JOURNAL OF MAGNETIC RESONANCE **130,** 217–232 (1998)
ARTICLE NO. MN971304

A Metropolis Monte Carlo Implementation of Bayesian Time-Domain Parameter Estimation: Application to Coupling Constant Estimation from Antiphase Multiplets

Michael Andrec and James H. Prestegard

Department of Chemistry, Yale University, New Haven, Connecticut 06511

Received March 11, 1997; revised October 14, 1997

The Bayesian perspective on statistics asserts that it makes sense to speak of a probability of an unknown parameter having a particular value. Given a model for an observed, noise-corrupted signal, we may use Bayesian methods to estimate not only the most probable value for each parameter but also their distributions. We present an implementation of the Bayesian parameter estimation formalism developed by G. L. Bretthorst (1990, *J. Magn. Reson.* **88, 533) using the Metropolis Monte Carlo sampling algorithm to perform the parameter and error estimation. This allows us to make very few assumptions about the shape of the posterior distribution, and allows the easy introduction of prior knowledge about constraints among the model parameters. We present evidence that the error estimates obtained in this manner are realistic, and that the Monte Carlo approach can be used to accurately estimate coupling constants from antiphase doublets in synthetic and experimental data.**

A common task in NMR spectroscopy is the accurate estimation of spectral parameters (such as splittings, linewidths, or intensities) and their uncertainties from time-domain data (FIDs). In addition, it is often possible to specify constraints among the parameters. For example, the presence of a known multiplet structure could lead to a significant reduction in the number of adjustable parameters, since one could specify constraints among the frequencies, phases, and intensities. Thus, the ideal quantitative NMR data analysis tool not only would estimate parameter values and uncertainties but also would allow the flexible specification of relationships among the model parameters.

We present here a new implementation of Bayesian parameter estimation based on the work of Bretthorst (*1*). Although it is certainly not the ideal quantitative NMR data analysis tool, it does possess many of the desirable attributes mentioned above. In addition, we hope to demonstrate that it is robust and flexible and that it can give reasonable estimates of the reliability of the extracted parameters. We have chosen to use the Metropolis–Hastings Monte Carlo algorithm (*2, 3*) to generate points in the parameter space according to the Bayesian posterior probability density. Using synthetic data with substantial added random noise, we demonstrate that the error estimates obtained using the Bayesian approach are consistent with a Monte Carlo error analysis in the case of two well-resolved Lorentzians, and that coupling constants can be reliably obtained from overlapping antiphase doublets in cases where the splitting is substantially smaller than one-half of the linewidth. An illustration of application to the extraction of coupling constants from experimental data containing an antiphase doublet with passive splittings is given.

THEORY

The Bayesian perspective on statistics asserts that a probability represents a degree of belief rather than a frequency of occurrence (*4*). In other words, it is possible to speak of the probability of a particular vector of parameter values among all possible parameter vectors in a statistical model. Thus, the process of parameter estimation and error estimation is intimately connected from a Bayesian point of view. For example, one could choose the best parameter estimate as the parameter vector Θ which maximizes the probability density $P(\Theta|\mathbf{D})$ given a data vector $\mathbf{D}$. Similarly, the uncertainty in Θ could be expressed by a credible interval, i.e., a hyperrectangle in Θ space that encloses a given fraction of the probability density $P(\Theta|\mathbf{D})$. Similarly, correlations between the model parameters can be determined from the covariance structure of $P(\Theta|\mathbf{D})$. The formal justification of such methods and a discussion of their relationship to classical point and interval estimation is beyond the scope of this paper, but has been amply discussed in the statistical literature (e.g., Refs. (*5, 6*)).

Previous use of Bayesian methods in the statistical analysis of NMR data was pioneered by Bretthorst, who developed the theoretical methodology (*1, 7–11*) and with co-workers evaluated its performance using synthetic NMR data containing well-resolved signals (*12, 13*). Applications of Bayesian methods by other workers have included the estimation of coupling constants from poorly resolved in-phase

1090-7807/98 $25.00

doublets (*14*) and the estimation of frequencies in multidimensional constant-time NMR data (*15*). In addition, a more theoretically oriented study using methodology similar to that of the work described here has recently appeared in the electrical engineering literature (*16*). In our implementation, we have chosen to use a subset of the theoretical results obtained by Bretthorst (*1*).

For convenience, we present a condensed version of the theoretical background below. Bretthorst's notation has been retained, with some minor changes to reflect the fact that we make no distinction between the linear and nonlinear parameters. For more details, the reader should consult Ref. (*1*). All of the relevant information for the parameter estimation problem is contained in the posterior probability density of the model parameters Θ given the data vector $\mathbf{D}$ and any prior information I: $P(\Theta|\mathbf{D}, I)$. Using Bayes' theorem, we find

$$P(\Theta|\mathbf{D}, I) = \frac{P(\mathbf{D}|\Theta, I)P(\Theta|I)}{P(\mathbf{D}|I)}. \quad [1]$$

By taking the prior probabilities $P(\Theta|I)$ to be uniform, uninformative improper distributions, we find that the posterior probability of the parameters is proportional to the likelihood of the data given a parameter vector:

$$P(\Theta|\mathbf{D}, I) \propto P(\mathbf{D}|\Theta, I). \quad [2]$$

Furthermore, we define our model functions to consist of linear superpositions of m exponentially damped sinusoids

$$f_R(t) = \sum_{j=1}^{m} U_j(t) = \sum_{j=1}^{m} I_j \cos(\omega_j t + \phi_j) e^{-\alpha_j t}$$
$$f_I(t) = \sum_{j=1}^{m} V_j(t) = \sum_{j=1}^{m} I_j \sin(\omega_j t + \phi_j) e^{-\alpha_j t} \quad [3]$$

for the real and imaginary channels, respectively. For a typical NMR problem, our model parameters Θ are then the resonance frequencies ω, the phases ϕ, the amplitudes I, and the relaxation rates α. If we define the noise to be that part of the data vector not fit by the model, and take the prior probability for a given noise vector to be a Gaussian with variance σ^2, then the likelihood of the data is given by

$$P(\mathbf{D}|\Theta, \sigma, I) = (2\pi\sigma^2)^{-N} \exp\left(-\frac{Q}{2\sigma^2}\right), \quad [4]$$

where Q is the sum of the squared residuals:

$$Q = \sum_{i=1}^{N} (d_R(t_i) - f_R(t_i))^2 + \sum_{i=1}^{N} (d_I(t_i) - f_I(t_i))^2$$
$$= (d_R \cdot d_R + d_I \cdot d_I) - 2 \sum_{j=1}^{m} (d_R \cdot U_j + d_I \cdot V_j)$$
$$+ \sum_{j=1}^{m} \sum_{k=1}^{m} (U_j \cdot U_k + V_j \cdot V_k). \quad [5]$$

In Eqs. [4] and [5], N is the number of complex data points, d_R and d_I are the real and imaginary data vectors, respectively, and the $(\cdot)$ operator denotes the scalar product of two vectors. If the noise variance is unknown, or if we do not desire to estimate it, then it can be eliminated by integration over a suitable prior probability (in this case, Jefferey's prior $P(\sigma|I) = \sigma^{-1}$):

$$P(\mathbf{D}|\Theta, I) \propto \int_0^\infty \frac{P(\mathbf{D}|\Theta, \sigma, I)}{\sigma} d\sigma \propto \frac{(N-1)!}{Q^N}. \quad [6]$$

Since the Metropolis–Hastings algorithm only considers ratios of probabilities from the same density function, the fact that the probability density is known only up to a constant factor is irrelevant to our implementation.

To obtain a practically useful implementation of the theoretical expressions shown above, it is imperative that the quantity Q be evaluated as efficiently and as accurately as possible. The first term in Eq. [5] corresponds to a projection of the data vectors onto themselves. Since this need be done only once during the calculation, its optimization with respect to speed is not crucial. For the second term in Eq. [5], we evaluate the projections of the data vectors onto the model basis functions U_j and V_j using a generalization of the Clenshaw recurrence formula for finite sums (see Appendix), which presumes that the data have been sampled uniformly in time. We have chosen this instead of Bretthorst's method of using the exponentially weighted Fourier transform (*10*) because we must be able to evaluate Q at arbitrary frequencies. Although the Fourier transform approach is valid, it can be used to evaluate Q only at discrete points determined by the degree of zero-filling. In addition, there is the very real possibility of inaccuracy due to truncation artifacts. In order to simplify the use of the recurrence formula, we perform all calculations using dimensionless units for the frequencies and damping factors,

$$\omega = \frac{2\pi f}{\mathrm{SW}}, \quad \alpha = \frac{b}{\mathrm{SW}}, \quad [7]$$

where f and b are the resonance frequencies and decay rates in hertz, respectively, and SW is the spectral width in hertz. To evaluate the elements of the ''interaction matrix'' $\mathbf{g}$

$$g_{ij} = U_i \cdot U_j + V_i \cdot V_j \quad [8]$$

appearing in the third term of Eq. [5], we make use of the closed-form expressions given by Bretthorst (Ref. (*10*), Eqs. [28] and [29]). Unlike some previous implementations of Bayesian and maximum-likelihood estimation (*7, 15, 17*), we make no assumptions concerning the orthogonality of the model basis functions, thereby ensuring that our evalua-

tion of Q will be accurate even if there is significant overlap of the Lorentzian lines.

We have chosen to use the Metropolis–Hastings Monte Carlo algorithm (*2, 3*) to sample directly from $P(\mathbf{D}|\Theta, \sigma, I)$ or $P(\mathbf{D}|\Theta, I)$ (Eqs. [4] and [6], respectively). In this way, we can be assured that we are obtaining an accurate picture of the posterior distribution, even if it is highly non-Gaussian. In addition, there is now a substantial body of literature on the use of Markov chain methods (such as Metropolis–Hastings) in statistical inference that deals both with theoretical aspects such as their convergence properties and with practical implementation issues (*18, 19*). The Metropolis algorithm for stochastically sampling a density function was first developed to solve problems in statistical thermodynamics (*2*), and was later generalized and introduced into the statistical literature by Hastings (*3*). For a general introduction to the Metropolis–Hastings algorithm and its mathematical underpinnings, the reader may consult Refs. (*18, 19*).

In our implementation we use the following version of the original Metropolis algorithm:

1. Let $q(\Pi)$ be a multivariate normal distribution with a covariance matrix $\mathbf{C}$ centered at a point Π in an M-dimensional parameter space (this space need not have the same dimensionality or parametrization as the space of all model parameters Θ defined in Eq. [3]).
2. Choose a starting point Π_0 and let $i = 0$.
3. Sample a point $\mathbf{Y}$ from $q(\Pi_i)$.
4. Sample a point U from a uniform (0, 1) distribution.
5. If $U \leq \dfrac{P(D|\mathbf{Y})}{P(D|\Pi_i)}$, then let $\Pi_{i+1} = \mathbf{Y}$, else, let $\Pi_{i+1} = \Pi_i$.
6. Let $i = i + 1$.
7. Go to 3.

It can be shown that no matter what the form of $q(\Pi_i)$, the above algorithm will converge to sampling from $P(D|\Pi_i)$ given a sufficient number of equilibration iterations, assuming certain regularity conditions.

We believe that the use of Monte Carlo sampling offers significant advantages in the Bayesian NMR parameter estimation problem. In particular, it allows the very straightforward incorporation of constraints among model parameters simply by a suitable choice of the Π space sampled by the Metropolis algorithm. Furthermore, it forces us to make very few assumptions about the shape of the posterior distribution. By contrast, the methods employed by Bretthorst at times approximate the posterior probability density function by expanding the argument of the exponential function in Eq. [4] in a Taylor series about the parameter vector which maximizes the posterior probability (*1, 8, 11*). This series expansion is then truncated at the quadratic term, which is equivalent to the assumption that the posterior is approximately Gaussian. Although this is true for many simple NMR parameter estimation problems, it need not be the case for some extreme but not unrealistic cases, as we will see below. Monte Carlo sampling is particularly advantageous when we wish to determine the limits of our ability to quantitatively interpret poor-quality data.

Although the Metropolis algorithm is very straightforward, it can be dangerous if misused. One insidious aspect of Monte Carlo-based sampling methods is the need to determine the degree of convergence. Two aspects of convergence must be considered when using Monte Carlo methods. The first is the convergence from the arbitrary starting point to the target distribution (in our case the posterior probability density), also known as the burn-in or equilibration period. The number of iterations required to achieve this convergence can be greatly influenced by the choice of starting point. For the types of data and models likely to be found in NMR data analysis, the most crucial parameters are the frequencies, as the posterior density is usually very sharply peaked in those dimensions. Fortunately, adequate starting values can be obtained from a simple Fourier transform spectrum. Assurance that convergence to a unique stationary distribution has been achieved can be obtained by performing multiple runs from several widely dispersed starting points (*20*).

The second important measure of convergence pertains to the efficiency with which the target distribution is sampled after the burn-in period. This often depends crucially on the sampling density $q(\Pi)$. A poor choice of $q(\Pi)$ can lead to extremely inefficient sampling, particularly if the target density is highly correlated. This could cause an inexperienced user to underestimate the uncertainty in the parameter estimates. In practice, inefficient sampling can usually be detected through the presence of low-frequency oscillations in the Monte Carlo points as a function of iteration number, and can be corrected by adjusting the sampling covariance matrix $\mathbf{C}$ so that $q(\Pi)$ better approximates $P(D|\Pi)$ (*19*). In our experience, $\mathbf{C}$ can be determined in an iterative fashion during the equilibration phase of the calculation. This can be done by performing several short Monte Carlo simulations and estimating the sample variances and correlation coefficients

$$K_{ij} = \frac{C_{ij}}{\sqrt{C_{ii}C_{jj}}} \qquad [9]$$

from the Monte Carlo samples. The variances C_{ii} are then scaled by a factor of 0.4–0.8 to keep the rejection rate from becoming too large, and the off-diagonal elements of $\mathbf{C}$ are recalculated from the correlation coefficients. This new sampling covariance matrix $\mathbf{C}$ is used for the next Monte Carlo run, the starting point of which is taken to be the last Monte Carlo sample from the previous run. This is repeated until plots of the parameter values vs iteration number have stabi-

lized, do not have low-frequency oscillations, and have rejection rates between 50 and 80%.

The argument could be made that the difficulties associated with the Monte Carlo sampling from highly correlated distributions could be avoided by a suitable reparametrization of the model (*19, 21*). In principle, a reparametrization could be found which could both alleviate the difficulties caused by correlation and restore approximate normality, allowing the use of Gaussian approximations such as those used by Bretthorst (*1, 7*). Although such reparametrizations could result in an improvement in computational efficiency, we feel that they are not helpful in the long run. In particular, reparametrization of the model functions could result in parameters with no obvious physical meaning, thereby depriving the analyst of an ability to judge parameter values on the basis of physical experience. Monte Carlo sampling based on the model in its spectroscopically natural parametrization has a simplicity and directness of interpretation and visualization that can assist rather than hinder the spectroscopist's intuition.

IMPLEMENTATION

The implementation outlined above was coded using the C programming language with an X/Motif-based graphical user interface in the package *X Rambo* (X-based Rigorous parAMeter estimation using Bayesian methOds). Output visualization and plotting were performed using the *xmgr* (P. J. Turner, v. 3.01) and *X Gobi* (Bellcore, 1995 release) software packages. Pseudorandom number generation was performed using the L'Ecuyer long-period uniform deviate generator and the Bays–Durham shuffling procedure given in (*22*). Independent multivariate normal pseudorandom numbers were generated from uniform deviates using the Box–Muller algorithm (*22, 23*). Correlated multivariate normal pseudorandom numbers were generated using the method of Barr and Slezak (*23*): given a vector Z of independent normal random numbers of mean zero and variance one, a vector X normally distributed with covariance matrix $\mathbf{C}$ can be calculated using

$$X = M + \mathbf{B}Z, \qquad [10]$$

where M is the vector of mean values, and $\mathbf{B}$ is the lower triangular matrix determined using the Cholesky factorization $\mathbf{C} = \mathbf{B}\mathbf{B}^{\mathrm{T}}$ (*24*). All data sets were analyzed using models of the form of Eq. [3] using Metropolis Monte Carlo sampling from $P(\Theta|\mathbf{D}, I)$ (Eqs. [2] and [6]). Means and the elements of the variance/covariance matrix of the Monte Carlo samples were estimated using Miller's updating method (*25*). Multivariate credible regions were determined from the Monte Carlo samples using the nonparametric method of Besag *et al.* (*26*).

RESULTS FOR A DATA SET CONTAINING TWO WELL-RESOLVED SIGNALS

To test our implementation and to provide some empirical justification for the use of Bayesian error estimates, we performed parameter estimation using a data set containing 256 complex points synthesized from the parameters shown in Table 1 (all phases ϕ_i were set to zero), and synthetic Gaussian noise with a standard deviation of 20 was added. The Fourier transform of the data set is shown in Fig. 1. The sampling covariance matrix $\mathbf{C}$ was determined using the iterative approach described above. A total of 60,000 Monte Carlo iterations were performed in the quantitative sampling phase of the calculation, of which every fourth value was stored to disk for subsequent analysis. The overall Metropolis rejection rate during the parameter estimation sampling was 70%. The final 60,000 iteration run required approximately 36 s of CPU time on a 150-MHz Silicon Graphics Indy R4400 computer, and the execution time was dominated by the evaluation of the projections of the model basis functions onto the data (i.e., the second term of Eq. [5]).

The Monte Carlo samples are shown graphically in Fig. 2, and summary statistics of the parameter estimates are given in Table 1. The fact that the parameter space is six-dimensional makes visualization of the output somewhat challenging. We simplify the visualization problem by plotting the Monte Carlo samples on 15 panels, corresponding to each of the pairwise combinations of the M independent model parameters. Therefore, each panel corresponds to a projection of the entire ''cloud'' of Monte Carlo samples onto a two-dimensional subspace corresponding to one of the $M(M-1)/2$ possible orthogonal ''viewpoints'' in the M-dimensional parameter space. Overall, the resulting figure gives a visual representation of the correlations between the model parameters. For example, it is immediately clear from Fig. 2 that all of the parameter estimates are uncorrelated, with the notable exceptions of positive linear correlation between the estimated linewidths and estimated intensities for each peak. This is not surprising, since the best fit to a set of observed FID data points given an underestimate of the decay rate would result in an underestimate of the $t = 0$ point of the FID, and *vice versa.*

In addition to the mean values and the standard deviations of the Monte Carlo samples, Table 1 also includes nonparametric estimates of the 70, 85, and 95% credible intervals. For comparison, given a *univariate* normal distribution, these credible intervals would correspond to approximately ± 1.0, ± 1.5, and ± 2.0 standard deviations about the mean, respectively. A nonparametric analysis of correlated normal pseudorandom numbers corresponding to the covariance matrix estimated from the Monte Carlo samples is consistent with the posterior probability being a multivariate Gaussian (data not shown). Since in this case the posterior distribution is approximately symmetrical, it can be assumed that the

TABLE 1
Results of Parameter Estimation for the Well-Resolved Two-Signal Data Set

	Peak 1			Peak 2		
	Frequency (Hz)[a]	Linewidth (Hz)[a]	Intensity	Frequency (Hz)[a]	Linewidth (Hz)[a]	Intensity
Correct values[b]	500.000	15.00	50.00	−100.000	5.00	40.00
Mean ± standard deviation of MC samples	499.81 ± 0.2	13.2 ± 2	46.6 ± 3	−100.11 ± 0.1	7.0 ± 2	42.6 ± 3
70% credible interval[c]	499.46	9.8	40.9	−100.38	4.0	37.4
	500.16	16.9	52.7	−99.84	10.1	48.1
85% credible interval[c]	499.39	9.1	40.0	−100.43	3.3	36.4
	500.23	17.5	53.8	−99.79	10.7	49.0
95% credible interval[c]	499.30	8.3	38.9	−100.48	2.6	35.3
	500.32	18.5	55.2	−99.73	11.6	50.5

[a] Parameter values have been converted from dimensionless units assuming a spectral width of 3000 Hz.
[b] Parameters used to create the synthetic data.
[c] Parameter values which define the corners of a hyperrectangle in the parameter space which encloses the indicated fraction of the posterior probability density.

means of the Monte Carlo samples will be good estimators for the maximum-likelihood parameters. It is clear from the results shown in Table 1 and Fig. 2 that the maximum-likelihood estimates determined from the Monte Carlo samples are consistent with the correct parameter values. In fact, all of the correct parameter values lie within the estimated 70% credible intervals (for this particular noise realization).

In order to convince ourselves that the credible intervals are indeed a realistic estimate of the uncertainties in the parameter values, we performed the following computational

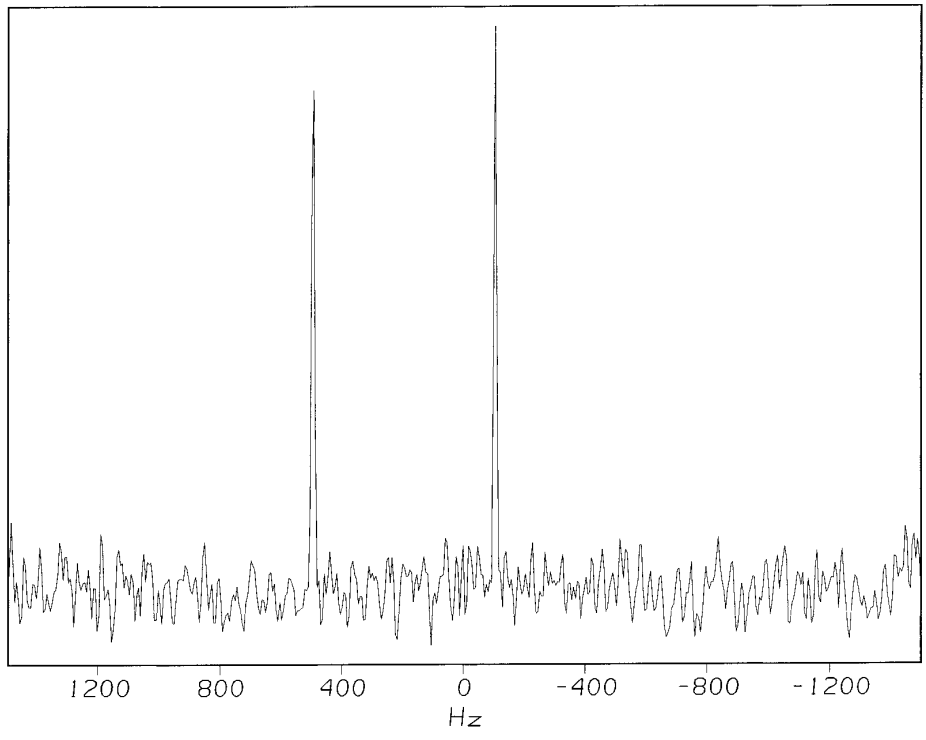

FIG. 1. Fourier transform of the well-resolved two-signal synthetic data set. The data were zero-filled to 512 points, and an 80° skewed sine-bell window function was applied prior to transformation.

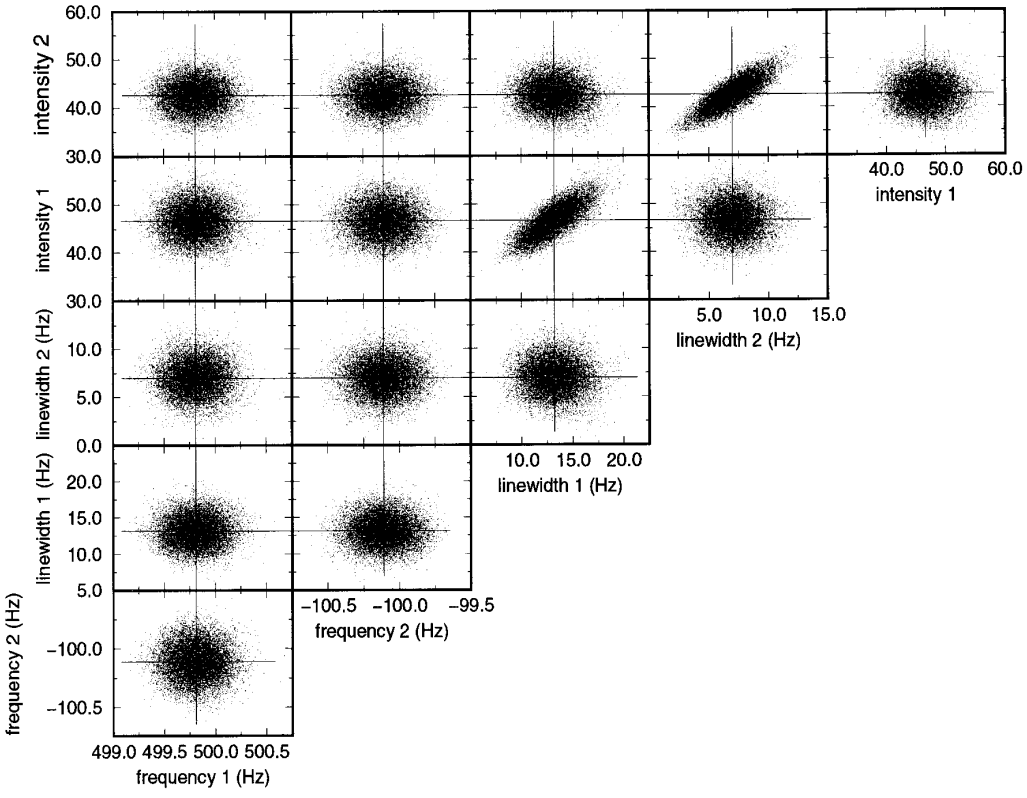

FIG. 2. Monte Carlo samples for the well-resolved two-signal synthetic data set. Each panel corresponds to a projection of the full set of Monte Carlo samples onto a plane corresponding to each pair of parameters. The solid lines indicate the maximum-likelihood estimates of each parameter.

experiment. For a given synthetic noise realization, we performed 40,000 iterations of Metropolis Monte Carlo sampling in order to determine the estimated parameters (the mean of the Metropolis Monte Carlo samples) and the estimated error (the standard deviation of the Metropolis Monte Carlo samples). We then repeated this process for 800 different noise realizations and determined the following four statistical parameters: (i) the mean of the estimated parameters, (ii) the standard deviation of the estimated parameters, (iii) the mean of the estimated errors, and (iv) the standard deviation of the estimated errors (Table 2). First of all, it is apparent from Table 2 that the means of the estimated parameters agree very well with the correct values shown in Table 1, indicating that the parameter estimates are unbiased in the classical sense. Furthermore, the standard deviations of the estimated errors are approximately an order of magnitude smaller than the means of the estimated errors, indicating that the estimated errors have negligible uncertainty (at least for the purpose of error estimation). Most importantly, the standard deviations of the estimated parameters and the means of the estimated errors are nearly identical, indicating that the Bayesian error estimates are empirically consistent

TABLE 2
Results of Parameter Error Validation for the Well-Resolved Two-Signal Data Set

	Mean of the estimated parameters	Standard deviation of the estimated parameters	Mean of the estimated errors	Standard deviation of the estimated errors
Peak 1 frequency[a]	500.03	0.20	0.21	0.02
Peak 1 linewidth[a]	15.3	2.0	2.1	0.06
Peak 1 intensity	50.5	3.4	3.4	0.18
Peak 2 frequency[a]	−100.00	0.14	0.14	0.01
Peak 2 linewidth[a]	5.67	1.6	1.7	0.15
Peak 2 intensity	41.1	2.7	2.9	0.18

[a] Parameter values have been converted from dimensionless units to hertz assuming a spectral width of 3000 Hz.

with the more classically motivated estimates. Therefore, we conclude that the width of the Bayesian posterior probability density is indeed a realistic estimate of the error in the parameter estimate.

RESULTS FOR COUPLING CONSTANT ESTIMATION FROM ANTIPHASE DOUBLETS

In order to explore the utility of our Monte Carlo-based Bayesian approach in more challenging problems, we used our program to estimate coupling constants from antiphase doublets in which the linewidth was significantly larger than the splitting. The estimation of coupling constants is obviously of great importance to practical NMR applications, because of the relationship of three-bond coupling constants to torsional angles about the central bond (*27*). For macromolecules, these coupling constants are often best measured from antiphase components of cross peaks in two- and three-dimensional correlation experiments (*27*). Double-quantum spectra used to separate signals from ^{13}C pairs as opposed to isolated ^{13}C sites (such as the INADEQUATE experiment) also return antiphase multiplet components (*28*). The peak-to-peak separation of an overlapped antiphase doublet is a complicated function of the actual splitting and the linewidth, and a naive measurement of the peak-to-peak separation will overestimate the coupling constant. More accurate estimates can be obtained by solving simultaneous equations involving the absorptive and dispersive peak-to-peak separations (*29*), nonlinear curve fitting in the frequency domain (*30–32*), or trigonometric manipulation of the time-domain data (*33*). It is generally accepted that these methods are unreliable when the linewidth is more than twice the magnitude of the splitting (*34*), but the estimation of the precision of the extracted splittings has been very difficult even under favorable conditions.

We generated three synthetic 256-point time-domain data sets, each containing a single antiphase doublet with a splitting of 4 Hz, and linewidths of 8, 16, and 28 Hz, respectively. Gaussian noise with standard deviation 20 was added to all three FIDs, resulting in frequency-domain signal-to-noise ratios of approximately 22:1, 9:1, and 4:1, respectively. The Fourier transforms of the data sets are shown in Fig. 3. All three data sets were analyzed using *X Rambo* in a manner similar to the well-resolved peak case described above, and the results are shown in Tables 3, 4, and 5. It was assumed that the intensities and linewidths of the two multiplet components were equal, and that the phases of the two components were 0° and 180°, respectively, thereby reducing the dimensionality of the parameter space Π to four from a potential maximum dimensionality of eight (the dimensionality of Θ). In addition, the frequencies of the two peaks of the multiplet were parametrized in terms of the frequency of the downfield multiplet component and the splitting. This allowed us to directly estimate the posterior distribution of the spectroscopically most interesting parameters.

Analysis of the 8- and 16-Hz linewidth data sets was performed using 60,000 Monte Carlo iterations, of which every fourth value was stored for analysis. The results are shown in Tables 3 and 4 and Fig. 4. It is clear that now all four parameters are linearly correlated. Both the mean and the maximum-likelihood parameter estimates are close to the correct values and the correct values are again within the 70% credible interval hyperrectangle. There is a clear asymmetry in the posterior distribution for the 16-Hz linewidth data set, as well as the appearance of nonlinear correlation in the frequency/linewidth, splitting/linewidth, and linewidth/intensity scatter plots, indicating that the posterior probability density can no longer be well approximated by a Gaussian distribution (Fig. 4b).

In order to test the limits of our ability to accurately measure coupling constants in noisy data, we attempted an analysis of an antiphase doublet with a 4-Hz splitting and 28-Hz linewidth. This is clearly an extreme example, the analysis of which would be nearly hopeless using standard methodologies. The resulting estimates are shown in Table 5 and Fig. 5. It is quite clear that the posterior distribution is wildly non-Gaussian. The distribution can be divided into two distinct regions: one which is in approximate agreement with the correct parameter values (which we will call region A), and one with intensities that are much larger and splittings that are much smaller than the correct values (region B). Since the posterior distribution has severe nonlinear correlations, it is not possible to choose a single covariance matrix **C** that would allow for efficient sampling. Instead, we chose to set the elements of the sampling covariance matrix corresponding to the frequency/intensity, splitting/intensity, and linewidth/intensity pairs to zero, and to perform 10.5 million Monte Carlo iterations, of which every 700th was saved for analysis. The simulation efficiency could be significantly improved by allowing **C** to depend on whether the current iteration point was in region A or B. Such a position-dependent sampling density could be easily incorporated using the Hastings generalization of the Metropolis algorithm (*3*), but has not yet been implemented in the current version of *X Rambo*.

As might have been expected given the linewidth and the signal-to-noise ratio of the data, the Monte Carlo Bayesian analysis indicates that only a vague quantitative interpretation is possible. For example, we could state with a reasonable degree of certainty that the coupling constant is less than 6 Hz based on the 95% credible interval (Table 5). A more precise determination of the splitting would require either data with a higher signal-to-noise ratio or further prior information about the other parameters. For example, we might know that an intensity greater than 1000 is physically impossible based on the molecular structure and the intensity of better-resolved doublets elsewhere in the spectrum. This

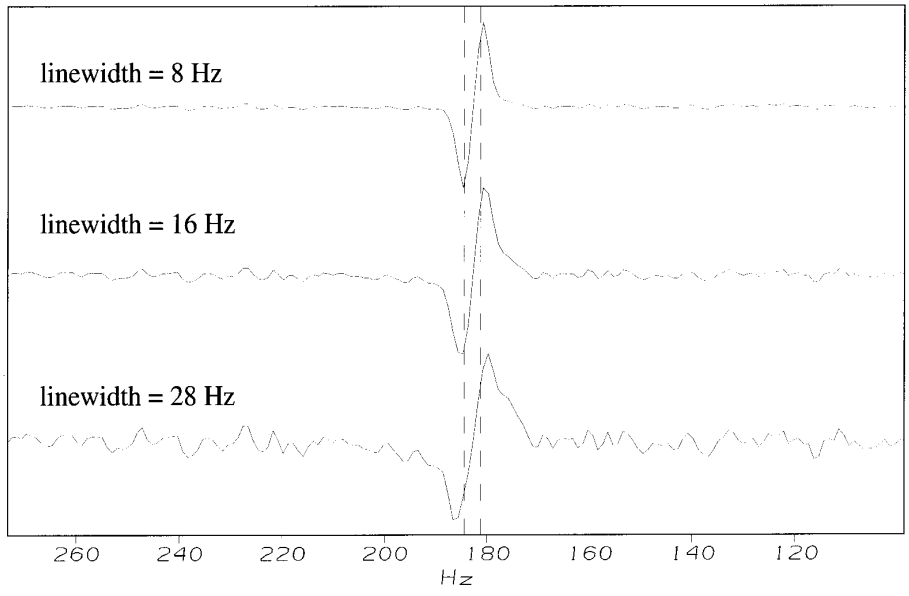

FIG. 3. Fourier transform of the antiphase doublet synthetic data sets, having a splitting of 4 Hz, and linewidths of 8, 16, and 28 Hz. The lines indicate the true 4-Hz splitting. The data were zero-filled to 512 points, and an 80° skewed sine-bell window function was applied prior to transformation.

would greatly reduce the posterior probability density in region B and allow the estimation of a realistic lower bound on the coupling constant. Such prior information could be incorporated in a quantitative fashion either in the form of equality constraints (such as intensity ratios predicted using the molecular structure) or as an informative prior probability $P(\Theta \mid I)$.

As a final example, we present the determination of coupling constants from antiphase doublets in actual experimental data. The data were obtained from a myristoylated peptide selectively labeled at several positions with ^{13}C and ^{15}N incorporated into field-oriented DMPC/DHPC micelles (*35*). This system spontaneously orients in magnetic fields, so that signals are split by residual through-space dipolar couplings to nearby spins, as well as through-bond scalar couplings to bonded spins. Accurate measurement of the splittings can provide valuable structural information (*36*). The experiment presented here is a one-dimensional ^{13}C INADEQUATE experiment (*37*) in which $^{13}C-^{13}C$ splittings appear as antiphase doublets. The data shown in Fig. 6a consist of

TABLE 3
Results of Parameter Estimation for the 8-Hz Linewidth Antiphase Doublet Data Set

	Frequency[a]	Splitting[a]	Linewidth[a]	Intensity
Correct values[b]	200.00	4.00	8.00	200.0
Mean ± standard deviation of MC samples	199.97 ± 0.04	4.07 ± 0.09	7.6 ± 0.7	190.5 ± 9
Maximum-likelihood estimate[c]	199.97	4.08	7.5	189.3
70% credible interval[d]	199.91	3.95	7.0	178.6
	200.03	4.17	8.2	203.0
85% credible interval[d]	199.90	3.92	6.8	175.6
	200.05	4.20	8.4	207.2
95% credible interval[d]	199.88	3.87	6.6	171.8
	200.07	4.23	8.6	212.1

[a] Parameter values have been converted from dimensionless units assuming a spectral width of 1100 Hz.
[b] Parameters used to create the synthetic data.
[c] The mean of 1% of the Monte Carlo samples having the largest likelihood.
[d] Parameter values which define the corners of a hyperrectangle in the parameter space which encloses the indicated fraction of the posterior probability density.

TABLE 4
Results of Parameter Estimation for the 16-Hz Linewidth Antiphase Doublet Data Set

	Frequency[a]	Splitting[a]	Linewidth[a]	Intensity
Correct values[b]	200.00	4.00	16.00	200.0
Mean ± standard deviation of MC samples	199.9 ± 0.2	4.0 ± 0.4	15 ± 1	192 ± 28
Maximum-likelihood estimate[c]	199.9	4.2	15.1	180
70% credible interval[d]	199.8	3.6	14.2	164
	200.2	4.4	17.2	230
85% credible interval[d]	199.7	3.4	13.8	158
	200.3	4.5	17.7	247
95% credible interval[d]	199.7	3.1	13.2	150
	200.5	4.6	18.4	278

[a] Parameter values have been converted from dimensionless units assuming a spectral width of 1100 Hz.
[b] Parameters used to create the synthetic data.
[c] The mean of 1% of the Monte Carlo samples having the largest likelihood.
[d] Parameter values which define the corners of a hyperrectangle in the parameter space which encloses the indicated fraction of the posterior probability density.

four such doublets, one of which is further split by a passive C–N coupling. There is also what appears to be an artifact arising from natural-abundance ^{13}C in the lipid chains at high field. The doublets arise from carbonyl carbon–alpha carbon pairs from a doubly labeled phenylalanine residue and a doubly labeled myristoyl chain in the molecule. For comparison with the results obtained using synthetic data above, the ratios of linewidths and coupling constants for the various doublets range from approximately 1:1 to 3:1, and the frequency-domain signal-to-noise ratios range from 5:1 to 15:1. Since each $^{13}C-^{13}C$ pair gives rise to two doublets with the same splitting, there are only two distinct active coupling constants in the spectrum, as shown in Fig. 6a. This prior information can be easily incorporated into the analysis using *X Rambo* by introducing equality constraints among the appropriate splitting parameters. No constraints involving the intensities and linewidths were used, and all phases were assumed to be 0° or 180°, resulting in a Π space having 30 dimensions (5 center frequencies, 3 splittings, 11 linewidths, and 11 intensities). The initial starting point was estimated from the Fourier transform spectrum, and convergence was achieved using the iterative method described above. The parameter estimation Monte Carlo was run for 60,000 iterations, of which every fourth value was stored to disk for subsequent analysis. The overall rejection rate was 75%.

A detailed visual inspection of the resulting samples using *X Gobi* revealed the posterior probability density to be well

TABLE 5
Results of Parameter Estimation for the 28-Hz Linewidth Antiphase Doublet Data Set

	Frequency[a]	Splitting[a]	Linewidth[a]	Intensity
Correct values[b]	200.00	4.00	28.00	200.0
Mean ± standard deviation of MC samples	201.0 ± 0.9	2 ± 2	30 ± 2	1020 ± 1060
Maximum-likelihood estimate[c]	199.9	4.3	26.8	174
70% credible interval[d]	199.6	0.3	25.6	150
	201.9	4.8	32.7	3037
85% credible interval[d]	199.4	0.2	24.0	129
	202.0	5.4	33.3	3486
95% credible interval[d]	199.1	0.2	22.3	112
	202.1	5.8	34.2	3818

[a] Parameter values have been converted from dimensionless units assuming a spectral width of 1100 Hz.
[b] Parameters used to create the synthetic data.
[c] The mean of 1% of the Monte Carlo samples having the largest likelihood.
[d] Parameter values which define the corners of a hyperrectangle in the parameter space which encloses the indicated fraction of the posterior probability density.

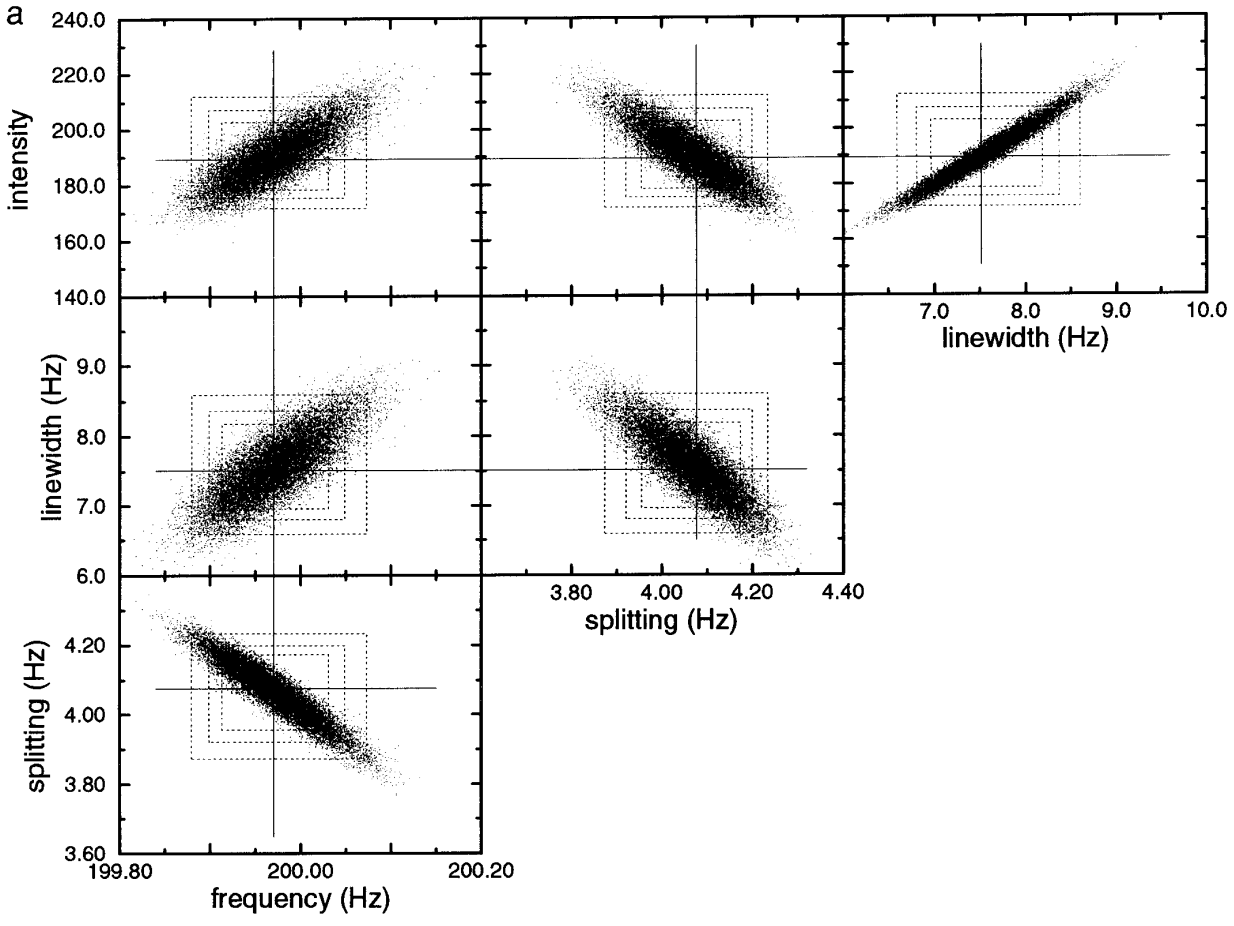
a
intensity
240.0
220.0
200.0
180.0
160.0
140.0
linewidth (Hz)
9.0
8.0
7.0
6.0
splitting (Hz)
4.20
4.00
3.80
3.60
199.80
200.00
200.20
frequency (Hz)
7.0
8.0
9.0
10.0
linewidth (Hz)
3.80
4.00
4.20
4.40
splitting (Hz)

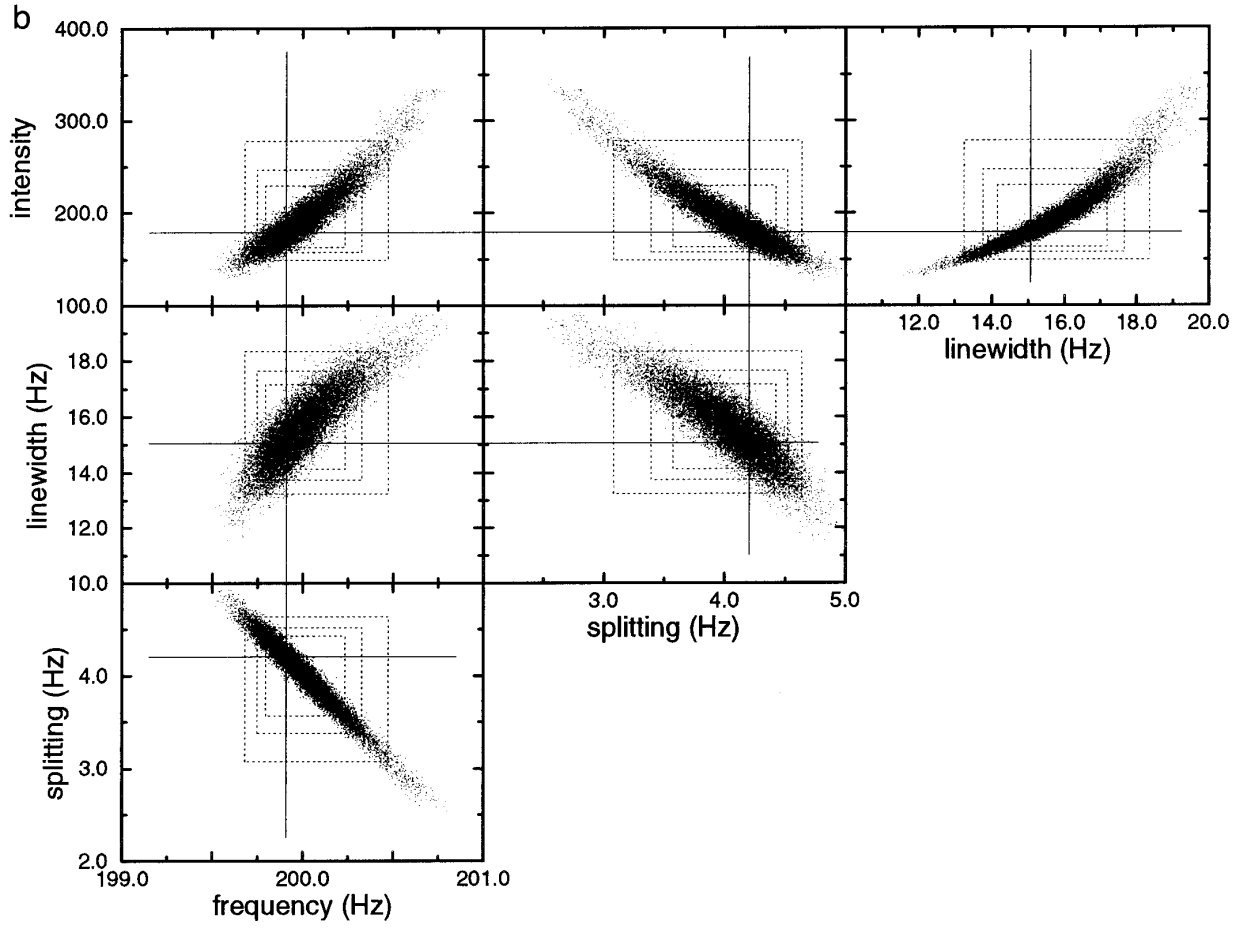
b
intensity
400.0
300.0
200.0
100.0
linewidth (Hz)
18.0
16.0
14.0
12.0
10.0
splitting (Hz)
4.0
3.0
2.0
199.0
200.0
201.0
frequency (Hz)
12.0
14.0
16.0
18.0
20.0
linewidth (Hz)
3.0
4.0
5.0
splitting (Hz)

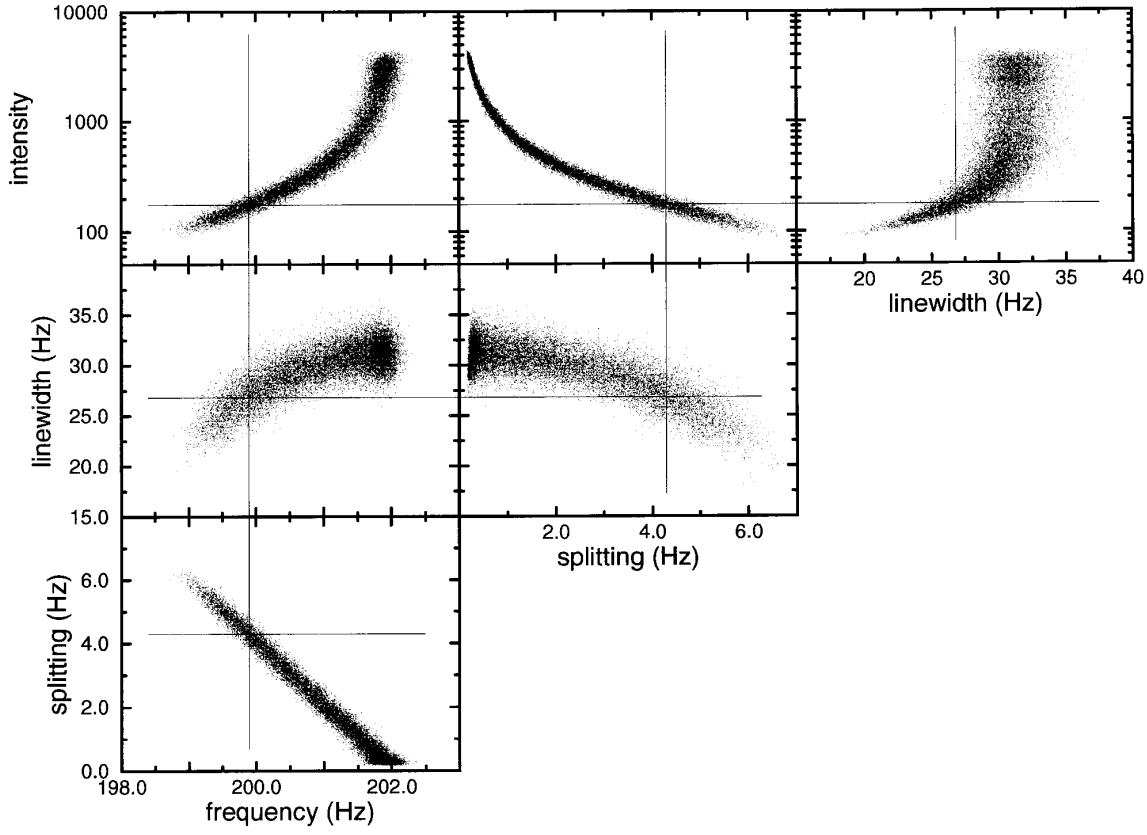

FIG. 5. Monte Carlo samples for the 28-Hz linewidth antiphase doublet data set. Each panel corresponds to a projection of the full set of Monte Carlo samples onto a plane corresponding to each pair of parameters. The solid lines indicate the maximum-likelihood estimates of each parameter. Note the logarithmic scale for the intensity parameter.

behaved and nearly Gaussian, with significant linear correlation between J_1 and J_2, as well as between linewidths and intensities within doublets. The resulting estimates of the posterior probability densities of the three coupling constants independent of the other parameters are shown in Fig. 6b. As might have been expected based on the relative linewidths, J_3 has the largest uncertainty of the three splittings, but it is clear that an estimate of J_3 and an estimate of the error are possible. It is clear that *X Rambo* is an effective tool for the determination of splittings from experimental data. Furthermore, we can easily determine splittings from multiplets containing both active and passive couplings, as well as impose equality constraints among the parameters.

DISCUSSION

The prominent feature of Monte Carlo sampling-based Bayesian parameter estimation is the ability to simultaneously obtain parameter and error estimates, while also providing an immediate assessment of any correlations among the model parameter estimates. Such an analysis of data with overlapping antiphase doublets, like that presented above, would have been difficult using the earlier approaches of Kim and Prestegard (*29*) or McIntyre and Freeman (*33*). In particular, the method of Kim and Prestegard cannot be used for the analysis of the 28-Hz linewidth antiphase data, as the signal-to-noise is insufficient to estimate the dispersive peak-to-peak separation. The low signal-to-noise also makes the method of McIntyre and Freeman difficult to apply, as the integral of the magnitude spectrum as a function of estimated J does not show a clean minimum. Estimation of the uncertainties in the estimated J is not possible using either method.

The presence of correlations among parameters in models for NMR data has in the past been a neglected source of information. It is clear from the simulated data sets

FIG. 4. Monte Carlo samples for the 8-Hz linewidth (a) and 16-Hz linewidth (b) antiphase doublet data sets. Each panel corresponds to a projection of the full set of Monte Carlo samples onto a plane corresponding to each pair of parameters. The solid lines indicate the maximum-likelihood estimates of each parameter. The boxes shown in dotted lines represent nonparametric estimates of the 70, 85, and 95% credible intervals (see text).

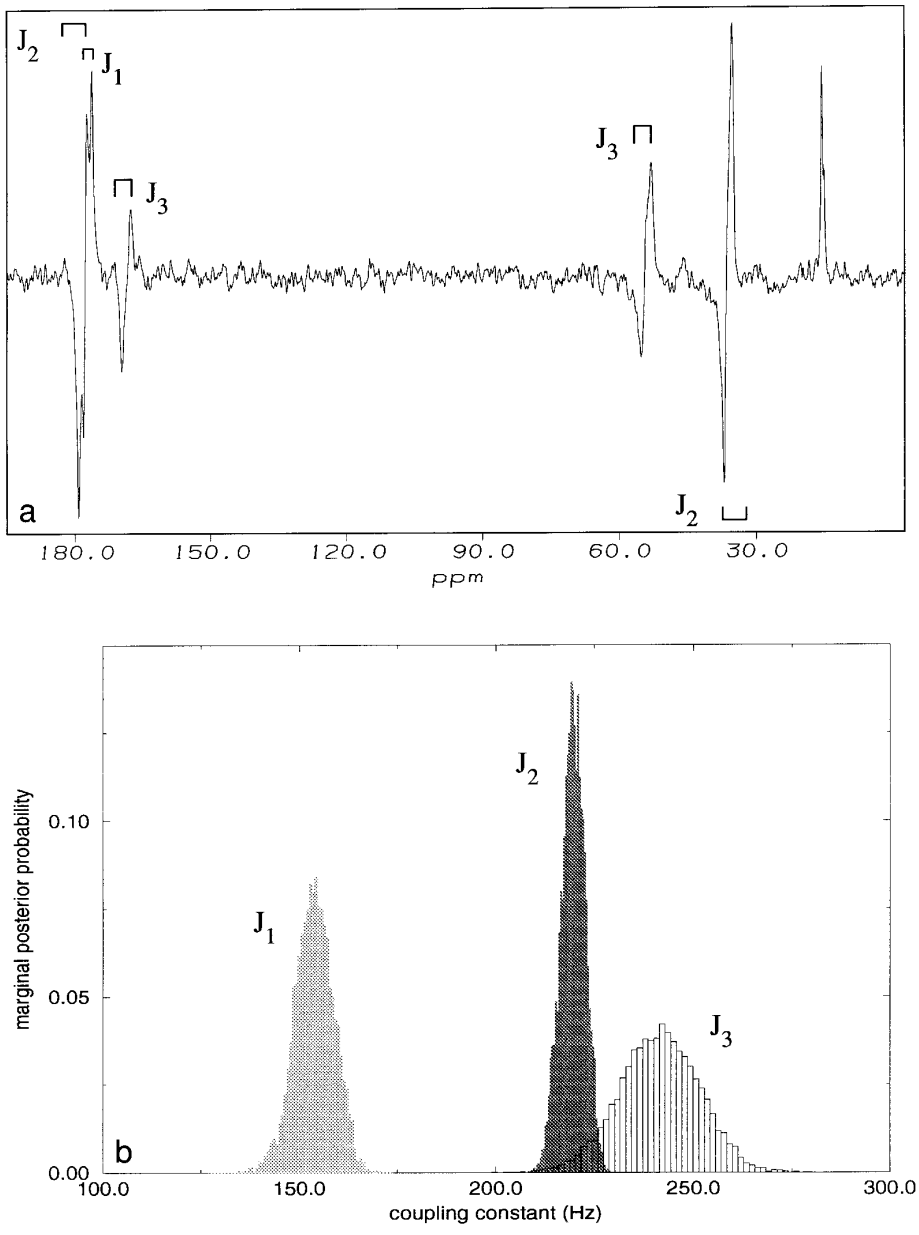

FIG. 6. (a) Fourier transform of a one-dimensional ^{13}C INADEQUATE data set obtained using the 15 amino acid myristoylated peptide Myr*-G†NIFANL*F†KGL*F†GKK containing backbone $^{13}C-^{13}C(*)$- and $^{15}N(\dagger)$-labeled residues incorporated into field-oriented DMPC/DHPC micelles (*37*). A 40-Hz exponential window function was applied prior to Fourier transformation. The upfield singlet is due to natural-abundance ^{13}C in the lipid chains. Splittings J_2 and J_1 have been assigned to the $^{13}C-^{13}C$ pair of the myristoyl group and the ^{15}N of the adjacent glycine, respectively. J_3 has been assigned to the $^{13}C-^{13}C$ pair of the second labeled leucine residue. No passive coupling is observed due to the averaging of the $^{13}C-^{15}N$ coupling to zero. The other labeled sites are not observed due to excessively broad linewidths and/or small splittings (*37*). (b) Histograms showing the estimates of the marginal posterior probability densities of the three coupling constants determined from the Bayesian Monte Carlo analysis. The maximum-likelihood estimates of J_1, J_2, and J_3 are 155, 219, and 239 Hz, respectively. The 85% credible intervals estimated using the full 30-dimensional joint density are 139–167, 211–227, and 212–268 Hz, respectively.

presented above that such correlations exist even for well-resolved Lorentzians, and that these correlations can become nonlinear for more complex models. Recognition of correlations among the parameter estimates could be used to reduce the uncertainty of parameter estimates by the introduction of prior information. One simple spectroscopically relevant application is the estimation of intensities as a function of a pulse sequence parameter, such as in a series of relaxation or pulsed field-gradient diffusion data sets. Since we do not expect the linewidths of the

peaks to change as a function of the relaxation delay or gradient strength, we could estimate a single linewidth parameter for each peak using all of the data sets simultaneously. This will presumably lead to a higher precision estimate of the linewidth, as we are taking advantage of a larger number of data points. Since linewidth and intensity estimates are correlated even for well-resolved Lorentzians, this will result in a concomitant decrease in the error in the intensity estimates. The ability to simultaneously analyze such groups of data sets has already been incorporated into *X Rambo*, along with its natural extension to the analysis of two-dimensional data.

The existence of strong nonlinear correlations between model parameters in some extreme cases also illustrates the dangers of the use of marginalized probability densities, which have been advocated as a mechanism for reducing the dimensionality of the parameter estimation problem (*1, 11*). Given a multivariate joint probability density function over all of the model parameters, one can determine a univariate or lower-dimensional multivariate marginal density function by integrating the product of joint density function and any prior density over the entire domain of the variables being eliminated. An example of marginalization can be seen in Eq. [6] above. Also, each panel in Figs. 2, 4, and 5 can be viewed as an approximate representation of the bivariate marginal density for that pair of parameters. Although the use of marginal densities can be convenient from a computational or visualization perspective, in general it represents a loss of information and can have an undesirable effect on the parameter estimates. For example, consider the posterior probability for the 28-Hz linewidth doublet data set shown in Fig. 5. If we had marginalized with respect to the intensity parameter, we would have obtained the density shown in the lower three panels of Fig. 5. It is clear that the maximum-likelihood estimate based on that marginalized density is in region B (splitting $\approx$ 0.4 Hz), whereas the maximum-likelihood estimate based on the full joint density is in region A (splitting $\approx$ 5 Hz). This seemingly paradoxical result arises from the fact that the probability mass corresponding to region B is spread out over an extremely large area in the intensity dimension. Similarly, any prior information about the intensity could have a very dramatic effect on the marginal density of the splitting. Therefore, marginalization can be a useful tool, but it should be used with caution.

One further point of concern is the effect of using an inappropriate model for the decay envelope of the signal. In all of the above analyses, we have assumed that the decay is a simple exponential, but this need not be the case for experimental data due to the presence of magnetic field inhomogeneity or multi-spin relaxation processes. Bretthorst has shown that the use of an incorrect decay model does not significantly impact the ability to accurately estimate the frequencies of the decaying sinusoids if the signals are well resolved (*7*). In order to investigate this effect in overlapped antiphase doublets, we generated 256-point noiseless synthetic FIDs containing two antiphase peaks at 200 and 208 Hz (at a spectral width of 1100 Hz), and a Gaussian decay envelope given by

$$\exp\left(\frac{-t^2}{2\beta^2}\right).$$

This choice of splitting and spectral width results in a single node at approximately the 138th data point (t = 125 ms) due to the beat frequency of the two antiphase components. From a data analysis perspective, estimation of the time at which this node occurs is equivalent to an estimation of the splitting. We then performed a Bayesian parameter estimation from this noiseless data with an exponentially decaying sinusoidal model using *X Rambo*. The deviation of the resulting estimate for the splitting from the correct value is shown in Fig. 7, and can be taken to be an estimate of the systematic error due to the use of the incorrect decay model.

As expected, the deviation is nearly zero if the frequency-domain linewidth is significantly smaller than the splitting. As the linewidth increases, however, the systematic error becomes significant. This is not surprising, as the signal envelope is now determined by the interference of the two antiphase components, the details of which are significantly influenced by the decay mode of the signal. For β values comparable in magnitude or longer than the time of the node point, the exponentially decaying sinusoidal fit will attempt to better fit the thinner tail of the Gaussian function by shifting the estimated node point to a later time, resulting in an underestimate of the splitting. As the lines become broader (i.e., β continues to decrease), the Gaussian envelope will cause a steeper decay *before* the node point, which the exponential model will attempt to fit by shifting the estimated node point to an earlier time, resulting in an overestimate of the splitting. Thus, the dependence of the systematic error of the splitting on β is not monotonic, but instead reaches an extremum when β is somewhat smaller than the position of the first node.

Systematic error due to an incorrect assumption of decay mode has also been observed by Yang and Havel in their frequency-domain analysis of COSY spectra (*31*). They recommended that a Gaussian lineshape be used in the analysis, and that the data should be processed in such a way as to make the peaks appear as nearly Gaussian as possible. Indeed, if it was known *a priori* that the decay mode was exponential with a rate a, one could use the window function $\exp(at)\exp(-bt^2)$ to perform a Lorentz–Gauss transformation (*34*). A nonlinear fit using a

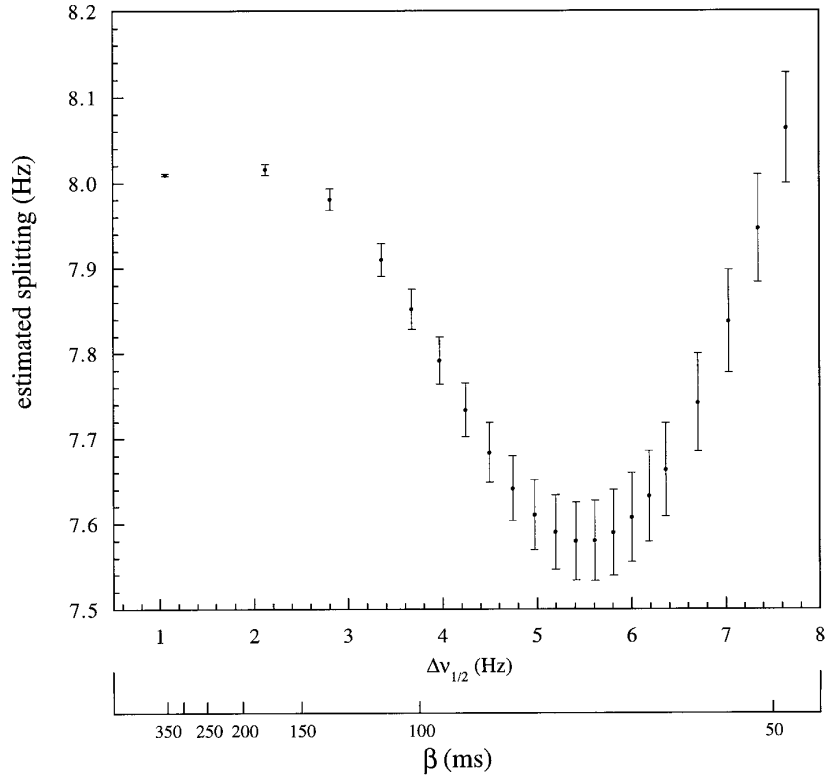

FIG. 7. An assessment of the systematic error in the estimate of a splitting from an antiphase doublet arising from the fitting of a noiseless Gaussian-decaying signal using an exponentially decaying model function (see text). The error bars correspond to the marginal standard deviations of the estimated splitting, and represent a range of splitting values consistent with the Gaussian-decaying data when fit with an exponentially decaying model. The estimated splittings are plotted as a function of the frequency-domain full width at half-height ($\Delta\nu_{1/2}$) and the time-domain Gaussian decay parameter β.

Gaussian model would then be free of systematic error, at least in principle. However, if the decay mode or rate is unknown, then the use of an arbitrarily chosen window function will not necessarily result in a fit which is free of systematic error, even if it results in peaks which appear to be Gaussian. If very high accuracy splitting estimates are required, the safest approach would be to fit the time-domain data without the use of any window function using a variety of possible decay models. The dependence of the splitting on the assumed model could then be assessed directly, and if significant differences are observed the best model could be chosen using statistical model selection methods.

Thus, we believe that Monte Carlo sampling-based Bayesian estimation represents a powerful tool for the quantitative analysis of NMR data. Prior information about constraints among model parameters can be easily incorporated, and the model can be parametrized in a manner which is consistent with the user's spectroscopic intuition. Parameter and error estimation is performed simultaneously, thereby providing an immediate assessment of the uncertainties in the estimated parameters. Very few approximations are made, thereby assuring the analyst that these uncertainties have been estimated in a realistic manner.

APPENDIX

A Generalization of the Clenshaw Recurrence Formula for Finite Sums

Consider a summation of the form

$$S = \sum_{k=0}^{N} d_k F_k, \qquad [11]$$

where the function F obeys the recurrence relation

$$F_{k+1} = a_k F_k + b_k F_{k-1}. \qquad [12]$$

If we define a new recurrence relation

$$y_k = a_k y_{k+1} + b_{k+1} y_{k+2} + d_k \qquad [13]$$

with

$$y_{N+1} = y_{N+2} = 0, \quad [14]$$

then the Clenshaw recurrence formula (*22, 38*) for the sum S is

$$S = F_1 y_1 + F_0 b_1 y_2 + d_0 F_0. \quad [15]$$

Equation [15] can be generalized to summations of the form

$$S = \sum_{k=0}^{N} d_k F_k G_k, \quad [16]$$

where the functions F and G obey the recurrence relations

$$F_{k+1} = a_k F_k + b_k F_{k-1}$$
$$G_{k+1} = c_k G_k. \quad [17]$$

If we define the recurrence relation

$$y_k = a_k c_k y_{k+1} + b_{k+1} c_{k+1} c_k y_{k+2} + d_k, \quad [18]$$

then the sum S can be written as

$$S = F_1 G_1 y_1 + F_2 G_0 b_1 y_2 + d_0 F_0 G_0. \quad [19]$$

Equation [19] can be easily verified by solving Eq. [18] for d_k and substituting into Eq. [16]. In our case,

$$G_k = e^{-\alpha_j k}, \quad c_k = e^{-\alpha_j}, \quad [20]$$

and

$$F_k = \begin{cases} \cos(\omega_j k) \\ \sin(\omega_j k), \end{cases} \quad [21]$$

where

$$a_k = 2 \cos \omega_j$$
$$b_k = -1 \quad [22]$$

for both sine and cosine modulation.

AVAILABILITY

Information on obtaining the software described in this paper can be found on the Prestegard lab WWW page accessed through http://www.ccrc.uga.edu.

ACKNOWLEDGMENTS

M.A. thanks Judit A. Losonczi for providing the myristoylated peptide data. This work was supported by grants from the National Institutes of Health: GM33225 and GM54160.

REFERENCES

1. G. L. Bretthorst, *J. Magn. Reson.* **88,** 533 (1990).
2. N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *J. Chem. Phys.* **21,** 1087 (1953).
3. W. K. Hastings, *Biometrika* **57,** 97 (1970).
4. H. Jeffreys, ''Theory of Probability,'' Oxford Univ. Press, London (1961).
5. J. C. Kiefer, ''Introduction to Statistical Inference,'' Springer-Verlag, New York (1987).
6. C. P. Robert, ''The Bayesian Choice: A Decision-Theoretic Motivation,'' Springer-Verlag, New York (1994).
7. G. L. Bretthorst, ''Bayesian Spectrum Analysis and Parameter Estimation,'' Lecture Notes in Statistics, Vol. 48, Springer-Verlag, New York (1988).
8. G. L. Bretthorst, *J. Magn. Reson.* **88,** 552 (1990).
9. G. L. Bretthorst, *J. Magn. Reson.* **88,** 571 (1990).
10. G. L. Bretthorst, *J. Magn. Reson.* **93,** 369 (1991).
11. G. L. Bretthorst, *J. Magn. Reson.* **98,** 501 (1992).
12. J. J. Kotyk, N. G. Hoffman, W. C. Hutton, G. L. Bretthorst, and J. J. H. Ackerman, *J. Magn. Reson.* **98,** 483 (1992).
13. J. J. Kotyk, N. G. Hoffman, W. C. Hutton, G. L. Bretthorst, and J. J. H. Ackerman, *J. Magn. Reson. A* **116,** 1 (1995).
14. K. S. Vines, R. F. Evilia, and S. L. Whittenburg, *J. Magn. Reson.* **100,** 195 (1992).
15. R. A. Chylla and J. L. Markley, *J. Biomol. NMR* **3,** 515 (1993).
16. P. Barone and R. Ragona, *IEEE Trans. Signal Proc.* **45,** 1806 (1997).
17. R. A. Chylla and J. L. Markley, *J. Biomol. NMR* **5,** 245 (1995).
18. W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, (Eds.), ''Markov Chain Monte Carlo in Practice,'' Chapman & Hall, London (1996).
19. J. J. K. Ó Ruanaidh and W. J. Fitzgerald, ''Numerical Bayesian Methods Applied to Signal Processing,'' Springer-Verlag, New York (1996).
20. A. Gelman, *in* ''Markov Chain Monte Carlo in Practice'' (W. R. Gilks *et al.,* Eds.), pp. 131–143, Chapman & Hall, London (1996).
21. W. R. Gilks and G. O. Roberts, *in* ''Markov Chain Monte Carlo in Practice'' (W. R. Gilks *et al.,* Eds.), pp. 89–114, Chapman & Hall, London (1996).
22. W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, ''Numerical Recipes in C: The Art of Scientific Computing,'' Cambridge Univ. Press, Cambridge (1992).
23. P. Bratley, B. L. Fox, and L. E. Schrage, ''A Guide to Simulation,'' Springer-Verlag, New York (1987).
24. G. H. Golub and C. F. Van Loan, ''Matrix Computations,'' The Johns Hopkins University Press, Baltimore (1989).
25. A. J. Miller, *J. Comput. Phys.* **85,** 500 (1989).
26. J. Besag, P. Green, D. Higdon, and K. Mengersen, *Stat. Sci.* **10,** 3 (1995).
27. D. A. Case, H. J. Dyson, and P. E. Wright, *Methods Enzymol.* **239,** 392 (1994).
28. R. R. Ernst, G. Bodenhausen, and A. Wokun, ''Principles of Nuclear

Magnetic Resonance in One and Two Dimensions,'' Clarendon Press, Oxford (1987).

29. Y. Kim and J. H. Prestegard, *J. Magn. Reson.* **84,** 9 (1989).

30. L. J. Smith, M. J. Sutcliffe, C. Redfield, and C. M. Dobson, *Biochemistry* **30,** 986 (1991).

31. J.-X. Yang and T. F. Havel, *J. Biomol. NMR* **4,** 807 (1994).

32. J.-X. Yang, A. Krezel, P. Schmieder, G. Wagner, and T. F. Havel, *J. Biomol. NMR* **4,** 827 (1994).

33. L. McIntyre and R. Freeman, *J. Magn. Reson.* **96,** 425 (1992).

34. J. Cavanagh, W. J. Fairbrother, A. G. Palmer, and N. J. Skelton, ''Protein NMR Spectroscopy: Principles and Practice,'' Academic Press, San Diego (1996).

35. C. R. I. Sanders, B. J. Hare, K. P. Howard, and J. H. Prestegard, *Prog. NMR Spectrosc.* **26,** 421 (1994).

36. B. A. Salvatore, R. Ghose, and J. H. Prestegard, *J. Am. Chem. Soc.* **118,** 4001 (1996).

37. J. A. Losonczi and J. H. Prestegard, *Biochemistry,* in press (1997).

38. C. W. Clenshaw, ''Mathematical Tables,'' Vol. 5, H. M. Stationery Office, London (1962).