

Saving Measurement Time in ^{13}C NMR Spectroscopy

Sharif D. Kunikeev[†] and Howard S. Taylor*

Department of Chemistry, University of Southern California, Los Angeles, California 90089-0482

Received: October 28, 2003

An algorithm to produce acceptably accurate Fourier transform NMR spectra using many fewer transients than commonly obtained is introduced and applied to ^{13}C chemical shift and one-dimensional INADEQUATE Lorentzian type spectra. The algorithm supplies criteria to recognize when to stop transient collection at a time too early for an acceptable spectrum to be produced by Fourier signal processing but not by harmonic inversion signal processing.

1. Introduction

In NMR spectroscopy the main method of signal processing is that of the Fourier transform (FT). In most experimental NMR studies noise is sufficiently prominent in the measured free induction decay (FID) time signal so as to make the Fourier processed spectrum unacceptable in the sense that the underlying spectrum cannot be extracted from the noisy one. The problem of reducing the noise level is approached in several ways. The most used method is the ubiquitous signal averaging over many transients or scans. In signal averaging the signal grows relative to noise quite slowly as $\sqrt{N_{\text{tr}}}$, N_{tr} being the number of measured transients. Moreover, there are often other considerations, such as sample concentration, field strength, and NMR machine design, that affect the signal intensity and cause the needed number of transients to be very high—so high in fact that some experiments are just not done. Additionally, when high spectral resolution is required, the Fourier methods need longer signals. As the ideal noiseless FID decays in time, the signal samples at longer times contain larger noise components and using them can be counterproductive.

Many NMR data processing methods^{1,2} try to improve this situation by attempting to maximize the extraction of signal information from noisy data using a priori knowledge regarding the underlying signal model. One general method, which is used in this paper, is the harmonic inversion method. This method comes in many versions^{3–10} and appears under such names as linear prediction,^{3,11} Pade approximation,⁴ filter diagonalization,^{5–8} regularized resolvent transform,⁹ and filter diagonalization method version 2000.¹⁰ The harmonic inversion method achieves a higher spectral resolution for a given signal length N than Fourier transform. This is due to the fact that the a priori information used here is that a linear combination of complex decaying exponentials is assumed to be a good model for an FID signal. The equations that relate the measured signal to the model when solved yield directly the position, width, and intensity (area under the line) of each resonance, and from them a spectrum can be reconstructed. The said equations are generically ill-conditioned, which means any solutions are extremely sensitive to small noise perturbations. To fix this, various regularization techniques are used^{11,12} which enable a

stable solution and spectrum to be found. When noise is not present or there is little noise, all these harmonic inversion and regularization methods work well and give similar results. However, there is a noise threshold above which they can give differing results, sometimes producing spectra with false or missing lines if a signal is “too” noisy; as such the method should not be used above this noise threshold. Thus, the problem is how to formulate such a harmonic inversion based signal processing scheme which will be able to produce acceptable spectra from noise reduced signals created by averaging a smaller number of transients than will be required by the Fourier method to get the same quality spectrum. Here, “acceptable” means, simultaneously, no fake or missing features, no undetermined “regularization” parameters, and small enough errors in the predicted Lorentzian parameters.

This paper suggests a windowed signal processing noise reduction strategy which is based on an eigenvector type signal processing theory to devise a conceptually simple and undemanding computational scheme to greatly reduce the number of measured transients (scans) required to obtain a Lorentzian type spectrum reconstructed from originally noisy data that have an acceptable level of errors in the Lorentzian parameters such as frequency position, width, and area under the Lorentzian. Our aim is to greatly reduce the machine time used to measure a spectrum or to enable new experiments to be done which presently require unacceptable amounts of machine time.

The strategy defines two qualitative criteria that depend on a signal constructed by averaging measured transients. The first criterion (section 2.2) is mathematical and uses the signal to construct a correlation matrix, the graph of whose eigenvalues (singular values) show a characteristic gap in their values when the criterion is satisfied. The second criterion (section 3) is the one common with the signal averaging plus Fourier analysis method, that is, the acceptance of the spectra when its prominent features are stable to further measurement. Satisfaction of the criteria determines when no further noise reducing signal averaging is needed for it to satisfy a model of the signal plus noise assumed by a mathematical noise reduction method, that of Cadzow.¹³ At this point it is shown that even though the Fourier transform of the signal will usually give an unacceptably noisy spectrum, Cadzow’s method can further reduce the noise so as to produce a “cleaned” signal.

Sections 2 and 3 represent a first application of the noise reduction methodology to the given systems. Because of this,

* Corresponding author. E-mail: taylor@usc.edu.

[†] Permanent address: Institute of Nuclear Physics, Moscow State University, 119899 Moscow, Russia. E-mail: kunikeev@usc.edu.

it at first appears that the method will always require for the experimentalist to work interactively with the measurement in order to determine the number of transients and ultimately machine time needed for the convergence of the spectra. That this is not true is explained in the discussion, section 4, where experiments are suggested to calibrate our methods so as to give tables of thresholds for the number of transients needed to satisfy our criteria in any generic class of experiments.

All examples presented in section 3 were chosen by our acknowledged colleagues to be challenging. The windowed noise reduction strategy method is first applied to ^{13}C chemical shift problem and then to the one-dimensional (1D) INADEQUATE experimental test problem. Here the INADEQUATE experimental example was analyzed because it was the most noisy spectra we could think of when done in one dimension. We are aware that it is better to do INADEQUATE in two dimensions using an algorithm such as perhaps FRED.¹⁴ In section 2, we give some of the mathematical details used in our signal processing scheme: (i) the windowing technique (section 2.1); (ii) the noise reduction preprocessor (section 2.2); and (iii) the harmonic inversion spectral estimator (section 2.3). Section 4 briefly summarizes our ideas. In the Appendix, we present results of a statistical error analysis of the parameters and the confidence spectra for the single Lorentzian obtained with the suggested noise reduction scheme.

2. Windowing

2.1. Windowing. All processing is done here by breaking the fast Fourier transform (FFT) spectrum into windows of 200–500 (300 is usual) Fourier grid points. The reasons for windowing as done here is that for the noise reduction part of the problem (i) without windowing the so-called singular value decomposition³ (SVD) graphs, to be discussed below, become too cluttered with signal and noise singular value points to be easily analyzed; (ii) certain windows will be much simpler to process than others and not windowing unnecessarily ties all features to the features most effected by noise; and (iii) without windowing the dimension $N/2$ of a Hermitian matrix arising in the Cadzow method (section 2.2) would be so large that a needed diagonalization could become too time-consuming. When harmonic inversion is used, windowing creates a smaller signal of length $N_d \ll N$, valid only in the window, which in turn determines the size, $N_d/2$, of the systems of equations to be solved. Since these equations have a rank less than their dimension $N_d/2$, they are ill-conditioned and unstable and the results may be very sensitive to small perturbations. For sufficiently small matrices which windowing ensures, the problem can be ameliorated by regularization procedures as the truncated SVD³ used here.

The edges of the windows are at Fourier grid points. Their placement ideally, based on prior knowledge or hints from the noisy FFT, surrounds regions of signal and begins and ends in regions of pure noise. In less than ideal situations a systematic windowing of the spectrum can be designed for all regions. If peaks, because of spectral density reasons, unavoidably appear at window edges where window induced distortions will occur, an additional window should be chosen so that the edge of the prior window falls interior to the new window. This is possible because windows do not know about each other and can overlap. Choosing windows is generally not a problem and becomes even less so with experience. Here the region of the spectrum where ^{13}C singlets appear is roughly known as is the fact that the spin–spin splittings are located about their base. This makes the choice of windows trivial. It is usually better to work one singlet at a

time, but if singlets are close, several can be studied in one window. Here for the singlet spectrum and the INADEQUATE spectrum N_d are set to 200 and 300 Fourier grid points, respectively.

At this point the “decimation” window program takes over and produces a signal of length N_d , called c_n^{bl} , “bl” for band limited, out of the measured signal of length N called c_n . This is inputted into the noise reduction scheme described in section 2.2. The production process, described with formulas in ref 3, is here discussed in words. In the usual Fourier spectrum of a signal made of a given number of scans, all intensities outside the window of N_d Fourier grid points are set to zero. The window spectrum is then shifted symmetrically about zero frequency and inverse Fourier transformed to produce a “new” signal. Since the original bandwidth was $2\pi/\tau$ (τ is the sampling or delay time) and now is reduced by the factor N_d/N , the new effective sampling or delay time will be $N\tau/N_d \equiv \tau_d$. Hence the band limited signal with n th element $c_n^{\text{bl}} \equiv c^{\text{bl}}(n\tau_d)$ is just the “new” signal element numbered $n = 0, 1, \dots, N_d - 1$. As $T \equiv N\tau = N_d\tau_d$, resolution is not affected by this signal length reduction. After all processing the real part of the frequencies must be shifted back to the original origin. Results near window edges are not reliable.

2.2. Noise Reduction Preprocessor. The noise reduction procedure takes a noisy signal and creates a low noise signal from it. This involves several steps detailed below, and we mark those sentences to distinguish them from the associated discussion.

Working with the N_d decimated time signal samples c_n (we drop the superscript “bl”), it is noted that $N_d - M + 1$, where M is roughly $N_d/2$ (small variations are not important), and linearly independent “measured signal” vectors $\vec{c}_n = (c_n, c_{n+1}, \dots, c_{n+M-1})$ can be created which define an M -dimensional vector space. The measured signal vectors can be assumed to be the sum of an “actual” signal vector \vec{x}_n and a random noise vector $\vec{\epsilon}_n$

$$\vec{c}_n = \vec{x}_n + \vec{\epsilon}_n \quad (1)$$

The harmonic model that fits much of NMR, and often ICR, further assumes that the noiseless time signal elements x_n are created from the sum of K damped harmonics; i.e., this spectrum is the sum of K complex Lorentzians. K is the “rank”, not always the number of observed spectral peaks as sometimes two or more Lorentzians underly one spectral peak. The model can be expressed vectorally and sample by sample respectively as

$$\vec{x}_n = \sum_{k=1}^K d_k z_k^n \vec{s}_k \quad \text{and} \quad x_n = \sum_{k=1}^K d_k z_k^n \quad (2)$$

where $(\vec{s}_k)^T = (1, z_k, z_k^2, \dots, z_k^{M-1})$, $z_k = \exp(-i\omega_k\tau_d)$, where d_k is the amplitude and ω_k is the complex frequency whose imaginary part is assumed negative. z_k is then a damped exponential in the time domain. The real part of ω_k is the frequency, and the modulus of imaginary part $|\text{Im}(\omega_k)|$ is twice the width of the Lorentzian of height $d_k/|\text{Im}(\omega_k)|$. The rank or dimension of the signal space K is not an input and is determined by the processing. Since the vectors \vec{s}_k are linearly independent, they form a basis set for a signal space in which must lie the signal vectors \vec{x}_n . The problem is to now find the signal space without first finding the vectors \vec{s}_k . Then the measured vectors \vec{c}_n will be projected onto the signal subspace and the projections taken as the “cleaned” signal vectors. In the field of signal processing this all is a textbook problem,^{3,16} and only the

prescription is given here. From the c_n the M -by- M Hermitian covariant correlation matrix is constructed (step 1) as

$$R_{ij} = \frac{1}{N_d - M + 1} \sum_{n=0}^{N_d - M} c_{n+i-1} c_{n+j-1}^* \quad (3)$$

Using a standard, computationally $O(M^3)$ diagonalization routine, R is diagonalized (step 2) and the real, nonnegative eigenvalues s_i , also called singular values, and eigenvectors \bar{u}_i , $i = 1, \dots, M$, are obtained and indexed so that $s_i > s_{i+1}$. A plot of $\ln s_i$ against index i is then produced. Now if K were known, the first K eigenvectors \bar{u}_i would be associated with the signal subspace. The complement space, orthogonal to the signal subspace, is the noise space. For an ideal white noise problem the rank K can easily be spotted as the number of low i eigenvalues that are widely separated relative to the remaining $M - K$ noise eigenvalues which are constant and equal to the mean square noise strength. They form a horizontal "string of pearls" to the lower right of the smallest "signal" point. The distance from the first noise point to the signal point is the "gap". Figure 1, bottom right, shows an example of a graph of $\ln s_i$ versus i for what is a high signal-to-noise (S/N) ratio damped exponential signal case. Note the "gap" between the noise and noiseless singular values. K is clearly 1. This "gap" appears even if the noise eigenvalues are not constant. In fact, they will be far from constant. Perfect "gaps" and constancy of the noise eigenvalues would require exclusively white noise, exclusively Lorentzian lines, and a very large N_d so as to create a representative ensemble average of R . This in turn requires a large number of vectors \bar{c}_n . The nonconstancy of the higher indexed eigenvalues is due at a minimum to the finite statistics and the use of a window which as explained above is generally necessary; nonwhite noise can also be a cause.

The "gap" in real cases can be estimated (step 3) by recognizing that they always appear at the "elbow" of curves as in Figure 1. When no gap exists, i.e., the method is not working because the S/N ratio is too low, one gets pictures as the two in rows 1 and 2 on the left in Figure 1. The qualitative change in the eigenvalue density and spacing (low on the left and high on the right) is an indicator of a gap. To understand this separation physically, we recall that perturbation theory indicates that eigenvalues that are well spaced are much less sensitive to perturbations than those in high-density regions; hence the identification of the former with signal and the latter with noise, and the larger the gap the better. Perturbation theory says that the noise subspace will have less influence on the signal subspace as the gap, which represents their eigenvalue differences, increases. In the worst cases this separation will pin the value of K down to within ± 1 . Taking the $+1$ case is safer and often leads to no change in the spectrum. When changes do occur or when no "gap" can be clearly estimated, then the number of transients used needs to be increased and the process restarted. *In fact, it is usually the appearance of the "gap" that determines the sufficiency of the number of transients for noise reduction to be implemented, for K to be estimated, for a local S/N ratio to be 1, and for transient collection to stop after roughly another 5–10% of the already collected transients are obtained. The extra number of transients is needed to test various convergences. Noise reduction preprocessing followed by FFT or harmonic inversion signal processing can now take over.* Any uncertainty in gap estimation is met by increasing the number of transients to test if the resulting spectrum is robust. Finding a converged spectra is the primary goal, and

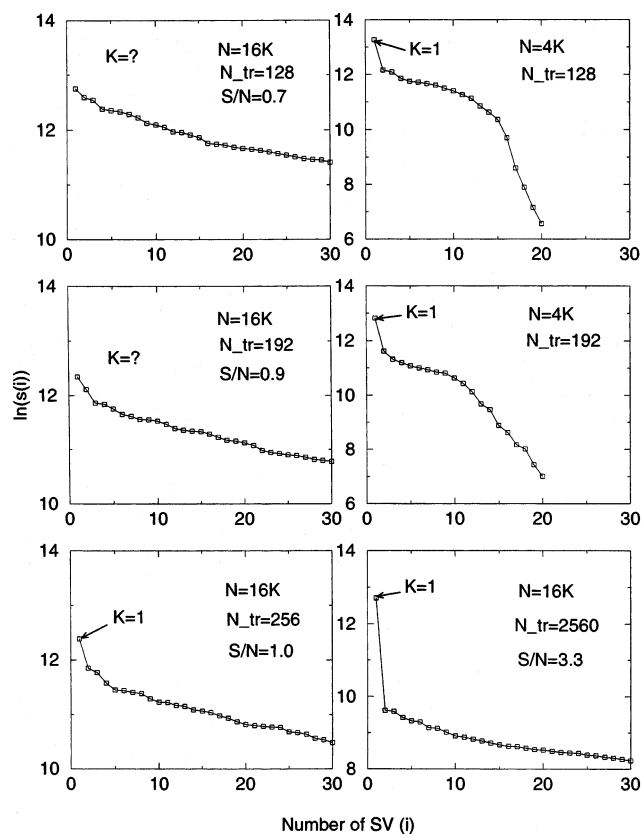


Figure 1. Distributions of logarithmic singular values $\ln s_i$ versus $i = 1, \dots, M$ associated with Figure 2 (carbonyl region, strychnine spectrum) and calculated for the correlation matrix eq 3. $M > 30$ not shown. N is the number of samples acquired; N_{tr} is the number of transients.

increased scanning is our tool to get there. Even with increases the saving in scanning relative to pure signal averaging should be large.

An obvious procedure that saves effort is to use any foreknowledge to pick the window with the least intense and/or the most narrowly separated features. An example is the carbonyl region in the ^{13}C singlet spectra. The number of scans here will be adequate for all remaining windows except when, as will be seen, high resolution is needed to resolve some only slightly less intense features. A most useful strategy is to, in parallel, process c_n , $n = 0, 1, \dots, N - 1$, and the same c_n but stopped at $n_{\max} = N/2$ or even $N/4$. The advantage of this is that there is less noise included in the lower n_{\max} case and K may become easier to estimate. When the K value estimated from the lower n_{\max} was used with the full length N signal, a better resolved but most often very similar spectrum was seen especially if the projected signal was subsequently fed into the harmonic inversion analyzer. Cutting the signal length decreases both resolution and number of \bar{c}_n vectors that can be formed and hence the statistics in the averaging of eq 3. Most of the time the high-resolution analysis then restores resolution. The first two rows of Figure 1, associated with Figure 2 (section 3), shows this effect. The bottom left and right graphs authenticate the choice of $K = 1$.

For now in this subsection we assume K can be determined and proceed to give the projection formula (step 4). In doing this we are basically following the work¹⁵ where this type of noise reduction method, which we adapted to the damped harmonic signal, was used for noise reduction in signals obtained from human speech. The formula is

$$\bar{c}_n^{\text{pr}} = \sum_{k=1}^K (\bar{u}_k^*, \bar{c}_n) \bar{u}_k \quad (4)$$

where (\cdot, \cdot) denotes the scalar, inner product of two vectors and the asterisk (*) stands for complex conjugate. Given the decimated signal vectors \bar{c}_n , we now have an estimate of the $N_d - M + 1$ projected signal vectors as $\bar{x}_n \cong \bar{c}_n^{\text{pr}}$. Since a particular $x_n = c_n^{\text{pr}} = c_n^{\text{NR}}$ element appears in multiple vectors, an arithmetic average is taken (step 5), the result of which restores the Hankel structure of the signal matrix as defined by $x_{nm} = x_{n+m}$. This is required by the model eq 1. As this averaging tends to cancel some of the corrections made in the projection, starting with the new signal the whole procedure is iterated (step 6) to convergence several times (here empirically five iterations are taken) to give the final denoised, cleaned signal.

Clearly a tradeoff is being made here. The first run of the SVD can be proven^{16,17} to give the best, in the least squares sense, rank K approximation to R (see the discussion of the minimum principle in refs 13 and 15), but it does not preserve the Hankel data matrix structure that is consistent with our model, eq 2. Averaging the projected vectors sacrifices a bit of the former to restore the latter. By iterating this diagonalization–projection–averaging procedure several times, we eventually get a K -rank Hankel matrix approximation to the original full rank noisy data matrix. Harmonic inversion methods are seen by us and others^{18,19} to give significantly improved results when fed signals resulting from the Cadzow iterative scheme. Much of the past harmonic inversion work omitted this step. This makes comparisons of our work and past experience with harmonic inversion method of uncertain value. Our method must stand on its own comparisons to experiments.

Fundamentally, the noise reduction is done so as to create a signal that is better represented by eq 2. Another way to look at the reason the method is effective starts by putting eq 1 into eq 4. This shows that \bar{c}_n^{NR} still contains noise as \bar{c}_n is not orthogonal to \bar{u}_k , $k = 1, \dots, K$. It also shows that any randomly oriented noise vector, which would have on the average equal weights on all M basis vectors \bar{u}_k , contributes only a factor proportional to $\sqrt{K/M}$ of its weight in the measured c_n to c_n^{NR} . Clearly one wants M as large as possible consistent with the statistics of eq 3, which improves with smaller M as more terms will appear in the sum.

Note that the model, eq 2, for the underlying signal is not the only one for which the noise reduction scheme could be applied. For applicability of the method it is sufficient for the underlying data matrix $x_{nm} = x_{n+m} = x(n+m)\tau$ to be a $K < M$ rank matrix. For example, this condition can be fulfilled for a polynomial function of order $P_k < M - 1$:

$$d_k^{P_k}(t) = \sum_{p=0}^{P_k} d_{kp} t^p \quad (5)$$

with $d_{k0} \equiv d_k$. Using Newton's binomial formula, one can write a data matrix in a factorizable form:

$$(d_k^{P_k})_{nm} = d_k^{P_k}((n+m)\tau) = \sum_{p=0}^{P_k} \sum_{l=0}^p d_{kp} C_p^l e_l(n) e_{p-l}(m) \quad (6)$$

where $C_p^l = p!/[l!(p-l)!]$ is the binomial coefficient and $e_l(n) = (n\tau)^l$. We wish to find the rank of this matrix. Since the vectors \bar{e}_l , $l = 0, \dots, P_k$, with components $\{e_l(n)\}_{n=0}^{M-1}$ are linearly independent, the rank of matrix (6) is $P_k + 1$. One can suggest a more general ansatz for the signal:

$$x(t) = \sum_{k=1}^K d_k^{P_k}(t) \exp(-i\omega_k t) \quad (7)$$

with the amplitude functions $d_k^{P_k}(t)$ approximated by a P_k -order polynomial (5) with coefficients d_{kp} . Since the data matrix constructed from the signal (7) has the rank $K_r = K + \sum_{k=1}^K P_k$, the above noise reduction scheme can also be applied to this type of signal if $K_r < M$. This might help to clean up those types of non-Lorentzian spectral features that can be well approximated by a Fourier transformed functional form (7). These possible extensions will be investigated in future papers.

2.3. The Harmonic Inversion Spectral Estimator. The c_n^{NR} , $n = 0, \dots, N_d - 1$, which are hopefully very similar to the exact noiseless signal samples x_n , can be Fourier transformed or subjected to a harmonic inversion analysis to give the spectrum $I(\omega)$ which can be presented in absorption, $\text{Re}[I(\omega)]$, or magnitude, $|I(\omega)|$, form. The infinite discrete Fourier transform (DFT) (or z -transform) of the signal described by eq 2 is given by

$$I(\omega) = \tau_d \sum_{n=0}^{\infty} x_n z^{-n} = \tau_d \sum_{k=1}^K \frac{d_k}{1 - z_k/z} = \tau_d \sum_{k=1}^K \frac{d_k}{1 - \exp(i(\omega - \omega_k)\tau_d)} \quad (8)$$

where $z = \exp(-i\omega\tau_d)$. The right-hand side of eq 8 obtained as a result of summing up an infinite series of signal points is the harmonic inversion spectral estimator expressed in terms of harmonic inversion parameters. If $|(\omega - \omega_k)\tau_d| \ll 1$, eq 8 reduces to a sum of complex Lorentzians. In the spectral regions far from resonance lines, $\text{Re}[I(\omega)] = (\tau_d/2) \sum_{k=1}^K d_k = \tau_d x_0/2$ (if the signal is properly phased, i.e., all d_k 's are real), so to get zero baseline the constant $\tau_d x_0/2$ should be subtracted from eq 8.

Let us construct from the signal eq 2 an infinite Hankel matrix $x_{nm} = x_{n+m} = x(n+m)\tau$, $n, m = 0, 1, \dots$. Then the matrix x_{nm} has a finite rank K and there exist K numbers $\alpha_1, \alpha_2, \dots, \alpha_K$ such that

$$x_q = \sum_{k=1}^K \alpha_k x_{q-k} \quad (q = K, K+1, \dots) \quad (9)$$

(see Vol. II, Chapter XV, Section 10, in ref 20 for proof). The linear prediction (LP) equations (9) enable one to calculate all the signal points knowing the first K signal points and the K LP equation coefficients; the LP coefficients in turn can be obtained as a solution of the system of K LP equations.

If the infinite matrix x_{nm} is of finite rank, then the DFT of the signal can be summed up to a rational function of z (Pade approximant):²⁰

$$I(\omega) = \tau_d \sum_{n=0}^{\infty} x_n z^{-n} = \tau_d \frac{P_K(z)}{Q_K(z)} \quad (10)$$

where $P_K(z) = \sum_{k=1}^K b_k z^{K-k+1}$ and $Q_K(z) = \sum_{k=0}^K a_k z^{K-k}$ are, respectively, numerator and denominator polynomials whose coefficients can be calculated from the following system of relations:

$$b_k = \sum_{r=0}^{k-1} a_r x_{k-1-r} \quad (k = 1, \dots, K) \quad (11)$$

$$\sum_{k=0}^K a_k x_{q-k} = 0 \quad (q = K, K+1, \dots, 2K-1) \quad (12)$$

Setting $\alpha_k = -a_k/a_0$, $k = 1, \dots, K$, we can write the relations eq 12 in the form eq 9. Therefore, the a_k coefficients can be obtained as a solution of a set of the LP equations,³ whereas the b_k ones can be obtained from the *explicit* relations eq 11.

The harmonic inversion parameters, z_k or ω_k , can be found, if necessary, by rooting the denominator polynomial. This problem can be effectively reduced to the diagonalization of the companion Hessenberg matrix.^{4,17} The parameters d_k are calculated via the residues of the Pade approximant eq 10 at the positions of the corresponding complex poles z_k .⁴

3. Examples

In this paragraph we first discuss a few “considerations” that face the user. (i) It is well to remember that higher N (length of signal), because the underlying signal is damped, means a higher fraction of noise is present in each new sample. As such, the gap and hence K will be seen first when a low N is used. If the signal had, say, $N = 32\,768$ samples, investigating $N = 16\,384$ incurs no extra experimental cost. Care must be taken when using this “helpful” idea. First, if the need to resolve very close peaks exists, the K gotten from a short N , which can fail to resolve two peaks, may be too small. On the other hand, lowering noise by noise reduction and not signal shortening always makes resolving close peaks easier. (ii) If N_{tr} is insufficient, extreme noise singular values can end up on the signal side of what appears to be a gap and signal singular values could (or could not) appear on the noise side. Respectively, this gives extra unphysical lines and missing lines. Closely spaced lines make everything more difficult. As will be seen in our examples, since these contrary considerations cannot be systematized and can be recognized as spectral nonconvergence, we will simply always increase N_{tr} and seek both to open a gap and to converge the spectrum with respect to increasing N_{tr} using the full given values of N . This will be our ultimate test and strategy.

If software were available, one could work interactively and stop the experiment when gaps opened and spectra stabilized for the estimated weakest features. Here we simulated working interactively by generally obtaining signals that had about 1/25 of the number of transients that signal averaging alone would have required for the case under consideration. The fraction 1/25 was chosen because as will be seen below it was what was needed for the noisiest, most difficult to resolve INADEQUATE splitting we worked with. Our transients came in pulses of eight scans, and we processed signals for increasing numbers of pulses until a gap opened and the spectrum stabilized for the weakest features or until our transients were consumed (the latter did not occur).

The chemical shift spectra of strychnine is our first test case. For the example of a 15 mM solution of strychnine in CDCl_3 , a spectrum showing all the chemically shifted singlet lines with the parameters obtained with an acceptable level of accuracy is achieved on a 20-year-old AM-360 MHz machine with a 5 mm diameter tube after 128 and then 256 transients, with $N = 16\,384$ samples, are measured and averaged. This represents a total acquisition time of 8 min 15 s and 16 min 30 s, respectively. Although times will vary among different machines, tube size, and concentration it should be noted that all our comparisons are between the same machines, tubes, and concentrations. Relative time savings should then be invariant to changes in machine.

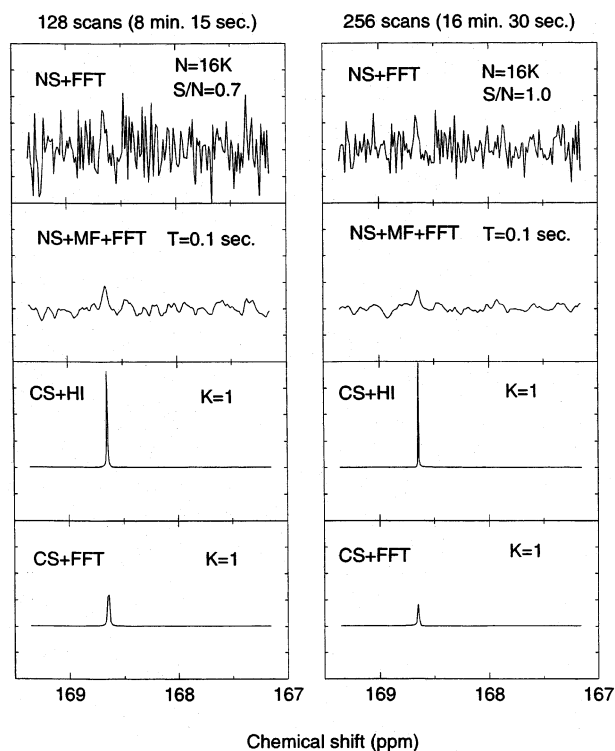


Figure 2. 1D ^{13}C NMR absorption spectrum of a 15 mM solution of strychnine in CDCl_3 for a carbonyl spectral region and with signal after 128 (first column) and 256 transients (second column) of length $N = 16\,384 \equiv 16\text{K}$ computed by, respectively, from top to bottom: NS + FFT – experimental noisy data; NS + MF + FFT – matched filter with exponential envelope, $\exp(-t/T)$, followed by using FFT; CS + HI method, i.e., windowing, then noise reduction followed by using high-resolution harmonic inversion spectral estimator; CS + FFT method, i.e., windowed noise reduction followed by using FFT. Pure signal averaging is estimated to use 16 h of scanning for similar results.

The given foreknowledge was that an amide carbonyl was in the molecule and that it was expected to give the lowest amplitude signal of all the carbons. We knew it must be in the 161–175 ppm region, so we placed two windows there each of 512 Fourier grid points (FG pts). Here, 74 FG pts is approximately 1 ppm. At 256 transients ($N_{\text{tr}} = 256$) a converged result was obtained and “the experiment was stopped”. For presentation purposes the gap opening and spectral convergence in this region was easier to view in the window of Figure 2. The second column is shown to demonstrate convergence of the processing. The top two (row order) sections of this figure work with the original noisy signal (NS) and the latter two work with the cleaned signal (CS) obtained after the Cadzow noise reduction preprocessing (section 2.2). All but the third use the fast Fourier transform (FFT). The second row entry exemplifies the use of the matched filter (MF) variant of Fourier noise reduction processing. The addition of letters such as “CS” or “NS + HI” or “FFT” tells which signal and which spectral estimator is used, respectively. The harmonic inversion spectrum, not being on a Fourier grid, is drawn on a 1000-point grid in the window of the spectrum in the figures.

Figure 1 shows the SVD analysis that goes with Figure 2. Starting with the full given $N = 16\,384$ signal, 256 transients were needed before the “string of pearls” with no gap obtained previously for 128 and 192 transients converted to a case with $K = 1$ and a clear gap. At this point the S/N ratio was 1. Most useful at times is when the first quarter of the signal was used for noise reduction at both 128 and 192 transients. A $K = 1$ gap showed up, as these signals had less noise. Now if the full

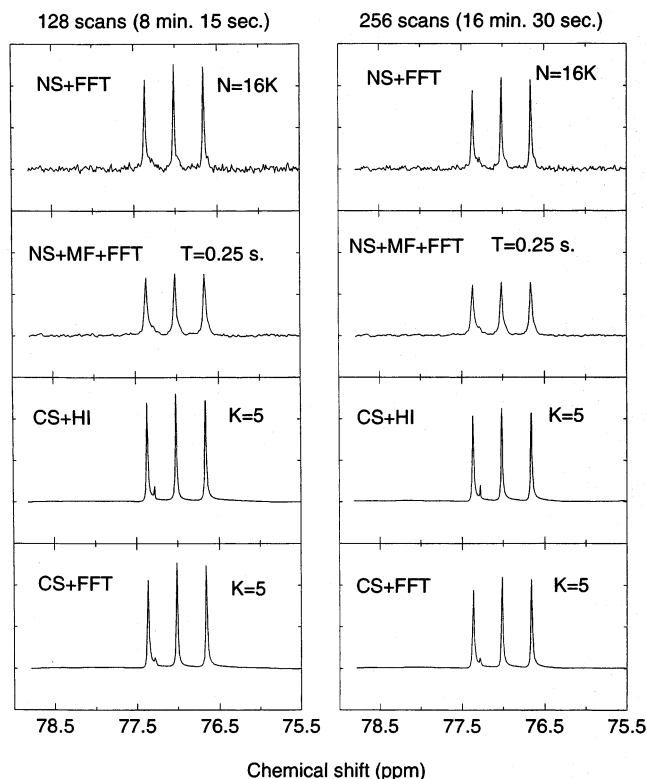


Figure 3. 1D ^{13}C NMR absorption spectrum of 15 nM solution of strychnine in CDCl_3 for deuterated chloroform triplet peak spectral region for noisy signal after 128 (left column) and 256 transients (right column) of length 16 K. The arrangement of plots is similar to that in Figure 2.

length signal for 128 and 192 transients was then assumed to have $K = 1$ and projection was done on the first eigenvector, Figure 2 shows excellent results were achieved. The quarter-length signal was too short to use throughout the analysis, but it was good enough to help in SVD analysis. The added information led us to accept the 256 transient result as “converged”. Last, since the method was new, for confidence an experiment with 2560 transients, which by signal averaging had a factor of about 3 higher S/N ratio, was done. Figure 1 shows it gave the same $K = 1$ result and the spectrum (not shown) except for a small line narrowing was the same as that for $N_{\text{tr}} = 256$. To study how the line position and width change with transients, calculations for $N_{\text{tr}} = 256, 384, 448,$ and 512 were done, and the deviation of the results of the first four from the fifth was taken. The average deviation was 0.001 ppm for the position at 168.646 ppm and 1.8×10^{-3} ppm for the width of 4.8×10^{-3} ppm. Such performance is consistent with the confidence spectrum and further analyses shown in the Appendix. The issue of line heights will be addressed there. Here we simply note that, for the type of line in Figure 2, the area is obtained with a precision of 2% and accuracy of better than 10%.

In routine work, not work aimed at demonstrating the various above ideas, in such cases as here, the user would just stick with $N = 16384$ until a gap opened up around $N_{\text{tr}} = 256$. A $K = 1$ would be read off. With no extra experimental work an $N = 8192$ or 4096 could be run to see that K is still 1. If it was not, the number of transients would be raised.

As a second example, a window was placed about solvent peaks as in Figure 3, because it was expected that if a chemical shift existed here it might be difficult to see due to the size and baseline of these peaks. The same 128 and 256 scan signal




Figure 4. Logarithmic SCD curves calculated for experimental data presented in left (128 scans) and right (256 scans) columns of Figure 3.

average sample data was used as before. We found a ^{13}C singlet at 77.3 ppm buried in the base of the deuterated chloroform triplet peaks. Figure 3 shows the results for the number of scans where a gap was evident (see Figure 4) and where the spectra converged (Figure 3) with respect to the strychnine peak. An important point should be made here. When using signal averaging at about eight scans, a gap corresponding to $K = 3$ emerges and the spectrum contains the three solvent peaks. At 128 scans two more break free and the shift peak appears in the spectrum. To test the convergence of the spectrum, 256 scans were processed to give a similar result. If lines of greatly different amplitudes exist, one cannot stop at the emergence of the first gap. The story is the same as with signal averaging, except that for the new method it all happens at many less scans. The spread of three high and two low singular values for $K = 5$ that can be seen in Figure 4 supports this argument.

Another point is illustrated here. The additional increase by two in number of signal singular values does not mean the peaks increase by two. The convergence of the spectrum also counts. In fact, if the number of transients increases by a factor of 10, we have seen that K can change to a value of 11 but the spectrum is essentially unchanged. The non-Lorentzian parts of the spectrum cause this phenomenon. By non-Lorentzian parts we mean the parts associated with the solvent. These peaks suffer from concentration and measurement effects and are not Lorentzian, especially in the overlapping solvent and sample peak regions and baseline regions. Our calculation of ω_k shows that the chemical shift peak is one pole. The non-Lorentzian solvent peaks require nine poles for representation. One extra pole is well away under the baseline.

The total spectrum is in Figure 5. For the rest of the spectrum using the $N_{\text{tr}} = 256$ signal, windows indicated in Figure 5 were put down and in each window a logarithmic singular value graph immediately showed a clear gap with a K value equal to the number of lines in each window as seen in Figure 5. The

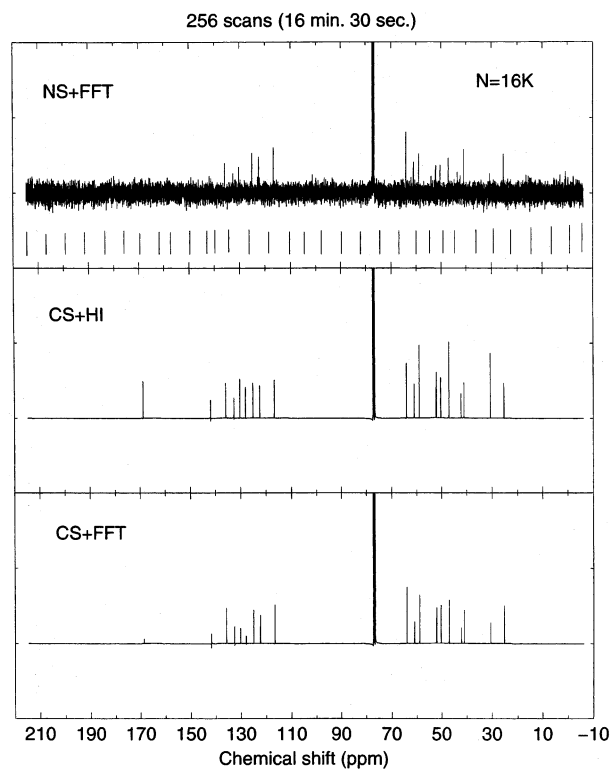


Figure 5. 1D ^{13}C NMR spectrum of 15 mM solution of strychnine in CDCl_3 with signal after 256 transients of length $N = 16\text{ K}$. Vertical lines on the top panel show window edges.

resulting windowed spectra were merged to form Figure 5. All 21 chemical shifts are now seen using 256 transients taken on our machine in 16 min 30 s. Windows where no signal showed up were recognized by the fact that *all* singular values dropped significantly in value, as would be predicted by perturbation theory, as more transients (and less noise) were used. When a window has both signal and noise as in Figures 1, 4, 7, and 9, the signal singular values stabilized and do not consistently and significantly drop in value as does the noise. This distinction is important because what appeared to be a gap with $K = 1$ did appear in one signal-free window. It was recognized as noise because it dropped in value precipitously between 128 and 256 transients.

To further demonstrate the method's ability to save transients, an FID signal from a 1D INADEQUATE experiment on 5% dicyclopentadiene in DCCl_3 was studied. Windows surrounding the fourth carbon at 54.75 ppm (see Figure 6) and the sixth carbon at 46.15 ppm (see Figure 8), in a number scheme that increases with decreasing ppm, were used as an example of a "typical" and a "difficult" case, respectively. Our foreknowledge was the position of the singlets and the order of magnitude of the splitting which made the choice of windows easy. We noted the fact that K could be slightly more than twice the number of splittings (which we pretended we did not know but which is not more than four) due to the fact that for the experiment done at our NMR facility the singlet central line was not totally wiped out and in fact could appear as several lines.

To start, we first tried 2176 scans (not shown) and both $N = 16\,384$ and $32\,768$ for the fourth carbon. The former gave $K = 7$; the latter gave $K = 6$. For the purposes of estimating K , we suspected the less noisy $K = 7$ was correct. To be sure, we raised number of transients N_{tr} to 2432 and 2560. Both gave a clear $K = 7$ at both half and full lengths and a converged spectrum as seen in Figure 6. Figure 7 shows the corresponding

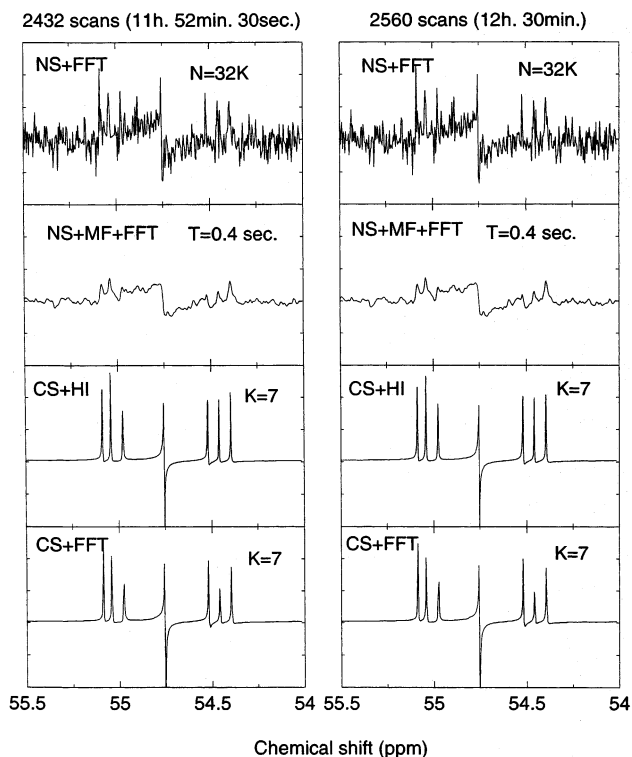


Figure 6. 1D INADEQUATE, spin-spin interaction ^{13}C - ^{13}C spectrum of 5% solution of dicyclopentadiene in CDCl_3 for the fourth carbon spectral region for a noisy signal after 2432 and 2560 scans of length $N = 32\text{ K}$. Pure signal averaging is estimated to use 2 weeks of scanning for similar results. The arrangement of plots in rows is similar to that in Figure 2.

SVD curves. The time used was 12 h as opposed to 2 weeks. The matched filter was of no use.

The sixth carbon, Figure 8, our last example (although all carbons were processed successfully) was confusing at $N_{\text{tr}} = 3712$ (our first estimate) because $N = 16\,384$ and $N = 32\,768$ not only had different rank, 7 and 8, respectively, they also had different spectra, with a fourth line on the right-hand side for $N = 32\,768$. As such, the number of transients was raised several times until at $N_{\text{tr}} = 4480$ the gap given by $N = 16\,384$ and $32\,768$ (Figure 9) was clearer and consistent. $N = 16\,384$ with less noise gave us confidence that the gap of $K = 7$ was correct for $N = 32\,768$. Here, as Figure 8 shows, spectral convergence was achieved and tested by using 4736 scans. For $N_{\text{tr}} = 4736$ the gap situation was the same as $N_{\text{tr}} = 4480$ (Figure 9).

The observed splitting of 28.9, 32.1, and 37.7 Hz is within 1.5 Hz (at worst) of that in a 95% solution spectra (which of course required quite a few less scans and may have concentration effects). It is quite hard to estimate the number of scans and the time needed for signal averaging plus FFT to resolve the middle line on the left-hand side, but for estimating the time, units of weeks would not have been out of place. The matched filter on the FFT could not resolve what turned out to be close peaks. It was unsatisfactory in general. This was as difficult an example as we and our colleagues could think of, as it was both "extremely noisy" and had lines that were difficult to resolve.

4. Discussion

The new processing methodology presented here envisions the use of signal averaging of transients until, for the case of an individual line or window, a gap appears and most importantly a spectrum becomes stable to increases in the

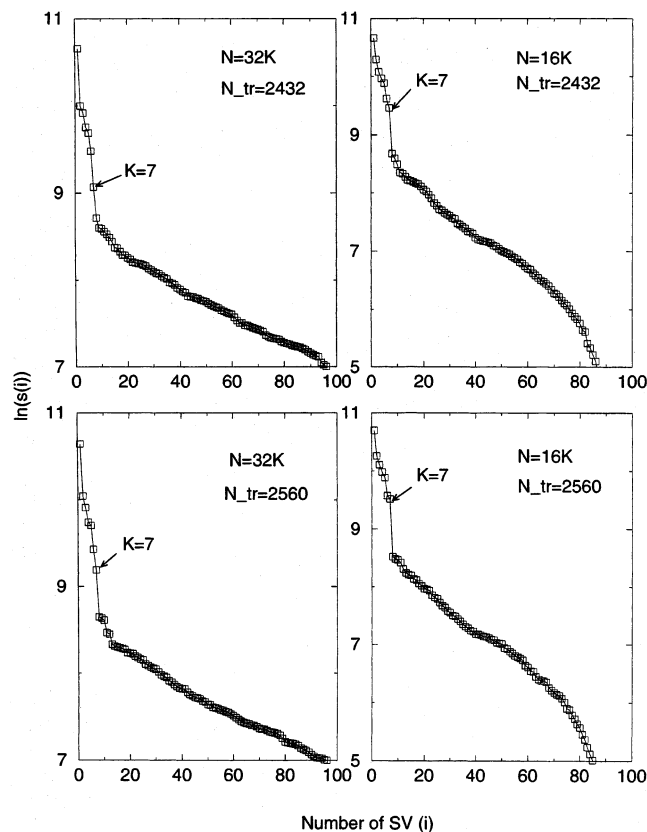


Figure 7. Logarithmic SVD curves calculated for experimental data presented in left (2432 scans) and right (2560 scans) columns of Figure 6.

number of transients. This says that the existence of a stable spectrum and of a gap will give a result of a quality that signal averaging would give with many more transients. Note K need not be constant when a Lorentzian line is sought in the presence of non-Lorentzian ones as our examples show.

The “cleaned” spectra reported here look unfamiliar in two senses. First, there is no baseline noise and, unlike when running an experiment using signal averaging followed by Fourier transform, no operator interactive way exists to know, based on the signal-to-noise ratio, when to stop the experiment and “accept” the spectrum. Second is that when harmonic inversion is also used (CS + HI in the figures) the lines seem unfamiliarly high and narrow. These are not real difficulties for the noise reduction or harmonic inversion methods but simply reflect that the information obtained is here obtained in different ways. For example, when harmonic inversion is used, the analysis allows the usual parameters of frequency position, width, area, and height of a Lorentzian peak (height = amplitude/width; amplitude = area/ π) to be read from the spectrum. The method also computes the frequency, width, and amplitude directly. As shown in the Appendix, the widths have the biggest error. When this error undervalues the width, the line height is too big. This is not a problem as the Appendix also shows that the directly extracted amplitude itself is significantly more accurate than the width. The amplitude is what is sought when areas under lines and peak intensities are studied.

The residual noise in the signal and the effect of artifacts are in the errors of the parameters and do not appear in the baseline spectra or line shape distortions. These errors are studied in the Appendix, where a confidence spectrum is given for the example of Figure 2. Interestingly, in the Fourier signal averaging methods artifacts are often ameliorated by a process that assumes a known reference signal and involves some steps common to

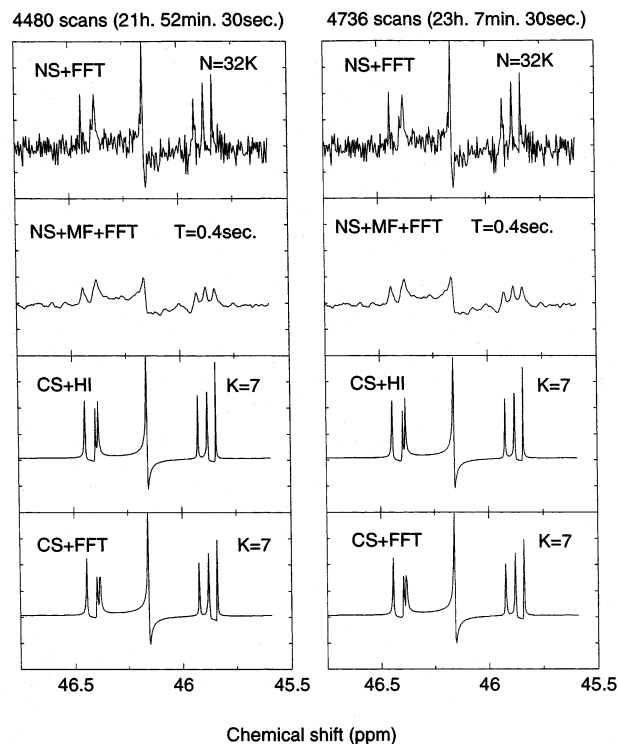


Figure 8. 1D INADEQUATE, spin–spin interaction ^{13}C – ^{13}C absorption spectrum of 5% solution of dicyclopentadiene in CDCl_3 for the sixth carbon spectral region for a noisy signal after 4480 (first column, $N = 32$ K) and 4736 scans (second column, $N = 32$ K). The arrangement of plots in rows is similar to that in Figure 2.

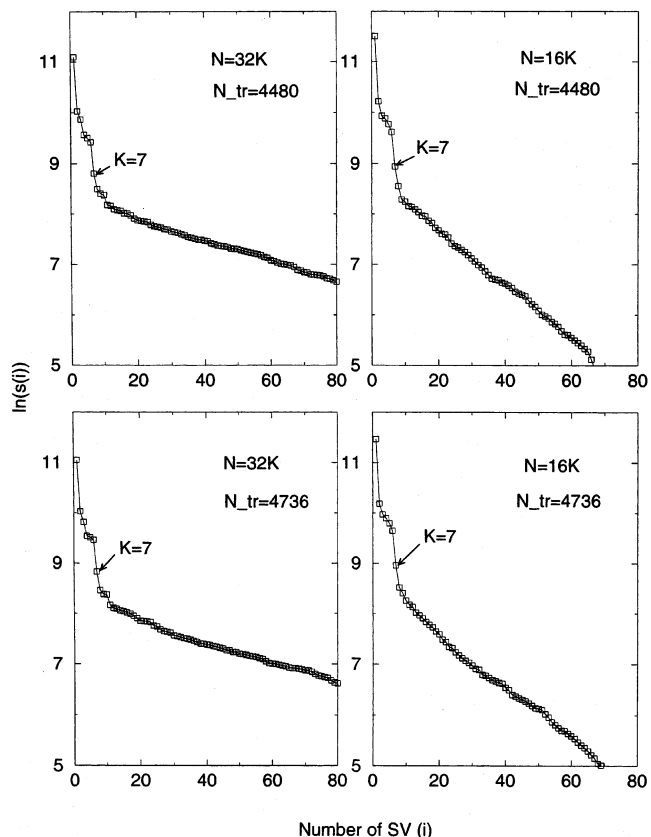


Figure 9. Logarithmic SVD curves calculated for experimental data presented in first (4480 scans, $N = 32$ K and 16 K) and second (4736 scans, $N = 32$ K and 16 K) columns of Figure 8.

our method as windowing and adjusting to an “ideal” experimental or model Lorentzian or Gaussian signal.²¹

The methodology used here breaks the spectral bandwidth into regions of smaller bandwidth (of the order of 200–500 Fourier grid points) called windows and carries out the analysis window by window, reassembling at the end the total or parts of the spectrum as required. This procedure reduces the size of matrices encountered in the analysis to at most 250×250 so that their diagonalization takes computing time on the order of seconds on today's desktop computers. Our experience is that the order of 30 s is needed per window because of an internal iteration inside the Cadzow procedure. The number of windows is on the order of, say, 50 for the total spectrum, so assuming, as certainly is possible, software is written to control the mechanics of applying the windowing procedure, a total analysis should take on the order of 30 min. Even when one considers the FFT processing to be essentially instantaneous, the tradeoff of using the procedures of this paper, which as the case may be will be seen to save from hours to weeks of machine time, can well be considered as favorable. Often, as in ^{17}O NMR spectra, only one window is studied and the total computing time is a minute or so.

Importantly when replacing the usual methods of spectral analysis by the new analysis for the case of experiments repeated under similar parameters such as sample concentration, signal length, and perhaps the experience based specification of the number of transients, user interaction can be altogether avoided after a few initial experiments. The initial experiments would follow the total algorithm of this paper to determine the number of transients needed to satisfy the threshold for cutoff criteria and use it without cycling in future experiments. For example, the fractional reduction in transients needed relative to the usual analysis using only signal averaging is, for ^{13}C chemical shift spectra, on the order of 1/60. If now one is willing to accept a fraction of 1/30, one could assume with certainty that the signal satisfies the criterion for cleaning by Cadzow regularization and thereby avoid user interaction. Similarly, if one knows in advance the weakest signal in a spectrum and works first in a window about it as is the case in Figure 2, the number of transients used in this first window to satisfy our criteria is guaranteed to be enough in other windows, again avoiding much of the cycling. To further eliminate the everyday need for user interaction, tables of N_{tr} cutoff threshold that depend on nuclear isotopes involved, molecular concentration, functional group, etc. are expected to follow acceptance of this methodology. In our experience, for threshold determination for a class of experiments, it is better to carry a converged Fourier analysis and to work backward to determine transient thresholds.

The factor 64 of time saving for the ^{13}C case comes from the observation that for single isolated lines, when the gap opens and the windowed noise reduction processing begins, the intensity of lines near true signal lines, whose position we eventually know, is roughly 1:1. $S/N = 8$ is desired, hence the factor 64 in experimental efficiency relative to any given signal.

Is it possible to miss a low-intensity line using the signal "cleaning" scheme? The answer is exactly as in signal averaging and is "yes" if you do not scan enough. Here we have a factor of roughly 64 advantage and can claim that the method introduced here will miss fewer lines than are missed in the signal averaging plus FFT processing scheme for a given number of scans. Of course, for a series of lines the above discussion holds for the line of smallest expected intensity.

Can "extra" and fake peaks appear as in the past when SVD-based methods were used? The answer is "no" when the criteria for the signal "cleaning" scheme are fulfilled and allow cycling to be stopped. This might occur before this point. For example,

this occurred as we converged to the sixth carbon in the INADEQUATE spectrum. The half-length signal for $N_{\text{tr}} = 3712$, which was below convergence in CS + HI, gave what turned out to be almost identical with our final converged spectra, but the full-length signal gave an increased K and an extra line. This all went away with increasing N_{tr} .

It is the open gap and the convergence of the spectrum as a function of signal averaging that matters most. Just as in pure signal averaging, any uncertainty is met by averaging more transients. Still, big savings are made: no great price is paid if a factor of 64 is reduced for confidence to a factor of say 50 or even 25. The big difference between how the harmonic inversion model was used in the past and how it is used here is a determination not to stick with the given signal and its S/N ratio but to converge the spectrum with respect to the number of transients once a gap opened; the gap in turn was "forced" to open by increased signal averaging and the increasing S/N ratio. *Essentially the signal was rejected until it was "cleaned enough"*. The fact remains that with all this increasing of signal averaging done here the spectrum is obtained with many times fewer transients than with pure signal averaging.

The noise reduction method may be able to be extended to other signals if a signal vector–subspace relationship such as eq 2 can be developed. This can be done for some multiple-dimension basically Lorentzian signals, and further development is being pursued in this laboratory. Here, as Mandelshtam and Shaka have shown,^{22,23} the high-resolution feature of harmonic inversion can contribute for many reasons with the most obvious of them being the ability to resolve frequencies in the "short acquisition direction".

The method used here gives the same type of savings in scans and improvements in resolution for proton NMR, but of course the real time savings are not as spectacular since these experiments do not take such long times in the first place.

The ability to shorten the time needed to produce a "cleaned" spectrum bodes well for carrying out experiments that only scan a short time. It also bodes well for observing spectra of those unstable species that survive the now shortened measurement time but do not survive the longer time. A lower "enrichment" of isotopes may be needed using the "cleaned" signal plus harmonic inversion method.

Even though spectra are Lorentzian, if the Lorentzians overlap significantly, parameter estimation from the spectrum is problematic. Since the methods of this paper should work in this case, the harmonic inversion method should yield the parameters directly without viewing the spectra. Some very preliminary work on ^{17}O – ^{13}C splittings attests to this.²⁴

Clearly the model, eq 2, requires $K \leq N_d/2$. If this is not true, the number of unknowns (d_k, ω_k) exceeds the number of givens, c_n^{NR} . Depending on field strength, this may not be so for big molecules such as proteins, and this could extend the limitations of our method to treating Lorentzian spectra for molecules with a quite high density of lines. Proteins may be out, but only further study will make this clear. At least care must be taken and experience gained if this method is to be applied to dense spectra. The happy thought is that despite the enormous current interest in proteins, the large majority of everyday NMR experiments in chemistry laboratories are of a type we can impact.

The most original aspect of this work is the recognition that using successively the signal averaging and the noise reduction preprocessor in conjunction with the "SVD gap" criteria of section 2.2 for switching between the two yields an acceptable spectrum with many fewer transients. The individual compo-

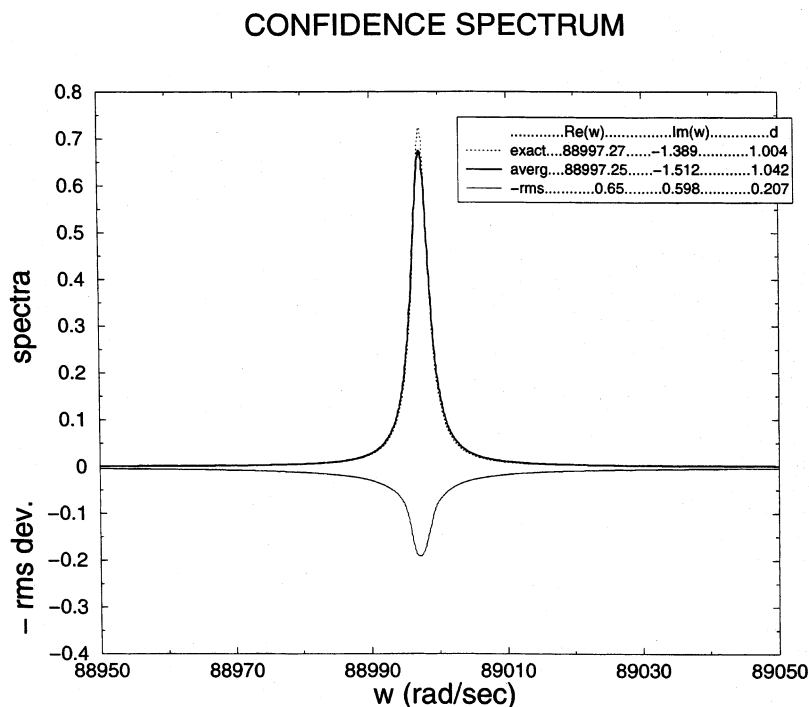


Figure 10. HI confidence spectrum obtained from 2000 independent white noise realizations added to a single Lorentzian peak with parameters shown in the insert and $S/N = 1.2$.

nents of the algorithm such as windowing, signal averaging, noise reduction preprocessing, and the harmonic inversion processing were all separately available. As is often the case for many advances, this advance in shortening the measurement time combines the components in a way they have not been previously combined to produce the reported algorithm. We know of no other algorithm that can achieve this time reduction for real experiments.

Acknowledgment. The authors are indebted to Prof. G. K. S. Prakash for guidance as to where our method could be of use in NMR. We thank A. Kershaw for supplying us with experimental data and advice. Dr. J. Hoch and Prof. C. McKenna are thanked for supplying certain test signals and advice. This work has been supported by NSF Grant PHY-0071742. H.S.T. thanks the Max Planck Institute for Complex Systems in Dresden, Germany, for making a visit possible during which Dr. Holgar Kantz made him aware of his work¹⁵ on noise reduction. Dr. E. Atilgan has helped us significantly with the study in the Appendix.

Appendix: Confidence Spectra for the Single Lorentzian

To model our process of analysis using the same bandwidth and signal length as in Figure 2, random white noise perturbations were added to a single damped exponential signal to form 2000 independent noisy realizations chosen to satisfy three criteria simultaneously. First, a gap as in Figure 1 had to be evident; second, the realizations with $S/N \approx 1.2$ had to show locally about where we knew the true peak to be, a roughly 1:1 “signal peak” to noise peak intensity; third, the resulting $\text{Im}(\omega_k)$, width parameter, had to be below -0.5 , i.e., $\Gamma \geq 1$ Hz, because when very narrow peaks occasionally occur, a small change in N_{tr} , something done in practice all the time in this method, wipes them out. That is, a too narrow peak is double checked and its width (not existence) is almost always not robust. The result for harmonic inversion, Figure 10, shows that obtaining the width is the main source of error and, since the maximum peak

height is $\text{Re}(d_k)/|\text{Im}(\omega_k)|$, causes the somewhat large root-mean-square (rms) deviations at the line center in Figure 10. The difficulties with the $\text{Im}(\omega_k)$ relative to those of the residue d_k and position $\text{Re}(\omega_k)$ are generic to all HI methods and also all methods in computational physics which compute for nontrivial examples, poles in the complex plane. Such methods include the direct calculation of resonance contributions to the cross sections in scattering theory and pure scattering calculations where resonances are then fit to Lorentzian parameters. What helps for harmonic inversion, as seen in the insert in Figure 10, is that d_k , the actual amplitude of the signal, is given to an accuracy of about 4%, albeit with a precision of 20%. Changing the S/N ratio from 1.2 to 1.7, equivalent to a doubling of the number of scans, improves these numbers to 1% and 15%, respectively.

For CS + FFT again the $\text{Re}(\omega_k)$ frequency is excellent, the imaginary part is poor, but the amplitude d_k , here the area/ π , has an accuracy of 10% and a precision of 2%. The accuracy estimates given here may be a bit pessimistic because our results are compared to the exact model parameters and not to parameters obtained by signal averaging. The latter may have canceling leakage and baseline estimation errors.

Harmonic inversion used on model signals that have no noise gives essentially perfect parameter values. As such, it should be no surprise that the rms error decreases with increasing S/N ratio, i.e., more transients. Future studies of this functional dependence may suggest the costs involved in terms of increasing numbers of scans needed to obtain various levels of confidence.

Different confidence levels will be associated with each model spectrum’s features; e.g., near degenerate $\text{Re}(\omega_k)$ and “small” Lorentzians at the base of large features will give lower confidence levels. In the future extensive confidence studies will be made of such cases as well of others that are commonly encountered.

Testing for confidence by decreasing in successive problems a high S/N ratio until the method works does not make sense

here, because our method essentially gives a threshold of detectability as roughly locally 1:1.

References and Notes

- (1) Hoch, J. C.; Stern, A. S. *NMR Data Processing*; Wiley-Liss: New York, 1996.
- (2) Stephenson, D. S. *Prog. NMR Spectrosc.* **1998**, *20*, 515–626.
- (3) Press, W. H.; Teukolsky, S. A.; Vetterling, W. T.; Flannery, B. P. *Numerical Recipes*, 2nd ed.; Cambridge University Press: Cambridge, 1992.
- (4) Main, J.; Dando, P. A.; Belkic, Dz.; Taylor, H. S. *J. Phys. A: Math. Gen.* **2000**, *33*, 1247–1263.
- (5) Wall, M. R.; Neuhauser, D. *J. Chem. Phys.* **1995**, *102*, 8011–8022.
- (6) Grozdanov, T. P.; Mandelshtam, V. A.; Taylor, H. S. *J. Chem. Phys.* **1995**, *103*, 7990–7995.
- (7) Mandelshtam, V. A.; Taylor, H. S. *J. Chem. Phys.* **1997**, *107*, 6756–6769.
- (8) Mandelshtam, V. A. *Prog. Nucl. Magn. Reson. Spectrosc.* **2001**, *38*, 159.
- (9) Chen, J.; Mandelshtam, V. A.; Shaka, A. J. *J. Magn. Reson.* **2000**, *147*, 129–137.
- (10) Chen, J.; Mandelshtam, V. A.; Shaka, A. J. *J. Magn. Reson.* **2000**, *146*, 363–368.
- (11) Koehl, P. *Prog. Nucl. Magn. Reson. Spectrosc.* **1999**, *34*, 257.
- (12) Hanke, M.; Hansen, P. C. *Surv. Math. Ind.* **1993**, *3*, 253.
- (13) Cadzow, J. A. *IEEE Trans. Acoustics, Speech. Signal Processing* **1988**, *36*, 49–62.
- (14) Dunkel, R. Correction and Automated Analysis of Spectral and Imaging Data. U.S. Patent 5,572,125, Nov 5, 1996.
- (15) Hegger, R.; Kantz, H.; Matassini, L. *IEEE Trans. Circuits Syst. I* **2001**, *48*, 1454–1461.
- (16) Depretterre, Ed. F. *SVD and Signal Processing: Algorithms, Analysis and Applications*; Elsevier/North-Holland: New York, 1988.
- (17) Golub, G. H.; Van Loan, C. *Matrix computations*, 3rd ed.; Johns Hopkins University Press: Baltimore, MD, 1996.
- (18) Totz, J.; Vandenboogaart, A.; Van Huffel, S.; Graverondemilly, D.; Dologlou, I.; Heidler, R.; Michel, D. *J. Magn. Reson.* **1997**, *124*, 400–409.
- (19) Dologlou, I.; Van Huffel, S.; Van Ormondt, D. *IEEE Trans. Signal Processing* **1997**, *45*, 799–803.
- (20) Gantmacher, F. R. *The Theory of Matrices*; Chelsea Publishing Co.: New York, 1974.
- (21) Morris, G. A. *J. Magn. Reson.* **1988**, *80*, 547.
- (22) Mandelshtam, V. A. *J. Magn. Reson.* **1998**, *144*, 343–356.
- (23) DeAngelis, A. A.; Hu, H.; Mandelshtam, V. A.; Shaka, A. J. *J. Magn. Reson.* **1998**, *144*, 357–366.
- (24) Work in progress with Prof. C. McKenna.