JOURNAL OF MAGNETIC RESONANCE **134,** 76–87 (1998)
ARTICLE NO. MN981516

Reference Deconvolution, Phase Correction, and Line Listing of NMR Spectra by the 1D Filter Diagonalization Method

Haitao Hu, Que N. Van, Vladimir A. Mandelshtam, and A. J. Shaka[1]

Chemistry Department, University of California, Irvine, California 92697-2025

Received April 21, 1998

We describe a new way to attack the problem of identifying and quantifying the number of NMR transitions in a given NMR spectrum. The goal is to reduce the spectrum to a tabular line list of peak positions, widths, amplitudes, and phases, and to have this line list be of high fidelity. In this context "high fidelity" means that each true NMR transition is represented by a single entry, with no spurious entries and no missed peaks. A high fidelity line list allows the measurement of chemical shifts and coupling constants with good accuracy and precision and is the ultimate in data compression. There are two parts to the problem. The first is to overcome common imperfections: the non-Lorentzian lineshapes that can arise whenever the magnetic field inhomogeneity is less than perfect, and nonzero time delays that cause frequency-dependent phase errors. The second is to fit the spectral features to a model of Lorentzian lines. We use the recently developed filter diagonalization method (FDM) to accomplish the reference deconvolution, the phase correction, and the fitting, and show good progress toward the goal of obtaining a high fidelity line list.

Key Words: **reference deconvolution; filter diagonalization method; FDM; NMR; line list; phase correction; linear prediction.**

INTRODUCTION

NMR spectroscopists spend a great deal of time analyzing Fourier transform (FT) NMR spectra. In some cases a "clean" spectrum with a high signal-to-noise (S/N) ratio can be used as evidence of chemical purity; in others it may be indicative of good spectrometer performance, as in a lineshape demonstration. Part of the strength of the FT *spectrum* is that it reveals a wealth of information and, as the limitations of the FT itself are well understood, is considered to be highly reliable. The weakness of the FT spectrum is that it is simply a graph of amplitude versus frequency: It must be analyzed somehow to pick out peaks, measure frequencies and linewidths, integrate multiplets, measure couplings, etc., which can occupy considerable operator time and may be prone to hidden errors that are less well understood.

For many spectra, a high fidelity tabular line list, in which each true peak was correctly identified and no spurious entries appeared, would offer a complete summary of the information content of the signal, excluding noise. The ultimate utility of such a line list depends to some extent on how well the spectrum can be modeled as a sum of discrete peaks rather than, say, just a function of frequency. In most cases, however, there are discrete peaks and so the utility of the line list will be high. In this article we describe significant progress made toward generating a high fidelity line list for liquid-state NMR spectra automatically, efficiently, and with little or no operator intervention. We show that by using the filter diagonalization method (FDM) it is possible to correct for common instrumental shortcomings such as even fairly large nonzero delays after the excitation pulse and imperfect shimming of the magnet. Furthermore, it is possible to estimate the reliability of the FDM line list.

THEORY

The Filter Diagonalization Method

Because FDM is new to NMR, we will briefly describe its structure and function. The idea of FDM was introduced by Wall and Neuhauser in 1995 as a method for spectral analysis of time correlation functions in quantum dynamics calculations (*1*). Very recently it has been improved (*2, 3*) and extended to the case of model multidimensional time signals (*4*) and large 2D NMR data sets (*5*). The object of 1D FDM is to fit a given complex time signal $c_n = C(t_n)$ defined on an equidistant time grid $t_n = n\tau$, $n = 0, 1, \ldots, N-1$, to the sum of exponentially damped sinusoids,

$$c_n = \sum_{k=1}^{K} d_k e^{-in\tau\omega_k} = \sum_{k=1}^{K} d_k e^{-2\pi in\tau f_k} e^{-n\tau\gamma_k} \qquad [1]$$

with a total of $2K$ unknowns: the K complex amplitudes d_k, and the K complex frequencies $\omega_k = 2\pi f_k - i\gamma_k$, which include damping. Even though the fitting problem of Eq. [1] is highly nonlinear, its solution can be obtained by pure linear algebra. The implicit assumption of a Lorentzian lineshape and uniform sampling rate allows this simplification, which is

[1] To whom correspondence should be addressed. E-mail: ajshaka@uci.edu.

1090-7807/98 $25.00

shared by other "high resolution" methods including linear prediction (LP) (*6–14*), the Prony method, MUSIC, and ESPRIT (*15, 16*). What makes FDM unique is to associate the complex signal c_n to be fitted by the form of Eq. [1] with the time autocorrelation function of a fictitious dynamical system described by an effective complex symmetric Hamiltonian operator $\hat{\Omega}$ with complex eigenvalues $\{\omega_k\}$ (*1*)

$$c_n = (\Phi_0|e^{-in\tau\hat{\Omega}}\Phi_0) \qquad [2]$$

so that the problem of fitting the observed time signal to the form of Eq. [1] becomes equivalent to diagonalizing the Hamiltonian $\hat{\Omega}$ or, equivalently, the evolution operator (*3*) over a single time step, $\hat{U} = \exp(-i\tau\hat{\Omega})$. In Eq. [2] a complex symmetric inner product is used, $(a|b) = (b|a)$ without complex conjugation, and Φ_0 is some "initial state." The key difference between FDM and, e.g., LP, is that in FDM we can choose a basis just as in any other problem in quantum mechanics. Furthermore, the basis can be chosen to make the procedure numerically robust and far more efficient, so that the numerical effort scales quasi-linearly with the size of the data set, making FDM comparable to the FFT itself in this important regard.

Suppose that there is a set of orthonormal eigenvectors, $\{Y_k\}$, that diagonalizes $\hat{U} = \exp(-i\tau\hat{\Omega})$. Then we can write

$$\hat{U} = \sum_k u_k|Y_k)(Y_k| = \sum_k \exp(-i\tau\omega_k)|Y_k)(Y_k| \qquad [3]$$

and inserting Eq. [3] into Eq. [2] we can obtain Eq. [1] if we let

$$d_k = (\Phi_0|Y_k)(Y_k|\Phi_0) = (Y_k|\Phi_0)^2. \qquad [4]$$

The *eigenvalues* thus determine *line positions and widths*, and the *eigenvectors* determine *line amplitudes and phases.* That is, one diagonalization of $\hat{U}$ yields all of the desired parameters. This is again in contrast to LP, in which a singular value decomposition (SVD) or similar procedure is used to find the LP coefficients, a polynomial rooting or equivalent procedure is then used to determine line frequencies and widths, and then *another* least squares problem, using these roots as input, must be solved to determine amplitudes and phases (*14*). Even when computation of the LP coefficients is made efficient (*12*), the overall algorithm is still unacceptably slow for large signals.

Even though $\hat{U}$ is not explicitly available, its matrix elements in an appropriately chosen basis are completely determined by the measured values c_n. The basis, which is not necessarily orthonormal, can be chosen in many different ways. Chosen naïvely, the number of basis functions would determine the size of the linear system which must be solved. For an FID of length N complex points, the maximum basis size is $N/2$: Each Lorentzian line requires two complex numbers to specify it, so that at most $N/2$ lines can be uniquely fitted to a signal of length N.

Consider the simplest basis, the Krylov vectors generated by the evolution operator: $\Phi_n = \hat{U}^n\Phi_0 = \exp(-in\tau\hat{\Omega})\Phi_0$. Based on Eq. [2] the matrix elements of $\hat{U}$ in this basis are trivial to obtain as

$$(\Phi_n|\hat{U}\Phi_m) = (\Phi_n|\Phi_{m+1}) = c_{m+n+1}, \qquad [5]$$

but since the basis is not orthonormal, the overlap matrix

$$(\Phi_n|\Phi_m) = (\hat{U}^n\Phi_0|\hat{U}^m\Phi_0) = (\Phi_0|\hat{U}^{m+n}\Phi_0) = c_{m+n}, \qquad [6]$$

which is again trivially related to the measured signal, must be assembled. It is convenient to adopt notation in which $\mathbf{U}^{(0)}$ is the $(M+1)\times(M+1)$ matrix representation of the overlap matrix, and $\mathbf{U}^{(1)}$ is that of $\hat{U}$. The fitting problem of Eq. [1] is then cast into a generalized complex symmetric eigenvalue problem

$$\mathbf{U}^{(1)}\mathbf{B}_k = u_k\mathbf{U}^{(0)}\mathbf{B}_k, \qquad [7]$$

whence the eigenvalues $u_k = \exp(-i\tau\omega_k)$ give the line positions and widths, and the eigenvectors $\mathbf{B}_k$ the intensities and phases:

$$d_k^{1/2} = \sum_{n=0}^{M} B_{nk}c_n, \qquad [8]$$

which follows from Eq. [4] by substituting

$$Y_k = \sum_{n=0}^{M} B_{nk}\Phi_n. \qquad [9]$$

Note that the eigenvectors $\mathbf{B}_k$ are normalized with respect to the overlap matrix $\mathbf{U}^{(0)}$, i.e.,

$$\mathbf{B}_k^{\mathrm{T}}\mathbf{U}^{(0)}\mathbf{B}_j = \delta_{kj}. \qquad [10]$$

A Single Line

It is instructive to follow through the calculation based on the simplest possible signal consisting of a single damped sinusoid of arbitrary initial phase with no noise, for which only two time points have been measured, $t = 0$ and $t = \tau$, so that

$$c_0 = Ae^{i\phi};\ c_1 = Ae^{i\phi}\exp(-i\omega_0\tau), \qquad [11]$$

where the frequency ω_0 is complex, $\omega_0 = 2\pi f_0 - i\gamma$, and $\gamma > 0$. There is just a single basis function Φ_0 so that the

eigenvalue problem of Eq. [7] boils down to a 1×1 matrix problem

$$(\Phi_0|\hat{U}\Phi_0)(\Phi_0|Y_0) = u_0(\Phi_0|\Phi_0)(\Phi_0|Y_0), \qquad [12]$$

which, using Eqs. [5] and [6], becomes

$$c_1 B_0 = u_0 c_0 B_0, \qquad [13]$$

showing that the complex eigenvalue u_0 is $c_1/c_0 = \exp(-i\omega_0\tau)$, from which ω_0 can be obtained if τ is known. The eigenfunction Y_0 is just proportional to Φ_0 itself, and the normalization condition, $B_0 c_0 B_0 = 1$, for the eigenvector $\mathbf{B}_0$ gives

$$B_0 = \frac{1}{\sqrt{c_0}}. \qquad [14]$$

Now using Eq. [8] we recover the expected result:

$$d_0 = (B_0 c_0)^2 = c_0 = Ae^{-i\phi}. \qquad [15]$$

This illustrative example does in fact reveal one interesting aspect: The line is perfectly identified and characterized even though only two time-domain data points have been measured. There are many situations in which this might prove useful, and in fact 2D spectra in which only two t_1 time points were used have already been obtained (*5*). A two-point FFT of course gives essentially no useful information: Even if zero-filled, some apodization of the second point is required! The output of FDM is the list of spectral parameters, in this case ω_0 and d_0, but an artificial or "ersatz" spectrum can be computed from them, with frequency points spaced as closely as one likes. This ability to apparently get something for nothing, or at least get more information from the measured data, is part of the attraction of all of the "high resolution" methods such as LP. The disadvantage of these methods is fairly obvious if one includes the effect of noise in the two measured points. The ersatz spectrum will still appear noiseless, but the parameters themselves will be in error. A slight noise spike on the c_1 point could even result in a "negative" linewidth. Unless there is some indication of how reliable the spectral parameters are, they cannot be completely trusted, and the experimentalist will, quite rightly, want to compute the FFT spectrum to get a feel for what the data "really" look like.

Many-Line Spectra and the Rectangular Window Fourier Basis

Suppose now that a long FID of N complex points has been acquired, which we want to fit by the form of Eq. [1]. The complex symmetric **U** matrices in the Krylov basis are very simple in structure, just consisting of successive rows shifted by a single data point,

$$\mathbf{U}^{(0)} = \begin{bmatrix} c_0 & c_1 & c_2 & \cdots & c_M \\ c_1 & c_2 & c_3 & \cdots & c_{M+1} \\ c_2 & c_3 & c_4 & \cdots & c_{M+2} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ c_M & c_{M+1} & c_{M+2} & \cdots & c_{M+M} \end{bmatrix};$$

$$\mathbf{U}^{(1)} = \begin{bmatrix} c_1 & c_2 & c_3 & \cdots & c_{M+1} \\ c_2 & c_3 & c_4 & \cdots & c_{M+2} \\ c_3 & c_4 & c_5 & \cdots & c_{M+3} \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ c_{M+1} & c_{M+2} & c_{M+3} & \cdots & c_{2M+1} \end{bmatrix}, \qquad [16]$$

and $\mathbf{U}^{(1)}$ is simply $\mathbf{U}^{(0)}$ with an entire column right-shifted. One can immediately recognize in these matrices the well-known "data matrices" of LP. As such, solving Eq. [7] might be reminiscent of some formulations of the LP-based methods. We note, however, that the real power of FDM will be due to the choice of a nonprimitive basis. We show the primitive basis formulation only to relate FDM to previous linear algebraic approaches, and because of its extreme simplicity.

Armed with Eqs. [7] and [16] the problem is certainly easy to set up, but it is unfortunately very difficult to solve. The **U** matrices are dense and far from diagonal. Furthermore, a routine 3 s acquisition time over a 10 ppm spectral width at 500 MHz would lead to $N \sim 15{,}000$ and $M \sim 7500$. The numerical effort for an eigenvalue problem generally scales as M^3, so that the computation takes forever and the memory requirements to set up the two 7500×7500 matrices are excessive. (Some simplification is possible taking into account the Toeplitz structure of the matrices; see Ref. *12*). If the true number of lines K happens to be far fewer than 7500, then the basis is overcomplete and $\mathbf{U}^{(0)}$ becomes singular (except for noise). The net result of these practical considerations is that the fitting problem formulated in this way is a huge and ill-conditioned linear system, with output that would be unreliable using double precision arithmetic even if computation time were immaterial. Note that exactly the same problems plaque parametric LP methods, where the filter length plays a similar role to M and it is known that, particularly when there is nonnegligible noise, the former should exceed the true (unknown) number of lines by a substantial margin (*14*) and where the matrix diagonalization is replaced by the determination of the M LP coefficients and the rooting of an Mth-order polynomial, followed by a huge linear least squares problem.

Fortunately, the primitive Krylov basis functions are not the only choice. Any linear combination of them will serve for a basis. In fact, the Fourier transform itself is just a recipe for taking linear combinations of data points in time to arrive at a data point in the frequency domain. Unlike the time points, the frequency points have little interaction when they are well separated, at least for linewidths narrower than the separation,

precisely because the spectrum contains "peaks" instead of "wiggles." Thus, a good choice is a Fourier basis (*1*), of which a particularly simple and efficient variant is the rectangular window Fourier basis (*2*):

$$\Psi(\varphi) = \sum_{n=0}^{M} e^{in\varphi}\Phi_n \equiv \sum_{n=0}^{M} e^{in(\varphi-\hat{\Omega}\tau)}\Phi_0, \qquad [17]$$

with a set of φ values taken to correspond to an evenly spaced frequency grid. The matrix elements of the operator $\hat{U}^p = \exp(-ip\tau\hat{\Omega})$, $p = 0, 1, 2, \ldots$, between any two functions $\Psi(\varphi)$ and $\Psi(\varphi')$, defined by Eq. [17], can also be evaluated only in terms of the measured signal c_n. This can be shown using the ansatz of Eq. [2]:

$$\begin{aligned}\mathbf{U}^{(p)}(\varphi',\varphi) &\equiv (\Psi(\varphi')|\hat{U}^p\Psi(\varphi)) \\ &= \sum_{n'=0}^{M}\sum_{n=0}^{M} e^{in\varphi}e^{in'\varphi'}(e^{-in'\tau\hat{\Omega}}\Phi_0|e^{-i(n+p)\tau\hat{\Omega}}\Phi_0) \\ &= \sum_{n'=0}^{M}\sum_{n=0}^{M} e^{in(\varphi-\varphi')}e^{i(n+n')\varphi'}c_{n+n'+p}. \qquad [18]\end{aligned}$$

Since for a finite Fourier series the basis $\{\Psi(\varphi)\}$ is not orthonormal, we must again compute the overlap matrix $(\Psi(\varphi)|\Psi(\varphi'))$, which corresponds to $p = 0$.

There is a huge difference in moving to the Fourier basis. The structure of $\mathbf{U}^{(0)}$, $\mathbf{U}^{(1)}$ is now diagonally dominant, with off-diagonal matrix elements that decay with a roughly sinc-like dependence along the minor diagonal, as shown schematically in Fig. 1. The diagonalization is thus far easier, and can be carried out, for example, in a block-diagonal fashion by simply ignoring all correlations between points in the frequency domain which are far away from each other. Thus, the eigenvalue problem is (i) diagonally dominant and (ii) of a small size (in practice $K_{win} \sim 3$–100) over each spectral window. In particular, K_{win} remains small even though M, the number of total lines used for the fit, and N, the length of the FID, are both huge. An FID that is twice as long can be analyzed by using twice as many windows, with matrices each of which take roughly the same computational effort to diagonalize, and which yield stable and reliable eigenvalues. One of the two summations in Eq. [18] can be evaluated analytically (*3*), and therefore the numerical effort to construct the matrix representations of the **U** matrices, which are obtained essentially from an FFT of the FID itself, scales as $N \log N$. Thus, the entire numerical effort to obtain the spectral parameters over the entire spectral width scales as $N \log N$ with a prefactor which is somewhat larger than that of the FFT itself. This change in scaling from cubic to quasi-linear makes FDM far preferable to any form of conventional LP and allows the fitting of huge signals containing myriads of lines.

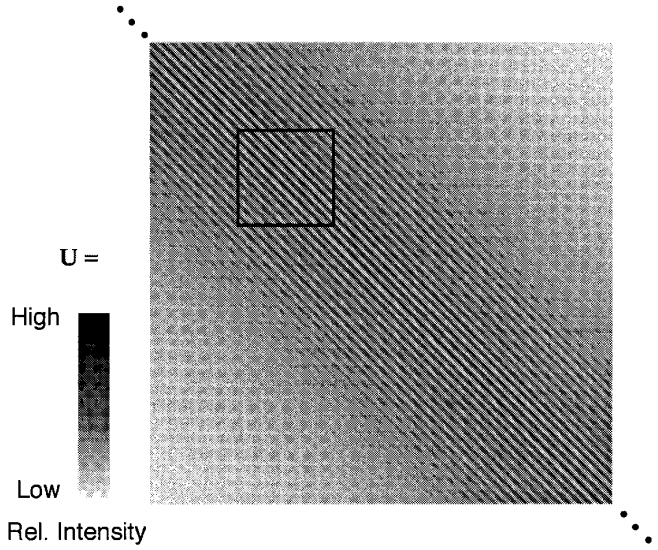

FIG. 1. A pictorial representation of the structure of the matrix representation of the evolution operator in the rectangular window Fourier basis. Elements along the minor diagonal decay with a roughly sinc-like dependence. This structure allows the diagonalization to be carried out over separate small windows, one of which is indicated by the box in the figure. Neglecting *all* off-diagonal elements gives, after solving the generalized eigenvalue problem, a spectrum reminiscent of the Fourier transform spectrum.

Although the formulation of the fitting problem as a generalized eigenvalue problem is compact and elegant, the Fourier basis is probably the most important aspect of FDM that makes it numerically superior to other linear algebraic methods of solving the fitting problem of Eq. [1]. The local nature of the Fourier basis allows solution for the frequencies ω_k in a small *a priori* chosen frequency interval, which operationally involves only small matrices corresponding to a small (and only locally complete) basis $\{\Psi(\varphi_j)\}$, $j = 1, 2, \ldots, K_{win}$. This greatly reduces the numerical effort for both matrix evaluation and for solution of the resulting small generalized eigenvalue problem. Interestingly, the LP-ZOOM method, first proposed in 1988, which is a modification of parametric LP using the z-transform to zoom in on a particular spectral region (*17*), theoretically shares many of the advantages of FDM for 1D spectral analysis. LP-ZOOM is numerically distinct from FDM, as it involves more adjustable parameters and the solution of more than one least squares problem, but is similar in spirit.

A detailed theoretical treatment of the version of 1D FDM that we employ can be found in the literature (*3, 4*). The practical numerical procedure can be summarized by the following steps:

1. Starting at the lowest frequency available according to the Nyquist criterion, choose a small enough frequency window $[f_{min}, f_{max}]$ for the spectral analysis of a given signal c_n

$= C(t_n)$, $n = 0, 1, 2, \ldots, N - 1$. "Small enough" is operationally defined by the resulting size of the linear algebraic problem, and could depend on the type of signal and/or the computational muscle available.

2. Set up an evenly spaced (if not otherwise specified) angular frequency grid $2\pi\tau f_{min} < \varphi_j < 2\pi\tau f_{max}$, $j = 1, 2, \ldots, K_{win}$, by choosing an appropriate value for K_{win}. A reasonable choice for K_{win} is $K_{win} = N(f_{max} - f_{min})/2SW$, where SW is the spectral width $1/\tau$. This is simply half the number of points that would appear in a conventional N-point complex Fourier transform of c_n within this same frequency window, which is consistent with the informational content of the data array c_n, i.e., this value of K_{win} is the largest number of lines that can give a unique fit to a signal of this length.

3. Evaluate three complex symmetric matrices $\mathbf{U}^{(p)}$ of the size $K_{win} \times K_{win}$ for, e.g., $p = 0, 1, 2$, using for the off-diagonal matrix elements of $\mathbf{U}^{(p)}_{jj'} \equiv \mathbf{U}^{(p)}(\varphi_j, \varphi_{j'})$

$$\mathbf{U}^{(p)}(\varphi,\varphi') = \frac{e^{-i\varphi}f_p(\varphi') - e^{-i\varphi'}f_p(\varphi) + e^{iM\varphi'}g_p(\varphi) - e^{iM\varphi}g_p(\varphi')}{e^{-i\varphi} - e^{-i\varphi'}} \quad [19]$$

where the functions f_p and g_p are the Fourier transforms of the (roughly) first and last half of the signal c_n with the time-domain data points shifted by p,

$$f_p(\varphi) = \sum_{n=0}^{M} e^{in\varphi}c_{n+p};\ g_p(\varphi) = \sum_{n=M+1}^{2M} e^{i(n-M-1)\varphi}c_{n+p}. \quad [20]$$

For the diagonal matrix elements the formula

$$U^{(p)}(\varphi,\varphi) = \sum_{n=0}^{2M} (M + 1 - |M - n|)e^{in\varphi} \quad [21]$$

is used.

4. Solve the generalized eigenvalue problem

$$\mathbf{U}^{(1)}\mathbf{B}_k = u_k\mathbf{U}^{(0)}\mathbf{B}_k \quad [22]$$

for the eigenvalues $u_k = \exp(-i\tau\omega_k)$ and eigenvectors $\mathbf{B}_k$ using, for example, a standard implementation of the QZ-algorithm (*15*).

5. Accept those eigenvalues u_k which satisfy an error criterion as, e.g.,

$$\|(\mathbf{U}^{(2)} - u_k^2\mathbf{U}^{(0)})\mathbf{B}_k\| < \epsilon \quad [23]$$

for a small enough ϵ. This is simply to make sure that none of the eigenvalues are spurious, for such spurious eigenvalues are very sensitive to any change in the parameters of the calculation. (Spurious eigenvalues are not those associated with noise. They arise from the use of only a locally complete basis, and do not occur often.)

6. Compute the complex amplitudes d_k using the expression

$$d_k^{1/2} = \sum_{j=1}^{K_{win}} B_{jk} \sum_{n=0}^{M} c_n e^{in\varphi_j} \quad [24]$$

or a substantially more accurate expression (*4*) which makes use of the entire measured signal rather than just the first $N/2$ points:

$$d_k^{1/2} = \frac{1 - e^{-\tau\gamma}}{1 - e^{-(M+1)\tau\gamma}} \sum_{j=1}^{K_{win}} B_{jk}U^{(0)}(\varphi_j, \tau(\omega_k + i\gamma)). \quad [25]$$

The adjusting parameter γ is chosen so that $U^{(0)}(\varphi_j, \tau(\omega_k + i\gamma))$ is numerically stable. One correct choice is $\gamma = -\mathrm{Im}\{\omega_k\}$ for $\mathrm{Im}\{\omega_k\} < 0$ and $\gamma = 0$ for $\mathrm{Im}\{\omega_k\} > 0$. In the limit $\gamma \to \infty$, Eq. [25] becomes Eq. [24].

7. Store the set of converged ω_k and d_k and use them as input for a spectral estimator which, for the absorption mode "ersatz" spectrum (*5*) corresponds to

$$A(F) = -\sum_k \mathrm{Im}\left\{\frac{d_k}{2\pi F - \omega_k}\right\}. \quad [26]$$

If $\mathrm{Im}\{\omega_k\} > 0$, then ω_k is replaced with ω_k^*. Because the form of the spectrum is dependent on the denominator of Eq. [26], it is sometimes economical to refer to each entry in the line list as a *pole.*

8. Choose the next frequency window.

A few comments are in order. First, note that in FDM the complex amplitudes are computed *without solving another linear least squares problem,* in contrast to all other linear algebraic approaches. Note that the formula for the amplitude and phase of each feature as given in Eq. [24] or [25] does *not* depend on the values of the other complex frequencies. In conventional LP methods, and in LP-ZOOM, the amplitudes are constructed by solving a linear system which uses, as part of its input, all the obtained frequencies. In a few cases ω_k might have a positive imaginary part resulting in a "negative" linewidth. These "bad" roots complicate the correct determination of the amplitudes of other features. One regularization procedure is simply to change the sign of the imaginary part. We adopt this procedure when constructing the ersatz spectrum using Eq. [26], replacing ω_k with ω_k^*. Note that in FDM this replacement does not affect the other spectral features. Sec-

ondly, the maximum number of possible lines, namely $N/2$ for a signal of length N, is used to fit the spectrum. Many of these ultimately turn out to be of low amplitude or wide linewidth and correspond to noise. However, by including them we improve the likelihood that the signal poles will be accurate. By adjusting the number of points of the entire FID that are used and repeating the calculation, it is possible to get a feel for the stability of the line list. Because FDM is very fast on any modern workstation or PC, it is possible to "play around" as one might routinely do with the FFT, e.g., using different weighting functions to get the best representation of the data. It is, of course, possible to prune the line list after the fact to weed out features of no further interest.

Other Fourier Bases

Experience with Fourier spectral analysis might suggest that the rectangular-window Fourier basis of Eq. [7] may not be the most efficient choice. For example, one could argue that the abrupt rectangular window might cause truncation artifacts analogous to those observed in an FFT of an unapodized FID. While such oscillations never actually occur in the FDM ersatz spectrum, they can be identified in the sinc-like structure of the off-diagonal matrix elements of **U** (see Fig. 1). It is not obvious that such oscillations can ever affect the final results of the generalized eigenvalue problem. In this section we address this question.

Just as when one transforms an FID to obtain the FFT spectrum, in the frame of FDM the truncation artifacts can be suppressed by implementing a window or apodization function g_n which, for convenience, we write as a function of the index n of the FID. Rewriting Eq. [17] as

$$\Psi_g(\varphi) = \sum_{n=0}^{M} g_n \exp(in(\varphi - \tau\hat{\Omega}))\Phi_0, \qquad [27]$$

we are free to experiment with g_n. The original formulation of FDM (*1*) used a Gaussian-window Fourier basis, which is reminiscent of the "pseudo-echo" filter that was often applied to early absolute-value 2D spectra (*19*). Such a window discriminates against the initial portion of the signal, and so results in a substantial reduction in S/N. Likewise, in FDM this window appeared to be inferior to the rectangular-window basis introduced later (*3*). This is clearly a result of the very inefficient sampling of the data with the Gaussian window. However, it might be the case that other more efficient windows could be superior to the simplest choice of the rectangular window. To this end, Chen and Guo have derived (*20*) an alternative expression to compute the **U** matrix elements in a Fourier basis (*1*) with a general filter g_n. Their expression unfortunately requires about a factor of K_{win} more CPU time than that for the special case of the rectangular window of Eq. [7]. It turns out, however, that a rectangular window is not the only one which can be implemented with an eye to numerical efficiency. There are at least two other window functions g_n with this property. The first obvious choice is the exponential window

$$g_n = \exp(-\gamma n), \qquad [28]$$

corresponding to replacing φ by $\varphi + i\gamma$ in Eq. [27] and consequently in the expressions for the matrix elements of **U** in Eq. [18]. Despite its simplicity, the exponential window is known to be a nonoptimal apodization function, as for a reasonably large γ, when it can have a noticeable effect, it starts to discriminate a large part of the Fourier series (excessive line broadening in the analogous case of an FID).

The simple cosine window (often referred to in the NMR literature as a sine-bell window (*21*) phase shifted by $\pi/2$),

$$g_n = \cos\left(\frac{\pi}{2}\frac{n}{M+1}\right), \qquad [29]$$

often leads to good results as its shape yields a good balance between efficient sampling on one hand and suppression of the Gibbs oscillations on the other. The cosine window can be implemented trivially, as inserting Eq. [29] into Eq. [27] leads (after dropping the unimportant factor of 1/2) to

$$\Psi_{\cos}(\varphi) = \Psi(\varphi - \alpha) + \Psi(\varphi + \alpha) \qquad [30]$$

with $\alpha = \pi/(2(M+1))$, which for the **U** matrices gives

$$\begin{aligned}\hat{U}^{(p)}_{\cos}(\varphi,\varphi') = {} & \hat{U}^{(p)}(\varphi + \alpha, \varphi' + \alpha) \\ & + \hat{U}^{(p)}(\varphi + \alpha, \varphi' - \alpha) \\ & + \hat{U}^{(p)}(\varphi - \alpha, \varphi' + \alpha) \\ & + \hat{U}^{(p)}(\varphi - \alpha, \varphi' - \alpha). \end{aligned} \qquad [31]$$

That is, the use of a cosine window corresponds to essentially the same numerical procedure and therefore is easily incorporated in the FDM code as an option. Similarly, virtually any trigonometric function g_n can be implemented efficiently.

To our disappointment, introducing the cosine window has not led to a significant numerical difference, except for a few cases where it perhaps showed a slightly better performance than the rectangular window. A simple explanation of why the use of the cosine window does not lead to significant improvement of FDM is the following: (i) FDM with a rectangular window is already very efficient and often shows nearly optimal performance (*3–5*); (ii) the use of another window g_n in Eq. [27] only affects the initial structure of the **U** matrices, having only a minor effect on the whole subspace formed by the K_{win} window basis functions: Eq. [31] simply means that

$\{\Psi_{\cos}(\varphi)\}$ is a certain linear combination of $\{\Psi(\varphi)\}$. Linear combinations of the $\{\Psi(\varphi)\}$ are taken in any event to diagonalize the **U** matrices, so that taking such combinations up front plays little additional role. We believe that this finding should also hold for all other reasonable types of window g_n.

Reference Deconvolution

The elegant idea of reference deconvolution as put forward by Morris (*22*) is simply to use a reference line, which could be TMS, as a lineshape or convolution reference as well as just a chemical shift reference. The known theoretical lineshape of TMS can be compared with the observed lineshape to assess the shimming and/or stability of the B_0 field. The observed response as a function of frequency $O(F)$ can be represented by a convolution of the perfect response $P(F)$ with the instrumental response $Q(F)$,

$$O(F) = P(F)*Q(F), \qquad [32]$$

which according to the convolution theorem means that the corresponding perfect FID $p(t)$ has been multiplied by the instrumental function $q(t)$, where $q(t)$ and $Q(F)$ form a Fourier transform pair. Clearly, $p(t) = o(t)/q(t)$, so that the observed data can be corrected if the instrumental function $q(t)$ can be obtained. Knowing the ideal lineshape for an isolated reference line, like TMS, allows $q(t)$ to be obtained as the ratio of the inverse FT of a small spectral region centered on the observed lineshape (and not containing other signals) to the inverse FT of the ideal lineshape:

$$q(t) = \frac{o_{TMS}(t)}{p_{TMS}(t)}. \qquad [33]$$

Then simply dividing the full FID by $q(t)$ accomplishes the desired deconvolution. If the shimming is poor, then $o_{TMS}(t)$ will decrease more rapidly than $p_{TMS}(t)$, so that $q(t)$ will eventually be small. Dividing by $q(t)$ thus blows up the noise in the tail of the FID just like a conventional resolution enhancement function. When Fourier spectral analysis is used, this may require some apodization, so that there is a limit to reference deconvolution which will depend on the S/N ratio. To the extent that there is nonnegligible noise in the baseline around the reference peak, there may also be some unwanted small features introduced by reference deconvolution with the noise. Other minor problems such as baseline discontinuity at the edges of the selected section have been successfully addressed (*23*).

Reference deconvolution is a powerful addition to FDM. The first weakness an experimentalist notices with Eq. [1] is the assumption of a Lorentzian lineshape. This is important to the algorithm because without it, the fitting becomes an essentially impossible large nonlinear optimization problem rather than a linear algebra problem. However, imperfect shimming, chemical exchange, and other dynamic processes can all lead to a distinctly non-Lorentzian lineshape. FDM will fit such a single line as a superposition of Lorentzians, but it would be preferable to have a *single* entry in the line list rather than a large number of closely clustered lines. This is especially important if one wishes to measure small splittings that may be almost unresolved but correspond to a J coupling that is of interest. For many situations it is possible to offer a surprisingly simple solution: reference deconvolution using the prescription laid out by Morris, but implemented using FDM itself.

The protocol is deceptively simple. FDM is used to fit the observed reference lineshape to a sum of Lorentzians by selecting a window centered on the reference line that is narrow enough to exclude the carbon-13 satellites. It is essential to have a long enough signal so that truncation effects are minimal, and so that the number of local basis functions is dense enough to result in a good match to the observed lineshape. The set $\{d_k, \omega_k\}$ for TMS then gives a very close analytic approximation to the function $o_{TMS}(t)$. The function $p_{TMS}(t)$ can be formulated as three closely spaced Lorentzians (to account for the Si-29 satellites), and the ratio of the two functions gives $q(t)$. Note that there is no baseline noise included in $q(t)$ except, perhaps, as slight errors in the set $\{d_k, \omega_k\}$. The FID is then preprocessed using $q(t)$.

The efficacy of the reference deconvolution can then be checked by processing the same TMS spectral window of the new FID. Ideally, only the three poles for the TMS lines should turn up, along with very small amplitude poles representing the baseline noise. If significant additional lines are needed, then further small corrections can be iteratively applied, in a procedure which converges providing the S/N is adequate and there are no other genuine signals within the window. When adequate results are obtained, the rest of the spectrum can be processed with confidence. Only lines from spins undergoing chemical exchange and the like should then require more than a single entry in the line list to achieve a good fit. It is possible to extend the idea further by using a model for a chemical exchange lineshape and then deconvolute this as well. Since FDM works over disjoint spectral windows, it is possible in principle to process the NH and aliphatic regions of the proton spectrum using different functions for the reference deconvolution. Note that reference deconvolution can be applied to any spectroscopy in which the observed lineshape can be represented as a convolution, extending the reach of FDM into many related areas in which the lineshape may be non-Lorentzian. It is even possible to convolute the Lorentzians found by FDM with the experimental lineshape to restore the original shape, should this be desired.

Phase Correction

Problems with the first few data points of the FID can arise from the transient response of the audio filters (digital or

analog) or probe ring-down and lead to baseline roll. Some pulse sequences, such as the 1–3–3–1 water suppression sequence (*24*), can introduce large linear phase rolls over the spectrum, and when a pulsed field gradient is applied immediately prior to acquisition, some stabilization delays usually need to be included. The usual linear phase correction routines applied to the FFT spectrum only work adequately when the corresponding delay is not too long, so that the phase roll over a typical linewidth is small.

FDM has good tolerance for this type of data imperfection. The usual LP approach is to discard the first few data points and then "backward predict" them from uncorrupted data to flatten the baseline (*25*). This approach will only work well when the LP backward extrapolation is reasonably good, i.e., when both the number of extrapolated points and the number of lines are not too large.

If the data is available only for $n = \tilde{n}, \tilde{n} + 1, \ldots, N$, then the frequencies ω_k extracted from such a signal by FDM should be the same, while the amplitudes $d_k(\tilde{n})$ will be modulated by the frequency-dependent factors

$$d_k(\tilde{n}) = d_k \exp(-i\tilde{n}\tau\omega_k). \qquad [34]$$

To correct an extracted amplitude $d_k(\tilde{n})$, one can thus simply multiply it by the factor $\exp(i\tilde{n}\tau\omega_k)$. Such a correction will be accurate if the extracted frequencies ω_k and amplitudes $d_k(\tilde{n})$ are reliable, which under certain conditions is often the case. Spectral estimator-based methods that attempt to correct the distorted baseline of the FFT spectrum by, e.g., fitting it to a smooth function are clearly inferior to FDM, as they require the solution of a nonlinear optimization problem, while the use of Eq. [34] is straightforward.

However, one should realize that the correction factor $\exp(i\tilde{n}\tau\omega_k)$, which is now a positive exponential, can be very large if the width of the kth pole is large. For a large time delay $\tilde{n}\tau$ even a very small error in determination of the amplitude $d_k(\tilde{n})$ could be amplified enormously if the width γ_k is large. The errors of the calculated spectral parameters, ω_k and $d_k(\tilde{n})$, are not the only source of instability, because the line list obtained by FDM does not only represent the desired signals. In fact, it always contains poles which represent the noise. For example, white uncorrelated noise results in poles with very large γ_k. However, such poles can be easily identified and removed from further analysis. Therefore, their presence will not affect the phase correction procedure. The problem will occur when such an obvious separation of the signal and noise poles is not possible. To this end Eq. [34] has certain limitations. That is, FDM will be a perfect method for correcting artifacts due to possibly very large time delays, but only for spectra with lines that are not too strongly overlapping and not buried in too much noise. We have found, however, that FDM is surprisingly robust on even complex spectra with considerable time delays as long as the spectrum is not too noisy.

As mentioned earlier, the amplitudes and complex poles of the broad spectral features are most sensitive to the parameters of the fit by FDM. This leads to instabilities when such poles are used as entries for the corrections of large time delays. In our numerical tests we found that the applicability of FDM for the phase correction can significantly be extended to longer time delays when the stabilization idea, originally proposed in a quite different context (*26*), is used in the frame of FDM. The idea is very simple and relies on the fact that the errors in determination of the spectral parameters have a random character and therefore can be reduced by averaging the FDM ersatz spectrum over the acquisition time or number of points, N. In particular, the phase of the unwanted features is sensitive to the exact length of the signal used in the calculation, whereas the signal poles are more stable. Since FDM is not time consuming, one can carry out such an averaging using many FDM calculations over long but slightly different length data records N_1, N_2, etc., extracted from the same FID. Although the averaging procedure does not seem to be very elegant, it turns out to work well when a single FDM result is unacceptable. As it is quite general, it can also be used in other contexts to stabilize the ersatz spectra computed by FDM.

Even when rather large delays are not at issue, FDM may still be quite useful. Note that the usual "automatic" zeroth- and first-order phase correction due to small delays and finite 90° pulse times can be applied extremely rapidly to the list of complex amplitudes d_k using any simple smooth function of frequency. There are other advantages: As long as the expected phase shift as a function of frequency can be modeled, for example, that which might be produced by a composite 180° pulse in a spin echo, then this nonlinear correction can be applied to the d_k as well, even though it may not be a simple function of frequency.

Matched Resolution Enhancement or "Cheating"

FDM can achieve an extremely good fit to most kinds of spectra that arise from NMR experiments in liquids. There is a difference, however, between having a good fit and having the true spectral parameters. Obviously, if there are thousands of overlapping Lorentzians within a small spectral window, then the length N of the FID may not be long enough to achieve an exact fit, as the number of poles in the window is limited to K_{win}. In this case the spectrum is intrinsically unresolved. The presence of noise poles is another issue. Nevertheless, in some cases the FDM line list is close to a perfect physical summary of the NMR transitions. By manipulating the spectral parameters directly, it is possible to present the spectral information in more informative ways.

For example, digital filtering for resolution enhancement is commonly applied in FT spectroscopy. One method is to multiply the FID by a positive exponential which matches some linewidth, and then apodize the end of the data to keep the amplification of the noise from being too severe (*27*). When

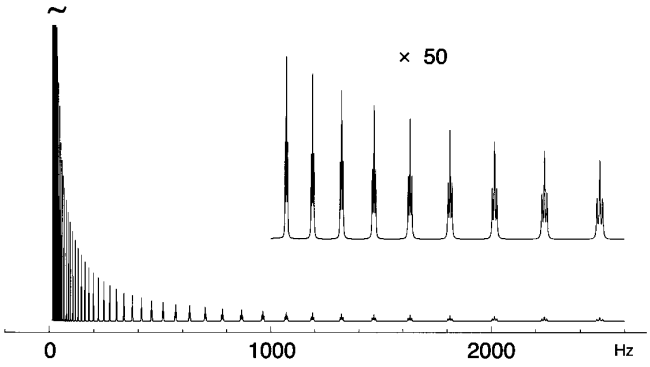

FIG. 2. The spectrum of triplets, "Jacob's ladder." A total of 50 triplets are present, beginning with a 5.0 Hz linewidth with 12.5 Hz coupling centered at 2487.5 Hz and finishing with a 28.8 mHz line width and 71.6 mHz couplings. A spectral width of 5 kHz is used, of which half is shown. The wide range of linewidths and splittings and the accumulation of lines near zero frequency makes this signal challenging for linear algebraic methods.

there are lines of differing width, some narrow features can be "overenhanced," giving undulations in the baseline around these signals. This is because the digital filter corresponds to *subtracting* a fixed width from each line. However, in FDM we can take a different approach and reduce the linewidth of each pole by a certain fraction of its original width. We call this procedure "cheating," because it can be quite dangerous, and because it will appeal to everyone.

Although cheating is certainly pushing the integrity of the line list to its limit, it can reveal buried couplings that were invisible in the FDM ersatz spectrum, or in the FT spectrum. Reference deconvolution, which tends to minimize the number of entries for each line, is certainly a good idea to use before cheating. Used cautiously, cheating can help reveal more of the information in the signal, as we show in the next section.

EXPERIMENTAL

A Numerical Example

To show the ability of FDM to identify and resolve purely Lorentzian lines of widely differing widths and spacings in a very long FID, we constructed a "Jacob's ladder" of 50 triplets, each of which has a splitting between the nearest neighbors of 2.5 times the full width at half height. The signal is constructed by assuming a spectral width of 5 kHz and beginning with a triplet with the right-hand line at Nyquist frequency of 2500 Hz and with a linewidth of 5.0 Hz and spacing of 12.5 Hz. This 1:2:1 triplet is then reproduced by multiplying both the frequency and linewidth by 0.9 a total of 49 times, keeping the amplitude and phase fixed, with an overall phase of zero. The damped sinusoids for these 50 triplets are coadded, normalized to an initial value $c_0 = 16{,}384$, sampled uniformly in time every 200 μs, and then integerized to mimic, approximately, an ADC by taking the closest integer value *less* than the true value. The test FID c_n before integerization is thus

$$c_n = 81.92^* \sum_{m=0}^{49} \exp(2\pi i(0.9)^m \times (2500.0 - 2.5i)(2.0 \times 10^{-4}n)) + 2\exp(2\pi i(0.9)^m(2487.5 - 2.5i)(2.0 \times 10^{-4}n)) + \exp(2\pi i(0.9)^m(2475.0 - 2.5i)(2.0 \times 10^{-4}n)) \quad [35]$$

and is heavily peaked at $t = 0$, diminishing rapidly as the lines get out of phase and the broader features decay, so that the narrow lines are digitized with only 6–8 bits. Integerization introduces colored noise into the baseline and is an important element of the test. The difference in line width between the broadest and narrowest feature is almost 175-fold, as is the difference in peak height. The integral of each multiplet should be the same, however. Figure 2 shows the spectrum of the test signal.

Figure 3 compares the FFT to FDM on the densest portion of the frequency range of the test signal, using 32K and 64K FIDs. The result of using the FFT, with appropriate apodization

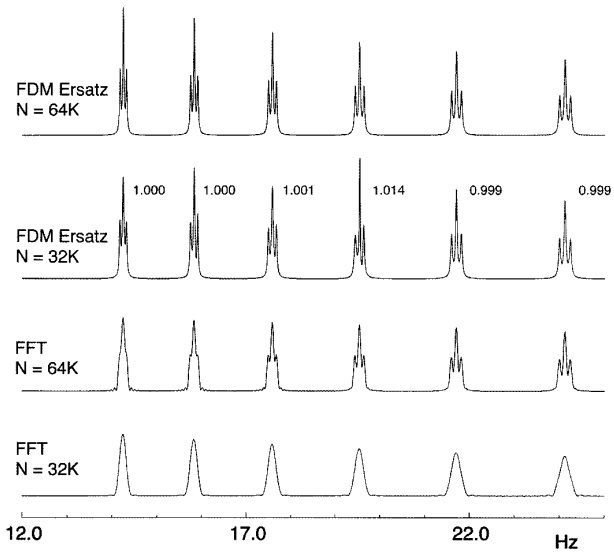

FIG. 3. A comparison of FFT and FDM on the densest part of the signal shown in Fig. 2. The relative integrals calculated by FDM are shown for the 32K FID. Note that a 32K point acquisition (6.553 s) is unable to resolve any of the fine multiplet structure if an FFT is used, even with zero-filling to 128K, because of the approximate resolution limit of 0.15 Hz according to the uncertainty principle. By contrast, FDM resolves all the multiplets with a 32K FID. The line frequencies extracted by FDM are highly accurate, and the relative integrals are within a few percent of the correct value for each multiplet. Some of the individual line intensity/width ratios are not as accurate, however, because the shorter FID does not have quite enough information content to pin down these parameters.

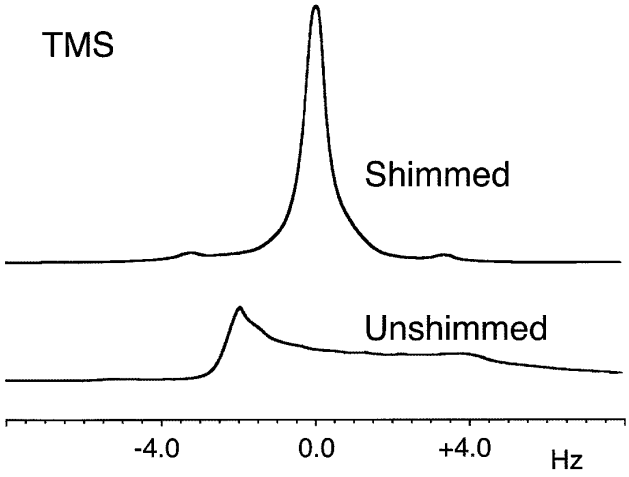

FIG. 4. A TMS reference line shimmed so that the Si-29 satellites are visible (top) and shimmed very poorly (bottom). FDM was used to fit both signals to a sum of Lorentzian lines and thereby construct the reference deconvolution function discussed in the text.

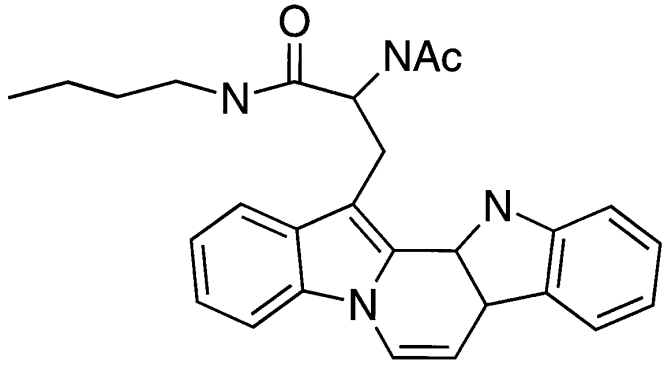

FIG. 5. The structure of a diindole reaction product, 12*H*-pyrido[1,2-*a*:3,4-*b*]diinole, used to test the reference deconvolution procedure. It was dissolved in $CDCl_3$ with added TMS.

to minimize, but not completely eliminate, the "sinc wiggles" that are found on the narrowest lines, and with zero-filling to 128K, shows that a 32K FID is unable to resolve any of the narrowest features, but that doubling the length of the FID to 64K resolves more of the multiplets. FDM does substantially better, although the peak heights of some of the narrow features are not correctly captured. The relative *integrals* of the multiplets, shown in the figure, are quite good, as are the frequency coordinates of the individual lines, but there is some error in the exact determination of the width/height individually. These errors arise from the incomplete information in the shorter FID, and disappear as the length of the signal is increased. The broader lines are correctly reproduced by both methods. The longest signal is processed, in its entirety, in a matter of tens of seconds with FDM on a low-end workstation.

Reference Deconvolution by FDM

Figure 4 shows a close-up of the TMS line in a solution of ~10 mg of the compound shown in Fig. 5 dissolved in $CDCl_3$, acquired on a 500 MHz Varian UnityPlus spectrometer with different instrumental linewidths by adjusting the shimming. In the second trace the line has been deliberately made asymmetric, broad, and very non-Lorentzian using the z^2 and z^4 shims. Rather than use a theoretical model for the TMS line, we simply used the FDM fit of the shimmed spectrum as an experimental standard. Thus, deconvoluting the poorly shimmed spectrum should yield results comparable to the properly shimmed spectrum, which is easily checked. The first two traces in Fig. 6 show a representative multiplet with the good and poor shimming, respectively. The effect of reference deconvolution on the second spectrum is shown in the third trace, giving results nearly indistinguishable from the correctly shimmed spectrum at top. Finally, the last trace shows the effect of matched resolution enhancement (cheating) by artificially multiplying the linewidth of all poles by a factor of 0.5. Some of the noise poles become more prominent as notches in the baseline, but in other respects the results are quite encouraging.

There is, of course, a limit to reference deconvolution by FDM. When signals that were broadened about twice as much as those shown were processed, the deconvolution function itself becomes unstable. The net result is that only the initial portion of the FID can be processed, and the number of poles in the frequency domain is not dense enough to achieve a good fit, or to uncover small splittings. Fortunately, it is fairly easy to tell visually whether the results are any good.

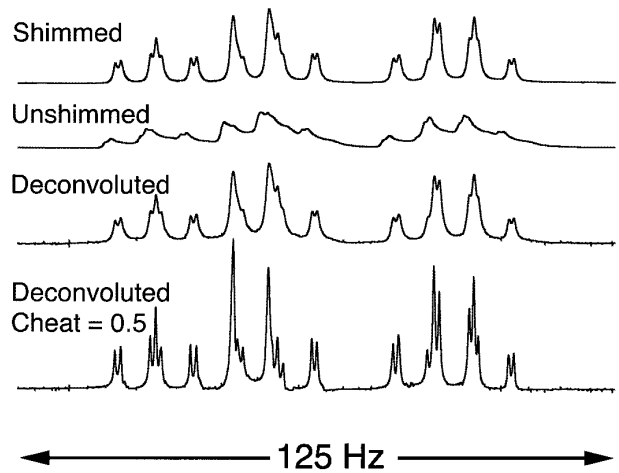

FIG. 6. An expanded view of a small portion of the 1D NMR spectrum of the molecule shown in Fig. 5. The first two traces correspond to same homogeneity as the two traces in Fig. 4. The third trace shows the result of reference deconvolution of the second trace with the deconvolution function obtained by analyzing the TMS signals in Fig. 4 by FDM. It agrees extremely well with the first trace, as it should. The last trace shows the effect of "cheating" and artificially making the linewidth of each feature in the deconvoluted spectrum only half as wide. Some small couplings are revealed. The FID had 32,768 complex points; the entire calculation including the reference deconvolution took about 15 s on a small SGI workstation (200 MHZ IP22 processor. FPU: MIPS R4000 Floating Point Coprocessor Revision: 0.0. CPU: MIPS R4400 Processor Chip Revision: 6.0, 64 Mb RAM.)

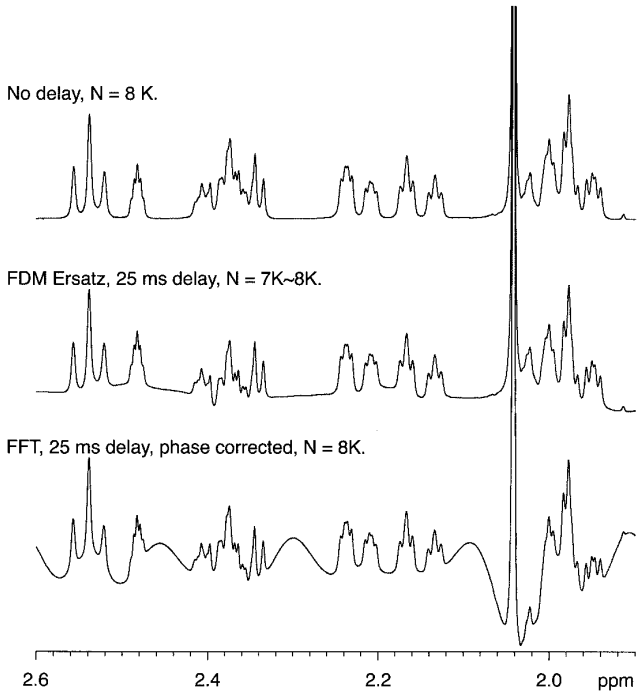

FIG. 7. An expanded view of the NMR spectrum of progesterone in the vicinity of the Me-19 resonance. *Top:* Conventional NMR spectrum, properly phased. *Middle:* Phase-corrected FDM ersatz spectrum obtained after deleting the first 100 data points. Twenty-five FDM ersatz spectra were computed, with signal lengths between 7K and 8K data points (an increment of 40 data points) and then summed. The calculation required about a minute of computer time (see the caption to Fig. 6). *Bottom:* FFT spectrum obtained by substituting the first 100 data points with zeroes, which is equivalent to applying an accurate linear phase correction to the spectrum. The phase roll is sufficiently violent that it would be next to impossible to pick out baseline points appropriately.

Phase Correction by FDM

Figure 7 shows a portion of the 1D spectrum of progesterone obtained over a 4 kHz spectral width using $N = 8192$ complex points. With a time delay of 25 ms, the first 100 data points are missing from the FID. This is quite a challenging phase correction problem because the spectrum is not simple, the FID has thousands of points, and a lot of data is missing. Nevertheless, the averaged ersatz FDM spectrum summing 25 different calculations with between 7K and 8K points shows very good agreement with the uncorrupted spectrum. By comparison, zeroing the first 100 points and then Fourier transforming the FID gives terrible baseline problems which also cannot be reliably distinguished from true signal. Even with 200 points missing, FDM can still extract a usable averaged ersatz spectrum (data not shown). Clearly, for milder situations in which a gradient stabilization delay needs to be included and a few milliseconds of data are missing, FDM will be very reliable.

SUMMARY

The filter diagonalization method is a powerful new addition to the linear algebraic arsenal for parameter and spectral estimation of NMR signals. FDM is numerically stable, lightning fast, and applicable to many different kinds of spectra. Some of the common instrumental shortcomings, such as poor shimming or nonzero dead times, can be conveniently overcome using FDM. In particular, reference deconvolution by FDM is a very useful technique. Like most parameter estimation methods, FDM is sensitive to noise and certainly cannot improve the intrinsic sensitivity of the NMR experiment by somehow digging signals out of noise. To the extent the signal matches the FDM ansatz, however, big improvements can be obtained, with the best case being sharp, well-resolved, purely Lorentzian lines. If the signal is nonstationary (e.g., t_1-noise along an indirect dimension of a multidimensional NMR experiment) or very noisy, or if the lines are very far from Lorentzian, a spectrum can be obtained, but the spectral parameters do not have the desired significance. However, because FDM can be implemented so efficiently, it could be used to control the experiment on-line, terminating acquisition under computer control once a stable line list is obtained. Time averaging is, of course, the only reliable method for improving *S/N* on a given instrument and sample. This is clearly an avenue worth exploring in the future. We note in passing that almost any computer is suitable for FDM. However, because the algorithm parallelizes trivially, parallel computers may be used for direct parameter estimation on huge multidimensional signals in less time than it takes to acquire the data. We will describe these exciting developments in a series of related publications.

ACKNOWLEDGMENTS

This material is based on work supported by the National Science Foundation CHE-9625674, and by a UC BioSTAR grant S97-18. Que Van acknowledges partial support from a Synthesis and Structure of Biological Macromolecules Training Grant T32 GM 07311-23. We are indebted to Prof. Benny Gerber for the use of an SGI workstation, and to Professor David Van Vranken and Mr. David S. Carter for supplying us with the molecule shown in Fig. 5.

REFERENCES

1. M. R. Wall and D. Neuhauser, Extraction, through filter-diagonalization, of general quantum eigenvalues or classical normal mode frequencies from a small number of residues or a short-time segment of a signal. I. Theory and application to a quantum-dynamics model, *J. Chem. Phys.* **102,** 8011–8022 (1995).
2. V. A. Mandelshtam and H. S. Taylor, Spectral analysis of time correlation function for a dissipative dynamical system using filter diagonalization: application to calculation of unimolecular decay rates, *Phys. Rev. Lett.* **78,** 3274–3277 (1997).
3. V. A. Mandelshtam and H. S. Taylor, Harmonic inversion of time signals and its applications, *J. Chem. Phys.* **107,** 6756–6769 (1997).
4. V. A. Mandelshtam and H. S. Taylor, Multi-dimensional harmonic inversion by filter diagonalization, *J. Chem. Phys.* **108,** 9970–9977 (1998).
5. V. A. Mandelshtam, H. Hu, and A. J. Shaka, Two-dimensional HSQC NMR spectra obtained using a self-compensating double

pulsed field gradient and processed using the filter diagonalization method, *Magn. Reson. Chem.* (in press).

6. H. Barkhuijsen, R. de Beer, W. M. M. J. Bovée, and D. van Ormondt, Retrieval of frequencies, amplitudes, damping factors, and phases from time-domain signals using a linear least-squares procedure, *J. Magn. Reson.* **61,** 465–481 (1985).
7. J. Tang and J. R. Norris, LPZ spectral analysis using linear prediction and the z-transform, *J. Chem. Phys.* **84,** 5210–5211 (1986).
8. H. Gesmar and J. J. Led, Spectral estimation of complex time-domain NMR signals by linear prediction, *J. Magn. Reson.* **76,** 183–192 (1988).
9. D. S. Stephenson, Linear prediction and maximum entropy methods in NMR spectroscopy, *Prog. NMR Spectrosc.* **20,** 515–626 (1988).
10. G. Zhu and A. Bax, Improved linear prediction for truncated signals of known phase, *J. Magn. Reson.* **90,** 405–410 (1990).
11. R. de Beer and D. van Ormondt, Analysis of NMR data using time domain fitting procedures, *NMR Basic Prin. Prog.* **26,** 201–248 (1992).
12. H. Gesmar and P. C. Hansen, Fast linear prediction and its application to NMR spectroscopy, *J. Magn. Reson. Series A* **106,** 236–240 (1994).
13. Y.-Y. Lin, P. Hodgkinson, M. Ernst, and A. Pines, A novel detection–estimation scheme for noisy NMR signals: Applications to delayed acquisition data, *J. Magn. Reson.* **128,** 30–41 (1997).
14. J. C. Hoch and A. S. Stern, "NMR Data Processing," Wiley-Liss, New York (1996).
15. S. Marple, Jr., "Digital Spectral Analysis with Applications," Prentice-Hall, Englewood Cliffs, NJ (1987).
16. R. Roy, B. G. Sumpter, G. A. Pfeffer, S. K. Gray, and D. W. Noid, Novel methods for spectral analysis, *Comput. Phys. Rep.* **205,** 109–152 (1991).
17. J. Tang and J. Norris, LP-ZOOM, a linear prediction method for local spectral analysis of NMR signals, *J. Magn. Reson.* **79,** 190–196 (1988).
18. C. B. Moler and G. W. Stewart, An algorithm for generalized matrix eigenvalue problems, *SIAM J. Numer. Anal.* **10,** 241–256 (1973).
19. A. Bax, R. Freeman, and G. A. Morris, A simple method for suppressing dispersion-mode contributions in NMR spectra: the "pseudo echo," *J. Magn. Reson.* **43,** 333–338 (1981).
20. R. Chen and H. Guo, Calculation of matrix elements in filter diagonalization: a generalized method based on Fourier transform, *Chem. Phys. Lett.* **279,** 252–258 (1997).
21. J. C. Lindon and A. G. Ferrige, Digitisation and data processing in Fourier transform NMR, *Prog. NMR Spectrosc.* **14,** 27–66 (1980).
22. G. A. Morris, Compensation of instrumental imperfections by deconvolution using an internal reference signal, *J. Magn. Reson.* **80,** 547 (1988).
23. A. Gibbs and G. A. Morris, Reference deconvolution. Elimination of distortions arising from reference line truncation, *J. Magn. Reson.* **91,** 77–83 (1991).
24. P. J. Hore, Solvent suppression in Fourier transform nuclear magnetic resonance, *J. Magn. Reson.* **55,** 283–300 (1983).
25. D. Marion and A. Bax, Baseline correction of 2D FT NMR spectra using a simple linear prediction extrapolation of the time-domain data, *J. Magn. Reson.* **83,** 205–211 (1989).
26. V. A. Mandelshtam, T. R. Ravuri, and H. S. Taylor, Calculation of the density of resonance states using the stabilization method, *Phys. Rev. Lett.* **70,** 1932–1935 (1993).
27. R. R. Ernst, G. Bodenhausen, and A. Wokaun, "Principles of Nuclear Magnetic Resonance in One and Two Dimensions," Clarendon Press, Oxford (1987).