ACADEMIC PRESS

Available online at www.sciencedirect.com

SCIENCE DIRECT®

Journal of Magnetic Resonance 163 (2003) 248–256

JMR Journal of Magnetic Resonance

www.elsevier.com/locate/jmr

Rapid construction of solid-state magnetic resonance powder spectra from frequencies and amplitudes as applied to ESEEM

Stefan Stoll and Arthur Schweiger*

Physical Chemistry Laboratory, ETH Zürich, CH-8093 Zürich, Switzerland

Received 27 December 2002; revised 6 March 2003

Abstract

In many Fourier-transform spectroscopies, such as pulse magnetic resonance (NMR, EPR), time-domain signals are acquired. Parameters are extracted from these signals by fitting numerical simulations to the experimental data. At present, simulations are often performed in frequency domain (FD). These computations generate a list of frequencies and amplitudes associated with the complex exponential components evolving during one or several variable time intervals. In order to compare simulations with experiments, this peak list is converted to a finite-length time-domain (TD) signal. This can be achieved either by directly evoluting the exponentials in time (direct method) or by rounding their frequencies and binning their amplitudes into a frequency-domain array (histogram method). The first approach is equivalent to a brute-force TD simulation and is slow for a large number of peaks. The second approach is a fast, but very crude approximation and is usually applied without considering in detail the errors involved. A third method introduced and illustrated here is based on the convolution and deconvolution of a short finite impulse response filter kernel. This convolution approach is much faster than the direct method and by orders of magnitude more accurate than the histogram method. For both TD and FD signals a detailed analysis of the errors and of the associated computational costs is presented. The convolution approach is applicable to any simulation problem where TD signals consist of a large number of complex exponentials. In particular, it is the method of choice for simulating 1D and 2D electron spin echo envelope modulation (ESEEM) spectra of disordered systems.

1. Introduction

Solid-state pulse EPR (electron paramagnetic resonance) and ESEEM (electron spin echo envelope modulation) in particular are established tools in magnetic resonance spectroscopy [1–3]. ESEEM allows the determination of the geometry and the electron spin density distribution in paramagnetic species such as transition metal complexes, radicals, and defect centers.

In ESEEM, as in many other Fourier-transform (FT) spectroscopic methods like NMR or ion cyclotron resonance, signals are acquired in the time domain (TD) and contain oscillating components in the form of complex exponentials. In the case of ESEEM, these exponentials arise from the time evolution of electron and nuclear coherences. For analysis, the signals are usually converted to the frequency domain (FD) by applying a discrete Fourier transform (DFT).

The information content of the TD and the FD signal is identical, but its representation in the two domains emphasizes different aspects. Whereas in TD amplitudes and phases can easily be identified, the FD representation exposes amplitudes and frequencies. A detailed analysis of these signals is rewarding both in TD and in FD. The TD signal thus gives easy access to the modulation depth, which apart from other information can yield the number of nuclei giving rise to the modulation. Dividing the TD signals from two different two-pulse ESEEM experiments allows the elimination of components common to both time traces. In FD, peak positions give access to magnetic parameters such as hyperfine and nuclear quadrupole couplings. For the interpretation of the spectra analytical formulas are available only for special and simple cases [1]. In general, however,

* Corresponding author. Fax: +41-1-632-1538.
E-mail address: schweiger@phys.chem.ethz.ch (A. Schweiger).

1090-7807/$ - see front matter © 2003 Elsevier Science (USA). All rights reserved.
doi:10.1016/S1090-7807(03)00120-4

investigated systems are so complex that numerical simulations are indispensable for reliably extracting magnetic parameters and structural information.

Since the two forms of the signal are interconvertible, computations are possible in both TD and FD. Much of the existing work relies on TD simulations [4–6]. Recently FD simulations have become more widespread, not only in ESEEM [7,8], but also in NMR [9–11]. FD methods have the potential of being significantly more efficient than TD methods [8,12]. Especially in ESEEM of disordered systems FD methods deserve special attention, since such simulations are particularly demanding due to the large number of spectra contributing to the overall powder spectrum.

In FD simulations the spectrum is computed in two steps. First the underlying quantum-mechanical equations are evaluated to obtain a list of peaks, each characterized by a frequency and an amplitude. Then the peak list is used to construct the actual ESEEM spectrum. In the case of a very large number of peaks ($>10^3$), this second step is crucial for the computational performance of the FD simulation. Existing methods are either too slow or too inaccurate. Using the TD evolution method [4–6], almost the whole time is spent on constructing the spectrum (see Table 1). The histogram method [7,8] introduces considerable errors (see example at the end of this contribution), as will be shown in Section 5.

In this work we first introduce the peak list (Section 2) associated with a TD signal and examine ways to construct the spectrum from it accurately (Section 3). Next approximate methods are introduced (Section 4) and the problems of the already known approximate histogram method are analyzed (Section 5). In Section 6, a new approach operating in FD superior in accuracy and speed is presented. Finally, the errors involved are examined, and the performance of the new approach is illustrated with an example from ESEEM spectroscopy.

2. TD signals and the peak list

The ESEEM TD signal of a disordered system is the sum of ESEEM TD signals of orientationally distributed paramagnetic species

$$s_{\mathrm{pow}}(t) = \sum_{q=1}^{Q} s(t, \mathbf{\Omega}_q), \tag{1}$$

where t is an interpulse time delay or a pulse length being incremented or decremented during the experiment. The quantity $\mathbf{\Omega}_q$ is specified by three Euler angles $\mathbf{\Omega} = (\alpha, \beta, \gamma)$ and describes the orientation of the molecule in the laboratory frame. The DFT of $s_{\mathrm{pow}}(t)$ gives the ESEEM powder spectrum.

In a powder sample, at least 10^{10} paramagnetic species with different orientations are present. The resulting spectrum can usually be modelled by a sum over less than 10^5 species. In simulations of experiments where a parameter in the spectral domain is varied, as in cw EPR (magnetic field) or ENDOR (radio frequency), analytical projection techniques [13–15] can reduce this number by a factor of 10–100. In ESEEM, however, spectra are obtained indirectly via the DFT of a TD signal, and projection methods are not applicable.

A 1D ESEEM time trace from a single orientation is the sum of P exponentials with frequencies ν_p, complex amplitudes A_p, and decay constants λ_p:

$$s(t, \mathbf{\Omega}_q) = \sum_{p=1}^{P} A_p(\mathbf{\Omega}_q) \exp\left[\mathrm{i}2\pi\nu_p(\mathbf{\Omega}_q)t\right] \exp(-\lambda_p t). \tag{2}$$

In FD, each exponential corresponds to a Lorentzian peak at frequency ν_p, and the sum gives the entire spectrum for the particular orientation $\mathbf{\Omega}_q$.

2D experiments can be described by the same formula by substituting t, ν_p, and λ_p by the vectors $\boldsymbol{t} = (t_1, t_2)$, $\boldsymbol{\nu}_p = (\nu_{1p}, \nu_{2p})$, and $\boldsymbol{\lambda} = (\lambda_{1p}, \lambda_{2p})$.

In ESEEM, like in most other solid-state magnetic resonance spectra of disordered solids, line shapes are determined by the orientation distribution. The decay of a TD signal is then not due to the decays of the single components, but rather to the rapid dephasing of the many exponentials contributing to the signal. The decay of the TD signal due to this inhomogeneous broadening is usually much faster than the relaxational decay described by λ_p; we therefore neglect λ_p:

$$s(t, \mathbf{\Omega}_q) = \sum_{p=1}^{P} A_p(\mathbf{\Omega}_q) \exp\left[\mathrm{i}2\pi\nu_p(\mathbf{\Omega}_q)t\right]. \tag{3}$$

Table 1
Computation times of a typical 2D HYSCORE spectrum for different spectrum construction methods

Method	Computation time (s)		
	Peak list generation	Spectrum construction	Total
TD evolution	13.27	2570.0	2583.3
Histogram	13.27	0.74	14.01
Convolution	13.27	3.05	16.32

Spin system: $S = 1/2$, $g = 2$, one ^{14}N nucleus ($I = 1$, $g_{\mathrm{n}} = 0.4038$), $a_{\mathrm{iso}} = 5\,\mathrm{MHz}$, $e^2qQ/h = 2.4\,\mathrm{MHz}$, $\eta = 0.5$. Pulse sequence: $\pi/2 - \tau - \pi/2 - t_1 - \pi - t_2 - \pi/2 - \tau -$ echo. Experimental parameters: $\tau = 136\,\mathrm{ns}$, pulse lengths 10/10/20/10 ns, $B_0 = 350\,\mathrm{mT}$, $\nu = 9.797369\,\mathrm{GHz}$. 256×256 points, $\Delta t = 50\,\mathrm{ns}$ in both dimensions. 4186 orientations. Computed on a standard Linux PC (866 MHz Pentium III, 128 MB RAM).

If homogeneous broadening is not negligible, it can be included by multiplying the final cumulative TD signal with a single exponential decay $\exp(-\lambda t)$.

For most ESEEM experiments, the peak positions ν_p depend only on the first two Euler angles (α, β). The third angle γ has only an impact for decoupling and nutation experiments, where the measured coherences evolve under microwave irradiation. Peak amplitudes A_p always depend on all three Euler angles.

Both ν_p and A_p can be computed from the usual quantum-mechanical equations describing the dynamics of spin ensembles. The frequencies ν_p are differences of eigenfrequencies of the propagation Hamiltonians, whereas the amplitudes A_p are products of elements from the representation matrices of pulse operators and detection operators [1]. In these computations, there is much room to improve efficiency [12], but this is beyond the scope of this contribution.

The "peak list" of (ν_p, A_p) values obtained from such computations already fully represents the spectrum and the TD signal.

3. Exact signals in time and frequency domain

Since the overall TD signal is a linear combination of exponentials, and the FT is a linear operation, it is sufficient to consider only a single exponential with frequency ν and amplitude 1. Here we restrict ourselves to the 1D situation, where ν is a scalar. Extension to 2D is straightforward, but would unnecessarily complicate the notation.

The TD exponential with frequency ν is given by

$$s(t) = \mathrm{e}^{\mathrm{i}2\pi\nu t}. \tag{4}$$

In an experiment, this signal is sampled at N points at times $n\Delta t$ with $0 \leqslant n \leqslant N-1$ and the dwell time Δt, giving

$$s[n] = s(n\Delta t) = \mathrm{e}^{\mathrm{i}2\pi\nu n\Delta t}. \tag{5}$$

The TD signal can be constructed by explicitly evaluating this formula for all n. However, it is faster to compute $s[n]$ by $N-1$ consecutive multiplications

$$s[n] = s[n-1]\mathrm{e}^{\mathrm{i}2\pi\nu\Delta t} \quad 1 \leqslant n \leqslant N-1 \tag{6}$$

with $s[0] = 1$.

The FT of the TD signal in Eq. (5) defined by

$$S(f) = \frac{1}{N}\sum_{n=0}^{N-1} s[n]\mathrm{e}^{-\mathrm{i}2\pi f n\Delta t} \tag{7}$$

is a periodic sinc function with a complex phase factor [16]

$$S(f) = \frac{1}{N}\mathrm{e}^{-\mathrm{i}\pi(f-\nu)(N-1)\Delta t}\frac{\sin[\pi(f-\nu)N\Delta t]}{\sin[\pi(f-\nu)\Delta t]}. \tag{8}$$

It can be considered as a back-folded aperiodic sinc function. When this function is sampled at the frequencies $m\Delta f$ with the increment $\Delta f = (N\Delta t)^{-1}$ and $0 \leqslant m \leqslant N-1$, the DFT of the original signal

$$S[m] = S(m\Delta f) = \frac{1}{N}\mathrm{e}^{-\mathrm{i}\pi(m-\kappa)\frac{N-1}{N}}\frac{\sin[\pi(m-\kappa)]}{\sin[\pi(m-\kappa)/N]} \tag{9}$$

is obtained, with the scaled frequency $\kappa = \nu/\Delta f$. Though feasible, evaluation of this expression for the impulse response is impractical due to its complexity compared to the TD signal.

The direct method (computation of the exponential in TD) needs one exponential computation and $N-1$ multiplications/peak. In the 2D case of an $N \times N$ signal, $(N-1)^2$ multiplications/peak are needed. The performance of Eq. (9) is even worse. The drawback of both methods lies in their slowness when many peaks have to be accumulated to construct a spectrum, as in Eq. (1).

4. Frequency-domain approximations

TD and FD signals of a single peak with amplitude 1 differ considerably in their character. The TD signal is distributed over the entire domain, and its magnitude is always 1. It seems not to be possible to devise any method that allows one to reconstruct the TD signal with less computational cost than the direct evaluation of the exponential according to Eq. (6). On the other hand, most of the FD signal (Eq. (9)) is concentrated around the frequency ν of the exponential. This compactness of the FD signal makes it possible to save computation time by computing only the important central part around ν and neglecting the small wings. Such approximate methods sacrifice accuracy for the sake of computational efficiency. So far, there is only one FD approximation method used in the literature [8]. We call it the "histogram method" and discuss it in the next section.

To assess the quality of an approximation, we define error measures for both TD and FD. An appropriate error function in TD is

$$\mu_{\mathrm{TD}}(\nu) = \max_{0 \leqslant n \leqslant N-1} |s^{(\nu)}_{\mathrm{approx}}[n] - s^{(\nu)}[n]|, \tag{10}$$

which describes the maximum of the magnitude difference between the approximate and the exact TD signal of a single exponential with amplitude 1 and frequency ν. We use the maximum instead of the sum of squares, because it correlates better with the visual difference of the two spectra. For FD the error measure corresponding to Eq. (10) is given by

$$\mu_{\mathrm{FD}}(\nu) = \max_{0 \leqslant m \leqslant N-1} |S^{(\nu)}_{\mathrm{approx}}[m] - S^{(\nu)}[m]|. \tag{11}$$

The average of μ over all possible frequencies ν is a frequency-independent overall error measure. It turns

out that for all cases studied μ in the above formulas is periodic in v with period Δv, so that the average can be limited to one period

$$\epsilon_{TD} = N\Delta t \int_0^{N\Delta t^{-1}} \mu_{TD}(v)\mathrm{d}v. \quad (12)$$

In FD, the corresponding error function

$$\epsilon_{FD} = N\Delta t \int_0^{N\Delta t^{-1}} \mu_{FD}(v)\mathrm{d}v \quad (13)$$

can be used. In practical computations, the integrals in the error functions have to be approximated by sums. In this work we integrate over 801 test frequencies uniformly distributed between two neighboring DFT frequencies. The error function varies slightly with N. It is minimal in the region $100 < N < 300$; at $N = 1024$ it is about 10% larger. We use a 128-point signal as standard.

The errors are visually noticeable if they are above 0.005. We thus can tolerate errors below a significance limit of $\epsilon_0 = 0.005$.

5. The histogram approximation

The histogram method is to the best of our knowledge the only FD approximation method reported [7,8,11]. This straightforward approach, which has not been studied in detail so far, takes a peak and rounds its frequency v to the nearest DFT frequency, i.e., to the nearest multiple of Δf:

$$k\Delta f = (\kappa - \beta)\Delta f = \lfloor \kappa + 1/2 \rfloor \Delta f = \lfloor v/\Delta f + 1/2 \rfloor \Delta f, \quad (14)$$

where $\lfloor x \rfloor$ indicates the largest integer not bigger than x. The peak is shifted by the frequency offset $\beta = \kappa - k$, which introduces a maximum rounding error of $\Delta f/2$. Then the value of the peak amplitude is added to bin k in the N-point FD vector. The periodic sinc of Eq. (9) is thus approximated by a slightly shifted periodic sinc, so that the discretization of the latter gives a non-zero value only in one bin of the FD vector (see Fig. 1). Finally, an inverse DFT (IDFT) generates the approximate TD signal.

The advantage of this method lies in the fact that for each peak only one rounding and one addition are needed. The speed-up compared to the direct evolution in the TD is enormous. However, the small frequency error introduced by the rounding can cause significant errors in both TD and FD.

In FD, the distortion is obvious (Fig. 1). The amplitude at the center is far from being correct, and the wings on both sides are completely missing. In the figure, the maximum magnitude error μ_{FD} is 1.04. In TD, the phase error increases with increasing n (see Fig. 2). If the rounding error is $\Delta f/2$, the TD error μ_{TD} is 2.

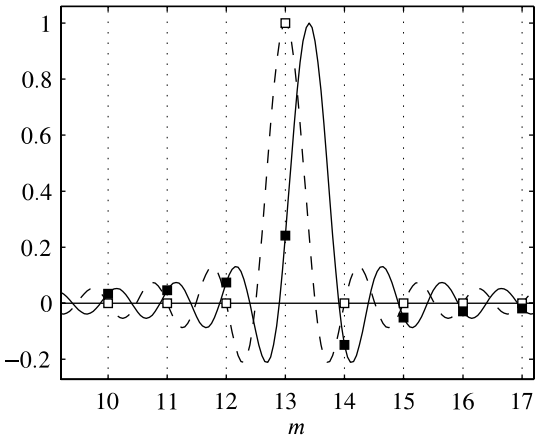

Fig. 1. The histogram method. (solid) Exact periodic sinc with $v = 13.4\Delta f$, (dashed) shifted periodic sinc with $v = 13\Delta f$, (■) correct discretization $S[m]$, and (□) $S_{approx}[m]$ in the histogram approximation. Only the real part is shown.

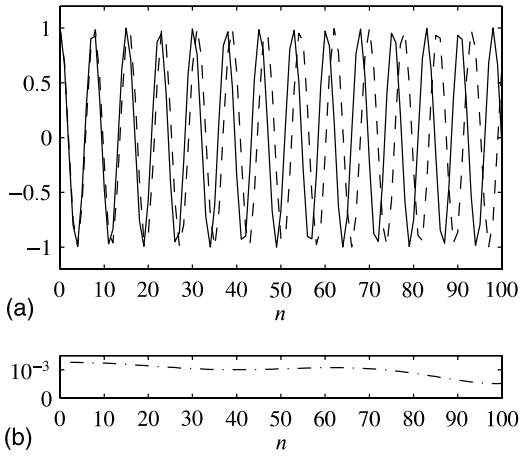

Fig. 2. (a) Real part of the exact TD signal (solid) for $\kappa = 13.4$ and of the histogram approximation (dashed, $k = 13$), $N = 101$ points. (b) Magnitude difference between the exact TD signal and the convolution approximation.

In the case of multiple peaks, the errors can become very disturbing due to interference effects. For example, two exponentials with opposite amplitudes 1 and -1 and frequencies $(k - 1/2)\Delta f$ and $(k + 1/2)\Delta f$ (with integer k) are shifted by the histogram method to $k\Delta f$ and $(k + 1)\Delta f$, respectively, and generate a real spectrum, whereas the correct spectrum should be imaginary with a different shape.

In another unfavorable case, two peaks with amplitudes 1 and -1 and frequencies $(k - \alpha)\Delta f$ and $(k + \alpha)\Delta f$ with $0 \leqslant \alpha < 1/2$ are both rounded to $k\Delta f$ and accumulated into the same bin. As a consequence, they cancel completely, although they should give a single broadened peak around $k\Delta f$. This error is maximum at $\alpha = 0.372$ with $\mu_{FD} = 1.44$.

To remedy these errors, an extended FD vector of length eN ($e > 1$) can be used, where the DFT frequency

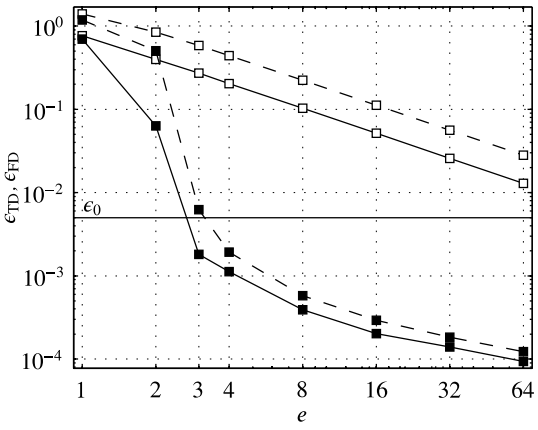

Fig. 3. ϵ_{FD} (solid) and ϵ_{TD} (dashed) for histogram (□) and the convolution method ((■) $M = 2$, α and γ are optimized separately for each e) as a function of the expansion factor e.

spacing is $\Delta f/e$. Increasing the resolution of the frequency axis by a factor of e reduces the maximum rounding error to $\Delta f/(2e)$. The IDFT gives a TD signal of length eN, where the first N points are an approximation of the exact TD signal.

Fig. 3 shows the dependence of the average TD and FD errors on the expansion factor e. With increasing e the error decreases, but obviously $e \geqslant 8$ is needed to reduce at least the mean FD error to an acceptable level. This is at the expense of a larger data array and a correspondingly slower IDFT, which might become prohibitive in 2D. A 256×256 TD signal requires an IDFT of at least a 2048×2048 FD array.

6. The convolution approach

The approximation method introduced in this work is based on convolution and deconvolution of the stick spectrum with a truncated and windowed sinc function (see the schematic illustration in Fig. 4).

Truncating the full kernel function in Eq. (9) to a width of $2M\Delta f$ around its center frequency f speeds up the computation, but introduces truncation errors. With decreasing M, the computation gets faster, but the truncation errors increase. M has thus to be chosen small enough to improve speed and large enough for the error to fall below a certain level of significance. The kernel function itself can also be varied to minimize the error.

The truncated kernel is best designed by apodizing a sinc function with a window function (Fig. 4, step (A)). This is a standard technique for filter design in digital signal processing [16]. The sinc function

$$\operatorname{sinc}(\gamma x) = \frac{\sin(\pi\gamma x)}{\pi\gamma x}, \quad -1 \leqslant x \leqslant 1, \tag{15}$$

contains a width parameter γ. For apodization we use the Kaiser window

$$K(\alpha, x) = \frac{\mathrm{I}_0\left(\alpha\sqrt{1-x^2}\right)}{\mathrm{I}_0(\alpha)}, \quad -1 \leqslant x \leqslant 1, \tag{16}$$

with the zeroth-order modified Bessel function I_0 and the width parameter α. K gives the best results among all classical windows. The resulting truncated kernel

$$R(y) = \operatorname{sinc}\left(\gamma\frac{y}{M}\right)K\left(\alpha, \frac{y}{M}\right), \quad -M \leqslant y \leqslant M, \tag{17}$$

is zero outside the interval $[-M; M]$ and depends on α and γ and on the half-width M. In contrast to Eq. (8), this function is real, and its evaluation is therefore twice as efficient as that of a complex truncated periodic sinc.

For each peak, the truncated kernel $R(y)$ is then sampled into the eN-point spectral vector $\tilde{S}$, centered at the position of the peak (see Fig. 4, step (1))

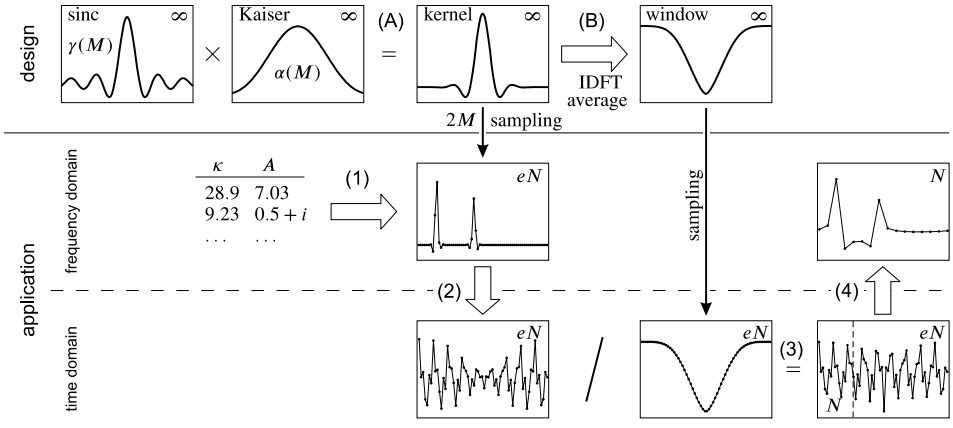

Fig. 4. Design and application of the convolution method. Design steps: (A) kernel as a product of sinc function and Kaiser window, (B) correction window as averaged IDFT. Application steps: (1) continuous convolution of peak list and kernel, eN-point discretization, (2) IDFT, (3) division by discretized correction window in TD (equivalent to circular convolution in FD), and (4) DFT of the first N points.

$$\tilde{S}[k+m] \leftarrow \tilde{S}[k+m] + R(m-\beta) \quad \begin{cases} \beta \neq 0, & -M+1 \leqslant m \leqslant M, \\ \beta = 0, & -M \leqslant m \leqslant M. \end{cases} \tag{18}$$

k and β are defined as in Eq. (14). In the case $\beta \neq 0$, i.e., when the peak frequency lies between two DFT frequencies, $2M$ points have to be evaluated. If the peak frequency exactly coincides with a DFT frequency ($\beta = 0$), $2M+1$ points have to be evaluated (see Fig. 5a). If the peak lies on the edges of $\tilde{S}$, the sampled kernel has to be wrapped around. Similar to the histogram method, an expansion factor e (default value $e = 4$) is used. Its influence on the performance is analyzed later. Mathematically, Eq. (18) corresponds to a continuous circular convolution of the stick spectrum T consisting of delta peaks at the given frequencies with the truncated kernel $R(y)$:

$$\tilde{S} = T \circledast R. \tag{19}$$

The result $\tilde{S}$ is a pseudo-spectrum with peaks at the correct positions, but with distorted line shapes.

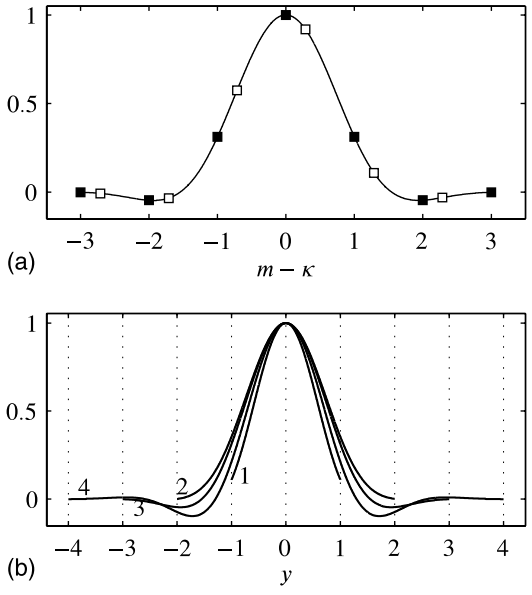

Fig. 5. (a) Sampling a $2M$ wide truncated kernel ($M = 3$). (■) $2M+1$ values for frequency offset 0, (□) $2M$ values for frequency offset $\neq 0$. (b) sinc/Kaiser kernels for different half-widths M.

The distortion can almost completely be removed by deconvoluting the pseudo-spectrum with the truncated kernel

$$S = \tilde{S} \circledast^{-1} R. \tag{20}$$

This circular deconvolution is performed in TD by dividing the IDFT of the pseudo-spectrum by a correction window W obtained by IDFT of the sampled kernel $R(m-\beta)$:

$$W_\beta[n] = \sum_{m=-eN/2+1}^{eN/2} R(m-\beta) \mathrm{e}^{\mathrm{i}2\pi nm/eN} \tag{21}$$

(Fig. 4, steps (2) and (3)). This correction window, however, is different for each peak and depends on its frequency offset β, since each β corresponds to a different set of $2M$ (or $2M+1$) values from the continuous kernel R, which in turn give different window functions in TD. If there is more than one peak in the pseudo-spectrum, the correction window cannot be optimal for all of them. But the mean error ϵ_{TD} can be minimized by using an average of the correction window W_β over all β:

$$\bar{W}[n] = \int_0^1 W_\beta[n] \mathrm{d}\beta \tag{22}$$

(Fig. 4, step (B)). If the continuous kernel is symmetric around its center, the average window $\bar{W}$ is a real function. The first N points of the TD signal obtained after application of $\bar{W}$ are an excellent approximation of the correct TD signal (Fig. 4, step (4)).

The entire procedure described above depends on M, α, and γ. M is the parameter which determines the computational cost. For a given M, α, and γ determine the quality of the approximation. α and γ are simultaneously optimized by a least-squares minimization of the average maximum time-domain deviation ϵ_{TD} (Eq. (12)). The resulting filter kernels for $M = 1$ to 4 are shown in Fig. 5b. Table 2 lists optimal α and γ values for different values of M. Both α and γ strongly depend on M. The minima in the error function are very flat in the region $\alpha_{\mathrm{opt}} \pm 1/2$. On the other hand, the error function is more sensitive to γ.

Numerically, quasi-continuous representations with ca. $5000M$ points of the kernel R and the average window $\bar{W}$ can be pre-computed and re-used (see Fig. 4, design steps). Instead of evaluating R in Eq. (18) anew for each peak, the values can be taken from its pre-computed representation. In this representation, peak

Table 2
Optimal parameters α and γ and the errors ϵ_{TD} and ϵ_{FD} for the convolution method depending on kernel half-width M

M	γ_{opt}	α_{opt}	ϵ_{TD}	ϵ_{FD}
1	1.6695	4.7255	0.05456	0.03368
2	1.4889	10.9444	0.00185	0.00111
3	2.5878	11.3687	0.00094	0.00036
4	3.8593	12.7798	0.00066	0.00018
5	5.2596	13.4354	0.00030	0.00012

frequencies are rounded to the next discrete frequency and are shifted by $\Delta f/10^4$ at most. The error introduced by this small shift does not contribute significantly to the overall error of the convolution method. To use the quasi-continuous correction window representation for the deconvolution of an eN-point signal, it has to be interpolated to eN points to obtain $\bar{W}[n]$.

7. Error analysis

Fig. 6 shows the errors ϵ_{TD} and ϵ_{FD} of the convolution method based on the optimized sinc/Kaiser kernel. As expected, both errors decrease with increasing M. ϵ_{TD} is two to three times larger than ϵ_{FD}. By increasing M from 1 to 3, the error compared to the histogram method drops by almost two orders of magnitude in both TD and FD domain. The kernel with $M = 2$ gives already satisfactory results and is the one of our choice.

Fig. 3 shows the error of the convolution method with the sinc/Kaiser kernel as a function of the expansion factor e. For $e = 1$, the new method is not much better than the histogram approach, but it gains rapidly when increasing e to 4, where both TD and FD error levels are no longer significant.

Fig. 7 shows the dependence of the errors μ_{TD} and μ_{FD} of the sinc/Kaiser kernel on the offset β of the peak from a DFT frequency. The TD error is periodic with β/e and is most sensitive to the offset around $n\beta/e$. This means that the averaged window used is not the best choice for peaks close to DFT frequencies, but it is still the best choice to minimize the overall mean error. The FD error exhibits a more complex behavior.

The error of the convolution method is different in character from the one of the histogram approach, which is a pure phase error caused by the frequency shift (Fig. 2a). In the convolution method only the amplitude

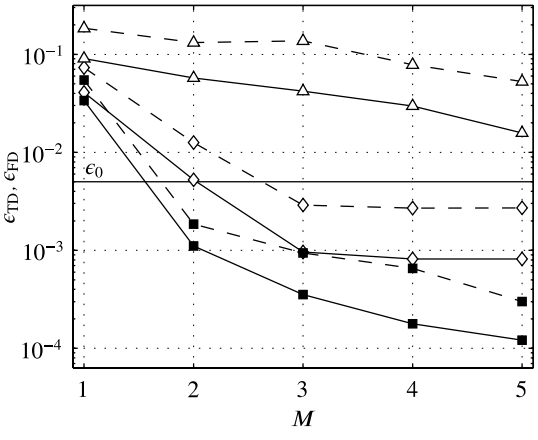

Fig. 6. ϵ_{TD} (dashed) and ϵ_{FD} (solid) for different kernels ($e = 4$) as a function of the kernel half-width M. (■) sinc/Kaiser, (◇) Gaussian, and (△) Lorentzian.

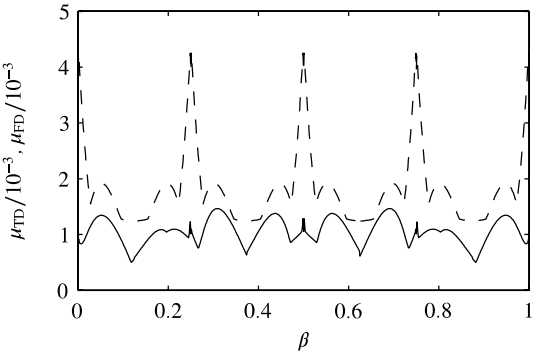

Fig. 7. μ_{FD} (solid) and μ_{TD} (dashed) as a function of the frequency offset β for the sinc/Kaiser kernel with $M = 2$, $e = 4$.

is incorrect and becomes more accurate with increasing time (see the residuals in Fig. 2b). The phase is exact within numerical accuracy.

The origin of the error is not obvious. The fact that an average window is used explains only part of it. For a single peak there would still be an error, if the matching correction window was used. In the aperiodic case, a finite FD signal like the sinc/Kaiser kernel always has an infinite inverse FT. Hence the correct deconvolution window would have infinitely wide wings that asymptotically approach zero in both directions on the time axis. But since the window is only represented in a finite TD, the tails fold back, and the window is aliased. The shorter the kernel in FD, the wider is its associated window in the TD, and the stronger is the effect of backfolding. Thus, it is the deconvolution of the pseudospectrum with an aliased version of the correct window that accounts for most of the error. The impact of this aliasing is already inherently minimized by the least-squares fit of α and γ to the TD error ϵ_{TD}.

The aliasing of the correction window also explains why the convolution method works well only for $e \geqslant 3$ (Fig. 3). For $e < 3$, significant parts of the window lie outside the finite TD region and are folded back.

Often physical lineshape functions, such as Lorentzians or Gaussians, with line widths γ are sampled into the FD, truncating the lineshape function at some distance $k\gamma$ from its center. Although this choice is obvious, it is not the best. We have examined these functions, using the kernels

$$R(y) = \left[1 + \frac{4}{3}\left(\frac{\gamma y}{M}\right)^2\right]^{-1}, \tag{23}$$

$$R(y) = \exp\left[-2\left(\frac{\gamma y}{M}\right)^2\right], \tag{24}$$

with $-M \leqslant y \leqslant M$ and optimizing γ for each given M as in the sinc/Kaiser kernel design. It turned out that even with this γ optimization Gaussians and especially Lorentzians perform significantly worse than the sinc/Kaiser design, as can be seen from Fig. 6.

If in an experimental signal different oscillating components decay at different rates, each peak in the spectrum will have a different linewidth. In this case, the explicit evaluation of a lineshape function for each peak is mandatory, and neither the histogram nor the convolution approximation are applicable.

However, lineshapes in ESEEM spectra of disordered systems are not determined by these relaxational broadenings, but rather by unresolved couplings resulting in an additional inhomogeneous broadening. In all practical cases it is therefore sufficient to convolute the spectrum with a Lorentzian or Gaussian lineshape after the powder spectrum has been constructed from the peak list.

8. Performance analysis

The convolution method is much faster than the direct TD evolution based on Eq. (6). The asymptotic computational costs for all three methods are listed in Table 3. All costs scale linearly with the number of peaks P, but only the TD evolution depends directly on the number of data points N. The dependence of the convolution method's performance on M and M^2 has little impact, since M is small. However, in the two FD methods, the additional cost for the IDFT [17] may be significant.

Fig. 8 illustrates the overall performance of the convolution method based on an analysis of the number of floating-point operations involved. For 1D experiments the convolution method is always faster than the TD evolution, and it performs exceedingly well in 2D experiments. The smallest gain is found in situations with $P < 10^2$, where the cost of the additional IDFT is significant, and in the case of $N < 2^5$, where the small value of N does not penalize the TD evolution too much. However, both situations are quite unusual. For ESEEM simulations, values of $N \approx 2^8$ and $P \approx 10^5$ are common. In this case, 1D peak accumulations are 20 times and 2D accumulations are 1000 times faster. The error is below the significance level ϵ_0 and 200 times smaller than the error in the histogram method.

As an illustrative example we apply the new method to the simulation of a typical HYSCORE spectrum (see Fig. 9). HYSCORE is a common 2D ESEEM experiment [1] and the primary method for determining weak interactions between electron spins and nuclei such as hydrogens and remote nitrogens. The powder spectrum is computed as the sum of single crystal spectra from 4186 orientations. The computation of the resulting list of 489 544 peaks takes 13.27 s on a standard PC (see Table 1). For the accumulation, the convolution method is only four times slower than the histogram method, but 840 times faster than direct TD evolution. In fact, with the latter approach the spectrum accumulation

Table 3
Asymptotic computational costs of TD construction in terms of arithmetic operations for various methods

Method	1D	2D
TD evolution	$O(PN)$	$O(PN^2)$
FD histogram	$O(P)$	$O(P)$
FD convolution	$O(PM)$	$O(PM^2)$

P, number of peaks; N, number of points along one dimension in TD (2D: total $N \times N$ points).

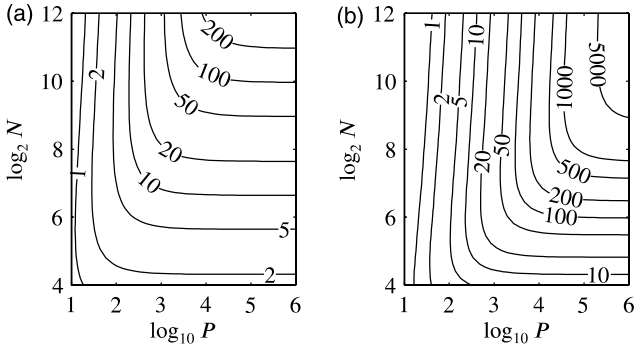

Fig. 8. Speed-up of the convolution method ($M = 2$, $e = 4$) compared to the direct TD evolution. P, number of peaks; N, number of points along one dimension in TD (2D: total $N \times N$ points). (a) 1D and (b) 2D.

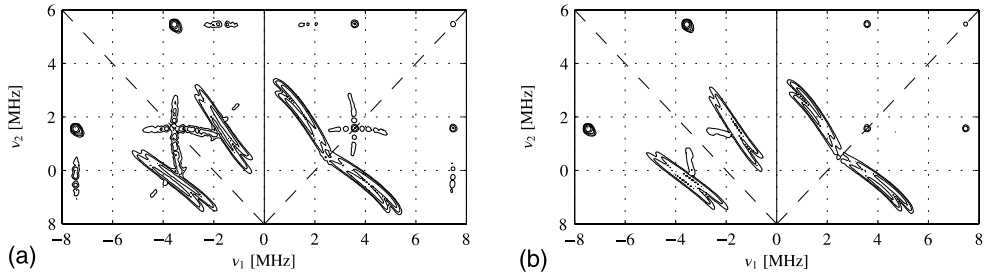

Fig. 9. Simulated ^{14}N HYSCORE spectra. (a) Histogram method and (b) convolution method. There is no visible difference between (b) and a spectrum accumulated with the exact TD evolution method. Spin system and experimental parameters as in Table 1. Logarithmic contour levels between 0.01 and 1 relative to the maximum intensity.

consumes 99.5% of the total computation time. As can be seen in Fig. 9, the histogram method generates significant artifacts with more than 1% relative intensity, which can lead to misinterpretation of the spectrum. These artifacts are absent in the spectrum computed with the convolution method, making it much more reliable than the histogram method.

All gains in computing time discussed so far refer only to the accumulation of peaks into a spectrum. In ESEEM simulations, the spin dynamical part of the simulations, i.e., the compilation of the peak list, can be optimized as well [12]. When the peak list is computed, positions and amplitudes can be interpolated. Unwanted and low-amplitude peaks can be thrown out. As a result, the overall performance of the simulation can be improved by at least another order of magnitude, depending on the nature of the spin system and the experiment. In contrast, TD methods cannot take advantage of these possibilities, because they usually do not compute the list of peaks explicitly.

9. Conclusions

The approximative FD convolution method presented in this work reduces the computational costs for spectrum construction dramatically compared to direct TD evaluations, and it is by orders of magnitude more accurate than the histogram approximation. For fitting experimental data with a low S/N ratio, this gain in accuracy is not always relevant, since the artifacts from the histogram method can be below noise level. Both FD approximation methods can then be accurate for peaks above the noise level. However, when using the histogram method in such cases, care has to exercised to assure that all artifacts are below noise level and all peaks above noise level are undistorted. The convolution method is thus generally preferred.

The approach has been developed for use in numerical simulations in ESEEM spectroscopy, but has not been tailored to specific experiments. It is generally applicable to all spectroscopic methods that acquire signals consisting of a large number of complex exponentials. Like the FD histogram method, it is not applicable if the exponentials have different decay constants.

Slight improvements of the accuracy for a given M might be possible, when a kernel more flexible than the sinc/Kaiser approach is used in the fitting process. Using a least-squares fitted weighted average for the deconvolution window might decrease error figures further, but the gain will be marginal.

Acknowledgments

This research has been supported by the Swiss National Science Foundation. The authors thank Kevin McGill from Stanford University and Marc Jolo from ETH for helpful discussions.

References

[1] A. Schweiger, G. Jeschke, Principles of pulse electron paramagnetic resonance, Oxford University Press, Oxford, 2001.
[2] Y. Deligiannakis, M. Louloudi, N. Hadjiliadis, Electron spin echo envelope modulation (ESEEM) spectroscopy as a tool to investigate the coordination environment of metal centers, Coord. Chem. Rev. 204 (1) (2000) 1–112.
[3] T. Prisner, M. Rohrer, F. MacMillan, Pulsed EPR spectroscopy: biological applications, Annu. Rev. Phys. Chem. 52 (2001) 279–313.
[4] Z. Mádi, S. Van Doorslaer, A. Schweiger, Efficient simulation of ESEEM spectra, J. Magn. Reson. 154 (2) (2002) 181–191.
[5] P. Nicholas, D. Fushman, V. Ruchinsky, D. Cowburn, The virtual NMR spectrometer: a computer program for efficient simulation of NMR experiments involving pulsed field gradients, J. Magn. Reson. 145 (2) (2000) 262–275.
[6] M. Bak, J.T. Rasmussen, N.C. Nielsen, SIMPSON: a general simulation program for solid-state NMR spectroscopy, J. Magn. Reson. 147 (2) (2000) 296–330.
[7] J.J. Shane, Electron spin echo envelope modulation spectroscopy of disordered solids, Ph.D. Thesis, University of Nijmegen, 1993.
[8] R. Szosenfogel, D. Goldfarb, Simulations of HYSCORE spectra obtained with ideal and non-ideal pulses, Mol. Phys. 95 (6) (1998) 1295–1308.
[9] M. Edén, Y.K. Lee, M.H. Levitt, Efficient simulation of periodic problems in NMR. Application to decoupling and rotational resonance, J. Magn. Reson. 120 (1) (1996) 56–71.
[10] T. Charpentier, C. Fermon, J. Virlet, Efficient time propagation technique for MAS NMR simulation: application to quadrupolar nuclei, J. Magn. Reson. 132 (2) (1998) 181–190.
[11] P. Hodgkinson, L. Emsley, Numerical simulation of solid-state NMR experiments, Prog. Nucl. Magn. Reson. Spectrosc. 36 (3) (2000) 201–239.
[12] S. Stoll, A. Schweiger, in preparation.
[13] D.W. Aldermanm, M.S. Solum, D.M. Grant, Methods for analyzing spectroscopic line shapes. NMR solid powder patterns, J. Chem. Phys. 84 (7) (1986) 3717–3735.
[14] H. Ebert, J. Abart, J. Voitländer, Simulation of quadrupole disturbed NMR field spectra by using perturbation theory and the triangle integration method, J. Chem. Phys. 79 (10) (1983) 4719–4723.
[15] A. Ponti, Simulation of one-dimensional magnetic resonance powder lineshapes reduced to area computation, Chem. Phys. Lett. 302 (3–4) (1999) 224–230.
[16] A.V. Oppenheim, R.W. Schafer, Discrete-Time Signal Processing, Prentice-Hall, Englewood Cliffs, NJ, 1989.
[17] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical Recipes in C: The Art of Scientific Programming, Cambirdge University Press, Cambridge, MA, 1992.