Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Magnetic Resonance 162 (2003) 176–188
JMR Journal of Magnetic Resonance
www.elsevier.com/locate/jmr

A statistical analysis of NMR spectrometer noise

Halfdan Grage[a,*] and Mikael Akke[b]

[a] Mathematical Statistics, Centre for Mathematical Sciences, Lund University, Box 118, Lund SE-221 00, Sweden
[b] Department of Biophysical Chemistry, Lund University, Box 124, Lund SE-221 00, Sweden

Received 28 August 2002; revised 13 January 2003

Abstract

Estimation of NMR spectral parameters, using e.g. maximum likelihood methods, is commonly based on the assumption of white complex Gaussian noise in the signal obtained by quadrature detection. Here we present a statistical analysis with the purpose of discussing and testing the validity of this fundamental assumption. Theoretical expressions are derived for the correlation structure of the noise under various conditions, showing that in general the noise in the sampled signal is not strictly white, even if the thermal noise in the receiver steps prior to digitisation can be characterised as white Gaussian noise. It is shown that the noise correlation properties depend on the ratio between the sampling frequency and the filter cut-off frequency, as well as the filter characteristics. The theoretical analysis identifies conditions that are expected to yield non-white noise in the sampled signal. Extensive statistical characterisation of experimental noise confirms the theoretical predictions. The statistical methods outlined here are also useful for residual analysis in connection with validation of the model and the parameter estimates.

Keywords: Correlation analysis; Frequency analysis; Gaussian noise; Maximum likelihood estimation; Statistical modelling; Residual analysis

1. Introduction

The statistical properties of the noise in a given signal have important implications for modelling of the signal and for estimation of its spectral parameters. As an example, the maximum likelihood (ML) method explicitly incorporates in the model the statistical distribution of the observations. In estimating model parameters, all relevant available information must be used—not only because this yields more precise and accurate parameter estimates [1], but also because arbitrary exclusion of some information would permit equally arbitrary conclusions to be drawn [2]. Therefore it is essential to include in the model any available information about the noise of the measured signal. ML methods have been applied successfully to estimate NMR spectral parameters [3–11], and is an alternative to Fourier transformation of the time domain data [1,12]. The complexity of the model required to describe the NMR signal is directly dependent on the statistical properties of the measurement noise, as outlined below (cf. Eqs. (2) and (3)). To the best of our knowledge, the literature on ML estimation of NMR spectra exclusively involves models based on the assumption that the complex signal includes white complex Gaussian noise. More specifically, the assumption is that the following conditions on the noise sequences in the real and imaginary channels are fulfilled: they are mutually uncorrelated; each of them is uncorrelated; they have equal variance; and they have a joint Gaussian distribution. The validity of this fundamental assumption is apparently taken for granted, and is very rarely discussed [4]. Statistical models appropriate for ML and non-linear least-squares estimation in cases where the noise is coloured have previously been considered briefly [11]. Also, Bayesian approaches have been presented that may allow the noise probability distribution to be quantified and included in the analysis in a more formal manner [13–15]. However, the notion that the noise is uncorrelated clearly dominates the literature.

In the field of MRI, the extent of noise correlation between multiple receiving coils has been addressed

* Corresponding author. Fax: +46-0-46-2224623.
E-mail addresses: Halfdan.Grage@matstat.lu.se, halfdan@maths.lth.se (H. Grage).

1090-7807/03/$ - see front matter © 2003 Elsevier Science (USA). All rights reserved.
doi:10.1016/S1090-7807(03)00038-7

[16,17], and a systematic empirical investigation of the statistical properties of the noise has been reported for this case [18]. High-resolution NMR spectroscopy is different from MRI in that the normal NMR probe configuration includes a single receiving coil, although a design using two orthogonal receiving coils has also been described [19]. Here, we limit our analysis to the case of a single receiving coil, for which quadrature detection is achieved by splitting the signal into two parallel channels (''real'' and ''imaginary''), and mixing these with two reference signals that are orthogonal in phase, but otherwise identical [20]. Since the same noise is given as input to the real and imaginary channels, it may not be immediately obvious that the noise sequences in the two channels should be uncorrelated after mixing with the high-frequency reference signal and subsequent low-pass filtering and sampling. Furthermore, it is not clear that the assumptions about the correlation properties of the noise will hold throughout the signal processing scheme outlined above, even if the noise in the signal emanating from the receiving coil may be adequately represented by a Gaussian distribution. Here we address the validity of these fundamental assumptions in some depth. Our analysis addresses the situation where the noise input is stationary. The complementary problem of the transient response to a noiseless signal has been analysed previously [21].

The interferogram is commonly modelled as a sum of exponentially damped complex sinusoids with additive Gaussian white noise:

$$S_n = m_n + Y_n = \sum_{k=1}^{p} A_k \mathrm{e}^{\mathrm{i}\phi_k} \mathrm{e}^{(-\alpha_k + \mathrm{i}2\pi f_k)t_n} + Y_n, \quad n = 1, 2, \ldots, N, \tag{1}$$

where $A_k > 0$ is the amplitude, ϕ_k the phase, $\alpha_k \geqslant 0$ the damping factor, and f_k the frequency of sinusoid number k. The variables Y_n are independent, identically distributed (IID) complex Gaussian random variables with zero mean and uncorrelated real and imaginary parts, $\mathrm{Cov}[\mathrm{Re}\, Y_n, \mathrm{Im}\, Y_n] = 0$, each with variance σ^2, so that $\mathrm{Var}[Y_n] = \mathrm{E}|Y_n|^2 = 2\sigma^2$, where Cov, Var, and E denote the covariance, the variance, and the mathematical expectation, respectively. Quite generally, if the noise sequence $\mathbf{Y} = (Y_1, \ldots, Y_N)^{\mathrm{T}}$, where the superscript T denotes matrix transposition, is complex Gaussian with mean $\mathbf{m}$ and covariance matrix $\mathbf{C} = \mathrm{E}[(\mathbf{Y} - \mathbf{m})(\mathbf{Y} - \mathbf{m})^H]$, the likelihood function can be written as

$$L(\mathbf{y}; \mathbf{m}, \mathbf{C}) = \frac{1}{\pi^N \det \mathbf{C}} \exp[-(\mathbf{y} - \mathbf{m})^H \mathbf{C}^{-1} (\mathbf{y} - \mathbf{m})], \tag{2}$$

where $\mathbf{y} = (y_1, \ldots, y_N)^{\mathrm{T}}$ is the sequence of experimental data points, the superscript H denotes the Hermitian transpose of a matrix, and $\det\mathbf{C}$ denotes the determinant of the matrix $\mathbf{C}$. Irrespective of the time-correlation properties of the noise, this likelihood model implies that the real and imaginary parts of Y_n are uncorrelated, that is $\mathrm{Cov}[\mathrm{Re}\, Y_n, \mathrm{Im}\, Y_n] = 0$ and $\mathrm{Cov}[\mathrm{Re}\, Y_m, \mathrm{Re}\, Y_n] = \mathrm{Cov}[\mathrm{Im} Y_m, \mathrm{Im} Y_n]$ for all m and n; in particular, this means that $\mathrm{Var}[\mathrm{Re}\, Y_n] = \mathrm{Var}[\mathrm{Im}\, Y_n]$ for all n [22–26]. This compact formulation of an N-dimensional complex Gaussian distribution with an Hermitian covariance matrix thus inherently restricts the applicability of the model to situations where the noise has this particular correlation structure [25,27]. In contrast, a $2N$-dimensional Gaussian distribution with real variables would allow any covariance structure of the noise to be modelled. The assumption that the noise is uncorrelated, i.e., $\mathbf{C} = 2\sigma^2\mathbf{I}$, further simplifies the likelihood function (2) to the familiar form

$$L(\mathbf{y}; \mathbf{m}, \sigma^2) = \frac{1}{\sigma^{2N}(2\pi)^N} \exp\left[-\frac{1}{2\sigma^2} \sum_{j=1}^{N} |y_j - m_j|^2\right]. \tag{3}$$

This simplified likelihood function provides the basis for the most commonly encountered ML algorithms used for estimation of model parameters in NMR spectroscopy, such as IQML [8,9], EM [4,5], FML [6], RELAX [10], and AMARES [7] (see also the review articles [1,12]). The fundamental assumptions regarding the statistical properties of the noise that lead to Eq. (3) are evaluated in this paper.

In the following, we first derive the covariance functions of the noise (in a system with a single receiving coil) before and after low-pass filtering, and study their dependence on the subsequent digital sampling of the analog signal (Section 2). Next, we present an experimental study of the statistical properties of measured NMR spectrometer noise (Section 3). The statistical methods used in Section 3 and outlined in Appendix A are also suitable for residual analysis during validation of the model used for estimation of the NMR parameters. These methods should be equally useful for empirical characterisation of noise in systems with multiple receiving coils, e.g., in MRI.

2. Theoretical analysis

The derivations presented below are based on a quadrature detection scheme, as described in [20]. After splitting the signal into two parallel channels, the real and imaginary parts of the noise will be perfectly correlated, since the splitting merely means that we have the signal $X(t) + \mathrm{i}X(t)$, where $X(t)$ is the thermal noise generated in the sample, coil, and preamplifier [28]. These are the primary sources of noise. Typically, the noise introduced by the remaining receiver components, e.g., the analog-to-digital converter, are minor in comparison [28,29]—although the situation may be different with cryogenic probes and very high static magnetic field strengths. Following cosine modulation of the real part and sine modulation of the imaginary part, the noise in continuous time can be written as

$$Y(t) = Y_R(t) + iY_I(t)$$
$$= X(t)\cos(2\pi f_0 t) - iX(t)\sin(2\pi f_0 t),$$

where f_0 denotes the carrier frequency. Assuming initially for the sake of simplicity that $X(t)$ is white noise with the Dirac impulse as covariance function, the modulated complex noise will have covariance function

$$\mathrm{Cov}[Y_R(t), Y_I(s)] = -\cos(2\pi f_0 t)\sin(2\pi f_0 s)\delta(s-t)$$
$$= -\frac{1}{2}\sin(4\pi f_0 t)\delta(s-t),$$

so that $\mathrm{Cov}[Y_R(t), Y_I(t)]$ oscillates as a sinusoid with frequency $2f_0$ (i.e., twice the carrier frequency), and the corresponding cross-correlation

$$\mathrm{Corr}[Y_R(t), Y_I(t)] = \mathrm{Cov}[Y_R(t), Y_I(t)]/\sqrt{\mathrm{Var}[Y_R(t)]\mathrm{Var}[Y_I(t)]}$$

will oscillate between +1 and −1 with the same frequency, as illustrated in Fig. 1. At this point in the analysis, it does not seem obvious that the real and imaginary parts of the noise are uncorrelated in either of the likelihood models (2) and (3). This assumption will have to be justified by an analysis of the effects of the subsequent filtering and sampling steps, as outlined next.

To see what happens to the noise in the real channel, where the signal is modulated by a cosine wave, we consider the more general case where $X(t)$ is stationary real-valued noise (but not necessarily Gaussian or white) with mean zero, autocovariance function $r_X(\tau) = \mathrm{Cov}[X(t), X(t+\tau)]$, and spectral density $R_X(f) = \int_{-\infty}^{\infty} r_X(\tau)e^{-i2\pi f\tau}\,d\tau$. Further, suppose that the audio filter has the real-valued impulse response $h(t)$ with transfer function $H(f)$, so that the output of the filter can be written as [27,30,31]

$$Z_R(t) = \int_{-\infty}^{\infty} h(\xi)Y_R(t-\xi)\,d\xi$$
$$= \int_{-\infty}^{\infty} h(\xi)\cos[2\pi f_0(t-\xi)]X(t-\xi)\,d\xi$$

which is a process with zero mean, so that the autocovariance function simplifies to $r_{Z_R}(t,s) = \mathrm{E}[Z_R(t)Z_R(s)]$, which, following the treatment outlined in [30], can be written as

$$r_{Z_R}(t,s) = \frac{1}{4}\int_{-\infty}^{\infty}[M^*(f,t) + M(-f,t)][M(f,s) + M^*(-f,s)]R_X(f)\,df, \quad (4)$$

where $M(f,t) \equiv e^{i2\pi(f-f_0)t}H(f-f_0)$, and the asterisk denotes complex conjugation. If $X(t)$ is a Gaussian process, then the output process $Z_R(t)$ will also be Gaussian, but since the input noise $Y_R(t)$ is not stationary we can, in general, not expect the output noise $Z_R(t)$ to be stationary. Expanding the part of the integrand that contains the function M, we obtain

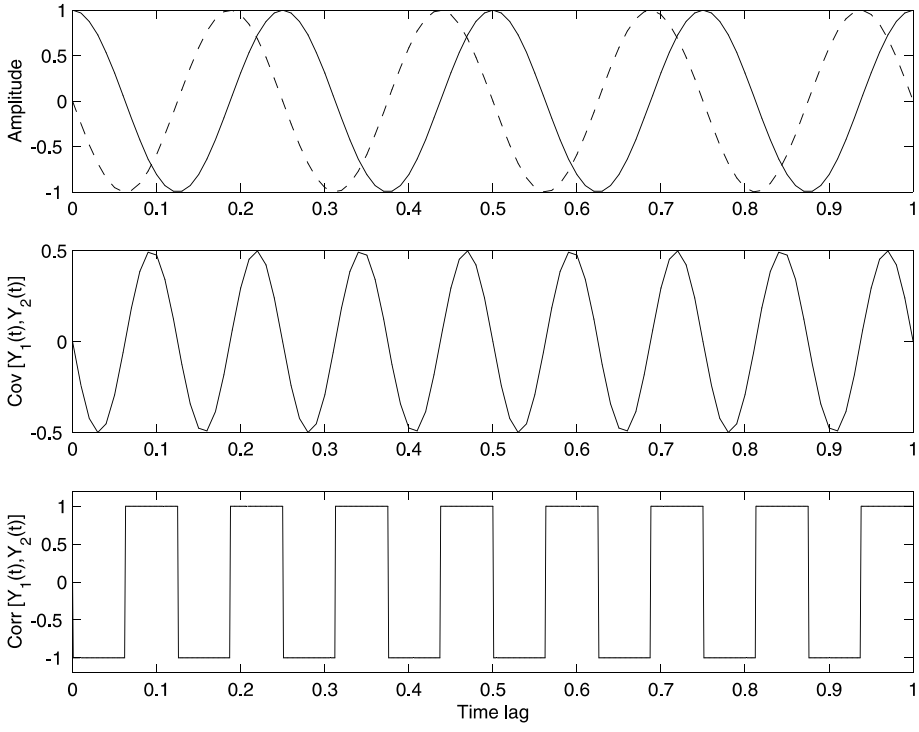

Fig. 1. The effect of amplitude modulation of the incoming noise on its correlation properties. Top: the modulating cosine (—) in the real channel and the corresponding modulating negative sine in the imaginary channel (– – –). Middle: the cross-covariance $\mathrm{Cov}[Y_R(t), Y_I(t)]$, where Y_R and Y_I denote the noise in the real and imaginary channels, respectively. Bottom: the cross-correlation $\mathrm{Corr}[Y_R(t), Y_I(t)]$.

$$
\begin{aligned}
&[M^*(f,t)+M(-f,t)][M(f,s)+M^*(-f,s)]\\
&\quad = \mathrm{e}^{\mathrm{i}2\pi(f-f_0)(s-t)}|H(f-f_0)|^2 + \mathrm{e}^{\mathrm{i}2\pi(f+f_0)(s-t)}|H(f+f_0)|^2\\
&\qquad + \mathrm{e}^{\mathrm{i}2\pi[f(s-t)+f_0(s+t)]}H^*(f-f_0)H(f+f_0)\\
&\qquad + \mathrm{e}^{\mathrm{i}2\pi[f(s-t)-f_0(s+t)]}H^*(f+f_0)H(f-f_0),
\end{aligned}
\tag{5}
$$

which consists of a (weakly) stationary part (the first two terms) and, if $f_0 \neq 0$, a non-stationary part (the last two terms). In the case that $f_0 = 0$, all terms are equal to $|H(f)|^2\mathrm{e}^{\mathrm{i}2\pi f(s-t)}$, which corresponds to the usual expression for filtered stationary noise [30]. However, in NMR spectroscopy f_0 is never zero, but typically in the radio-frequency range.

Assume that the filter is such that $H(f) = 0$ when $|f| > f_{\mathrm{sc}}$, where f_{sc} is the stop-band cut-off frequency of the filter [32,33], with $f_{\mathrm{sc}} < f_0$; typically f_{sc} is in the audio range. Then

$$H^*(f-f_0)H(f+f_0) = H^*(f+f_0)H(f-f_0) = 0$$

for all frequencies f, so that only the stationary part passes through the filter. The integral in (4) can then be simplified (by changing variables and using the fact that $|H(f)|^2$ and $R_X(f_0-f)+R_X(f_0+f)$ are even functions of f, while $\sin[2\pi f(s-t)]$ is an odd function of f) to yield

$$
\begin{aligned}
r_{Z_{\mathrm{R}}}(t,s) = \frac{1}{4}\int_{-f_{\mathrm{sc}}}^{f_{\mathrm{sc}}} &\cos[2\pi f(s-t)]|H(f)|^2[R_X(f_0-f)\\
&+ R_X(f_0+f)]\,\mathrm{d}f.
\end{aligned}
\tag{6}
$$

Let us make the further assumptions that the noise $X(t)$ is white and band-limited with bandwidth $f_w \leqslant f_0$ (corresponding to the bandwidth of the coil resonance circuit), and has a spectral density which is symmetric about f_0,

$$R_X(f) = \begin{cases} R_0, & f_0 - f_w \leqslant |f| \leqslant f_0 + f_w,\\ 0, & \text{otherwise.}\end{cases} \tag{7}$$

If it holds that $f_w > f_{\mathrm{sc}}$ (i.e., the bandwidth of the noise is comparably large; this requirement may not be met in all cases, e.g., with high-Q probes and large oversampling ratios), so that $R_X(f_0+f) = R_X(f_0-f) = R_0$ for $|f| < f_{\mathrm{sc}}$, then the integral in (6) can be simplified to

$$r_{Z_{\mathrm{R}}}(t,s) = \frac{R_0}{2}\int_{-f_{\mathrm{sc}}}^{f_{\mathrm{sc}}} \cos[2\pi f(s-t)]|H(f)|^2\,\mathrm{d}f. \tag{8}$$

If we assume that the audio filter is ideal, with impulse response $h(t)$ and transfer function

$$H(f) = \begin{cases} H_0, & |f| \leqslant f_{\mathrm{sc}},\\ 0, & |f| > f_{\mathrm{sc}},\end{cases} \tag{9}$$

then we can easily solve the integral in (8) to get

$$
\begin{aligned}
r_{Z_{\mathrm{R}}}(t,s) &= \frac{R_0|H_0|^2}{2}\int_{-f_{\mathrm{sc}}}^{f_{sc}} \cos[2\pi f(s-t)]\,\mathrm{d}f\\
&= R_0|H_0|^2\frac{\sin[2\pi f_{\mathrm{sc}}(s-t)]}{2\pi(s-t)}.
\end{aligned}
\tag{10}
$$

This is the autocovariance function of the noise in continuous time. The autocovariance of the sampled noise is obtained by inserting the discrete values of s and t as determined by the sampling scheme. Since the time lags $s-t$ between sampling times depend on the sampling frequency, this means that the correlation structure of the sampled signal will depend on the sampling frequency (see Fig. 2). For example, if the sampling frequency is equal to twice the filter cut-off frequency, $f_s = 2f_{\mathrm{sc}}$, then the autocorrelations of the sampled signal will all be zero (except at lag zero, where the value is one). In contrast, deviations from this condition are expected to introduce a non-negligible correlation structure in the noise.

A common assumption in NMR spectroscopy is that the noise $X(t)$ is thermal white Gaussian noise (so-called Johnson noise) [28,34] with spectral density

$$R_0 = 2k_{\mathrm{B}}T\Re, \tag{11}$$

where k_{B} is the Boltzmann constant, T is the absolute temperature, and $\Re$ is the resistance of the coil. If the values of T, $\Re$, and H_0 are known, the theoretical autocovariance function (10) can be calculated explicitly. In particular, the variance of the noise will be

$$\sigma^2_{Z_{\mathrm{R}}} = 2k_{\mathrm{B}}T\Re|H_0|^2 f_{\mathrm{sc}}. \tag{12}$$

A similar analysis as above shows that the sine-modulated noise in the imaginary channel, $Y_{\mathrm{I}}(t) = -\sin(2\pi f_0 t)X(t)$, has exactly the same autocovariance function as the cosine-modulated real part. This is expected because the input noise $X(t)$ is stationary, and the two modulating waves differ only in phase.

As an example of a non-ideal filter, we consider the Butterworth filter, for which explicit formulae are still manageable. The frequency response of a qth order Butterworth low-pass filter is defined by the squared magnitude transfer function [32,33]

$$|H(f)|^2 = H(f)H(-f) = \frac{1}{1+(f/f_{\mathrm{hc}})^{2q}}, \tag{13}$$

where f_{hc} is the 3-dB cut-off frequency, at which $|H(f)|^2$ has reached half its original value. Thus, the Butterworth filter is completely specified by the filter order and the 3-dB cut-off frequency (see Fig. 2 for an illustration). The $2q$ poles of the analytic continuation of (13) into the complex plane are given by [33]

$$s_k = 2\pi f_{\mathrm{hc}}\mathrm{e}^{\mathrm{i}\pi(q+2k+1)/2q}, \quad k = 0,1,\ldots,2q-1. \tag{14}$$

Using residue calculus to evaluate the integral (8), and neglecting the contributions from frequencies in the stop-band (since for the Butterworth filter the condition that $H(f) = 0$ when $|f| > f_{\mathrm{sc}}$ does not hold exactly), the autocovariance function of the cosine-modulated white noise filtered through a Butterworth filter of order q and 3-dB cut-off frequency f_{hc} can be found to be approximately

$$r_{Z_{\mathrm{R}}}(t,s) = (-1)^q\pi R_0 f_{\mathrm{hc}}\sum_{k=0}^{q-1}\mathrm{e}^{s_k(s-t)}\prod_{j=0,j\neq k}^{2q-1}\frac{2\pi f_{\mathrm{hc}}}{s_k - s_j}. \tag{15}$$

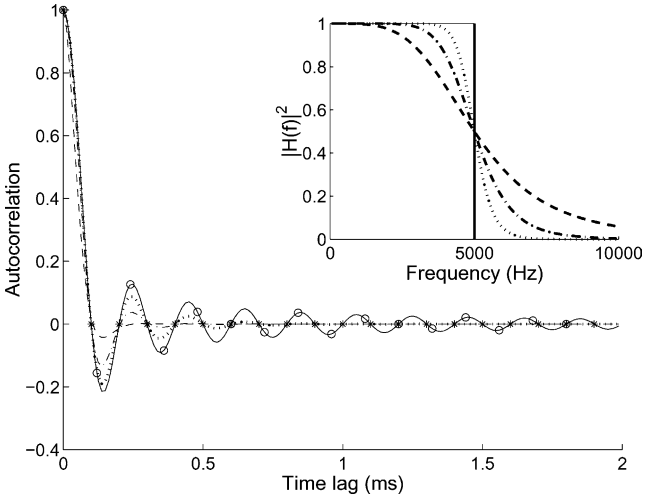

Fig. 2. The theoretical autocorrelation function when the filter 3 dB cut-off frequency is $f_{hc} = 5$ kHz (for the ideal filter this is simply the cut-off frequency). For the ideal filter (—), values taken when the sampling frequency is equal to twice the filter cut-off frequency, $f_s = 2f_{sc}$, are denoted by $*$. Commonly, the filter width is set to be 10–20% larger than the spectral width, and for the case with $f_s = 2f_{sc}/1.2$, denoted by $\circ$, we have almost a "worst case scenario." Also shown are the autocorrelation functions of white noise filtered through Butterworth filters of order two (– –), four (·—), and eight (···). Inlay: squared magnitude of transfer functions for the same filters.

This can also be expressed as a sum of exponentially damped sinusoids with damping ratio $2\pi f_{hc}\cos[\pi(q+2k+1)/2q]$ and frequency $2\pi f_{hc}\sin[\pi(q+2k+1)/2q]$, where the index k runs from zero to the nearest integer larger than or equal to $q/2-1$. This is in accord with other cases where a stochastic process is represented as filtered white noise [23,35]. Although the zeros of the autocovariance function of the output noise from a Butterworth filter do not coincide with the zeros of the output noise from the ideal filter, the difference seems to decrease with increasing filter order, and will be negligible (at least for most practical purposes) already for $N = 4$, as shown in Fig. 2. Similarly, the difference between the autocovariance functions of an ideal filter and a quasi-elliptical 8-pole filter (see Section 3.1) is expected to be small.

Treating the cross-covariance function along the same lines, and letting ψ denote the difference in phase between the modulating sinusoids in the two channels, the expression corresponding to (4) becomes

$$r_{Z_R,Z_I}(t,s) = \frac{1}{4}\int_{-\infty}^{\infty}[M^*(f,t)+M(-f,t)][e^{-i\psi}M(f,s) + e^{i\psi}M^*(-f,s)]R_X(f)\,df. \quad (16)$$

Under the same assumptions concerning the filter as those made to arrive at Eq. (6), only the stationary part of the noise passes through the filter, and we obtain

$$r_{Z_R,Z_I}(t,s) = \frac{1}{4}\cos\psi\int_{-f_{sc}}^{f_{sc}}\cos[2\pi f(s-t)]|H(f)|^2 \times [R_X(f_0+f)+R_X(f_0-f)]\,df + \frac{1}{4}\sin\psi\int_{-f_{sc}}^{f_{sc}}\sin[2\pi f(s-t)]|H(f)|^2 \times [R_X(f_0+f)-R_X(f_0-f)]\,df. \quad (17)$$

Thus, provided that the spectral density of the noise is symmetric around f_0 in the interval $[f_0-f_{sc}, f_0+f_{sc}]$, which was part of the assumption leading to (8), the second integral vanishes and we get

$$r_{Z_R,Z_I}(t,s) = \frac{1}{2}\cos\psi\int_{-f_{sc}}^{f_{sc}}\cos[2\pi f(s-t)]|H(f)|^2 \times R_X(f_0+f)\,df, \quad (18)$$

which we recognise as a factor $\cos\psi$ times the autocovariance function. In the ideal case of a phase difference $\psi = \pi/2$ between the modulating sinusoids in the two channels also this term vanishes, and the real and imaginary parts of the noise will be uncorrelated at all times: $r_{Z_R,Z_I}(t,s) = 0$, for all t and s. Notably, this is true irrespective of the form of the filter transfer function $H(f)$ within the interval.

Modern high-resolution spectrometers enable the use of digital filters. The additional signal processing steps involved, including oversampling, filtering and subsequent "decimation," require additional consideration beyond what has been outlined above. Unfortunately, the details of the digital filter design on commercial spectrometers are proprietary information, preventing us at this time from performing an in-depth theoretical treatment of the noise response to these filters. Of course, the statistical methods outlined below (Section 3.2 and Appendix A) may still be applied in order to investigate empirically the characteristics of digitally filtered noise.

3. Analysis of experimental noise

3.1. Data acquisition

Experimental noise was acquired at room temperature on a Varian Unity Inova 600 MHz spectrometer, equipped with a conventional inverse probe. No sample was inserted into the probe. The receiver gain was adjusted to optimise the dynamic range for sampling of the noise. Data acquisition was performed by gating the receiver on and sampling a single series of 1024 complex points in quadrature. The carrier frequency f_0 was 599.89 MHz, and the spectral width (i.e., the sampling frequency f_s) was 10,000 Hz. The filter width was either 5000 or 6000 Hz, i.e., the filter 3 dB cut-off was set to either $2f_{hc} = f_s$ or $2f_{hc} = 1.2f_s$, the latter of which corresponds to commonly used settings in routine spec-

trometer useage. A conventional analog 8-pole quasi-elliptical filter was used (the standard analog filter on this type of spectrometer).

3.2. Statistical methods

A large number of tests is available for testing hypotheses concerning correlation, independence, and normality. The tests used in the following data analysis (see Appendix A) have been chosen with the aim of providing a systematic approach to the validation of the assumptions concerning the random noise on which the statistical model is built. The analysis presented here may equally well be applied to residual analysis in connection with model validation.

The tests described in Appendix A can be divided into tests of a certain hypothesised distribution, tests of randomness based on combinatorial arguments, and tests of correlation properties. The latter category can be subdivided into tests in the time domain that are based on correlation function estimates, and tests in the frequency domain that are based on spectrum estimates. The following tests were performed (see Appendix A for explanations): test of joint normal distribution; test of marginal normal distribution; tests of randomness; tests of autocorrelation functions; tests of auto-spectra; test of cross-correlation function; and test of cross-spectrum.

The computations were performed using Matlab[1] (The MathWorks), including the Matlab Statistics Toolbox [36], the Matlab System Identification Toolbox [37], and the Matlab Signal Processing Toolbox [38]. No adjustments were made to account for the fact that multiple tests were performed on the same data, and it should be noted that some of the tests rely on the results of some of the other tests for their validity. For instance, if the real and imaginary parts of the noise are correlated, then conclusions drawn only from separate analyses of the two parts may be misleading.

3.3. Results

A first picture of the raw data is given in Fig. 3 (the results in this section are for the case with $2f_{\mathrm{hc}} = f_s$ unless otherwise is explicitly stated). If the complex observations have a complex Gaussian distribution—which, as noted in Section 1, means that the real and imaginary parts are uncorrelated and have equal variance—with zero mean, then the cloud of points in Fig. 3 should be fairly circularly symmetric with the largest density around the origin. It seems that the cloud has its centre somewhat towards the southwest, indicating a non-zero mean that arises because the channels are not perfectly balanced. This rather common effect is also discernible in the separate plots of the real and imaginary parts of the time series in Fig. 3. The observed offset suggests that the model interferogram (1) should include this term, unless it is removed by a suitable phase cycling scheme, e.g., CYCLOPS [39].

The probability plots in Fig. 4 can be used to evaluate the assumption of a Gaussian distribution. If the data are normally distributed, but perhaps with a non-zero mean, then—after centering of the data—the absolute values of the observations should follow a Rayleigh distribution, and the angles of the complex observations should follow a uniform distribution over the interval $[-\pi, \pi]$. The two probability plots in Fig. 4 do not show any striking deviations from what would be expected under the hypothesis of normally distributed data. Table 1 presents the p-values from the tests described in Appendix A.1 applied to the experimental data. It can be seen that the Kolmogorov–Smirnov test of the hypothesis that the absolute values follow a Rayleigh distribution gives no reason to reject the hypothesis on the 5% level of significance. Neither does the test of the hypothesis that the angles follow a uniform distribution give any reason to reject this hypothesis.

Figs. 5 and 6 offer a possibility to evaluate if there is any correlation between the real and imaginary parts of the signal. The estimated cross-correlation functions in Fig. 5 do not show such systematic deviations that there is reason to reject the hypothesis that the real and imaginary parts of the signal are uncorrelated (Appendix A.6). Since 39 lags are included in the plot, we would expect about two (i.e., 5% of 39) cross-correlation values to fall outside the 5% critical bounds. The squared coherency spectrum estimate (Appendix A.7) in Fig. 6 does not show any signs of any systematic connection between the real and imaginary parts of the signal—the smoothed estimate lies well below the estimated 5% critical bound. If we consider it sufficiently established that the real and imaginary parts of the signal are independent, we can proceed to study them separately.

The autocorrelation function estimates for the case when $2f_{\mathrm{hc}} = f_s$ in Fig. 5 do not give any reason to reject the hypothesis that each of the two time series (the real part and the imaginary part of the signal, respectively) is a series of uncorrelated observations (Appendix A.4). Since the plots cover 19 lags, we should expect about one autocorrelation value to fall outside the 5% critical bounds. A portmanteau test comprising the 32 first lags of the autocorrelation function for the real and imaginary parts, respectively, does not give any reason to reject the hypothesis that the noise in each of the two channels consists of independent identically distributed observations (see Table 1). For the case when $2f_{\mathrm{hc}} = 1.2f_s$ the autocorrelation function estimates in Fig. 5 show a significant and systematic deviation from zero for the first lag, which is in accord with the theoretical prediction (see Fig. 2).

[1] M-files are available from HG on request.

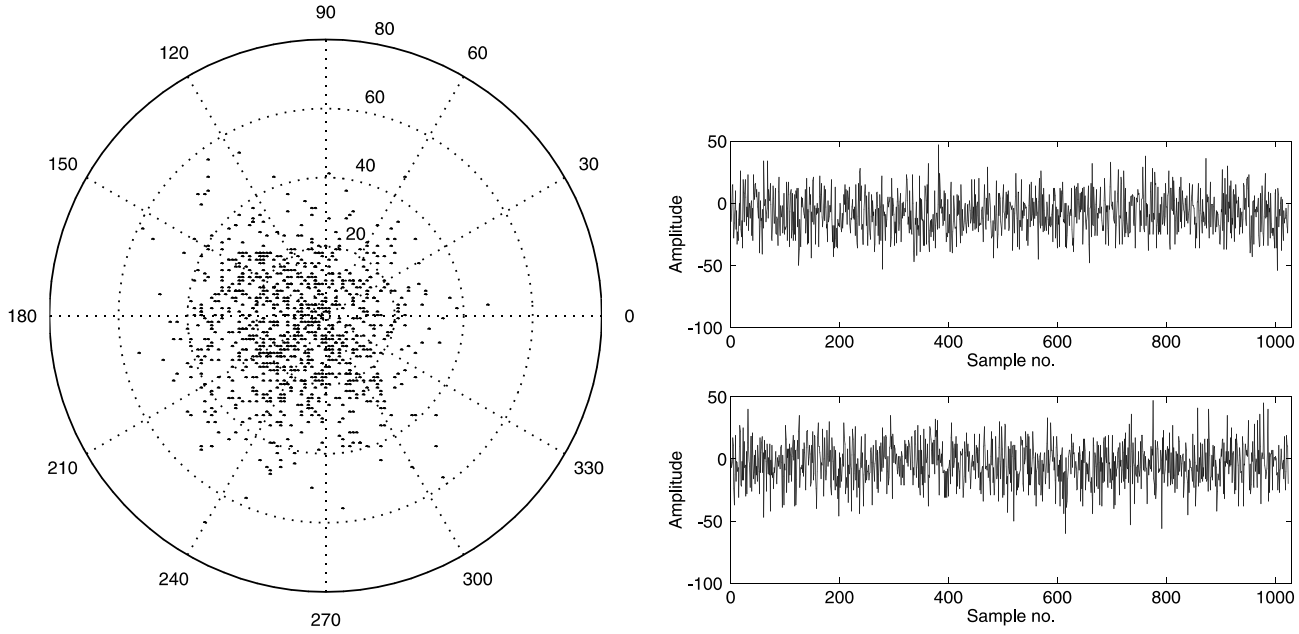

Fig. 3. Left: the complex observations represented in polar coordinates. Top right: the real part of the signal. Bottom right: the imaginary part of the signal.

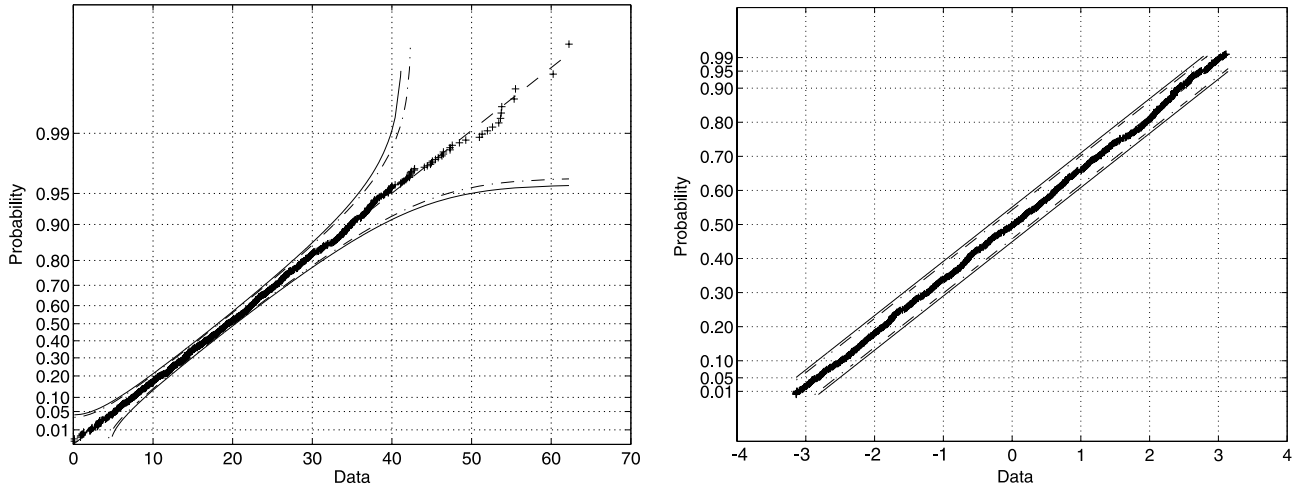

Fig. 4. Left: Rayleigh probability plot of the absolute values of the centred complex observations. Right: uniform probability plot of the angles of the centred complex observations. Critical bounds for $\alpha = 0.05$ and $\alpha = 0.01$ are also shown in the figures.

Concerning the randomness of the sequences, the turning points test and the difference-sign test (Appendix A.3) applied to the real and imaginary parts give p-values well above any reasonable level of significance for the case with $2f_{hc} = f_s$, but for the case with $2f_{hc} = 1.2f_s$ the turning points test indicates that the observed number of turning points significantly exceeds the number expected for an IID sequence. To avoid ties, in the case with $2f_{hc} = f_s$, 20 observations were removed from the real part of the signal and 15 observations were removed from the imaginary part, and in the case with $2f_{hc} = 1.2f_s$, 11 and 17 observations, respectively, were removed.

The histogram and normal probability plot for the imaginary part of the noise in Fig. 7 do not show any large deviations from the normal distribution, and the Kolmogorov–Smirnov test of the normality hypothesis (Appendix A.2) does not show significance (Table 1). Similar graphs and test results are obtained for the real part of the noise. The p-values from t tests of the hypothesis that the mean of the signal in each channel is zero are given in Table 1. The conclusion is that both means are significantly different from zero. Also, the hypothesis that the two means are equal must be rejected. Bartlett's test of the hypothesis that the two

Table 1
Summary of test results

Test		p-value	
		f_s	$1.2f_s$
Joint normal distribution	Rayleigh distribution radii	0.99	0.76
	Uniformly distributed angles	0.87	0.93
Portmanteau	Real part	0.81	0.25
	Imaginary part	0.40	0.20
Turning points	Real part	0.14	0.01
	Imaginary part	0.37	0.05
Difference sign	Real part	0.93	0.89
	Imaginary part	0.90	0.95
Marginal normal distribution	Real part	0.85	1.00
	Imaginary part	0.84	0.99
Cumulative periodogram	Real part	0.28	0.00
	Imaginary part	0.06	0.00
Zero mean	Real part	0.00	0.00
	Imaginary part	0.00	0.00
Equal means		0.00	0.00
Equal variances		0.46	0.08

The column headed f_s refers to the case with $2f_{hc} = f_s$ and the column headed $1.2f_s$ refers to the case with $2f_{hc} = 1.2f_s$.

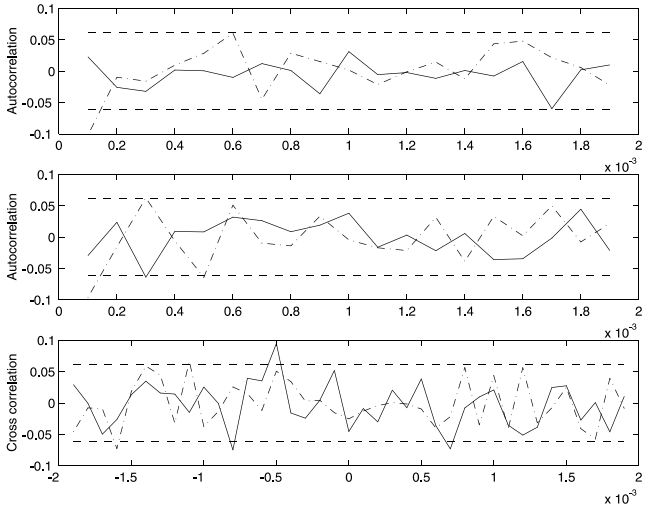

Fig. 5. Estimated correlation functions pertaining to the two channels of the quadrature detection. The sampling interval is $\Delta t = 0.1$ ms. The autocorrelation functions have been estimated for a span of the time lag up to 2 ms, and the cross-correlation function has been estimated for a lag span of 4 ms. (—) Represents the case with $2f_{hc} = f_s$ and (– · –) represents the case with $2f_{hc} = 1.2f_s$. Top: autocorrelation function for the real part of the signal. Middle: autocorrelation function for the imaginary part of the signal. Bottom: cross-correlation function for the real and imaginary parts. Approximate point-wise critical bounds (- - -) for $\alpha = 0.05$ are also shown.

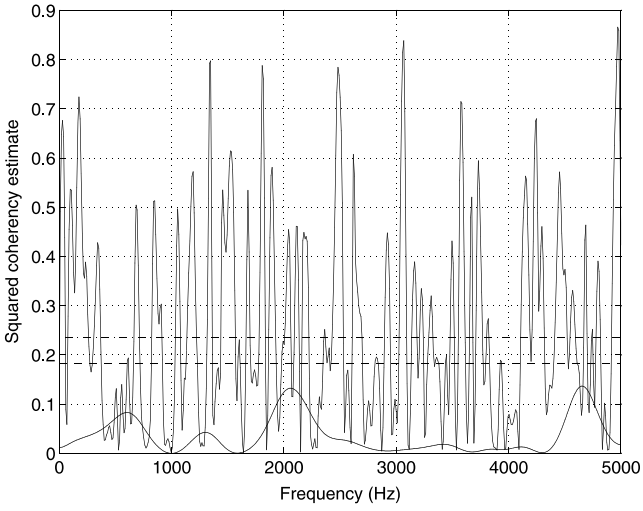

Fig. 6. Plot of the estimated squared coherency spectrum for the real and imaginary part of the signal. The estimates have been smoothed with a Hamming window of lag size 256 and 30, respectively. The critical bounds for the more smoothed estimate are also shown, with $\alpha = 0.05$ and $\alpha = 0.01$.

normal distributions have the same variance does not lead to rejection on the 5% level.

The spectral estimates for the imaginary part of the noise for the case when $2f_{hc} = f_s$ in Fig. 8 indicate that it has a rather constant spectral density, and the cumulative periodogram with critical bounds based on the Kolmogorov–Smirnov test (Appendix A.5) does not give any reason to reject the hypothesis of white Gaussian noise on the 5% level (the p-values are shown in Table 1). On the other hand, the corresponding plots in Fig. 9 for the case when $2f_{hc} = 1.2f_s$ do give reason to reject the hypothesis of white Gaussian noise. Similar graphs and test results are obtained for the real part of the noise.

4. Conclusions

The theoretical and experimental analyses presented here address the common assumption that the noise in a

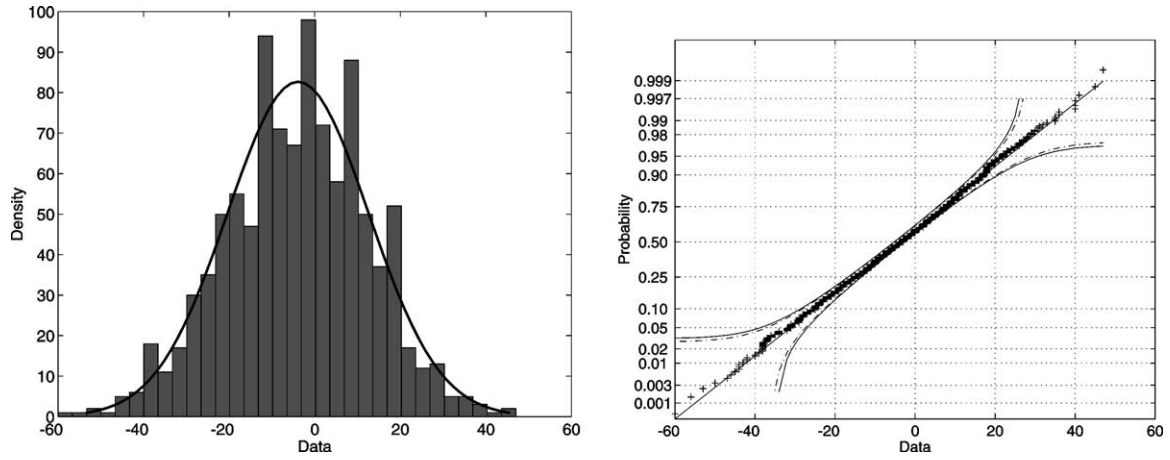

Fig. 7. Plots for the imaginary part of the signal. Left: histogram and estimated Gaussian density. Right: normal probability plot with critical bounds for $\alpha = 0.05$ and $\alpha = 0.01$.

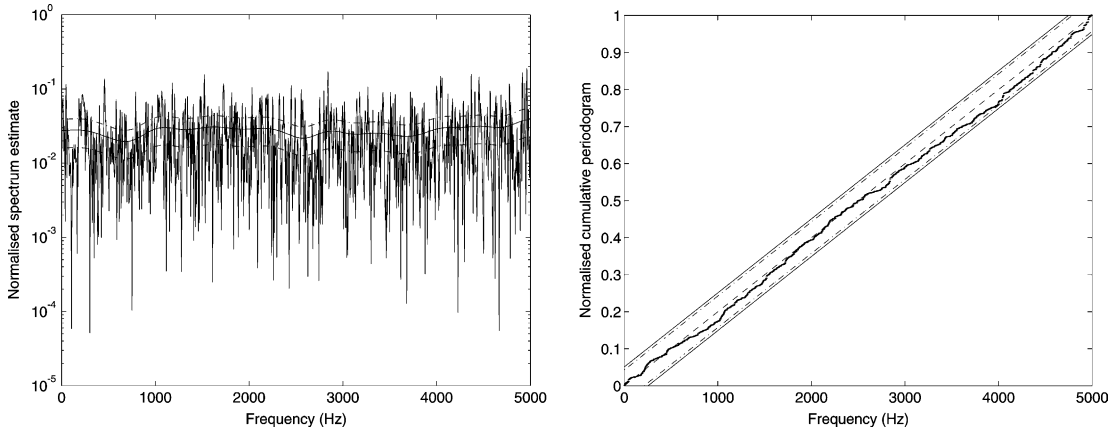

Fig. 8. Plots for the imaginary part of the signal for the case when $2f_{hc} = f_s$. Left: raw periodogram and smoothed spectral estimate with 2σ-bounds. Right: normalised cumulative periodogram with critical bounds for $\alpha = 0.05$ and $\alpha = 0.01$.

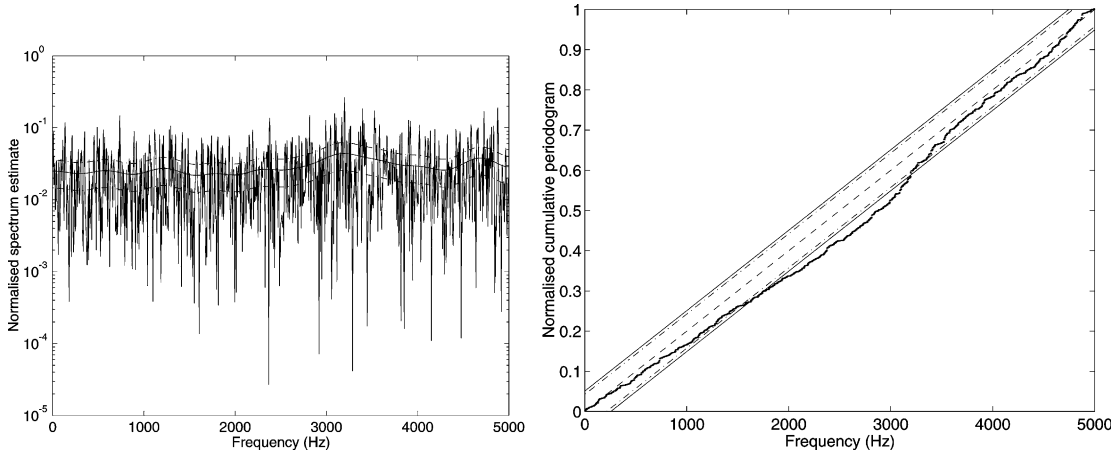

Fig. 9. Plots for the imaginary part of the signal for the case when $2f_{hc} = 1.2f_s$. Left: raw periodogram and smoothed spectral estimate with 2σ-bounds. Right: normalised cumulative periodogram with critical bounds for $\alpha = 0.05$ and $\alpha = 0.01$.

complex NMR signal acquired from a single receiving coil using quadrature detection is uncorrelated (white) and Gaussian. The theoretical analysis is based on the assumption that the dominant noise source is the thermal noise generated by the sample, coil, and preamplifier, and that it can be adequately modelled as white noise. By analysing the signal processing steps prior to digital sampling of the signal, we show that certain conditions have to be met for the sampled complex noise to be white with uncorrelated real and imaginary parts. If the noise is white, with a symmetric spectral density centred at the mixing frequency, and band-limited with a band-width greater than the stop-band cut-off frequency of the audio filter, the noise in the two channels will be uncorrelated at all times, i.e., $\mathrm{Cov}[\mathrm{Re}\, Y_m, \mathrm{Im}\, Y_n] = 0$ for all m and n, provided that the phase difference between the two modulating sinusoids is exactly $\pi/2$. Under the same assumptions, the autocovariance function for the noise in each channel will be a sinc function in the case of an ideal filter. The shape of this sinc function implies that sampling with a frequency of $f_s = 2f_{\mathrm{hc}}/n$, where n is an integer, yields uncorrelated noise in each channel. In contrast, any other ratio between the sampling and filter frequencies is expected to introduce correlations in the noise sequences. When a Butterworth filter of even order is used, the autocovariance function of the filtered noise can be well approximated by a sum of exponentially damped sinusoids which will have zeros in good agreement with the sinc function. Statistical analyses of experimental noise obtained using a standard analog filter (an 8-pole quasi-elliptical filter) bear out the theoretical predictions: noise sampled at twice the filter cut-off frequency (i.e., $f_s = 2f_{\mathrm{hc}}$) is shown to be white; in contrast, noise sampled using a 20% larger filter width (i.e., $1.2f_s = 2f_{\mathrm{hc}}$) is shown to be significantly non-white. Of note, it is common practice to use a filter width that reaches 10–20% outside of the spectral width when analog filters are used. The analysis presented here suggests that if the application in question requires that the sampled noise is white, it would be better to use a filter cut-off frequency of $2f_{\mathrm{hc}} = f_s$, while adjusting f_s such that distortions due to the filter are avoided in the spectral region of interest.

Regarding parameter estimation using maximum likelihood methods, the ML estimator will still be a least-squares estimator, and hence well motivated from a statistical point of view, even if the noise is not perfectly uncorrelated. However, in case the noise is actually correlated, an ML estimator based on the assumption of uncorrelated noise will be less efficient (i.e., have larger variance) than an ML estimator that explicitly models the noise correlation. A more serious problem is that parameter estimate confidence intervals computed under a false assumption of uncorrelated noise may be misleading, and hence—in the worst case—underestimate the uncertainty of the estimates. The consequences of such correlations will have to be investigated for the individual case. Model validation in connection with NMR signal estimation may be performed by analysing the residuals; an adequate model is expected to yield residuals that have the same statistical characteristics as the noise specified in the model. The procedures outlined here enable efficient analysis of residuals in this respect, and should be of general use in ML-based methods for parameter estimation of NMR signals.

Acknowledgments

M.A. acknowledges support from the Swedish Research Council and the Swedish Foundation for Strategic Research. The authors would also like to thank professors Jan Holst and Tobias Rydén, Department of Mathematical Statistics at Lund University, for valuable comments during the writing of this article.

Appendix A. Statistical tests

A.1. Test of joint normal distribution

A zero mean bivariate normal variable $Z = (X, Y)$ with X and Y uncorrelated with equal variances, $\sigma_X^2 = \sigma_Y^2 = \sigma^2$ can be expressed in polar coordinates as $Z = (R\cos\theta, R\sin\theta)$ where the radius R has a Rayleigh distribution with scale parameter σ and the angle θ is uniformly distributed on the interval $[-\pi, \pi]$ [27]. This can be used to check if the real and imaginary parts of the noise are independent and jointly Gaussian: if the hypothesis that the angles of the complex observations are uniformly distributed in the interval $[-\pi, \pi]$ is true, then the empirical cumulative distribution function should follow a straight line in a uniform probability plot, without any systematic or extreme deviations. (For the general concepts of probability plotting and its use for examining goodness-of-fit see, e.g. [40].) Since the distribution is fully specified by the null hypothesis, critical bounds based on the Kolmogorov–Smirnov test can be applied directly and the p-value can be calculated [41]. These bounds are constructed so that there is a probability $1 - \alpha$ that the empirical distribution function lies entirely within the bounds.

In a similar way, a probability plot can be constructed to test the hypothesis that the absolute values of the complex observations come from a Rayleigh distribution. In this case, since the scale parameter σ of the distribution is not specified by the null hypothesis, a modification of the Kolmogorov–Smirnov test will have to be applied, using the ML estimate of the unknown parameter and appropriate critical values [42]. The critical bounds thus obtained are interpreted in the same way as in the previous case. In this case there is no

formula available for the calculation of the p-value, so it has been estimated using the bootstrap method with 1000 bootstrap samples [43].

A.2. Test of marginal normal distributions

The assumption that the real and imaginary parts of the signal each are independently and normally distributed sequences can be checked further with a normal probability plot for each part. Since the mean and variance of the distribution are not specified by the null hypothesis, a Kolmogorov–Smirnov–Lilliefors test with the ML estimates of the unknown parameters can be used to obtain critical bounds for the empirical distribution function [41,44]. The p-values have been estimated using bootstrap with 1000 bootstrap samples [43].

A.3. Tests of randomness

One way to test the randomness of a sequence is to count the number of peaks and troughs, collectively called turning points, of the series. If T is the number of turning points of the sequence then, under the null hypothesis of an IID sequence, the expected number of turning points μ_T and the variance σ_T^2 of T can be computed. If the observed value of T is much greater than μ_T this indicates that the series is fluctuating more rapidly than would be expected for a random series. A value of T much smaller than μ_T would indicate a more slowly varying series with positive correlations between neighbouring points. This can be tested more formally with a test based on the fact that T is asymptotically normal with parameters μ_T and σ_T^2 when the null hypothesis is true [22,45].

Since the test based on the turning points has a poor performance as a test against trend, it could preferably be accompanied by, for instance, a difference-sign test. If S is the number of points of increase of the series, i.e., the number of positive first differences of the series, then, for an IID sequence, the expectation μ_S and variance σ_S^2 of S can be computed, and it can be shown that S is asymptotically normal with μ_S and σ_S^2 as parameters [22,45]. This implies that the null hypothesis of an IID sequence can be formally tested. If the observed value of S is much greater than μ_S, this indicates an increasing trend in the data, while a value of S much smaller than μ_S indicates a decreasing trend.

The turning points and difference-sign tests presuppose that the observations have a continuous distribution, so that the probability that two or more consecutive observations are equal (ties) is zero. When the observations are digitised, they will also be discretised, which means that the conditions for these two tests are not satisfied in a strict sense. But as long as the number of ties is small compared with the total number of observations, it seems reasonable to apply the tests, if the ties are removed from the data (i.e., if two or more consecutive observations are equal, all of these observations except the first one are removed).

A.4. Tests of autocorrelation functions

If the data is a realisation of an IID noise process, approximately 0.95 of the sample autocorrelations should lie between the critical bounds $\pm 1.96/\sqrt{N}$, based on a large sample normal approximation [22]. As a complement to this point-wise procedure, where each correlation is checked individually, a portmanteau test can be used, considering instead a single overall statistic which depends on all estimated correlations $\hat{\rho}(i)$ in a given interval $1 \leqslant i \leqslant m$. If the hypothesis of IID noise is correct, this test statistic will have an asymptotic (as $N \to \infty$) χ^2 distribution with the number of degrees of freedom equal to the number of correlations considered [22,35]. If the mean is found to be significantly different from zero and is subtracted from the series before the test, this should be adjusted for by reducing the number of degrees of freedom by one.

A.5. Tests of auto-spectra

In the complementary frequency analysis, the raw periodograms have been computed as the absolute square of the FFT of the time series, normalised by multiplication with the sampling interval [26]. The spectral estimates are based on the estimated covariance functions, which have been multiplied by a Hamming window of lag size 30 and then Fourier transformed [26]. An even more effective means for the detection of periodicities in the data may be provided by the cumulative periodogram, which can be combined with the Kolmogorov–Smirnov test to check the hypothesis of Gaussian white noise. The cumulative periodograms have been calculated for the maximal number of frequencies, i.e., $N/2 = 512$. There is a probability $1 - \alpha$ that the estimated normalised cumulative periodogram lies entirely within the critical bounds based on the Kolmogorov–Smirnov test [22,31,35].

A.6. Test of cross-correlation function

An approximate test for the independence of the two time series constituting, respectively, the real and the imaginary part of the complex signal, can be obtained by prewhitening one of them before estimating the cross-correlations and then comparing these with the same point-wise critical values as for the above autocorrelation analysis [22]. For the prewhitening procedure an auto-regressive (AR) filter of order 10 has been used [22,26,31,35].

A.7. Test of cross-spectrum

In a frequency analysis of the relation between two time series, one possibility is to inspect the squared coherency spectrum, which is defined as

$$\kappa^2(f) = \frac{|R_{Z_R,Z_I}(f)|^2}{R_{Z_R}(f)R_{Z_I}(f)}, \tag{A.1}$$

where $R_{Z_R}(f)$ and $R_{Z_I}(f)$ are the spectral densities of the real and imaginary parts of the signal, respectively, and $R_{Z_R,Z_I}(f)$ is the corresponding cross-spectral density. The coherency spectrum can, roughly speaking, be viewed as a correlation coefficient in the frequency domain [22,26,31]. The squared coherency spectrum has been estimated using Welch's averaged periodogram method [26], applying a Hamming window of lag size 256 as well as a more smoothed version with a window of lag size 30. For the more smoothed estimate critical bounds have been estimated by a Monte Carlo simulation with 1000 replicates, based on the hypothesis that Z_R and Z_I are uncorrelated and each consists of white Gaussian noise. Thus, if this hypothesis is correct, there is approximately probability $1-\alpha$ that the estimated squared coherency spectrum lies entirely below the bound.

References

[1] R. de Beer, D. van Ormondt, Analysis of NMR data using time domain fitting procedures, in: M. Rudin (Ed.), In Vivo Magnetic Resonance Spectroscopy I: Probeheads and Radiofrequency Pulses, Spectrum Analysis, Springer Verlag, Berlin, 1992, pp. 201–248.

[2] J.K. Lindsey, Parametric Statistical Inference, Clarendon Press, Oxford, 1996.

[3] D. Spielman, P. Webb, A. Macovski, A statistical framework for in vivo spectroscopic imaging, Journal of Magnetic Resonance 79 (1988) 66–77.

[4] M.I. Miller, A.S. Greene, Maximum-likelihood estimation for nuclear magnetic resonance spectroscopy, Journal of Magnetic Resonance 83 (1989) 525–548.

[5] S.C. Chen, T.J. Schaewe, R.S. Teichman, M.I. Miller, S.N. Nadel, A.S. Greene, Parallel algorithms for maximum-likelihood nuclear magnetic resonance spectroscopy, Journal of Magnetic Resonance, Series A 102 (1993) 16–23.

[6] S. Umesh, D.W. Tufts, Estimation of parameters of exponentially damped sinusoids using fast maximum likelihood estimation with application to NMR spectroscopy data, IEEE Transactions on Signal Processing 44 (1996) 2245–2259.

[7] L. Vanhamme, A. van den Boogaart, S. Van Huffel, Improved method for accurate and efficient quantification of MRS data with use of prior knowledge, Journal of Magnetic Resonance 129 (1997) 35–43.

[8] G. Zhu, Y. Hua, Quantitative NMR signal analysis by an iterative quadratic maximum likelihood method, Chemical Physical Letters 264 (1997) 424–428.

[9] G. Zhu, W.Y. Choy, B.C. Sanctuary, Spectral parameter estimation by an iterative quadratic maximum likelihood method, Journal of Magnetic Resonance 135 (1998) 37–43.

[10] Z. Bi, A.P. Bruner, J. Li, K.N. Scott, Z.-S. Liu, C.B. Stopka, H.-W. Kim, D.C. Wilson, Spectral fitting of NMR spectra using an alternating optimization method with a priori knowledge, Journal of Magnetic Resonance 140 (1999) 108–119.

[11] P. Stoica, T. Sundin, Exact ML estimation of spectroscopic parameters, Journal of Magnetic Resonance 145 (2000) 108–114.

[12] L. Vanhamme, T. Sundin, P. Van Hecke, S. Van Huffel, MR spectroscopy quantitation: a review of time-domain methods, NMR in Biomedicine 14 (2001) 233–246.

[13] G.L. Bretthorst, in: Bayesian Spectrum Analysis and Parameter Estimation, Lecture Notes in Statistics, vol. 48, Springer, New York, 1988.

[14] G.L. Bretthorst, Bayesian spectrum analysis on quadrature NMR data with noise correlations, in: J. Skilling (Ed.), Maximum Entropy and Bayesian Methods, Proceedings of the 8th MaxEnt Workshop held at St. John's College, Cambridge, England, August 1–5, 1988, Kluwer, Dordrecht, 1989, pp. 261–273.

[15] G.L. Bretthorst, Bayesian analysis. IV. Noise and computing time considerations, Journal of Magnetic Resonance 93 (1991) 369–394.

[16] J.W. Carlson, Power deposition and noise correlation in NMR samples, Magnetic Resonance in Medicine 10 (1989) 399–403.

[17] A. Jesmanowicz, J.S. Hyde, W. Froncisz, J.B. Kneeland, Noise correlation, Magnetic Resonance in Medicine 20 (1991) 36–47.

[18] E.R. McVeigh, R.M. Henkelman, M.J. Bronskill, Noise and filtration in magnetic resonance imaging, Medical Physics 12 (1985) 586–591.

[19] C.-N. Chen, D.I. Hoult, V.J. Sank, Quadrature detection coils—a further $\sqrt{2}$ improvement in sensitivity, Journal of Magnetic Resonance 54 (1983) 324–327.

[20] H.D.W. Hill, Spectrometers: a general overview, in: D.M. Grant, R.K. Harris (Eds.), Encyclopedia of Nuclear Magnetic Resonance, Wiley, New York, 1996, pp. 4505–4518.

[21] D.I. Hoult, C.-N. Chen, H. Eden, M. Eden, Elimination of baseline artifacts in spectra and their integrals, Journal of Magnetic Resonance 51 (1983) 110–117.

[22] P.J. Brockwell, R.A. Davis, Time Series: Theory and Methods, second ed., Springer, New York, 1991.

[23] A. Papoulis, Probability Random Variables and Stochastic Processes, third ed., McGraw-Hill, New York, 1991.

[24] A. van den Bos, The multivariate complex normal distribution—a generalization, IEEE Transactions on Information Theory 41 (1995) 537–539.

[25] P. Stoica, R.L. Moses, Introduction to Spectral Analysis, Prentice-Hall, Upper Saddle River, NJ, 1997.

[26] L. Ljung, System Identification. Theory for the User, second ed., Prentice-Hall, Upper Saddle River, NJ, 1999.

[27] K.S. Miller, Complex Stochastic Processes, Addison-Wesley, Reading, MA, 1974.

[28] D.I. Hoult, Sensitivity of the NMR experiment, in: D.M. Grant, R.K. Harris (Eds.), Encyclopedia of Nuclear Magnetic Resonance, Wiley, New York, 1996, pp. 4256–4266.

[29] R.A. Beckman, E.R.P. Zuiderweg, Guidelines for the use of oversampling in protein NMR, Journal of Magnetic Resonance A 113 (1995) 223–231.

[30] K.S. Miller, Multidimensional Gaussian Distributions, Wiley, New York, 1964.

[31] M.B. Priestley, Spectral Analysis and Time Series, Academic Press, San Diego, 1981.

[32] R.W. Hamming, Digital Filters, third ed., Dover, Mineola, NY, 1998.

[33] W.E. Higgins, D.C. Munson Jr., Infinite impulse response digital filter design, in: S.K. Mitra, J.F. Kaiser (Eds.), Handbook for Digital Signal Processing, Wiley, New York, 1993, pp. 279–335 (Chapter 5).

[34] D.I. Hoult, R.E. Richards, The signal-to-noise ratio of the nuclear magnetic resonance experiment, Journal of Magnetic Resonance 24 (1976) 71–85.
[35] G.E.P. Box, G.M. Jenkins, G.C. Reinsel, Time Series Analysis: Forecasting and Control, Prentice-Hall, Englewood Cliffs, NJ, 1994.
[36] B. Jones, MATLAB Statistics Toolbox, Version 2.2. The Math Works, Inc., Natick, MA, July 1998.
[37] L. Ljung, System Identification Toolbox, Version 4.0.5. The Math Works, Inc., Natick, MA, April 1998.
[38] The Math Works, Inc., Natick, MA. Signal, Processing Toolbox. Version 4.2, July 1998.
[39] D.I. Hoult, R.E. Richards, Critical factors in the design of sensitive high resolution nuclear magnetic resonance spectrometers, Proceedings of the Royal Society, A 344 (1975) 311–340.
[40] R.B. D'Agostino, M.A. Stephens (Eds.), Goodness-of-Fit Techniques, Marcel Dekker, New York, 1968.
[41] A. Stuart, J. Keith Ord, Kendall's Advanced Theory of Statistics, vol. 2, fifth ed., Edward Arnold, London, 1991.
[42] M. Chandra, N.D. Singpurwalla, M.A. Stephens, Kolmogorov statistics for tests of fit for the extreme-value and Weibull distributions, Journal of the American Statistical Association 76 (1981) 729–731.
[43] J.S. Urban Hjort, Computer Intensive Statistical Methods, Chapman and Hall, London, 1994.
[44] H.W. Lilliefors, On the Kolmogorov–Smirnov test for normality with mean and variance unknown, Journal of the American Statistical Association 62 (1967) 399–402.
[45] M. Kendall, A. Stuart, J. Keith Ord, The Advanced Theory of Statistics, vol. 3, fourth ed., Charles Griffin & Co, London, 1983.