Journal of Magnetic Resonance 204 (2010) 64–75

Contents lists available at ScienceDirect

Journal of Magnetic Resonance

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

Cardinal series to filter oversampled truncated magnetic resonance signals

Stéphane Rodts, Dimitri Bytchenkoff *, Teddy Fen-Chong

Institut Navier — Université Paris-Est[1], 2 allée Kepler, 77420 Champs sur Marne, France

ARTICLE INFO

Article history:
Received 25 November 2009
Revised 5 February 2010
Available online 19 February 2010

Keywords:
Digital filter
Bayesian statistics
Inverse problem
Truncation
Oversampling

ABSTRACT

Digital low pass filters are routinely used to improve the signal-to-noise ratio of NMR signals, e.g. FID or echoes, when pass band widths of the available analogue filters do not correspond to the spectral width of the signals. Applying digital filters will always necessitate an oversampling of the signal to filter.

The digital filters with which the commercial spectrometers are nowadays equipped and most of those known to date from literature were designed to be applied to signals in the time domain. Nevertheless, most of them are aimed at optimising the filtering of signals in the frequency domain and tend to distort them in the time domain, especially when applied to truncated signals.

Herein we propose a low pass filter that preserves all the features of the signal in both domains. The method consists in fitting raw NMR data with a finite sum of truncated cardinal sine functions and requires nothing but the signal being a band-limited function. We devised sensible and, in practice, hardly restrictive rules for setting parameters of the filter and applied it to various computer-simulated and experimentally measured truncated data sets to demonstrate its success in filtering both FID and echo signals.

1. Introduction

An NMR-signal is inevitably corrupted with electronic noise, an intrinsic characteristic of the electronic circuits of any NMR-spectrometer or -tomographer. Extra noise may result from interactions between the NMR-tomographer and bulky conductive samples, e.g. in medical MRI. Nowadays, with higher-field superconducting magnets and low temperature units becoming available, the signal-to-noise ratio can be, as stated by the Curie law, most often improved by working at higher static magnetic fields or cooling the sample to increase the signal itself. Moreover, cryogenically cooled probes can be used to lower noise due to thermal fluctuations in RF coils. Nevertheless, there are situations in which it turns out more advantageous or even necessary to work at lower magnetic fields or higher temperatures. This includes coherence transfer via weak scalar couplings and measurements of dipolar relaxation rates [1] in presence of high chemical shift anisotropy (CSA) in liquids, gradient NMR and NMR-imaging of intrinsically inhomogeneous systems such as concretes [2,3] or foams [4], as well as certain relaxometry studies [5]. Furthermore, there are NMR hardware peculiar to in situ studies, e.g. GARField [6] and the NMR MOUSE [7], whose constant and radio frequency magnetic field sources as well as receiver are externalised. This leads to higher field inhomogeneity and lower sensitivity than can usually be achieved at the conventional spectrometers. In such situations, measured signals will necessarily have to be filtered.

To process a signal with a digital filter, it must first be passed trough an analogue filter and then digitised by analogue-to-digital converters (ADC). The sampling frequency must be set larger than the pass band width of the analogue filter to avoid that high frequency noise folds on to the spectrum of the signal when it is digitised. As any NMR spectrometer has only a limited number of analogue filters, satisfying this condition often amounts to automatic setting the sampling frequency up to several orders of magnitude higher than that set by the spectroscopist to accurately record all the harmonics in the spectrum only, were the signal noise-free.

In one-dimensional and the direct dimension of multi-dimensional experiments, the signal can also be deliberately oversampled since this can, as will later transpire, be used for further signal-to-noise improvement. The oversampling, though useful and often necessary, results in a tremendous increase in the size of data to process with the digital filter.

There are numerous NMR experiments, e.g. CPMG and PGSE, in which information about physically or chemically relevant quantities are extracted direct from signals in the time domain. Filtering them from high frequency noise will improve the signal-to-noise ratio at those moments where the values of the signal are of interest, echoes maxima or beginnings of FID-signals and thus increase

[1] Institut Navier encompasses civil engineering laboratories associated with the Laboratoire Central des Ponts et Chaussées (LCPC), Ecole Nationale des Ponts et Chaussées (ENPC) and Centre National pour la Recherche Scientifique (CNRS).

* Corresponding author. Fax: +33 (0) 1 40 43 54 50.
E-mail addresses: Stephane.Rodts@lcpc.fr (S. Rodts), Dimitri.Bytchenkoff@lcpc.fr (D. Bytchenkoff).

1090-7807/$ - see front matter © 2010 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmr.2010.02.002

the precision with which the physical quantities will be determined. On the other hand, in spectroscopic and imaging experiments, in which signals measured as a function of time are meant to be Fourier-transformed, digital filtering will solely contribute to minimising [8] the volume of data without loss of information.

Three approaches [9] have so far been widely taken to damp high frequency noise. The first consists in processing and downsampling oversampled signals with one of in-built digital filters of the spectrometer software at the very moment as they are being measured and digitised. The output data of such filters is a signal freed of high frequency noise, digitised at the frequency that corresponds to its spectral width. The spectrometers are equipped with causal, finite (FIR) and infinite (IIR) impulse response filters. The FIR filter calculate each output sample as a linear combination of those samples that have already been to that point digitised by the ADC of the spectrometer. Each output sample of the IIR filter is a linear combination of both the samples that have been digitised by the ADC and those previously supplied by the filter itself. Here decimation by a factor N is achieved simply by retaining only one of N samples.

Another approach, known as low pass convolution filter, consists in acquiring an entire data set and then convoluting it in the time domain with an appropriate band-limited function. Convolution filters are mathematically similar to, though, significantly better than causal filters, as acquisition and filtering now take place sequentially, providing greater freedom in choosing the kernel.

Applying either type of filter to an FID signal will result in distortions, as in their algorithms, signals are supposed to be continuous functions of time, equal to zero at the boundaries of the acquisition interval. The non-zero value of the FID signal at $t = 0$ is erroneously interpreted by the filters as a discontinuity of the signal or, alternatively, spurious high frequency harmonics. The filters will eliminate them and thus distort several first samples of the signal in the time domain. Moreover, operating in real time, those filters will tend to dephase various signal harmonics and thus artificially postpone the beginning of signals by a so-called group delay in the time domain. All this will result in dephased spectra with distorted base-line or faulty images. Contrariwise, convolution filters with certain, rather sophisticated [10,11], kernels do give well phased spectra at the expense of signal-to-noise ratio. Extra noise, though, will appear outside the spectrum, which makes convolution filters perfectly suitable for spectroscopic experiments. On the other hand, they still fail utterly in the time domain when applied to truncated signals (see below).

The third approach consists in filtering and possibly restoring missing parts of the time-dependent signal by fitting it with model functions [9]. However, the choice of those functions may depend on the type of experiments and will require certain prior knowledge of the system under study. In particular, various linear prediction (LP) methods [12–14] proved successful in both improving spectral resolution, which in theory can always be turned into increase in signal-to-noise ratio, and alleviating phase and amplitude distortions due to, respectively, delayed start of acquisition and truncation of the signal. Nevertheless, those methods rely on the assumption that the spectral lines are Lorentzian or Gaussian and thus are not applicable to many physically relevant systems, e.g. those in which transverse magnetisation is known to decay non exponentially [15].

This paper reports on a low pass filter that we have designed to process truncated oversampled NMR signals in the time domain. The method consists in fitting raw experimental data with a series of cardinal sine functions and requires nothing but the signal to be a band-limited function. Such expression thoroughly allows for that part of the original noise-impaired signal whose spectrum fits into the cardinal series band width, while the rest of the signal will be largely overlooked. The band width of the series can be set to just over that of the signal by an appropriate choice of cardinal sine functions, thus making of it an excellent low pass filter for NMR signals.

The paper is structured as follows: We first remind some basic properties of the band-limited functions and convolution low pass filters and emphasise limits of the former for processing truncated data sets. Second, we introduce a low pass filter based on the use of cardinal series. Third, we derive sensible rules for setting parameters of the filter by numerically optimising them for two computer-simulated noise-free data sets. Forth, we examine tolerance of our filter towards signal truncations by filtering various experimental and computer-simulated signals and compare it with that of a convolution filter with Lanczos kernel. Finally, we examine performance of our method by applying it to data sets collected in CPMG experiments.

2. Theory

2.1. Properties of band-limited functions

If a continuous time–function $x(t)$ is band-limited in the low pass sense with a bandwidth Ω_o, i.e.

$$\int_{-\infty}^{\infty} x(t)e^{-i\omega t}dt = 0 \text{ for } |\omega| > \frac{\Omega_o}{2} \tag{1}$$

then, according to the sampling theorem [16,17], it can be expressed with no loss of information as an infinite cardinal series

$$x(t) = \sum_{m=-\infty}^{\infty} x_m sinc\frac{\Omega}{2}(t-\tau_m) \tag{2}$$

where *sinc* stands for the cardinal sine function (see also Fig. 1A)

$$sinc\ \alpha = \frac{\sin\alpha}{\alpha}, \tag{3}$$

time t is arbitrary ($-\infty < t < \infty$), coefficients $x_m = x(\tau_m)$ are discrete samples of the function at a countable yet infinite number of moments $\tau_m = 2\pi m/\Omega$ and the sampling frequency Ω fulfils the Nyquist criterion [18], i.e. it must not be less than the function band width Ω_o

$$\Omega = \frac{2\pi}{\tau_{m+1}-\tau_m} \geqslant \Omega_o \tag{4}$$

The band width Ω_o was defined here so that $\Omega_o/2\pi$ corresponds to the sweep width in Hertz (SWH) in the customary usage of NMR spectroscopists. The reader is referred to Table 1 for a list of all frequency parameters.

In the most general case the series in Eq. (2) is mathematically proved [19] to exhibit only a mean square convergence, though, in practice it most often converges uniformly. Furthermore, Eq. (2) reduces into an identity when $t = \tau_m$, which implies that the series will pass through all the samples.

Quite clearly, samples x_m become interdependent, when the function $x(t)$ is over-sampled, i.e. $\Omega > \Omega_o$, and given the assumed band-limited character of the function, Eq. (2) can be generalised [19] as

$$x(t) = \frac{\Omega'}{\Omega''}\sum_{m=-\infty}^{\infty} x_m sinc\frac{\Omega'}{2}(t-\tau_m)$$

$$\text{where } \Omega_o < \Omega' < 2\Omega'' - \Omega_o \text{ and } \Omega'' = \frac{2\pi}{\tau_{m+1}-\tau_m} \tag{5}$$

It should be mentioned that the expression of the function $x(t)$ as an infinite cardinal series is no longer unique when $\Omega' \neq \Omega''$, i.e. there

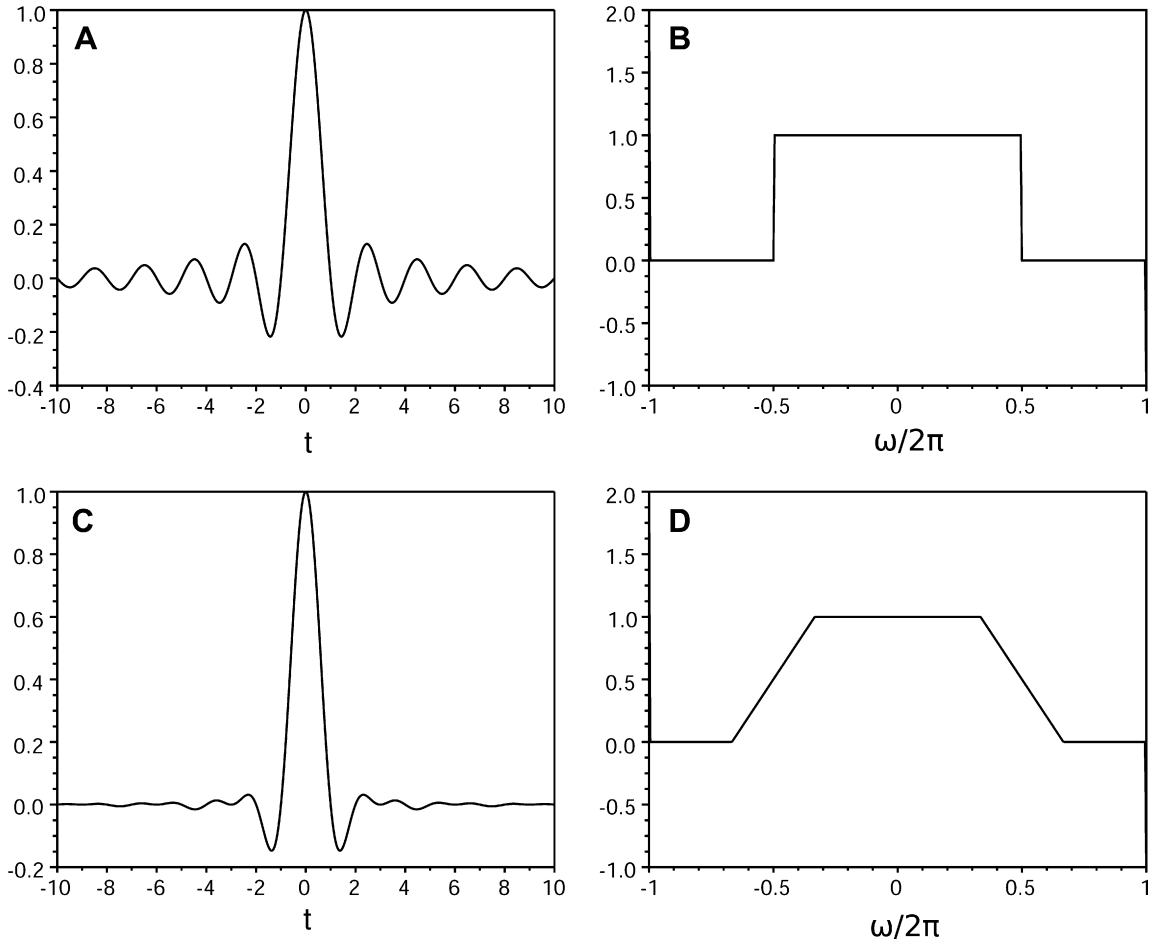

Fig. 1. (A) Cardinal sine function $sinc(t) = sin(\pi t)/(\pi t)$. (B) FT (in arbitrary units) of the cardinal sine function. (C) Lanczos kernel $x(t) = sinc(\pi t)sinc(\pi t/3)$. (D) FT (in arbitrary units) of the Lanczos kernel.

Table 1
List of frequency parameters.

Symbol table	
Ω_o	Signal band width
Ω	ADC sampling frequency
Ω'	Cardinal series band width
Ω'_o	Cardinal kernel band width
$\Omega'' = 2\pi/(\tau_{m+1} - \tau_m)$	Where τ_m mth cardinal sine maximum location
Ω_f	Analogue filter pass band

is an infinity of sets of coefficients a_m other than actual samples $x_m = x(\tau_m)$ of the function that would still satisfy Eq. (5).

In practice, NMR-signals are either band-limited or can reasonably be viewed as such and are passed first through an analogue filter with a pass band width Ω_f ($\Omega_f \geqslant \Omega_o$). They are then sampled at a frequency Ω that must be higher than the band width of the analogue filter ($\Omega \geqslant \Omega_f$) to avoid the folding of high frequency noise on to the signal band width. This usually leads to an oversampling of the signal, i.e. sampling at a much higher frequency ($\Omega \gg \Omega_o$) than it would be required by the Nyquist criterion to just accurately record all frequency components of the signal. Finally, the analogue-to-digital converters (ADC) of the spectrometer can output only a finite number of complex regularly spaced noise-impaired samples $x(t_n) + \delta x(t_n)$ with $1 \leqslant n \leqslant N$, of which one would most often like to restore an entire and preferably further filtered signal $s(t)$ with $t_1 \leqslant t \leqslant t_N$. To do so, one can, unless the signal is truncated, apply to it a convolution filter, whose basic principle we shall briefly remind in the next section. Otherwise, we propose, further below, a low pass filter based on using a finite cardinal series.

2.2. Convolution filter

The convolution low pass filter consists in convoluting a band-limited signal $x(t)$ of band width Ω_o impaired with white noise $\delta x(t)$ with a likewise band-limited kernel $f(t)$ of bandwidth Ω'_o subject to $\Omega_o \leqslant \Omega'_o$.

$$s(t) = \int_{-\infty}^{\infty} \{x(t') + \delta x(t')\} f(t - t') dt' \tag{6}$$

This amounts to the multiplication of the spectrum of the signal by the Fourier transform (FT) of the kernel

$$\tilde{s}(\omega) = 2\pi\{\tilde{x}(\omega) + \tilde{\delta x}(\omega)\}\tilde{f}(\omega) \tag{7}$$

or in more detail

$$\tilde{s}(\omega) = 2\pi\{\tilde{x}(\omega) + \tilde{\delta x}_l(\omega) + \tilde{\delta x}_h(\omega)\}\tilde{f}(\omega) \tag{8}$$

where $\delta x_l(\omega)$ and $\delta x_h(\omega)$ stand for noise at frequencies in intervals $|\omega| \leqslant \Omega_o/2$ and $\Omega_o/2 < |\omega| < \Omega_f/2$, respectively. We assumed here that the signal was sampled in accord with the Nyquist criterion (Eq. 4) and long enough for it to come to zero by the end of the acquisition; and thus used an integral with infinite boundaries rather than a truncated series. Theoretically the ideal type of kernel is a cardinal sine function

$$f(t) = \frac{\Omega_o}{2\pi} sinc\left(\frac{\Omega_o}{2} t\right) \tag{9}$$

whose FT equals one when $|\omega| \leqslant \Omega_o/2$ and zero elsewhere (see also Fig. 1B). Taking such a kernel will simplify Eq. (7) to

$$\tilde{s}(\omega) = \tilde{x}(\omega) + \tilde{\delta x}_l(\omega) \tag{10}$$

Thus, the filter with rectangular-shaped spectrum rids raw data of that part of white noise whose frequencies exceed the band width of the signal and thus increases the signal-to-noise ratio by a

$$\text{factor of } \frac{\sigma}{\sigma'} = \sqrt{\frac{\Omega_f}{\Omega_o}} \tag{11}$$

where σ and σ' are standard deviations of the raw and filtered signals, respectively, i.e.

$$\sigma = \sqrt{\langle|\delta x^2(t)|\rangle} \text{ and } \sigma' = \sqrt{\langle|\delta s^2(t)|\rangle} \tag{12}$$

The factor σ/σ' reflects a uniform noise reduction in the time domain, whilst in the frequency domain the same, owing to Parseval's theorem, factor corresponds to noise alleviation outside the pass band of the filter. Within this band, noise is left intact by the filter.

2.3. Cardinal series filter

To filter a truncated signal, we propose to fit the N available raw data samples $x(t_n) + \delta x(t_n)$ with a truncated cardinal series

$$x(t_n) + \delta x(t_n) \approx \sum_{m=M_{inf}}^{M_{sup}} a_m sinc\frac{\Omega'}{2}(t_n - \tau_m) \text{ for } 1 \leqslant n \leqslant N \tag{13}$$

Only an infinite cardinal series would give an exact expression for the noise-free samples $x(t_n)$ of the band-limited function $x(t)$. Nevertheless, fast convergence of the series allows, in practice, to limit it to a reasonable number of members. The number $M_{sup} - M_{inf} + 1$ of cardinal sine functions to include into the truncated series as well as their band width Ω' and maxima locations τ_m are constant parameters to optimise, subject to conditions [19]

$$\Omega_o \leqslant \Omega' \tag{14}$$

$$\Omega'' \geqslant \frac{\Omega' + \Omega_o}{2}, \text{ where } \Omega'' = \frac{2\pi}{\tau_{m+1} - \tau_m} \tag{15}$$

One then looks for complex coefficients a_m that achieve the best fit of the samples. According to Bayesian analysis [20], finding maximum likelihood values of such coefficients for normally distributed noise, amounts to minimising

$$min_{a_m(M_{inf} \leqslant m \leqslant M_{sup})} \times \sum_{n=1}^{N} \left| x(t_n) + \delta x(t_n) - \sum_{m=M_{inf}}^{M_{sup}} a_m sinc\frac{\Omega'}{2}(t_n - \tau_m) \right|^2 \tag{16}$$

This can be written in a matrix form as

$$min_A = |X - MA|^2 \tag{17}$$

where M is a non-square matrix

$$M = \begin{pmatrix} sinc\frac{\Omega'(t_1 - \tau_{M_{inf}})}{2} & \dots & sinc\frac{\Omega'(t_1 - \tau_{M_{sup}})}{2} \\ \vdots & & \vdots \\ sinc\frac{\Omega'(t_N - \tau_{M_{inf}})}{2} & \dots & sinc\frac{\Omega'(t_N - \tau_{M_{sup}})}{2} \end{pmatrix} \tag{18}$$

and X and A are column vectors

$$X = \begin{pmatrix} x(t_1) + \delta x(t_1) \\ \vdots \\ x(t_N) + \delta x(t_N) \end{pmatrix} \text{ and } A = \begin{pmatrix} a_{M_{inf}} \\ \vdots \\ a_{M_{sup}} \end{pmatrix} \tag{19}$$

Eq. (17) is formally solved by the pseudo-inverse method as

$$A_{opt} = (M^t M)^{-1} M^t X. \tag{20}$$

To calculate A_{opt}, we chose to compute the matrix M^tM and vector M^tX explicitly, as this, in our opinion, requires less computer memory than other algorithms [21] when dealing with large data sets. However, M^tM may turn out singular, when the number of elements in A exceeds that in X, or ill-conditioned. The former happens when the problem of finding the minimum in Eq. (16) turns out ill-posed, i.e. there are several distinct sets of coefficients a_m that provide an equally good fit of the data samples $x(t_n) + x\delta\ (t_n)$ with the truncated series of Eq. (13). Computational stability can nevertheless be restored by Tikhonov regularisation [20,21], which consists in substituting Eq. (20) by

$$A'_{opt} = (M^t M + \lambda I)^{-1} M^t X \tag{21}$$

where λ is a small real positive constant and I, an identity matrix. This amounts to replacing Eq. (17) by

$$min_A = |X - MA|^2 + \lambda|A|^2 \tag{22}$$

and does not noticeably compromise exactitude of the fit in Eq. (13).

Thus found optimum coefficients A'_{opt} can then be used to calculate the complex signal at an arbitrary moment t as

$$s(t) = G^t(t) A'_{opt} \tag{23}$$

where G is a column vector

$$G(t) = \begin{pmatrix} sinc\frac{\Omega'}{2}(t - \tau_{M_{inf}}) \\ \vdots \\ sinc\frac{\Omega'}{2}(t - \tau_{M_{sup}}) \end{pmatrix} \tag{24}$$

In particular, one may want to calculate $s(t)$ at the same moments $t = t_n$, $(1 \leqslant n \leqslant N)$ at which the signal $x(t) + \delta x(t)$ was originally sampled by the ADC of the spectrometer.

If $\delta x(t_n)$ is white Gaussian noise of standard deviation

$$\sigma = \sqrt{\langle|\delta x^2(t_n)|\rangle} \text{ for } 1 \leqslant n \leqslant N \tag{25}$$

that of the Gaussian noise of the signal $s(t_n)$ processed as described above can be expressed as

$$\sigma'(t_n) = \sqrt{\langle|\delta s^2(t_n)|\rangle} = \sigma\sqrt{V^t(t_n)V(t_n)} \quad \text{for} \quad 1 \leqslant n \leqslant N \tag{26}$$

where V stands for a column vector

$$V(t) = M(M^t M + \lambda I)^{-1} G(t) \tag{27}$$

Thus, noise reduction σ'/σ depends only on the parameters of the cardinal series and sample moments and is independent of the raw data samples themselves.

This improvement of the signal-to-noise ratio stems from lessening noise of frequencies higher than $\Omega'/2$. As could be expected from the sampling theorem, the band-limited signal $x(t)$ with the band width Ω_o is accurately described by even the finite series of truncated cardinal sine functions. In theory, the series is itself a band-limited function with a band width Ω'. In practice, the FT of the series determined within $[t_1,t_N]$ may turn out to have a band width somewhat larger than Ω', owing to this truncation. On the other hand, the cardinal series can't adequately take account of the part of the white electronic noise that falls out of its band width. Thus, the coefficients A'_{opt} and likewise calculated signal

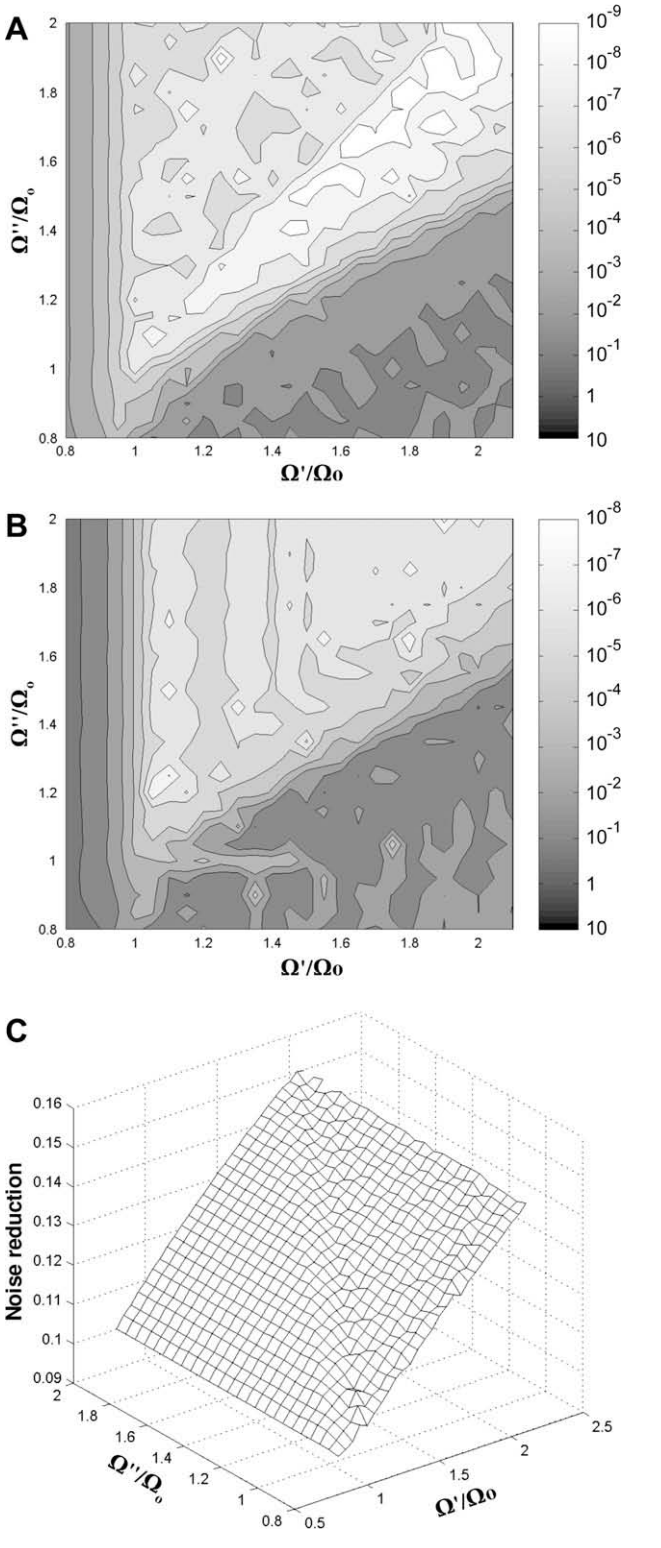

Fig. 2. Two-dimensional charts of the cardinal series filter exactitude $|s(0)-1|$ calculated as a function of Ω' and Ω'' for the computer-simulated data sets modelled as noise-free samples of the functions of (A) Eq. (28) and (B) (29), where $\omega_1 = 2\pi$, $\omega_2 = \pi$, $a = 0.6$, $b = 0.4$ and $t_o = 1$. The functions were sampled in an interval $[t_1,t_N]$, $t_1 = -1.4$, $t_N = 1.4$ at a normalised sampling frequency $\Omega/\Omega_o = 100$. The calculations were carried out for 27 and 25 values of Ω'/Ω_o and Ω''/Ω_o incrementally varied from 0.8 to 2.1 and from 0.8 to 2.0, respectively. Ten additional cardinal sine functions with maxima beyond either boundary of the sampling area were provided for a further possible improvement in the filter performance. (C) Three-dimensional diagram of the noise reduction σ'/σ at $t = 0$ that can be achieved by applying the filter to the data sets described above.

$s(t)$ are expected to hold all the information about the raw data samples of the band-limited signal $x(t)$. On the contrary, the information about electronic noise at frequencies higher than $\Omega'/2$ will

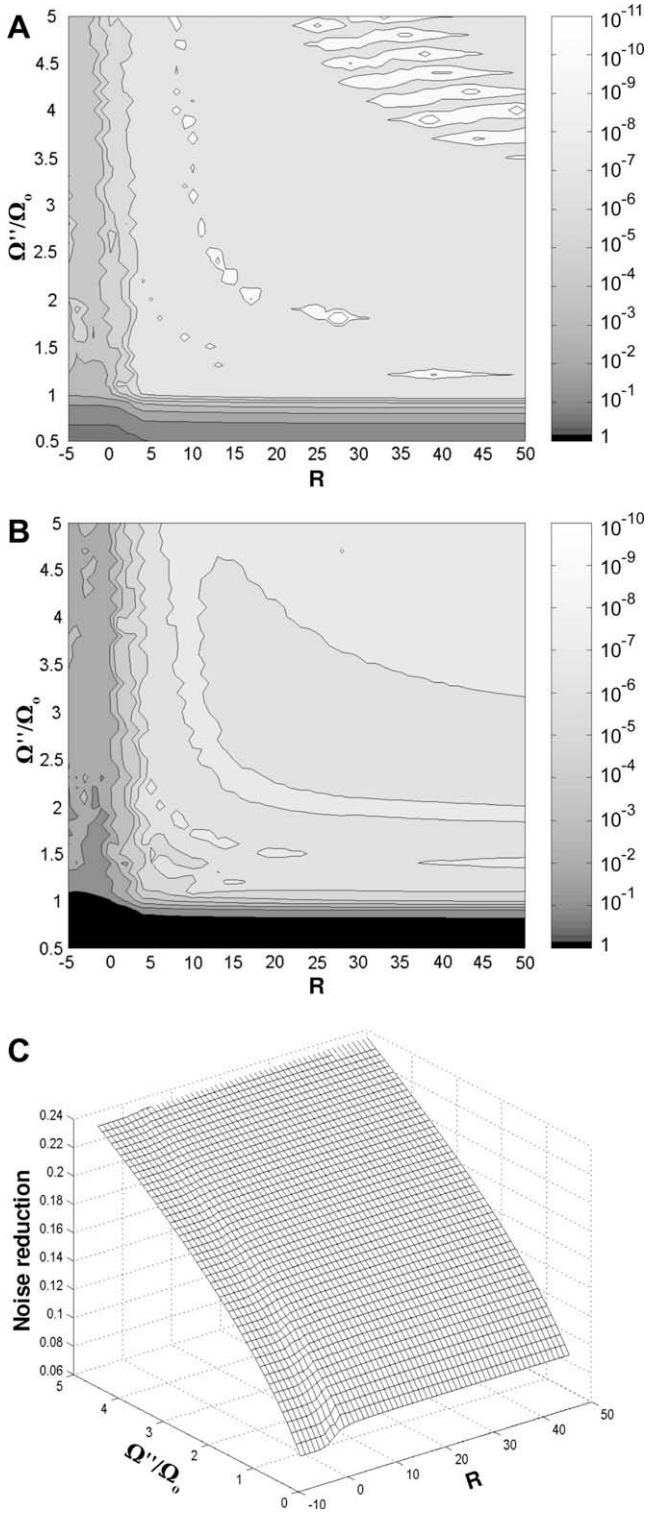

Fig. 3. Two-dimensional charts of the cardinal series filter exactitude $|s(0)-1|$ calculated as a function of R and Ω'' for the data sets of (A) Eq. (28) and (B) (29), respectively. The calculations were carried out for 56 and 90 values of R and Ω''/Ω_o incrementally varied from -5 to 50 and from 0.5 to 5, respectively. (C) Three-dimensional diagram of the noise reduction σ'/σ at $t = 0$ that can be achieved by applying the filter.

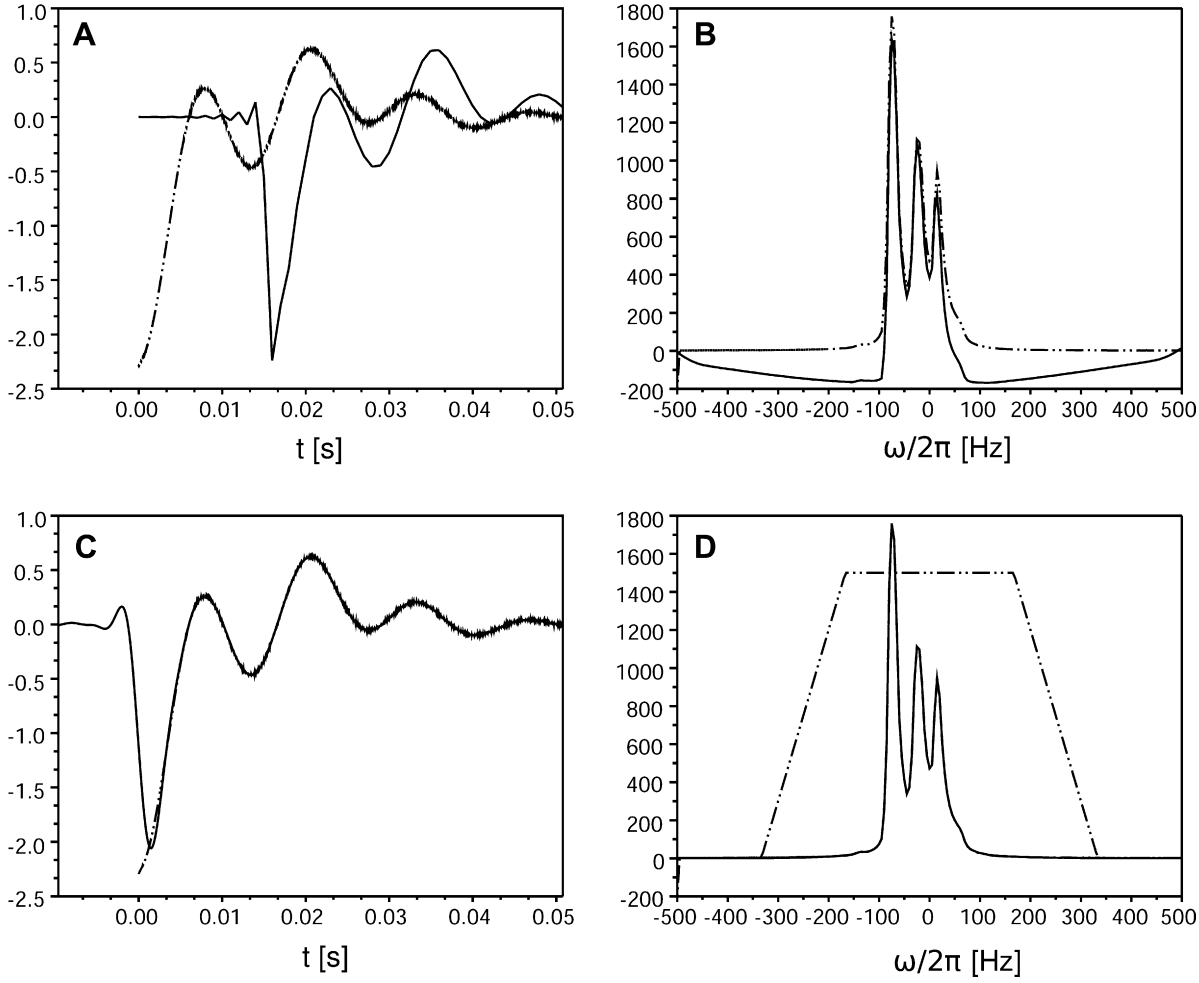

Fig. 4. (A) real part of an FID signal acquired as described in Section 5 before (dashed line) and after (solid line) it was processed with a Bruker spectrometer in-built digital filter with a pass band of 1 kHz; (B) uniformly phase-corrected spectrum (dashed line) of the raw signal and uniformly and linearly phase-corrected spectrum (solid line) of the filtered signal of Fig. 2A; (C) real part of an FID signal before (dashed line) and after (solid line) it was processed with a digital Lanczos filter with a pass band of 333 Hz; (D) uniformly phase-corrected spectrum (dashed line) of the raw signal and uniformly and linearly phase-corrected spectrum (solid line) of the signal processed with the Lanczos filter; FT (also dashed line) in arbitrary units of the Lanczos kernel. The spectra appear identical to each other.

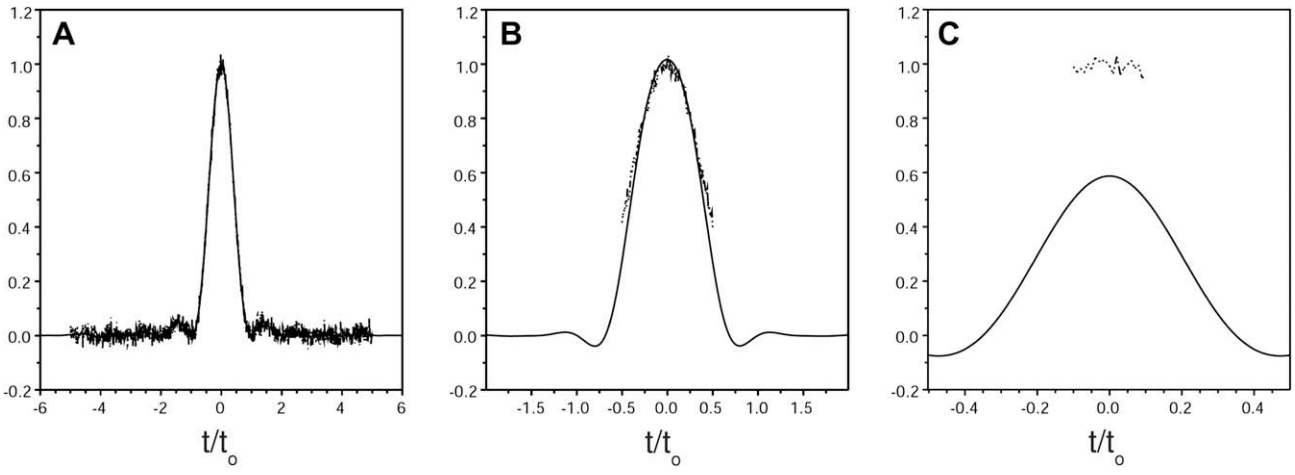

Fig. 5. Signals modelled as a $sinc^2(\pi t/t_o)$ function defined for (A) $-5 \leqslant t/t_o \leqslant 5$, (B) $-0.5 \leqslant t/t_o \leqslant 0.5$ and (C) $-0.1 \leqslant t/t_o \leqslant 0.1$ and impaired with computer-simulated white Gaussian noise of standard deviation 0.02 before (dotted line) and after Lanczos (continuous line) they were processed with a Lanczos filter. The pass band width of the filter was set equal to the band width $\Omega_o = 4\pi/t_o$ of the signals to filter.

be largely lost, making of the cardinal series an excellent low pass filter for NMR-signals.

3. Results and discussion

3.1. Setting the cardinal series filter parameters

Exactitude, precision and calculation efficiency of the 'cardinal series filter' depends on the choice of values of the parameters Ω', Ω'', M_{inf} and M_{sup} in Eq. (16). To derive sensible rules for setting these values, we optimised them by processing data sets modelled as N noise-free samples of two band-limited functions with a band width $\Omega_o = 4\pi/t_o$

$$x(t) = sinc^2\left(\pi\frac{t}{t_o}\right) \tag{28}$$

$$x(t) = a\cdot\cos(\omega_1 t) + b\cdot\cos(\omega_2 t), \quad \text{with} \quad \omega_1 = 2\pi/t_o, \omega_2 = \pi/t_o, a = 0.6 \text{ and } b = 0.4 \tag{29}$$

where the constant unit t_o was set to one in either function. The functions have shapes somewhat reminiscent of spin echoes with a maximum at $t = 0$. The spectrum of the function of Eq. (28) can be viewed as a broad line centred at $\omega = 0$, while that of the function of Eq. (29) as four infinitely narrow lines at $\omega = \pm\omega_1$ and $\omega = \pm\omega_2$. Such spectra seam to us to correspond to two distinct situations that most often occur in practice. The functions were sampled in an interval $t_1 \leqslant t = t_n \leqslant t_N$, $t_1 = -1.4t_o$, $t_N = 1.4t_o$ at a sampling frequency Ω set to 100 times the band width Ω_o, which corresponds to a oversampled truncated signal and would be a usual situation in one-dimensional experiments or in the direct dimension of multi-dimensional experiments.

Analysis of the filter parameters on an exponentially damped harmonically oscillating function (not shown here) proved less insightful, as such FID-like signal is not band-limited, strictly speaking. Nevertheless, the filter with the parameters optimised on the echo-like signals of Eqs. 28 and 29 will perform beautifully (see below) when applied to the FID-like signals.

First, we examined the exactitude $|s(0) - x(0)|$—also known as 'aliasing error' [22,23]—with which the filter calculates the signal at its maximum $x(0) = 1$ as well as the noise reduction σ'/σ that can be thus achieved, had the samples been noise-impaired. Fig. 2A and B show two-dimensional charts of the $|s(0) - 1|$ values calculated for the data sets mentioned above as a function of Ω' and Ω''. The calculations were carried out for 27 and 25 values of Ω'/Ω_o and Ω''/Ω_o incrementally varied from 0.8 to 2.1 and from 0.8 to 2.0, respectively. We also provided (see Eq. 13) 10 additional cardinal sine functions whose maxima appear beyond either boundary of the sampling area, i.e. τ_m satisfied

$$t_1 - 10\frac{\pi}{\Omega''} \leqslant \tau_m \leqslant t_N + 10\frac{\pi}{\Omega''} \tag{30}$$

for a further possible improvement in the filter performance. Clearly, the method fails when $\Omega' < \Omega_o$ and $2\Omega'' < \Omega_o + \Omega'$, as then the Nyquist criterion is not met. Elsewhere, the filter introduced a bias of order 10^{-6} to 10^{-7}, which is rather marginal compared to the accuracy with which NMR signals can be collected. We believe that the bias may correspond to a round off error in M singular values that occurs when the matrix product M^tM in Eq. (21) is explicitly calculated in a 64-bit double precision format. Certain roughness of the surfaces in Fig. 2A and B, as well as those in Fig. 3A and B and Fig. 6A and B (see below) was found to depend on the form of the signal to be processed and is, we reckon, irrelevant to the subject discussed herein. Fig. 2C shows a two-dimensional chart of the noise reduction factor σ'/σ at $t = 0$ as a function of Ω' and Ω''. Unlike the exactitude, this factor depends on the series constant parameters only and is independent of

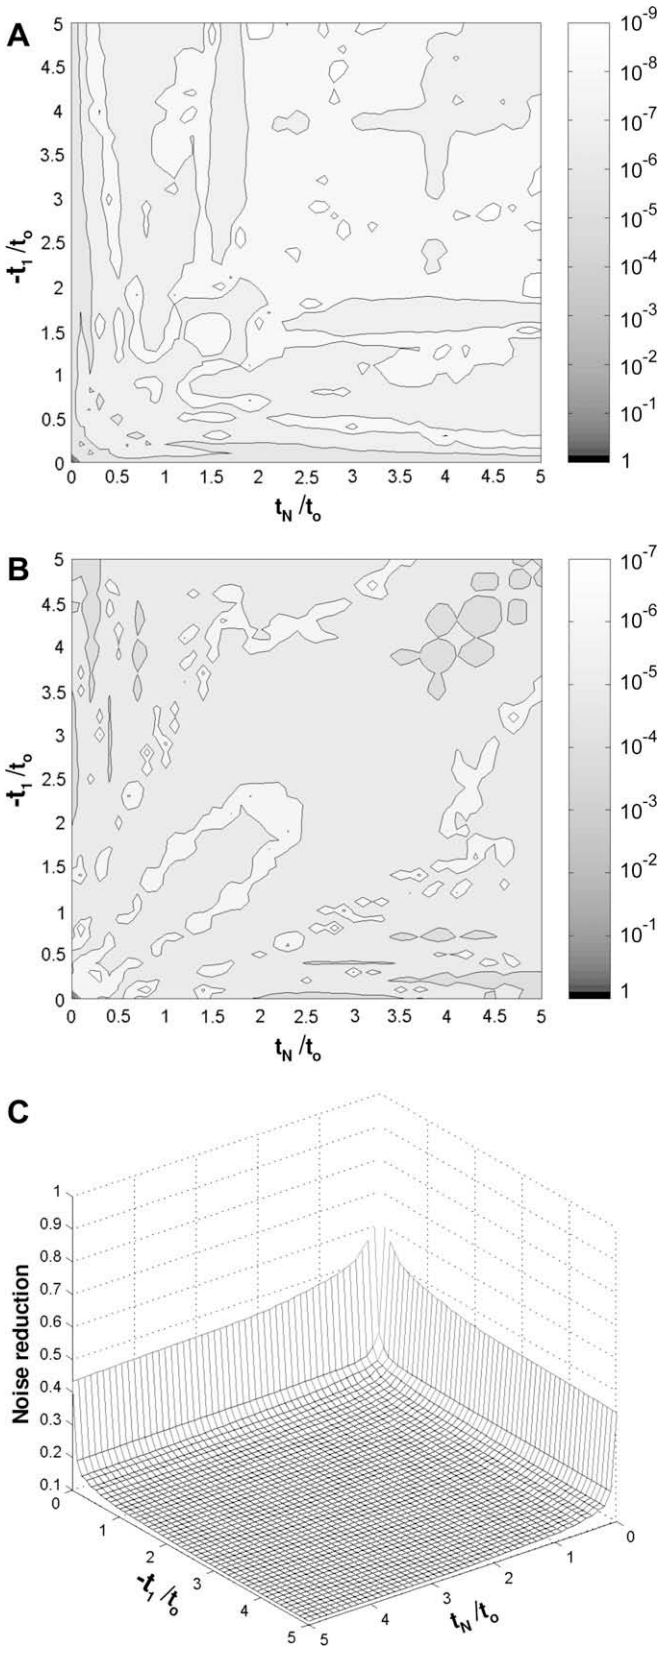

Fig. 6. Two-dimensional charts of the $|s(0) - 1|$ values calculated as a function of the lower t_1 and upper t_N boundaries of the sampling area for the data sets of (A) Eq. (28) and (B) (29), respectively. The calculations were carried out for *51* values of either parameter independently and incrementally varied from $-5t_o$ to *0* and from *0* to $5t_o$, respectively. (C) Three-dimensional diagram of the noise reduction σ'/σ at $t = 0$ that can be achieved by applying the filter.

sampled signals (see Eq. 26). The noise reduction proved quite insensitive to Ω'' and, as it could be expected, decreases monotonously when Ω' diminishes. Similar results were obtained when the model functions of Eqs. 28 and 29 were further truncated. This allowed us to infer that Ω' and Ω'' can be reasonably set equal to one another and just over the band width Ω_o of a signal to filter. In what will follow $\Omega'' = \Omega'$.

Second, given poor performance of cardinal series in extrapolating functions, the interval $[\tau_{M_{inf}}, \tau_{M_{sup}}]$ where cardinal sine functions reach their maximum should encompass the sampling area $[t_1, t_N]$, i.e.

$$\left[\tau_{M_{inf}}, \tau_{M_{sup}}\right] = \left[t_1 - R\frac{2\pi}{\Omega''}, t_N + R\frac{2\pi}{\Omega''}\right] \tag{31}$$

where an integer R stands for the number of cardinal sine functions, to include in the series, with maxima beyond each of the sampling area boundaries. The number R to optimise is related to the parameters M_{inf} and M_{sup} as

$$M_{inf} = int\left(\frac{\Omega'' t_1}{2\pi}\right) - R \quad \text{and} \quad M_{sup} = int\left(\frac{\Omega'' t_N}{2\pi}\right) + R \tag{32}$$

Fig. 3A and B show two-dimensional charts of the $|s(0) - 1|$ values calculated for the data sets mentioned above as a function of R and $\Omega'' = \Omega'$. The calculations were carried out for 56 and 90 values of R and Ω''/Ω_o incrementally varied from -5 to 50 and from 0.5 to 5, respectively. When the Nyquist criterion $\Omega'' > \Omega_o$ is fulfilled, varying Ω'' any further has little impact on the exactitude with which the filter predicts the value of the signal at its maximum. From now on, we set $\Omega'' = \Omega' = 1.1\Omega_o$. The exactitude proved reasonable for all explored positive values of R and was particularly good for $R > 4$, where the filter introduced, what we believe to be, a round off error of order 10^{-7}. The noise reduction at $t = 0$ (see Fig. 3C) proved almost independent of R and decreased monotonously with decreasing Ω''. To be on the safe side without compromising on computational efficiency, we decided to set $R = 6$ hereafter.

3.2. Tolerance of the cardinal series filter towards signal truncations

Performance of the convolution filter will depend crucially on the period $T = t_N - t_1$ during which the signal is sampled. When it becomes too short to allow for an adequate spectral resolution $\delta\omega = 2\pi/T$ of the signal, the filter distorts it very badly. This makes the convolution filter unsuitable for processing truncated signals.

Clearly, an infinitely long acquisition interval $-\infty < t < \infty$, as required by Eq. (6), cannot be realised in practice. Nevertheless, distortions introduced by the filter will be negligible when the signal is measured during such an interval that it equals zero at both of the acquisition interval boundaries, which will often be experimentally impossible. FID signals reach their maximum at the lower boundary of the acquisition interval and often cannot be measured long enough for nuclear spin relaxation to reduce them to zero. Nor can echoes always be recorded between two sequential moments at which spin magnetisation is refocused. In such situations, the filter will distort the signal everywhere, distortions being particularly severe near the discontinuities of the signal owing to the truncation. Neither artificial setting of values of the signal $x(t)$ to zero outside the acquisition interval $[t_1,t_N]$ nor treating the signal as a periodic function with a period equal to the acquisition interval can alleviate the truncation problem.

Application of a digital filter with a pass band of 1 kHz to a raw FID signal (dashed line in Fig. 4A) of ethanol first passed trough an analogue filter with a pass band of 100 kHz and oversampled at the frequency of 100 kHz gives a signal (solid line in Fig. 4A) downsampled to 1 kHz, distorted near $t_1 = 0$ and shifted as a whole with respect to the raw signal along the time axis. A frequency-independent and linear, to compensate for the time shift, phase correction and FT of the filtered signal gives a spectrum (solid line in Fig. 4B) similar to the uniformly phase-corrected spectrum of the raw signal (dashed line in Fig. 4B), though, with a twisted base line.

Fig. 4C shows the raw signal of Fig. 4A extended by setting its value to zero outside the interval $[0, t_N]$ before (dashed line) and after (solid line) it was processed with a convolution filter with the Lanczos kernel (see also Fig. 1C and D)

$$f_{Lan}(t) = \frac{\Omega}{2\pi} sinc\left(\frac{\Omega}{2}t\right) sinc\left(\frac{\Omega}{6}t\right), \tag{33}$$

which has been recognised by many as a filter providing the best compromise between filtering performance and limiting truncation-related distortions [10,11], and pass band width $\Omega/3\pi$ set to 333 Hz. The convolution filters tends to smooth the discontinuity of the signal at $t = 0$ and by doing so distorts the signal, as can be seen in Fig. 4C, for time close to zero. The uniformly phase-corrected FT of these signals, on the other hand, appear identical (see Fig. 4D); and base line is now perfect. Thus, the Lanczos filter proves perfectly adequate for processing truncated signals for spectroscopic applications. It leaves, though, much to be desired when truncated signals are to analyse in the time domain. Success

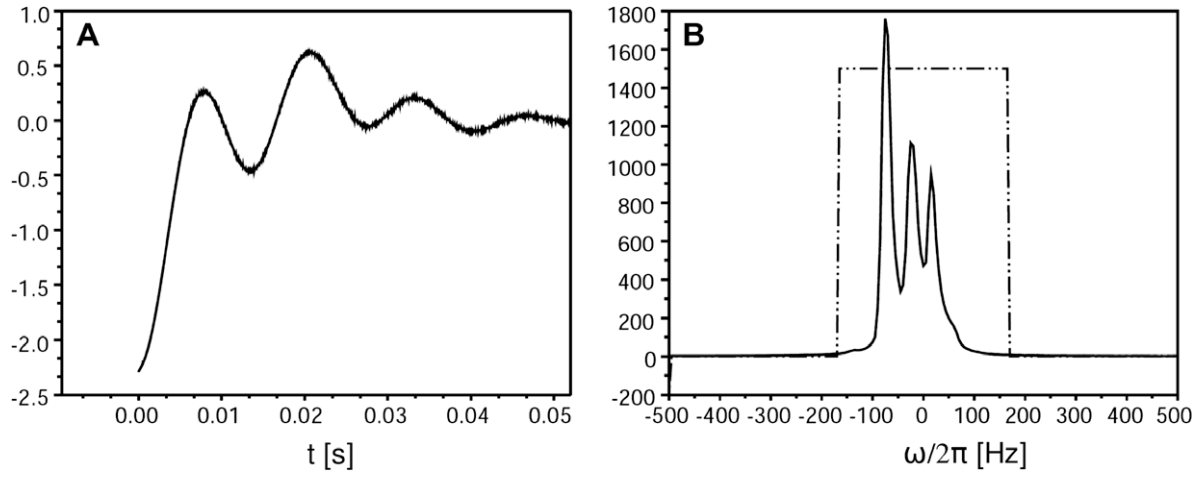

Fig. 7. (A) Real part of an FID signal acquired as described in Section 5 before (dashed line) and after (solid line) it was processed by the cardinal series filter with a pass band of 333 Hz; (B) uniformly phase-corrected spectra of the raw (dashed line) and filtered (solid line) signals of Fig. 7A, FT (also dotted line) in arbitrary units of the truncated cardinal sine function. The raw and filtered signals coincide perfectly in either domain.

in the application of a convolution filter to a band-limited signal depends crucially on its spectral resolution $\delta\omega$. To ensure this, the signal must be sampled long enough for $t_N - t_1 \gg 2\pi/\Omega_o$. Lowering the resolution worsens greatly the performance of the filter and it fails utterly when applied to truncated signals, i.e. when $t_N - t_1 \leqslant 2\pi/\Omega_o$.

To illustrate this, we applied a Lanczos convolution filter to signals (dashed line in Fig. 5A–C), modelled as a function

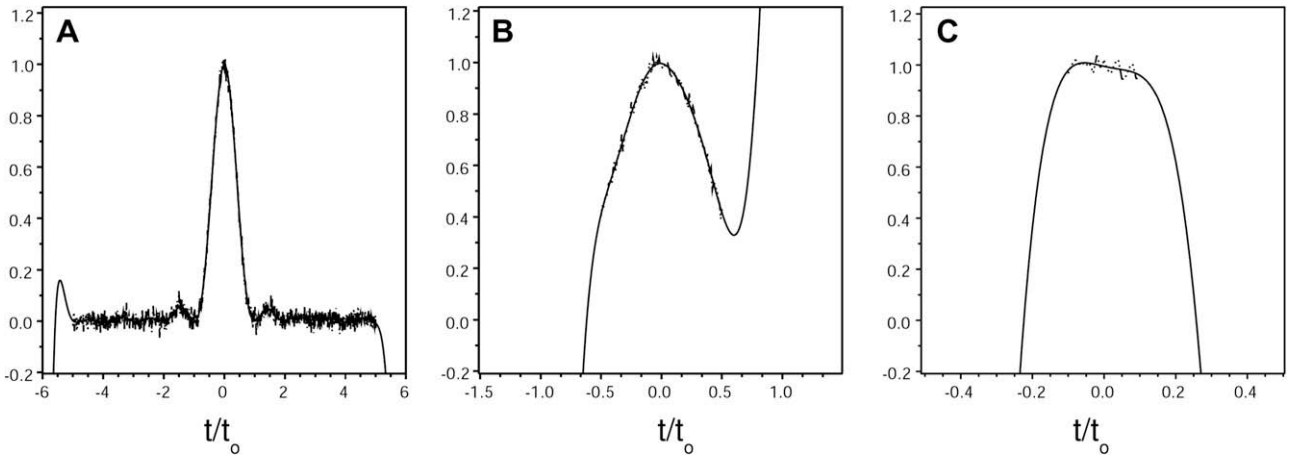

Fig. 8. Signals of Fig. 5 before (dotted line) and after (continuous line) they were processed with a cardinal series filter. The pass band width of the filter was set equal to the band width $\Omega_o = 4\pi/t_o$ of the signals.

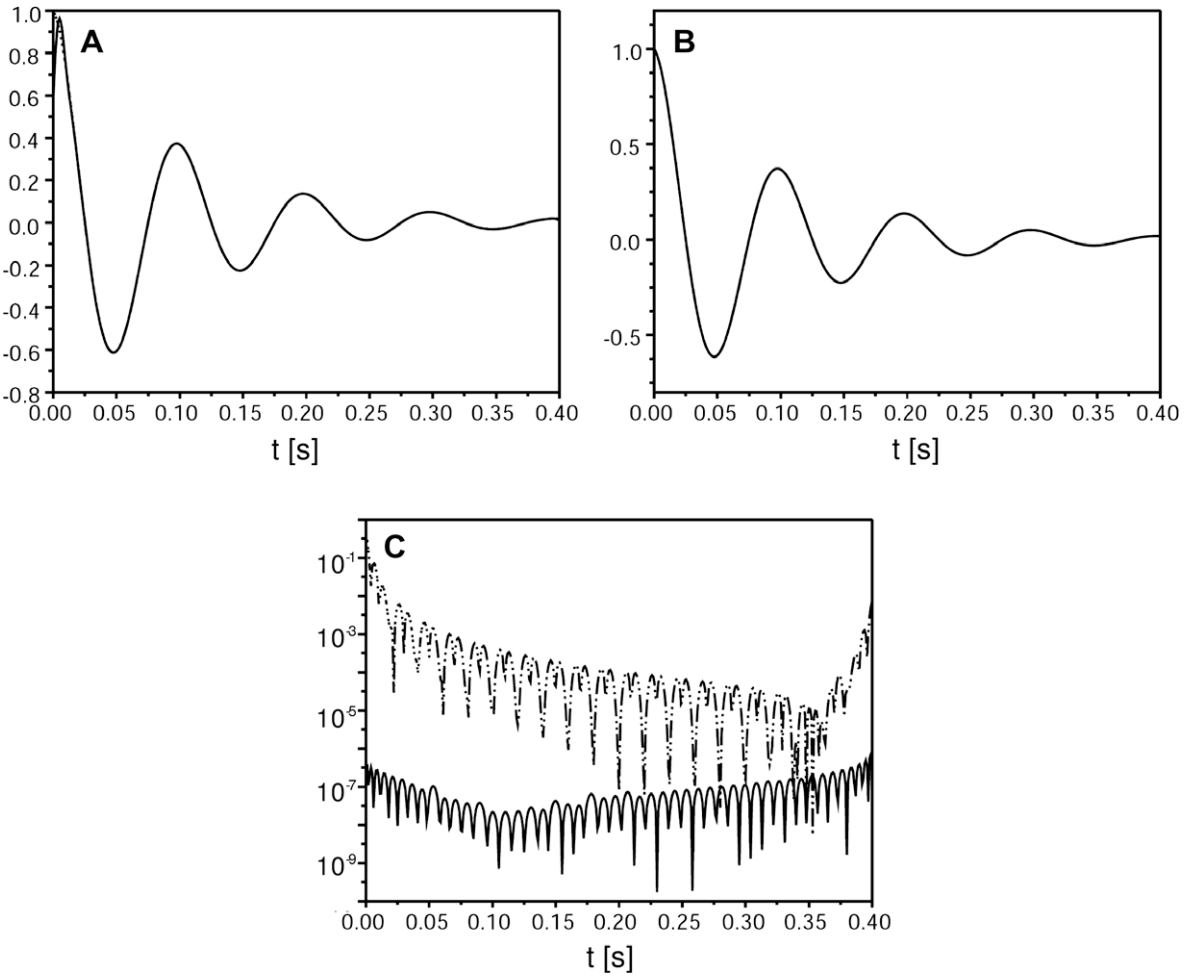

Fig. 9. (A) Noise-free FID-like signal modelled as the function in Eq. (35) before (dashed line) and after (continuous line) it was processed with a Lanczos filter. (B) Noise-free FID-like signal modelled as the function in Eq. (35) before (dashed line) and after (continuous line) it was processed with a cardinal series filter; (C) absolute value of the difference between the noise-free FID-like signal modelled as the function in Eq. (35) and the output signal of the Lanczos (dotted line) and cardinal series (solid line) filter. The pass band width of both filters was set to 100 Hz.

$$x(t) = sinc^2\left(\pi\frac{t}{t_o}\right), \quad \text{where } t_o \text{ stands for an arbitrary time unit,} \tag{34}$$

sampled at regular intervals $\delta t = t_o/100$ within (a) $-5t_o \leqslant t \leqslant 5t_o$, (b) $-0.5t_o \leqslant t \leqslant 0.5t_o$ and (c) $-0.1t_o \leqslant t \leqslant 0.1t_o$ and impaired with computer-simulated white normally distributed noise of standard deviation *0.02*. The filter pass band width was set equal to the band width $\Omega_o = 4\pi/t_o$ of the signals to filter. Applying the filter to the signal for which $t_N - t_1 = 20 \cdot (2\pi/\Omega_o)$ results in a significantly cleaner signal (continuous line in Fig. 5A) that preserves all the features of the original data set. Applied to the truncated signal of Fig. 5B, for which $t_N - t_1 = 2 \cdot (2\pi/\Omega_o)$, the filter gives a cleaner signal (continuous line in Fig. 5B) that differs, however, markedly from the original data set.

Applying the same filter to the even more truncated signal for which $t_N - t_1 = 0.4 \cdot (2\pi/\Omega_o)$, gives a clean, albeit completely disfigured signal (continuous line in Fig. 5C).

Convolution filters with fancier spectra, e.g. cardinal sine or Gaussian functions, do hardly better [24]. The filter based on using truncated cardinal series, contrariwise, will prove (see Fig. 8) suitable for processing truncated signals.

Once all the parameters were set as described above, i.e. $\Omega' = \Omega'' = 1.1\Omega_o$ and $R = 6$, we examined how the performance of the filter is affected by the size of data sets. Fig. 6A and B show two-dimensional charts of the $|s(0) - 1|$ values calculated for the data sets of Eqs. 28 and 29 as a function of the lower t_1 and upper t_N boundaries of the sampling area. The calculations were carried out for 51 values of either parameter independently and incrementally varied from $-5t_o$ to *0* and from *0* to $5t_o$, respectively. Whatever the size of data sets, the inexactitude is under 10^{-6}. Such an inexactitude can be safely ignored when it deals with processing NMR data. Fig. 6C shows a three-dimensional diagram of the noise reduction. When both $-t_1$ and t_N increase, the noise reduction σ'/σ tends quickly to 0.105, which corresponds to the value predicted by Eq. (11) for the convolution filter with the cardinal sine kernel. It should be pointed out that the expression on the right hand side in Eq. (11) gives the exact value of the signal-to-noise gain factor σ/σ' for signals determined on infinite supports and stands for a lower bound of its best achievable limit when it deals with truncated signals. When one or both of t_1 and t_N tend to zero, the noise reduction is still fairly good.

To further contrast the two types of filter, we applied a cardinal series filter to the experimentally acquired proton FID signal of ethanol (see also Fig. 4A). The input and output, shown in Fig. 7A, appear identical to one another everywhere along the time axis and at $t = 0$, where the convolution filter suffered a major

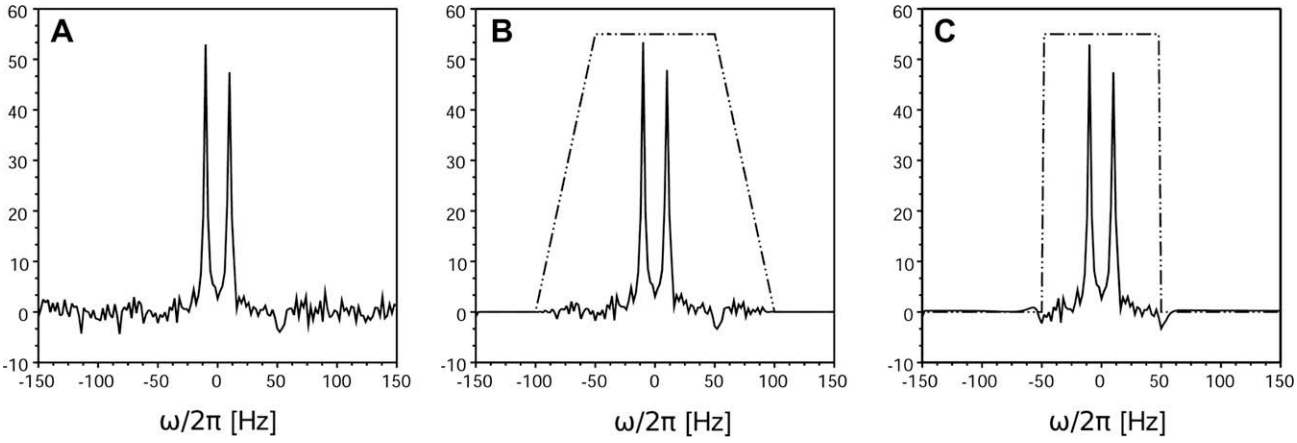

Fig. 10. (A) Uniformly phase-corrected FT of the FID-like signal of Fig. 9A impaired with computer-simulated white normally distributed noise with 0.05 standard deviation; (B) uniformly phase-corrected FT of the noise-impaired signal of Fig. 9A processed with a Lanczos filter; (C) uniformly phase-corrected FT of the noise-impaired signal of Fig. 9A processed with a cardinal series filter. The pass band width of both filters was set to 100 Hz. The dotted line in (B) and (C) indicates the FT in arbitrary units of the Lanczos kernel and cardinal series in Eq. (13), respectively.

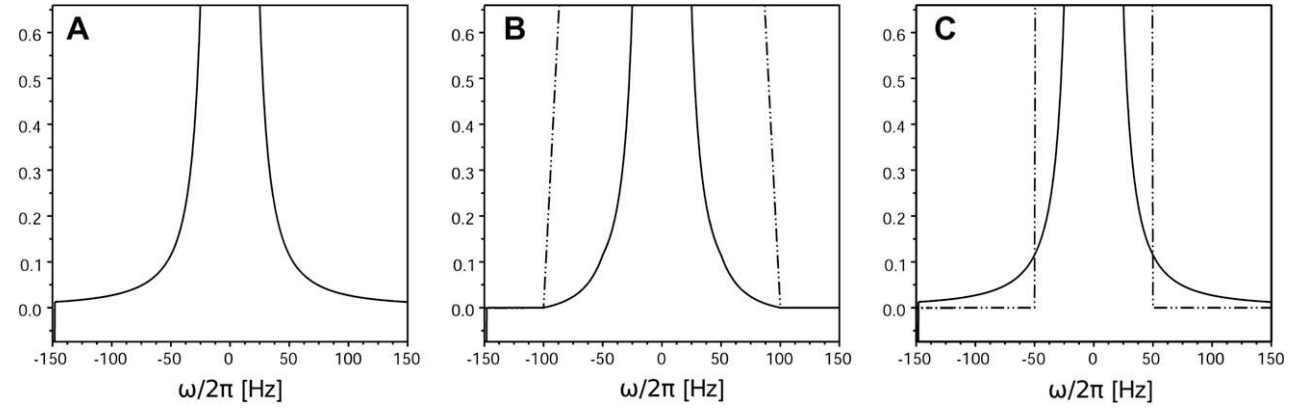

Fig. 11. (A) Uniformly phase-corrected FT of the noise-free signal of Fig. 9A; (B) uniformly phase-corrected FT of the noise-free signal of Fig. 9A processed with a Lanczos filter; (C) uniformly phase-corrected FT of the noise-free signal of Fig. 9A processed with a cardinal series filter. The pass band width of both filters was set to 100 Hz. The dotted line in (B) and (C) indicates the FT in arbitrary units of the Lanczos kernel and cardinal series in Eq. (13), respectively.

set-back, in particular. Thus the cardinal series filter proves much more tolerant to signal discontinuities. And so do their FT, shown with a dashed and continuous line, respectively, in Fig. 7B.

We also applied a cardinal series filter to the computer-simulated truncated echo-like data sets of Fig. 5A–C. Yet again not only do the filter output data sets, shown in Fig. 8A–C respectively, have higher signal-to-noise ratios but also preserve all the features of the original signal for all three input data sets.

To investigate in more detail how the cardinal series filter performs being applied to signals that can only approximately be viewed as band-limited, we applied it to a data set modelled as an exponentially damped simple harmonic oscillation

$$x(t) = \cos(\omega t) \cdot \exp\left(-\frac{t}{T_2}\right), \tag{35}$$

with a angular frequency $\omega/2\pi = 10$ Hz and decay constant $T_2 = 100$ ms sampled in an interval $[t_1, t_N]$ $t_1 = 0$, $t_N = 400$ ms at a sampling frequency 1 kHz. This signal was processed with the Lanczos (Fig. 9A) and cardinal series (Fig. 9B) filter with pass band width set to 100 Hz. The FT of such a noise-free FIDlike signal is an, infinitely wide, Lorentzian. In practice, though, only a very limited a part of the whole spectrum will be of interest. Fig. 9C shows the absolute value of the difference between the input and output signals of the Lanczos (dotted line) and cardinal series (solid line) filter. The error introduced by the Lanczos filter varies greatly as a function of time, is particularly large—about 40 per cent—at $t = 0$ and stays less than 10^{-3}–10^{-4} for a major of the rest of the sampling interval. The error introduced by the cardinal series filter is, contrariwise, rather uniform and is less than 10^{-6}–10^{-7} for the whole sampling interval. The precision of the cardinal series filter is superior not only to that of the Lanczos filter but also to the best mathematical estimations yet reported [22,23] of the error introduced in approximating the signal with truncated cardinal series; and may be accounted for by the fact that the coefficients a_m were left unspecified and sought for to only ensure the best possible agreement between the cardinal series and signal to filter $x(t)$, rather than identifying them with the signal samples $x(\tau_m)$.

Computer-simulated white normally distributed noise of standard deviation 0.05 was added to the signal (dashed line of Fig. 9A), whose FT is shown in Fig. 10A, and processed with the Lanczos and cardinal series filter. The FT of the output signals of either filter, shown in Fig. 10B and C, respectively, appear identical to the input signal within the pass band width of the filters. Thus both types of filter preserve the noise-impaired signal within their pass band, i.e. within $|\omega/2\pi| \leqslant 50$ Hz. The Lanczos filter has a fine roll-off and, thus, only attenuates noise within $50\text{ Hz} \leqslant |\omega/2\pi| \leqslant 100$ Hz. The cardinal series filter, contrariwise, appears, at least at first sight, to have an infinitely narrow roll-off and eliminates noise whose frequencies exceed the spectral band width of the truncated cardinal sine function that constitute the series of Eq. (13), i.e. $|\omega/2\pi| \leqslant 50$ Hz.; and thus performs very much as would the convolution filter with cardinal sine kernel and bass band width of 100 Hz, should it have been applicable to truncated signals.

Nevertheless, a closer look at the processed noise-free signals, shown in Fig. 11B and C, respectively, reveals a subtle difference between them. While the output signal of the Lanczos filter has a perfectly flat base line outside the interval $|\omega/2\pi| \leqslant 100$ Hz, that of the cardinal series filter preserved the wings of the input signal, shown in Fig. 11A, in the intervals $|\omega/2\pi| \geqslant 50$ Hz. This suggests to us that high frequency noise in Fig. 10C was indeed removed whilst a true high frequency components of the signal—resulted from its truncation and essential for a proper description of the signal at $t = 0$ in the time domain—was preserved by the filter.

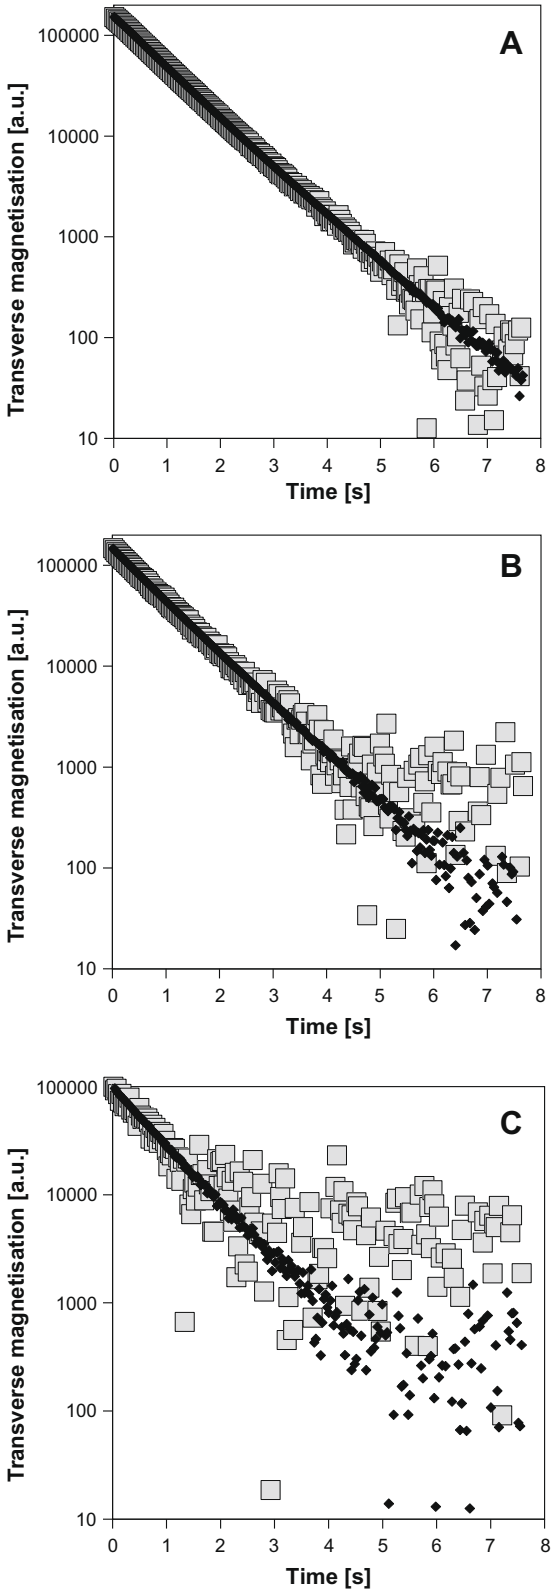

Fig. 12. Signal amplitudes, before (bigger grey squares) and after (smaller black diamonds) an application of the cardinal series filter to echoes acquired in CPMG-experiments on (A) 50-, (B) 5- and (C) 0.5-gramme samples prepared as described in part 5.

3.3. Application of the cardinal series filter for processing experimental NMR signals

The cardinal series filter was applied to data sets obtained in CPMG experiments on a proton system in 50, 5 and 0.5 gramme samples of a blend described in Experimental to rectify the amplitude of spin echoes summits—the only quantity of interest that could be extracted from the signal in the time domain. The band width $\Omega_o/2\pi$ of the spectrum was found to be about 300 Hz, whereas the narrowest analogue filter had a pass band width $\Omega_f/2\pi = 100$ kHz. The acquisition interval was deliberately set to 3 ms, i.e. $t_N - t_1 = 0.9 \cdot (2\pi/\Omega_o)$, to truncate the signal. The ADC sampling frequency was set to 100 kHz, to oversample it.

Fig. 12A–C show the experimentally measured (big grey squares) and filtered (small black squares) amplitudes of the spin echoes maxima for the three above mentioned samples, respectively, which corresponded to the signal-to-noise ratio varying from decent, in Fig. 12A, to quite poor, in Fig. 12C.

In particular, Fig. 12C shows that the signal-to-noise ratio could be raised from 10 for the experimentally determined raw spin echoes amplitudes (big grey squares) to about 150 for those (small black squares) that were obtained by filtering the entire data set as suggested above. This improvement is close to its theoretical maximum of 18.3 (see Eq. 11) under present experimental conditions (see in Section 5).

4. Conclusions

We proposed a digital low pass filter for truncated band-limited NMR signals. It consists in fitting raw data in the time domain with a finite series of truncated cardinal sine functions. In contrast with those widely used so far, our filter is essentially model-free, requires no prior knowledge about the system under study and preserves all the features of the signal in both the time and frequency domain.

Sensible rules were derived for setting the parameters of the filter.

The filter is applied after acquisition of the oversampled signal was completed, thus requiring more random access memory (RAM) than causal filters. Nevertheless, its performance is such that the acquisition interval can deliberately be shortened, should there be shortage of RAM.

5. Experimental

All numerical simulations, calculations and data-processing algorithms were coded in Fortran 95 programming language [25,26]. Where necessary, matrices were inverted using LDLT decomposition.

All NMR spectra were acquired at a vertical wide-bore Bruker 24/80 Avance DBX spectrometer equipped with a 20 cm birdcage RF coil and operating at 0.5 Tesla. One-dimensional proton NMR signals of a sample of ethanol were registered after a single 90° excitation pulse. The water proton CPMG experiments were carried out on 50, 5 and 0.5 gramme samples of an emulsion (a blend of CaOH, water and dodecane). In each two-minute experiment 300 complex data points were acquired for each of (odd) 256 collected of 512 generated echoes. The band width $\Omega_o/2\pi$ of the echoes was about 300 Hz. The ADC sampling frequency $\Omega/2\pi$ and analogue filter pass band width $\Omega_f/2\pi$ were both set to 100 kHz.

Acknowledgments

This work was supported by the Agence Nationale de la Recherche (ANR) of France through the ANR-06-JCJC-0106 project. A need for developing algorithms to filter truncated NMR signals first occurred to us in talking to Dr. Paméla Faure and Dr Julie Magat.

References

[1] D. Bytchenkoff, G. Bodenhausen, J. Magn. Reson. 165 (2003) 1–8.
[2] P.J. Prado, B.J. Balcom, S.D. Beyea, T.W. Bremner, R.L. Armstrong, P.E. Grattan-Bellew, Chem. Conc. Res. 28 (2) (1998) 261–270.
[3] P.F. Faure, S. Caré, C. Po, S. Rodts, J. Magn. Reson. Imaging 23 (2005) 311–314.
[4] S. Rodts, J.-C. Baudez, P. Coussot, Europhys. Lett. 69 (4) (2005) 636–642.
[5] P.J. McDonald, J.-P. Korb, J. Mitchell, L. Monteilhet, Phys. Rev. E 72 (2005) 011409.
[6] P.J. McDonald, P.S. Aptaker, J. Mitchell, M. Mulheron, J. Magn. Reson. 185 (2007) 1–11.
[7] B. Blümich, S. Anferova, S. Sharma, A.L. Segre, C. Federici, J. Magn. Reson. 161 (2003) 204–209.
[8] B. Leonard, J. Audio Gen. Soc. 124 (2008) 7398.
[9] J.-B. Poullet, D.M. Sima, S. van Huffel, J. Magn. Reson. 195 (2008) 134–144.
[10] A.S. Glassner (Ed.), Graphics Gems, Academic Press, 1990. pp. 147–165.
[11] J. Blinn Jim, Blinn's Corner: Dirty Pixels, Morgan Kaufman Publishers Inc., 1998. pp. 26–27.
[12] J. Tang, J.R. Norris, J. Magn. Reson. 69 (1986) 180–186.
[13] J. Tang, J.R. Norris, J. Magn. Reson. 79 (1988) 190–196.
[14] Y. Zeng, J. Tang, A. Bush, J.R. Norris, J. Magn. Reson. 83 (1989) 473–483.
[15] A. Guillermo, J.-P. Cohen Addad, D. Bytchenkoff, J. Chem. Phys. 113 (2000) 5098–5106.
[16] V.A. Kotelnikov, On the transmission capacity of 'ether' wire in electrocommunication, Material for the first all-Union conference on questions of communication, Izd. Red. Upr. Svyazzi RKKA, Moscow (1933).
[17] C.E. Shannon, Mathematical theory of communication, Bell Syst. Tech. J. 27 (1948) 379–423. pp. 623–656.
[18] H.T. Nyquist, Trans. AIEE 47 (1928) 617–644.
[19] R.J. Marks II, Introduction to Shannon sampling and interpolation theory, Springer-Verlag, New York, 1991.
[20] J. Idier (sous la direction) Approche bayésienne pour les problèmes inverses, Hermès Science Publications, Paris (2001) G.L. Bretthorst, J. Magn. Reson. 88 (1990) 533–551.
[21] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical recipes in Fortran—The art of scientific computing, Cambridge University Press, 1986.
[22] F.J. Beutler, IEEE Trans. Inform. Theory 22 (5) (1976) 568–573.
[23] A. Ya. Olenko, T.K. Pogány, J. Math. Anal. Appl. 324 (2006) 262–280.
[24] D. Bytchenkoff, S. Rodts, P. Moucheront, T. Fen-Chong, J. Magn. Reson. 202 (2010) 147–154.
[25] International standards organisation, ISO/IEC 1539: 1997, Information technology—Programming languages—Fortran (Fortran 95) Geneva (1997).
[26] J.C. Adams, W.S. Brainerd, J.T. Martin, B.T. Smith, J.L. Wagener, The Fortran 95 Handbook—Complete ISO/ANSI Reference, MIT Press, Cambridge, Massachusetts, 1997.