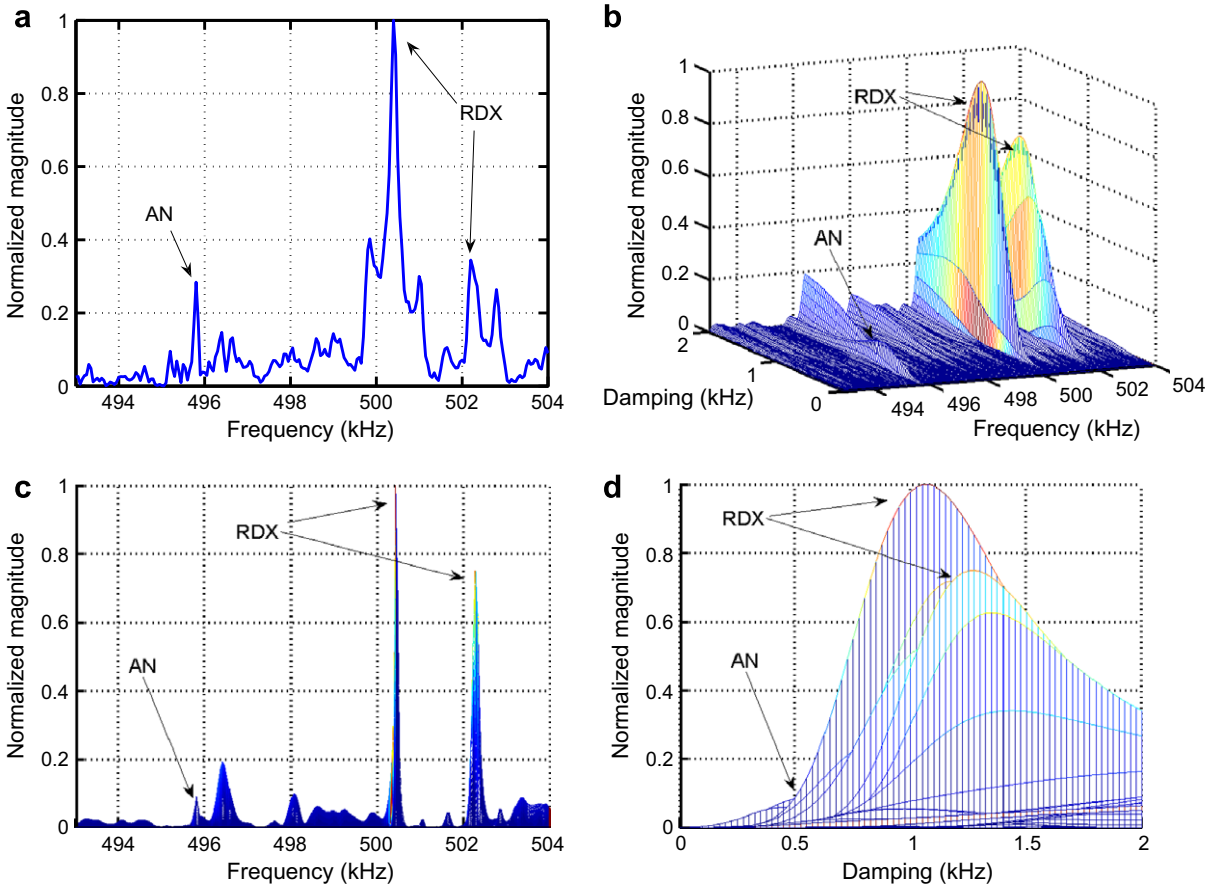

Fig. 8. (a) PerLS and (b) dIAA spectrum of a sample containing both AN and RDX. (c) and (d) show the magnitude-frequency and magnitude-damping views, respectively, of the dIAA spectrum in (b). Note that there are spurious peaks in both spectra but they are particularly strong for PerLS.

dBESP were shown, as expected, to outperform the nonparametric methods. However, we note that dBESP is limited to data that are sampled in blocks of regularly spaced samples. The simulations show that among the two nonparametric methods compared in this paper, dIAA outperforms the least-squares periodogram, PerLS, and that the sampling scheme can create spurious peaks in the PerLS spectrum but not in the dIAA spectrum. For larger data sizes, however, the differences between PerLS and dIAA are smaller and the artifacts caused by the sampling decrease, especially when the gap between two consecutive blocks diminishes.

When the methods were applied to experimental data, dIAA produced spectra where the peaks were clearly visible, whereas PerLS contained spurious peaks or artifacts. The dIAA spectrum for a mixed sample containing both AN and RDX also showed some spurious peaks that were likely due to interferences in the data as simulations for this type of data, which in the interest of brevity are not presented in this paper, showed no such spurious peaks. Furthermore, these simulations also showed that some leakage in damping might occur also for dIAA, possibly causing a small bias in the damping estimates.

Finally, we remark on the fact that issues such as line splitting and spurious peaks, that may affect parametric methods such as LP, have not been observed for dIAA. Generally, all nonparametric methods will exhibit some fluctuations in the spectral estimate due to the variance in the estimate, as well as peaks from interference components and artifacts in the data, but this is as expected; the method is only yielding an accurate depiction of the spectral content of the observed data.

Acknowledgments

Some of the computations were performed using UPPMAX resources under Project p2007002. The authors would like to thank Mr. P. Babu for insightful discussions.

References

[1] P. Stoica, R. Moses, Spectral Analysis of Signals, Prentice Hall, Upper Saddle River, NJ, 2005.
[2] Y. Wang, J. Li, P. Stoica, Spectral Analysis of Signals — The Missing Data Case, Morgan & Claypool, 2005.
[3] F.A. Marvasti (Ed.), Nonuniform Sampling: Theory and Practice, Kluwer Academic/Plenum Publishers, 2001.
[4] Y. Hua, T.K. Sarkar, Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise, IEEE Trans. Acoust. Speech Signal Process. 38 (5) (1990) 814–824, doi:10.1109/29.56027.
[5] P. Stoica, T. Sundin, Nonparametric NMR spectroscopy, J. Magn. Reson. 152 (2001) 57–69, doi:10.1006/jmre.2001.2377.
[6] I. Marshall, J. Higinbotham, S. Bruce, A. Freise, Use of Voigt lineshape for quantification of in vivo 1H spectra, Magn. Reson. Med. 37 (5) (1997) 651–657, doi:10.1002/mrm.1910370504.
[7] L. Vanhamme, T. Sundin, P. van Hecke, S. van Huffel, MR spectroscopy quantitation: a review of time-domain methods, NMR Biomed. 14 (4) (2001) 233–246, doi:10.1002/nbm.695.
[8] S. Mierisová, M. Ala-Korpela, MR spectroscopy quantitation: a review of frequency domain methods, NMR Biomed. 14 (4) (2001) 247–259, doi:10.1002/nbm.697.
[9] D.-K. Yang, J.E. Atkins, C.C. Lester, D.B. Zax, New developments in nuclear magnetic resonance using noise spectroscopy, Mol. Phys. 95 (1998) 747–757.
[10] D.-K. Yang, D.B. Zax, Bandwidth extension in noise spectroscopy, J. Magn. Reson. 135 (1998) 267–270.
[11] R.H. Pursley, J. Kakareka, G. Salem, N. Devasahayam, S. Subramanian, R.G. Tschudin, M.C. Krishna, T.J. Pohida, Stochastic excitation and Hadamard correlation spectroscopy with bandwidth extension in RF FT-EPR, J. Magn. Reson. 162 (2003) 35–45.

[12] M.-Y. Liao, D.B. Zax, Analysis of signal-to-noise ratios for noise excitation of quadrupolar nuclear spins in zero field, J. Phys. Chem. 100 (1996) 1483–1487.
[13] S.D. Somasundaram, A. Jakobsson, M.D. Rowe, J.A.S. Smith, N.R. Butt, K. Althoefer, Robust detection of stochastic nuclear quadrupole resonance signals, IEEE Trans. Signal Process. 56 (9) (2008) 4221–4229.
[14] P. Stoica, E.G. Larsson, J. Li, Adaptive filterbank approach to restoration and spectral analysis of gapped data, Astron. J. 120 (2000) 2163–2173.
[15] Y. Wang, P. Stoica, J. Li, T.L. Marzetta, Nonparametric spectral analysis with missing data via the EM algorithm, Dig. Signal Process. 15 (2) (2005) 191–206, doi:10.1016/j.dsp.2004.10.004.
[16] H. Gesmar, J.J. Led, Two-dimensional linear-prediction NMR spectroscopy, J. Magn. Reson. 83 (1) (1989) 53–64, doi:10.1016/0022-2364(89)90291-6.
[17] V.A. Mandelshtam, H.S. Taylor, A.J. Shaka, Application of the filter diagonalization method to one- and two-dimensional NMR spectra, J. Magn. Reson. 133 (2) (1998) 304–312, doi:10.1006/jmre.1998.1476.
[18] R. Bro, PARAFAC. Tutorial and applications, Chemomet. Intel. Lab. Syst. 38 (2) (1997) 149–171, doi:10.1016/S0169-7439(97)00032-4.
[19] J.C. Hoch, A.S. Stern, NMR Data Processing, Wiley, New York, 1996.
[20] A.S. Stern, K.-B. Li, J.C. Hoch, Modern spectrum analysis in multidimensional NMR spectroscopy: comparison of linear-prediction extrapolation and maximum-entropy reconstruction, J. Am. Chem. Soc. 124 (9) (2002) 1982–1993, doi:10.1021/ja011669o.
[21] S.G. Hyberts, G.J. Heffron, N.G. Tarragona, K. Solanky, K.A. Edmonds, H. Luithardt, J. Fejzo, M. Chorev, H. Aktas, K. Colson, K.H. Falchuk, J.A. Halperin, G. Wagner, Ultrahigh-resolution 1H–13C HSQC spectra of metabolite mixtures using nonlinear sampling and forward maximum entropy reconstruction, J. Am. Chem. Soc. 129 (16) (2007) 5108–5116, doi:10.1021/ja068541x.
[22] A.S. Stern, D.L. Donoho, J.C. Hoch, NMR data processing using iterative thresholding and minimum l1-norm reconstruction, J. Magn. Reson. 188 (2) (2007) 295–300, doi:10.1016/j.jmr.2007.07.008.
[23] E. Gudmundson, A. Jakobsson, S.D. Somasundaram, On the reconstruction of gapped sinusoidal data, in: 2008 IEEE International Conference on Acoustics, Speech and Signal Processing, Las Vegas, Nevada, 2008.
[24] T. Yardibi, J. Li, P. Stoica, M. Xue, A.B. Baggeroer, Source localization and sensing: a nonparametric iterative approach based on weighted least squares, IEEE Trans. Aerospace Electron. Syst., in press.
[25] W. Roberts, P. Stoica, J. Li, T. Yardibi, F.A. Sadjadi, Iterative adaptive approaches to MIMO radar imaging, IEEE J. Select. Topics Signal Process., in press.
[26] P. Stoica, J. Li, H. He, Spectral analysis of nonuniformly sampled data: a new approach versus the periodogram, IEEE Trans. Signal Process. 57 (3) (2009) 843–858, doi:10.1109/TSP.2008.2008973.
[27] N.R. Butt, A. Jakobsson, Coherence spectrum estimation from non-uniformly sampled sequences, IEEE Signal Process. Lett., in press.
[28] P. Stoica, J. Li, J. Ling, Missing data recovery via a nonparametric iterative adaptive approach, IEEE Signal Process. Lett. 16 (4) (2009) 241–244, doi:10.1109/LSP.2009.2014114.
[29] E. Gudmundson, A. Jakobsson, Efficient algorithms for computing the Capon and APES filters, in: 41st Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA, 2007.
[30] J.A.S. Smith, M.D. Rowe, R.M. Deas, M.J. Gaskell, Nuclear quadrupole resonance detection of landmines, in: International Conference on Requirements and Technologies for the Detection, Removal and Neutralization of Landmines and UXO, vol. 2, Vrije Universiteit Brussel, Brussel, Belgium, 2003, pp. 715–721. ISBN 9080826111.
[31] P. Stoica, N. Sandgren, Spectral analysis of irregularly-sampled data: paralleling the regularly-sampled data approaches, Dig. Signal Process. 16 (6) (2006) 712–734, doi:10.1016/j.dsp.2006.08.012.
[32] Y. Yao, S.P. Pandit, Cramér-Rao lower bound for a damped sinusoidal process, IEEE Trans. Signal Process. 43 (4) (1995) 878–885.