Available online at www.sciencedirect.com
ScienceDirect
Journal of Magnetic Resonance 184 (2007) 302–314
JMR Journal of Magnetic Resonance
www.elsevier.com/locate/jmr

Fitting of deuterium quadrupole echo spectra with multiple motional models

Margaret A. Eastman [a,*], Mark A. Nanny [b]

[a] Department of Chemistry, Oklahoma State University, Stillwater, OK 74078-0447, USA
[b] School of Civil Engineering and Environmental Science, University of Oklahoma, Norman, OK 73019-0631, USA

Received 24 August 2006; revised 1 November 2006
Available online 28 November 2006

Abstract

Dynamic information is generally extracted from deuterium quadrupole echo spectra by matching a spectrum calculated for a particular motional model to the experimental spectrum. In this work, a set of computer programs has been written to facilitate fitting of calculated spectra to experimental spectra that represent from one to five motional models. The fitting program requires pre-calculated libraries of spectra for the models of interest, and accomplishes the fitting either by a systematic method or by simulated annealing. The systematic method is convenient for fitting with one or two motional models, but the simulated annealing method is faster for two or more models, if the libraries are made up of hundreds of spectra. The parameter Q, with the standard deviation of the spectral points estimated as the standard deviation of the baseline noise, provides a stringent measure of goodness of fit. Acceptable fits of experimental data as judged by this criterion have not been found, even in the case of ring flip motion in phenylalanine-d_5 in which the fit may be judged acceptable by eye. An example of fitting with isotropic and methyl rotation motional models of alanine-d_3, which have distinct spectral patterns, shows that it is possible to obtain reasonably accurate estimates of the relative amounts of deuterium representing the different models, even from poorly fitted spectra.

Keywords: Quadrupole; Deuterium; Fitting; Dynamics; Annealing

1. Introduction

Deuterium quadrupole echo spectra have been used to study dynamics of many systems including macromolecules [1], small molecular complexes [2], and small molecules in association with substrates such as zeolites [3], silica [4] or clay [5,6]. The shape of the quadrupole echo spectrum is generally sensitive to motions with a range of rates from 10^3 to 10^7 s^{-1} [7].

Kinetic information cannot be extracted directly from spectra. Instead, an experimental spectrum must be matched with a calculated spectrum, generated for the intermediate range of rates noted above by diagonalization of the matrix **A**, defined by

$$\mathbf{A} = i\boldsymbol{\Omega} + \mathbf{K}, \tag{1}$$

where **K** is a matrix of jump rates between sites in the motional model, and **Ω** is the diagonal matrix of NMR frequencies for each of the sites in the model [8,9]. A detailed motional model, quadrupole parameters, jump rates, and acquisition parameters must be supplied for this calculation. Calculated spectra are usually matched by eye with experimental spectra taken at various temperatures, allowing kinetic rates to be assigned to each temperature, and activation energies of the motion to be calculated [10].

In our work on dynamics of aromatic hydrocarbons in association with natural organic matter, we encounter situations in which an experimental spectrum suggests the presence of more than one motional model, as has also been observed, for example, for benzene on montmorillonite clay [5] and silica [4]. We have embarked on the work described here to facilitate and partially automate fitting

* Corresponding author. Fax: +1 405 744 6007.
E-mail address: meastman@chem.okstate.edu (M.A. Eastman).

1090-7807/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmr.2006.11.003

of calculated spectra with multiple (or single) motional models to experimental spectra. We have also attempted, with limited success, to establish a goodness of fit parameter that can be used to evaluate and compare different fits.

Aliev and Harris [11] automated fitting of quadrupole echo spectra for a single motional model using the method of simulated annealing [12], with a cost function R,

$$R = \frac{1}{n}\sum_{i=1}^{n}(y_i^{\text{exp}} - y_i^{\text{calc}})^2, \tag{2}$$

where y^{exp} is the experimental spectrum to be fit, y^{calc} is the calculated spectrum, and n in our work is the number of points in the spectral pattern. Simulated annealing originated by analogy to physical annealing of a solid, in which the solid is melted and then slowly cooled such that local high-energy configurations are avoided and the lowest energy state is reached [12,13]. In this procedure, a cost function (analogous to energy) is minimized while the value of a control parameter (analogous to temperature) is lowered. From a randomly chosen starting configuration (set of variables) a random step is made; if the step lowers the value of the cost function it is always accepted, otherwise, the step is accepted with probability

$$P = \exp(-\Delta R/c), \tag{3}$$

where c is the control parameter and ΔR is the difference between the current and previous values of the cost function. To simulate slow cooling, the control parameter is lowered in a series of steps according to some cooling schedule, and at each step of the control parameter a number of steps making random changes in the variables are carried out. The non-zero probability of accepting an uphill step makes it possible to escape from local minima, increasing the likelihood of finding the global minimum of the cost function.

The approach of Aliev and Harris [11] was to calculate spectra during the fitting process. Whether spectra can be conveniently calculated during the fitting, or are better calculated beforehand, depends upon how quickly a spectrum can be calculated, which depends upon the number of sites in the motional model. For our benzene models with 30 or 32 sites, we have chosen to first calculate large libraries of spectra, with each library representing one motional model with a particular range of kinetic rates, and then have the fitting program read in these libraries and combine them as requested to find the best fit. We have utilized simulated annealing, with the modification that the best fits found during the process are saved and used to facilitate the fitting as described in the Methods section. The calculation of libraries before the fitting also allows a brute force systematic approach in which all possible combinations of the spectra from chosen libraries are made. We compare these methods for test spectra with one to three motional models made up of combinations of spectra from calculated libraries, and also demonstrate the simulated annealing technique on four- and five- model test spectra.

While R is an excellent indicator of the best match of calculated to experimental data for a single experimental spectrum, R is not a good parameter to compare for fits of different experimental spectra, because it depends upon the noise and any other error levels in each spectrum. So, R cannot be used as a general goodness of fit parameter that provides a distinct criterion by which fits can be judged as acceptable. For this, we turn to the parameter $P\chi(\chi^2,\nu)$ [14] or Q [15], which provides a statistical measure of goodness of fit. Q is the incomplete γ-function, $Q(0.5\nu,0.5\chi^2)$ [15], where ν is the number of degrees of freedom and

$$\chi^2 = \sum_{i=1}^{n}\left\{\frac{(y_i^{\text{exp}} - y_i^{\text{calc}})}{\sigma_i}\right\}^2. \tag{4}$$

For an acceptable fit, the value of Q should be about 0.5, while a value greater than 0.001 is acceptable on occasion and very low values are indicative of a poor fit [14,15]. χ^2 resembles R but includes the standard deviation of the experimental data (σ_i). We propose to calculate σ_i from the experimental spectrum itself as the standard deviation of baseline points. We demonstrate how Q calculated with this σ_i can be used as a goodness of fit measure, although it will be likely to judge many fits of experimental data as unacceptable.

2. Computational methods

2.1. Program development

2.1.1. DMS

While the FORTRAN program MXQET [8] is available for quadrupole echo spectral calculations, for flexibility we have written a program in the C programming language, the Deuterium Motional Simulation Program (DMS). Like MXQET, DMS calculates solid-state deuterium quadrupole echo NMR spectra for samples that are not spinning, using the methodology presented by Greenfield et al. [8], and has the option of applying the approximate corrections for finite pulse width, exchange during pulses, and the virtual free induction decay (FID) described by Barbara et al. [16]. The NMR signal is calculated in the time domain as a FID, and Fourier transformed to give a spectrum. The simulation is of an echo sequence of the type

$$\pi/2 - \tau_1 - \pi/2 - \tau_2 - \text{acquire}.$$

DMS requires as input the quadrupole parameters ($\chi = e^2qQ/h$ and η), particulars of the simulated data acquisition (such as number of points, dwell time, and pulse width), and a motional model or series of motional models with kinetic parameters. In the ideal calculation the $\pi/2$ pulses are delta functions with no width, and the echo top occurs at $\tau_2 = \tau_1$. When corrections for finite pulse width are made the echo will not appear at $\tau_2 = \tau_1$, and τ_2 must be made shorter than τ_1, beginning the acquisition early so that the echo top is not missed. Then, cubic spline interpolation [8,17,18] is used to locate the echo top and

left shift so that the echo top is the first point in the FID prior to Fourier transformation.

Evaluation of the FID requires diagonalization of the complex non-Hermitian matrix **A** (Eq. 1), accomplished in DMS by incorporation of the CLAPACK routine called ZGEEV. CLAPACK is a C translation of the FORTRAN LAPACK routines for solving systems of linear equations, linear least squares problems, eigenvalue problems and singular value problems [19], available on the world wide web at http://www.netlib.org.

DMS does not have all the features of MXQET. It does not handle composite π/2 pulses, always uses equal a priori probabilities for the sites in the kinetic model, and is limited to no more than two separate angle sets corresponding to transformations between different frames that determine the details of the motion.

The calculation of the FID also requires an average of the signal over all orientations with respect to the magnetic field. The method we use has been discussed by Koons et al. [20] as the Conroy–Wolfsberg method, and by Bak and Nielsen [21] as the ZCW method. Conroy's method of evaluating multidimensional integrals [22] represents the approximate value of the integral as a sum over M values which depend upon the rational fractions p_i/M. The $\{p_i\} = \{p_1,p_2,\ldots,p_k\}$ are odd integers selected by Conroy's scheme to minimize the error in the approximated integral, where k is the number of variables in the function being integrated. Thus, we need Conroy sets of M,p_1,p_2, for $k = 2$, since we are generating two angles. A set of M Euler angles γ and β (equivalent to polar angles ϕ and θ, respectively) are generated by DMS using the formulae reported by Bak and Nielsen [21] for the Spherical ZCW method,

$$\gamma = 2\pi \bmod f\left[\frac{p_1(i-1)}{M}\right]$$

and

$$\beta = \cos^{-1}\left(1 - 2 \bmod f\left[\frac{p_2(i-1)}{M}\right]\right),\ i = 1\ldots M, \qquad (5)$$

where mod$f$ is a function that returns the fractional part of the argument. The method is called "spherical" because the angles form a spherical set, meaning that they each subtend the same amount of solid angle and so do not require any weighting function to be applied in the averaging. A separate program has been written for the generation of Conroy sets with lowest possible error that searches for p_1 and p_2 given M, either exhaustively by trying all possibilities, or randomly if M is large. For each $\{M,p_1,p_2\}$ the integral of the test function f_2 is calculated as Conroy's sum [22], and the error is evaluated as the difference between this sum and the known value of the integral calculated analytically. Attempts to further optimize the Conroy sets specifically for the deuterium spectrum by evaluating a static deuterium spectrum using the best Conroy sets output for a particular M value resulted in the conclusion that all good Conroy sets determined from the evaluation of f_2 were equally good in the spectral calculations.

The DMS program allows all sites within a frame to exchange, or allows the kinetic matrix to be entered explicitly, as a list of rates in units of s^{-1} between pairs of sites in the model. Entering a kinetic matrix in detail allows nearest neighbor (NN) exchange or any other selection of exchange possibilities. If all sites exchange, two possibilities may be chosen: equal rates (in units of s^{-1}) (AEE) or unequal rates (AEU), in which all sites exchange with the same rate in units of rad/s, but individual jumps have different rates in units of s^{-1}, depending upon the jump angle between sites. The jump angle is defined as the angle between the z axes of the quadrupole principle axis systems, meaning the angle between the C–D bonds in the two sites. In the AEE and AEU exchange patterns, in the case of two motional frames, exchange does not occur between sites that have different sub-sites in both of the frames, i.e., a jump is made only within one frame.

2.1.2. LCP

Since we have chosen to calculate libraries of spectra prior to fitting, the Library Creation Program (LCP) has been written to facilitate this. LCP produces files describing the library and listing the jump rates, and an input file for either DMS or MXQET. Running the appropriate spectral calculation program with this input causes all the spectra for the library to be calculated. A library may contain spectra for more than one π/2 pulse width and more than one τ value, but these parameters are not varied in the fitting. The kinetic rates and in some cases a single angle in the Euler angle sets defining the motional model are varied in a library, and are then the variables in the fit. Kinetic rates are varied between an input minimum and maximum in logarithmic steps, while an angle, if varied, is incremented in linear steps. Allowed DMS kinetic modes are AEE and AEU, and a limited case of explicit entry of the rate matrix in which all rates for the single motional frame are identical. For MXQET "all sites exchange" and "hard collision = Y" are always used. Thus the kinetic file only has to store one or two lists of kinetic rates, rather than a complete rate matrix for each spectrum in the library.

The rationale behind the choice of which parameters to treat as fixed and variable in the fitting is as follows. It is assumed that the quadrupole parameters are known, and that the acquisition and processing parameters are determined by the experimental setup and data processing. Therefore, these parameters are not variables in the fitting. It is convenient to construct a library with several π/2 pulse widths and τ values, so that otherwise similar experimental spectra collected with differing values of these parameters may be fit using the same library. In this case, the fitting program described below will find and use only those spectra from a library with pulse width and delay that match those of the experimental spectrum, ignoring the others. Once the motional model is chosen, the main parameters to be identified by a fit are the rates of jumping between sites, so these rates are the main variables in the fitting; every rate associated with the model is always a library

variable. We also find that for some more complicated models (such as small- and large-angle wobble, see below) a certain Euler angle in the model may take on any of several values, making it desirable to search for the best value in fitting. This choice of variables is conservative, as it does not allow parameters such as the $\pi/2$ pulse width and line broadening, which can legitimately only be made equal to the experimental values, but in practice might be varied to give a better looking fit, to vary in the fitting. Of course, any parameter may be varied in the search for a good fit by constructing several libraries with different values of the parameter.

2.1.3. DFP

The Deuterium Fitting Program (DFP) takes in experimental spectra and calculated spectral libraries, each representing a different motional model, and attempts to find the best (lowest R (Eq. 2)) combination of one spectrum from each of the requested libraries to fit to each experimental spectrum. A best-fit spectrum, y^{calc}, is a sum:

$$y^{\mathrm{calc}} = \sum_i a_i S_{i,j(i)}, \tag{6}$$

where the a_i are coefficients and the $S_{i,j(i)}$ are library spectra. The spectrum $S_{i,j(i)}$ from the ith library will have a certain index, $j(i)$, identifying its position in the library. In the systematic method all possible combinations of one spectrum from each library are tried, and least squares fitting (employing the DGELS routine from CLAPACK) is used to determine the coefficients multiplying the spectra. This method is limited to fits that require no more than three libraries. In the simulated annealing (SA) method a block of a fixed number of simulated annealing steps starting with a random configuration is carried out and the best fits found during this process are saved. Blocks of SA steps are iterated up to a fixed number of times entered by the user, unless the entered stop criterion for R is met, in which case the calculation terminates. At the end of the calculation, the overall best fit found, with lowest R, is reported. In addition to reporting the single best fit found, DFP can provide an average over a group of saved fits that have values of R no greater than a designated factor times the lowest R found.

2.1.4. Simulated annealing

Several choices must be made for simulated annealing, resulting in many possible ways of implementing the method. Using spectral libraries calculated prior to the fitting calculation, each library used in the fit brings two variables, the spectral index in the library and the coefficient multiplying the spectrum. At first, we attempted simulated annealing varying all of these variables at each step. Reasonable ranges for the variables must be provided, and the more closely these can be specified, the more quickly an acceptably low value of R can be attained, because of the random nature of the search. Since it is difficult to know beforehand the ranges that are reasonable for the coefficient variables, we calculated ranges by taking a large number of random spectral combinations and obtaining the coefficients with least squares fitting; the minimum and maximum coefficients defined the starting range. This method worked well for fits requiring two libraries of the size described here, but failed when three libraries were needed, apparently because as the number of variables increases wide starting ranges make location of the correct fit more difficult. With this consideration, we have chosen to vary only the spectral index variables in simulated annealing, and to find the coefficients at each step with least squares fitting. This removes the need to specify any initial range for the coefficients.

A block of simulated annealing (SA) steps as mentioned above consists of 100 steps over which the control parameter is lowered, at each of which 1000 steps are taken in which a random change is made in each of the spectral index variables. These numbers of steps are fixed for an executable of DFP, but could be changed by recompiling the program. The initial control parameter, c_0, is selected using an acceptance ratio [13] of 0.95. A routine for determination of c_0 tries several values of control parameter, taking 1000 SA steps and determining empirically the acceptance ratio as the percent of accepted steps. When two control parameter values are found with acceptance ratios bracketing the desired value, the average of these two is taken as c_0. Within a block of SA steps, the control parameter is reduced with a simple cooling schedule, according to

$$c_n = \alpha c_{n-1}, \tag{7}$$

where c_n is the control parameter at the nth control-parameter step after the first, in which c_0 is used [12,13]. The cooling schedule parameter α is set to 0.95.

At each step, a random change is made in each spectral index variable, by taking a random number from 0 to 1, multiplying this by the total allowed range for the variable divided by two, and multiplying again by a variable factor f, such that the maximum possible step is one half the size of the range. Bounds checking is done, so that a step that puts the variable outside the initial allowed range is not permitted; instead the step takes the variable to its minimum or maximum value. The variable factor f starts out as one at the beginning of each SA block, and is reduced at each control parameter step in the same manner as the control parameter, by multiplication by α. Thus, this factor acts to reduce the size of steps taken as the control parameter is reduced. An additional critical aspect of the random step in our algorithm is that its size must be permitted to be no smaller than one, such that each step moves at least to the next neighboring spectrum in the library. At each step the new configuration found by the random change in variables is either accepted or rejected, according to the usual rules for SA given in the Introduction section, using R as the cost function [11].

The program tests at each accepted step whether the value of R of the new configuration is lower than the largest

value of R of a saved group of best configurations, and saves the new configuration in this group if it is. It also tests whether the criterion for R entered by the user has been met, and terminates the calculation if it has. At the end of each control parameter step, before moving on to the next lower value of the control parameter, we have added a feature not usual to SA, which is that the configuration with the lowest R found in the array of best configurations saved so far is used as the starting configuration for the next control parameter step. This action seems to significantly improve the speed of achieving a best fit.

Some fits require only one block of SA steps, but if more than one block is required, restricting the range over which steps can be taken is critical to rapidly converging to a correct fit. To do this, two separate arrays of configurations with best R values are saved, one of which is initialized at the start of each block of SA steps (local best array), and one of which is saved over the entire calculation (overall best array). The local best array is updated at each accepted step and utilized to find the starting configuration for each control parameter step as noted above. At the end of each SA block, any entry in the local best array with lower R than the entry of highest R in the overall best array replaces that highest-R entry in the overall best array. The overall best array is then searched to find the minimum and maximum values for each spectral index variable and 5% of the current allowed range is added to the maximum and subtracted from the minimum to give the new allowed range. This new range is used only for the purposes of selecting a random starting configuration and making random changes in the variables in the next SA block. The original total range is always used for bounds checking, so that if the new restricted range does not contain the best fit, stepping outside that range toward the best fit is possible.

2.1.5. Spectral centering and baseline adjustment in DFP

Because the fitted and experimental spectra are compared point by point, a library must have the same total number of points, N, as the experimental spectra. Also, it is essential that both experimental and calculated spectra be centered at the same frequency point, and that the baseline be positioned at an amplitude of zero. Before use in fitting, all library spectra are adjusted to move the baseline close to zero by averaging those 5% of the total N points on each end of the spectrum and subtracting this average from all points. The library spectra are then collectively normalized so that the highest peak in all of the spectra has amplitude one. As a group, all experimental spectra are also subjected to this baseline adjustment and normalization. Collective normalization preserves the relative heights of the spectra.

R, χ^2, and Q are calculated only for points in the spectral pattern (n in number), not baseline points, requiring that points belonging to the baseline be identified. The centering algorithm also uses just the spectral pattern points. The number of baseline points is determined by moving from the outer edge of the spectrum to the center on both ends, attempting to find the points where the amplitude increases significantly from the level of the baseline.

The centering algorithm finds all the peaks in the experimental spectral pattern with height at least 50% of the largest peak height using finite difference derivatives to locate maxima. It then constructs a center of mass of the pattern, assigning a height of one to each peak to eliminate effects of asymmetry, and shifts the center of mass to the point $N/2$. The code corrects the center of mass in a few cases were the above simple algorithm gets into difficulty, but it depends upon the spectrum not being too asymmetrical, or too far off-center initially.

2.1.6. Goodness of fit and error estimates in DFP

The calculation of χ^2 and Q requires a standard deviation of each point in the spectral pattern. We propose the following form for σ_i, the standard deviation for spectral point i:

$$\sigma_i = \sigma_N, \tag{8}$$

where the noise standard deviation, σ_N, is calculated for each spectrum as the standard deviation of the amplitude of all baseline points. Completion of the calculation of Q requires the quantity v, equal to $n-M'$, where M' is the number of parameters in the fit. Since both the index of the spectrum from each library and each coefficient are varied in the fitting procedure, this total number of parameters is taken as M'. If the libraries were not used, the rates and angles would have to be varied independently in the fit, so they would contribute to M'.

For the overall best-fit spectrum, errors in the coefficients a_i Eq. (6) are determined from an error bound calculation routine provided in the LAPACK Users' Guide [19], using an input error of twice the noise standard deviation σ_N in place of the machine precision. Percentages of each motional model are calculated from the coefficients, and the error in the percentages is obtained from the coefficient errors by propagation [14]. Errors in jump rate(s) and the value of the varied angle are estimated as the largest difference between the particular value and the adjacent value for the next higher or lower spectrum in the library. For an average fit spectrum, the standard deviation of each rate or angle parameter serves as the error estimate. The standard deviation of the coefficients is compared to the average of the coefficient errors determined for the individual fits with the error bound routine, and the larger of the two is taken as the error. Errors in percentages of each motional model are again propagated from the errors in the average coefficients.

2.2. Motional models and calculations

Four basic motional models have been used in the calculations presented here. Isotropic motion (ISO) is modeled with 32 sites in which the C–D bond is oriented along an axis from the center to each one of the 20 face centers

and 12 vertices of an icosahedron [5,6]. This model was originally used with the AEE kinetic mode [6], but we also try AEU and NN modes. Methyl rotation (MeRot) is modeled with three sites with angle of 70.5° between the C–D bond and the threefold rotation axis [10]. Small-angle and large-angle wobble (SAW, LAW) for benzene are models with two frames, one for 6-fold rotation with an angle of 90° between the C–D bond and the rotation axis, and one for wobbling of the rotation axis in a cone [5,8]. Following Xiong et al. [5] we use a 30-site model with 10 wobble sites and three sites representing the sixfold rotation, but we implement the AEU kinetic mode instead of AEE. Phenyl ring flip (RF) is modeled with two sites with Euler angles $(0, 60, 0)$ and $(180, 60, 0)$ [8].

Fitting calculations were done on a computer with dual Intel Xeon 2.4 GHz processors and 1 GB of random-access memory, running the Linux operating system.

3. Results

Table 1 summarizes the libraries of calculated spectra used here, including libraries for benzene-d_6 small-angle wobble (SAW), large-angle wobble (LAW), and isotropic (ISO), alanine-d_3 methyl rotation (MeRot) and ISO, and phenylalanine-d_5 180° ring flip (RF) motional models. For fits of the ISO motional model to aqueous alanine-d_3 spectra, we found that kinetic modes in which all sites exchange with unequal rates (AEU), all sites exchange with equal rates (AEE), and only nearest neighbor sites exchange (NN) were indistinguishable (not shown). We selected the AEU mode for both ISO and SAW models in all fitting shown here.

Test spectra for demonstrating the systematic and simulated annealing (SA) methods were made by taking spectra from the benzene-d_6 libraries, multiplying them by different factors, and combining them. For all these multiple-model spectra, separate features representing each motional model were visually identifiable in the spectral shape. These test spectra were fit with DFP using the libraries from which they originated to compare the efficacy and time taken for each method, using one- to five- library fits. Fig. 1 shows one of these spectra, which were all correctly fitted with most values of R less than 1.0×10^{-30} and none greater than 1.0×10^{-14}. Fig. 2 shows the relative time taken for each fit plotted against the number of least squares fits done, which is the total number of possible combinations of spectra from the libraries used for systematic fitting, or the total number of steps taken in simulated annealing. For libraries of the size used here (891 spectra for SAW and LAW models included in all fits, and 50 spectra for the ISO model included in fits with three or more libraries) the systematic method was faster only for a one-library fit. The simulated annealing method was faster for two-library fits, and SA fits requiring three to five libraries were all completed substantially more efficiently than the single three-library systematic fit. The number of iterations of simulated annealing (number of blocks of SA steps) was 1–2 (6–26, 33) for the three- (four-, five-) library fits.

A few tests were done to explore the use of Q as a goodness of fit parameter. Table 2 and Fig. 3 show fits of artificial spectra made by combining an alanine MeRot library spectrum and real experimental noise. In this case, σ_N is the only source of error in the spectrum. Q values are within the acceptable range (>0.001 and <1.0). The rates given by the fits are not identical to the test spectrum rate of $1.009 \times 10^8\ s^{-1}$, indicating that library spectra with rates in the range 0.93×10^8 to $1.5 \times 10^8\ s^{-1}$ are very close in appearance and essentially indistinguishable. Fig. 4 shows fits with the 180° ring flip model of an experimental spectrum of phenylalanine-d_5. Systematic fits of this data with χ values in the range of 168–170 kHz and η values from 0.03 to 0.04 are acceptable as judged by eye, while Q is zero for all fits. Flip rates for the best fits indicate that the rate is not well defined but is greater than 1×10^7 (Table 3).

Spectra of combinations of two alanine-d_3 samples (solid and aqueous) held in separate containers were taken to compare percentages of different motional models, calculated from the coefficients multiplying library spectra, to the known percentages. Spectra of the individual samples showed only a pattern consistent with ISO motion in the aqueous alanine-d_3, while spectra of the solid appeared to represent MeRot motion only. Spectra of the combined samples were fit with ISO and MeRot libraries with quadrupole coupling constant, χ, of 167.0 kHz [10] and the correct experimental pulse with of 1.9 μs in two-library systematic fits. However, the best-fit spectrum did not match the experimental pattern well at the shoulders and horns (Fig. 5a). To remedy this, a fit with χ of 164.0 kHz and pulse width of 3.0 μs was tried; this too did not provide a good fit as judged by eye (Fig. 5b). Better matching of the somewhat rounded horns and sloping shoulders of the experimental pattern was obtained by combining two or three MeRot libraries with different χ values in three- and four-library SA fits. One spectrum from the three-library fitting is shown in Fig. 5c for comparison to the other models used, and Fig. 6 shows further examples of superimposed experimental and three-library fitted spectra. All of the experimental spectra are poorly matched by calculated three-library spectra as judged by the parameter Q, which is zero. Fig. 7 shows a graph of the percent of each motional model from the overall best fit three-library spectra plotted against the known percent. Percentages from DFP differ by 0.1 to 4.1 (0.5 to 9.4) percentage points from the known percentages for the spectra with τ of 40 μs (100 μs). Errors provided by DFP vary from 5 to 37 percentage points for MeRot and 3 to 21 percentage points for ISO. The form of this graph was similar for the other fits tried; even fits that were poorer as judged by eye gave similar accuracy. The trend toward overestimation of the ISO component as the percent of ISO increases is more pronounced in the data from spectra with τ of 100 μs, indicating the shorter τ value of 40 μs, which gives spectra that

Table 1
Parameters of spectral libraries

Library name	Model	Kinetic mode	No. Rates	No. Angles	Minimum rate	Maximum rate	Minimum angle	Maximum angle	No. points	90° pw (μs)	LB (Hz)	χ (kHz)	η	No. PA angles
Benzene-d_6 Libraries, τ = 20, 40 μs														
BnzSAW	SAW	AEU	9 wobble/9 C_6	11	5E3 wobble, 1E6 C_6	5E5 wobble, 1E8 C_6	4°	24°	1024	2.25	318	183.0	0.04	1154
BnzLAW1	LAW	AEU	9 wobble/9 C_6	11	5E3 wobble, 1E6 C_6	1E6 wobble, 1E8 C_6	25°	45°	1024	2.25	318	183.0	0.04	1154
BnzLAW2	LAW	AEU	9 wobble/9 C_6	11	5E3 wobble, 1E6 C_6	1E6 wobble, 1E8 C_6	46°	66°	1024	2.25	318	183.0	0.04	1154
BnzISO	ISO	AEU	50	11	1.0E+05	5.0E+08	–	–	1024	2.25	318	183.0	0.0	1154
Alanine-d_3 Libraries, τ = 40, 100 μs														
AlaMeRot	MeRot	AEE	150	–	2.0E+06	1.0E+09	–	–	1024	1.9	1000	167.0	0.0	9644
AlaMeRot167	MeRot	AEE	150	–	2.0E+06	1.0E+09	–	–	8192	1.9	1000	167.0	0.0	9644
AlaMeRot160	MeRot	AEE	10	–	2.0E+06	1.0E+09	–	–	8192	1.9	1000	160.0	0.0	9644
AlaMeRot164	MeRot	AEE	10	–	1.0E+07	1.0E+11	–	–	8192	3.0	1000	164.0	0.0	9644
AlaISO167	ISO	AEU	100	–	5.0E+06	1.0E+09	–	–	8192	1.9	1000	167.0	0.0	1154
AlaISO164	ISO	AEU	10	–	1.0E+07	1.0E+11	–	–	8192	3.0	1000	164.0	0.0	1154
Phenylalanine-d_5 Libraries, τ = 40 μs														
PheRF180	RF	AEE	200	–	1.0E+05	5.0E+09	–	–	512	1.8	500	180.0	0.05	100004
PheRF170	RF	AEE	20	–	1.0E+07	1.0E+11	–	–	512	1.8	500	170.0	0.05	100004
PheRF168	RF	AEE	20	–	1.0E+07	1.0E+10	–	–	512	1.8	500	168.1	0.036	100004

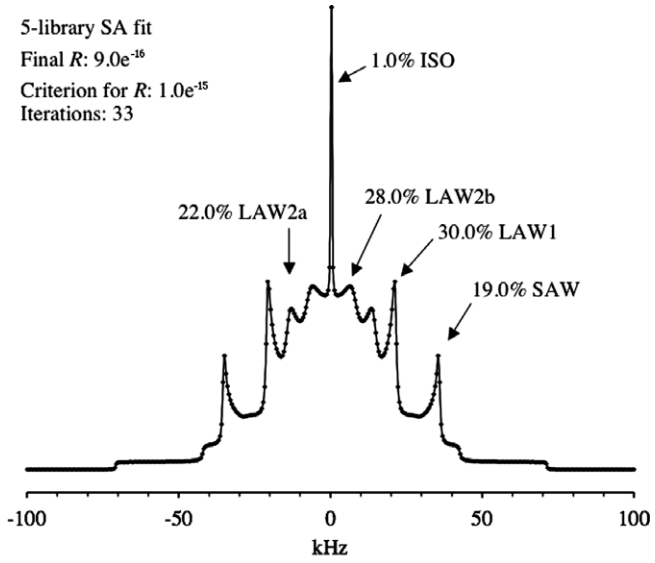

Spectrum	Library	Wobble Angle	Wobble Rate (rad/s)	C_6 Rate (rad/s)
SAW	BnzSAW	4°	5.0 X 10^3	1.0 X 10^8
LAW1	BnzLAW1	39°	5.2 X 10^5	1.0 X 10^8
LAW2a	BnzLAW2	50°	2.7 X 10^5	1.0 X 10^8
LAW2b	BnzLAW2	60°	1.0 X 10^6	1.0 X 10^7
ISO	BnzISO	ISO Rate = 8.8 X 10^7		

Fig. 1. Five-library simulated annealing (SA) fit of a test spectrum made up from combinations of the spectra from the benzene-d_6 libraries of Table 1. Solid line is the calculated spectrum, and symbols are the test spectrum. Both fitted and test spectra have the composition written on the plot. The R value for the fit, stop criterion and number of iterations actually taken are listed.

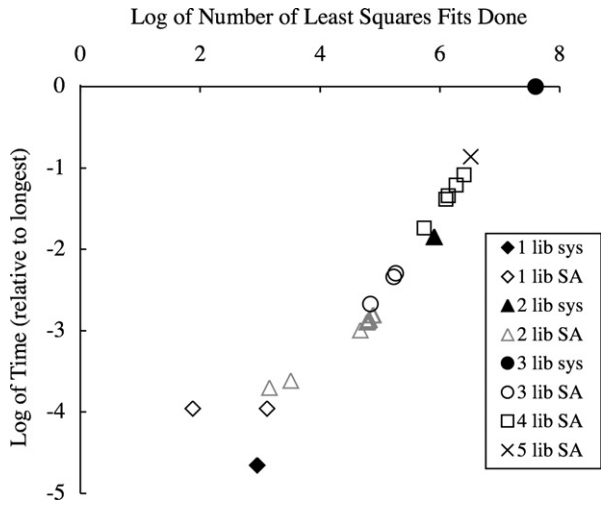

Fig. 2. Plot of log of time taken for each fit of test spectra made up from library spectra vs. log of the number of least squares fits done during the fitting calculation. Times are relative to the longest time taken, which was 12.47 h for the three-library systematic fit. Fits using the same number of libraries are labeled with the same symbol, with the filled symbol representing the systematic fit (sys) and the open symbol representing the simulated annealing fit (SA). The libraries used are for benzene (Bnz) undergoing small-angle wobble (SAW), large-angle wobble (LAW), and isotropic (ISO) motions (Table 1). Included are two one-library fits of spectra from BnzSAW and BnzLAW1, 11 two-library fits combining spectra from BnzSAW and BnzLAW1 in increments of 10% from 0 to 100%, and three three-library fits, five four-library fits and one five-library fit (shown in Fig. 1), all combining spectra from BnzSAW, BnzISO, and one or more of the LAW libraries. The same fits were done with both methods, except that only one three-library combination was fit with the systematic method, and no four- and five-library systematic fits were done.

are better fit as judged by eye, provides more accurate percentages of the components.

4. Discussion

Matching of calculated deuterium quadrupole echo spectra to experimental spectra is an effort that cannot be fully automated, because judgment must always be applied in interpreting the experimental spectrum to select the appropriate motional model(s). Regardless of the method used for fitting, all physically reasonable alternative models must be considered, and reasonable ranges of parameters for those models must be explored in the search for a best fit. We have demonstrated the success of the group of programs presented here as a tool to facilitate this search when spectra represent from one to five motional models. The Deuterium Fitting Program (DFP) group of programs can be obtained from http://chem.okstate.edu/~dfp, along with documentation and sample inputs.

Our examples of fitting have employed only spectra representing fast motions. The appropriate models have rates of 1×10^7 or greater, and the rates are not well defined, because quadrupole echo spectra are not expected to distinguish rates this fast. Although we have not demonstrated it by example, the DFP group of programs is designed to extract rates from spectral fits, since the rate(s) (and possibly also an angle in the model) are what is varied in a library. The averaging feature may be very useful for arriving at an error estimate for rates, by allowing fits that are close enough to be considered indistinguishable to be

Table 2
Parameters for fits to ideal spectra with experimental noise added

Spectrum	Noise	σ_N	ν	Coef.	Rate (s^{-1})	$R \times 10^5$	χ^2	Q
Fig. 3a	Air	0.004	196	1.07 ± 0.02	9.3×10^7	1.5	199	0.4
Fig. 3b	Glass	0.008	188	1.06 ± 0.03	1.2×10^8	7.6	243	0.004
Fig. 3c	Teflon	0.014	180	1.03 ± 0.06	1.5×10^8	19.8	175	0.6

Parameters for fits to ideal noisy spectra made up from AlaMeRot (alanine methyl rotation) spectrum with rate of 1.009×10^8 s^{-1} and noise from three experimental spectra. Spectra are plotted in Fig. 3. Particulars of the AlaMeRot library are in Table 1.

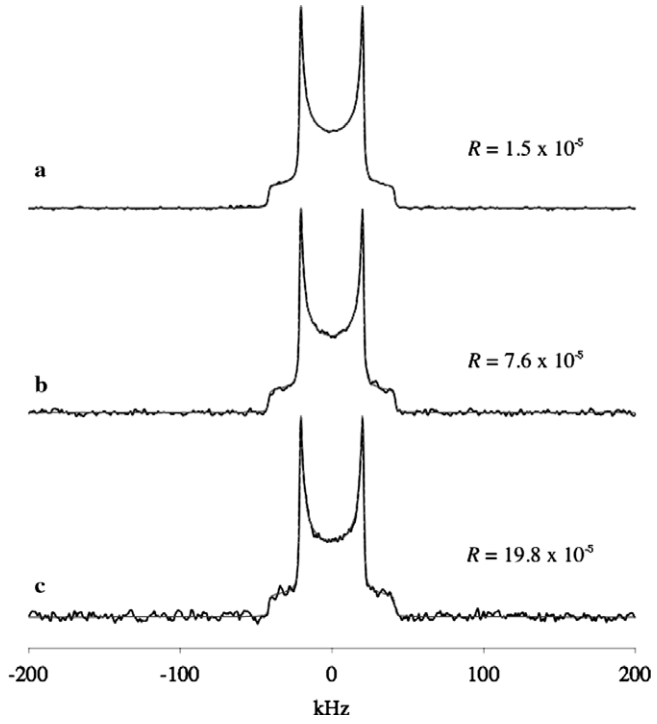

Fig. 3. Fit spectra (dashed gray lines) and artificial experimental spectra (solid black lines) made up of a AlaMeRot (alanine methyl rotation, Table 1) library spectrum with jump rate of $1.009 \times 10^8\ s^{-1}$ and experimental noise spectra obtained for (a) no sample, (b) glass, and (c) Teflon tape. Noise spectra were scaled by different amounts to produce differing signal-to-noise ratios. Fits are systematic, using the AlaMeRot library. Table 2 lists parameters of the spectra and fits.

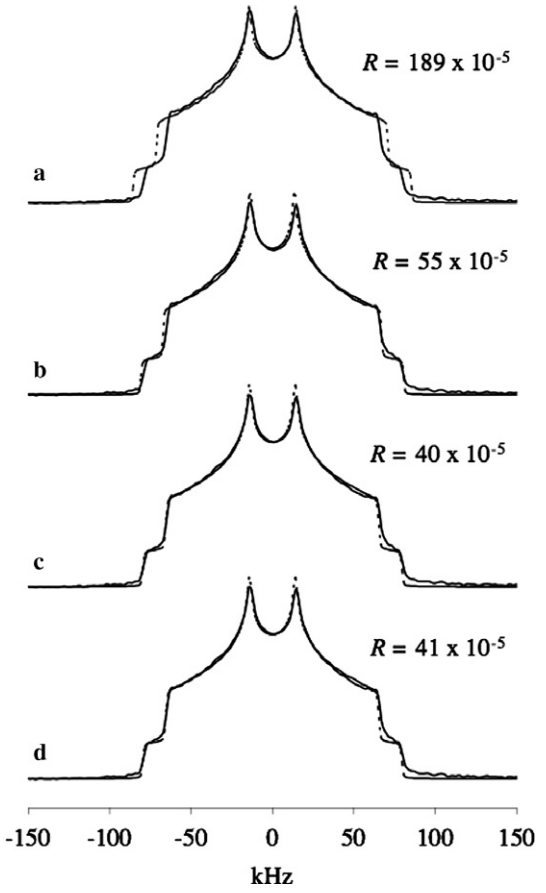

Fig. 4. Fitted 180° ring flip (RF) model spectra (dashed lines) and experimental phenylalanine-d_5 spectrum (solid lines). Libraries used were (a) PheRF180, (b) PheRF170, (c and d) PheRF168. All are best overall fits except d, which is an average of 20 spectra with R no greater than 10 times the lowest R. Values of R are written on the plots, and other parameters are in Table 3.

averaged together. Some trial and error is required to determine the step size for rates in the library, and the criterion for which fits to average.

Both the simulated annealing and systematic methods are capable of providing correct results for test spectra whose composition is exactly known. These fits of artificial spectra were essential to develop the fitting program and verify the efficacy of the methods. For two-library fits, the systematic approach is most convenient even though it is slower, because it does not require the input of a limit on the number of iterations and the stop criterion for R. For fits with three or more motional models with large numbers of library spectra, simulated annealing is the best choice due its greater speed. Fits of test spectra may be done to determine the appropriate number of iterations, which can be overestimated to avoid stopping the fit too early. Our experience is that values of R are rarely below 1×10^{-7} for experimental spectra, so a stop criterion equal to or below this will often be a good choice.

We have shown that Q can in principle be an effective goodness of fit parameter. Q values show that the fits of ideal spectra with added noise (Table 2 and Fig. 3) are acceptable, as expected because the added noise is known to be the only contribution to the error in the spectral amplitude. Values of Q in the acceptable range were not found for fits of the ring flip motional model to real experimental phenylalanine-d_5 data (Fig. 4), even though the fits of Figs. 4c and d might be judged as acceptable to the eye and the quadrupole parameters for these fits correspond to those found by Gall et al. [23] for the short T_1 component of phenylalanine-d_5. Q provides a very stringent criterion for an acceptable fit. Hence when experimental spectra deviate in a manner easily visible to the eye from the calculated spectra, Q is zero, indicating a poor fit. This problem, evidenced in Figs. 4–6, may be due in part to a lack of the proper estimate of experimental errors that go into the calculation of σ_i. Perhaps a larger σ_i value that includes all sources of random experimental error could be obtained by taking the standard deviation averaged over only the spectral points for several spectra of the same sample acquired at different times. We have experimented with this (results not shown), but rejected this procedure because it is inconvenient to take multiple spectra of every sample and can be very expensive in experiment time when spectra must be acquired for long times. The failure to find fits to experimental spectra with acceptable Q may also be due to experimental artifacts not entirely known and not taken into account in the quadrupole echo spectral

Table 3
Fits of ring flip model to phenylalanine-d_5 spectrum

Spectrum	Library	e^2qQ/h	η	Rate $\times 10^{-9}$ (s^{-1})	$R \times 10^5$	χ^2	Q
Fig. 4a	PheRF180	180	0.05	5.0 ± 0.3	189	30410	0
Fig. 4b	PheRF170	170	0.05	0.11 ± 0.07	55	8851	0
Fig. 4c	PheRF168	168.1	0.036	0.06 ± 0.03	40	6522	0
Fig. 4d	PheRF168	168.1	0.036	2 ± 3	41	6575	0

Parameters for fits shown in Fig. 4 of phenylalanine-d_5 spectrum with $\sigma_N = 0.004$ and $v = 290$. Particulars of the 180° ring flip (RF) libraries are in Table 1.

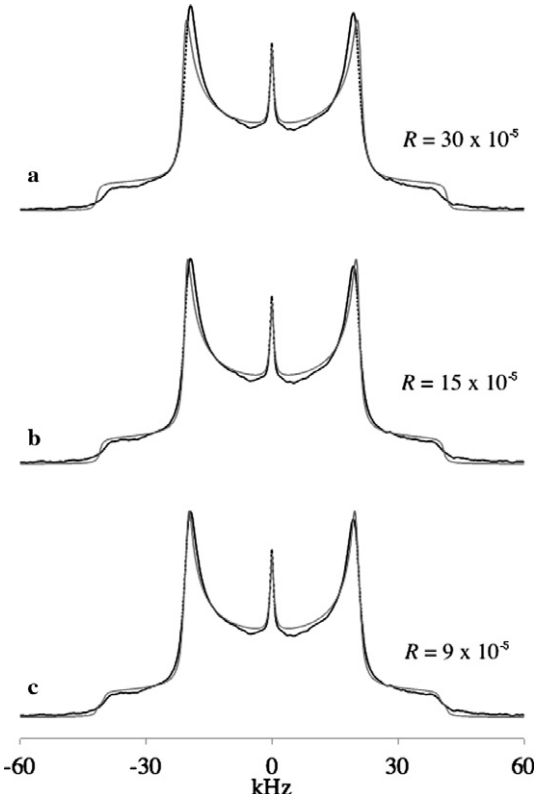

Fig. 5. Fitted spectra (gray lines) and experimental combined solid and aqueous ($1.7 \pm 0.1\%$) alanine-d_3 spectrum (black lines) with $\tau = 40$ μs. Fit spectra are best overall from (a) systematic fit with libraries AlaMeRot167 and AlaISO167, $\chi = 167$ kHz, 90° pulse width = 1.9 μs, (b) systematic fit with libraries AlaMeRot164 and AlaISO164, $\chi = 164$ kHz, 90° pulse width = 3.0 μs, and (c) simulated annealing fit with libraries AlaMeRot167, AlaMeRot160, and AlaISO167. The alanine (Ala) methyl rotation (MeRot) libraries are described in Table 1.

calculation. One well-known artifact is the asymmetry that arises from the long rise and fall times of the radiofrequency pulses relative to their length and cannot be reliably avoided, as described by Vold [24]. Spectra can be artificially made symmetric by setting the imaginary part to zero before Fourier transformation or by averaging corresponding points in the two halves of the spectrum, but it is our experience with the phenylalanine-d_5 data that these procedures do not lead to an acceptable Q value, probably because the true spectrum, which would be obtained with perfect square pulses, is not well represented by the artificially symmetric spectrum. Another source of the problem of unacceptable fits is approximations in the spectral calculation method that lead to slight errors in the calculated spectrum. Also, arriving at an exactly correct model may be difficult even if an individual calculated spectrum is correct, for example, if a range of quadrupole parameters is appropriate. Any of these latter difficulties will make the calculated spectral shape a poor match to the experimental, and will result in low Q even if the estimate of the standard deviation is adequate. In cases where experimental imperfections or limitations of the theoretical calculations are not addressed, or if the estimate of σ_i cannot be amended to include any other error sources, Q will not be of practical use.

Since Q as implemented here has limited usefulness, it may be tempting to compare values of R. However, as stated in the Introduction, it is not valid to compare R values for fits of different spectra, such as those shown in Figs. 3 and 6. R values may be compared for different fits of a single spectrum, as in Figs. 4 and 5, and they show the expected trend of lower R as the fit becomes better as judged by eye. However, R is not a substitute for Q, because we do not know what value of R constitutes an acceptable fit for any given spectrum. The best that can be done is to search for fits with smaller and smaller values of R.

Fig. 7 shows that even when fits are poor, as judged both by the value of Q and by eye, quantitative percentages of motional models can be extracted, with an accuracy of less than five percentage points for the example given. This is encouraging to the extraction of quantitative parameters in cases where the imperfections and limitations mentioned above cannot be readily overcome. A favorable outcome may have resulted here because the models can be well distinguished by eye: two models with very similar powder pattern shapes that both give poor fits might cause the percentages to have greater errors. Perhaps variation of the angle between the C–D bond and the threefold rotation axis in the methyl rotation library, or the use of a different motional model in which the C–D bond wobbles, would make more sense physically than our choice of multiple χ values, and also lead to a better fit, resulting in better accuracy and precision than shown in Fig. 7. The use of a liquid sample, certainly inappropriate for quadrupole echo spectroscopy, needs to be explained. It was chosen in part as a convenient sample for combined spectra of known composition, and in part because the isotropic (ISO) motional model is of interest to us. The verification of percentages obtained from fitting accomplished with the combined spectra is a useful exercise because in some applications

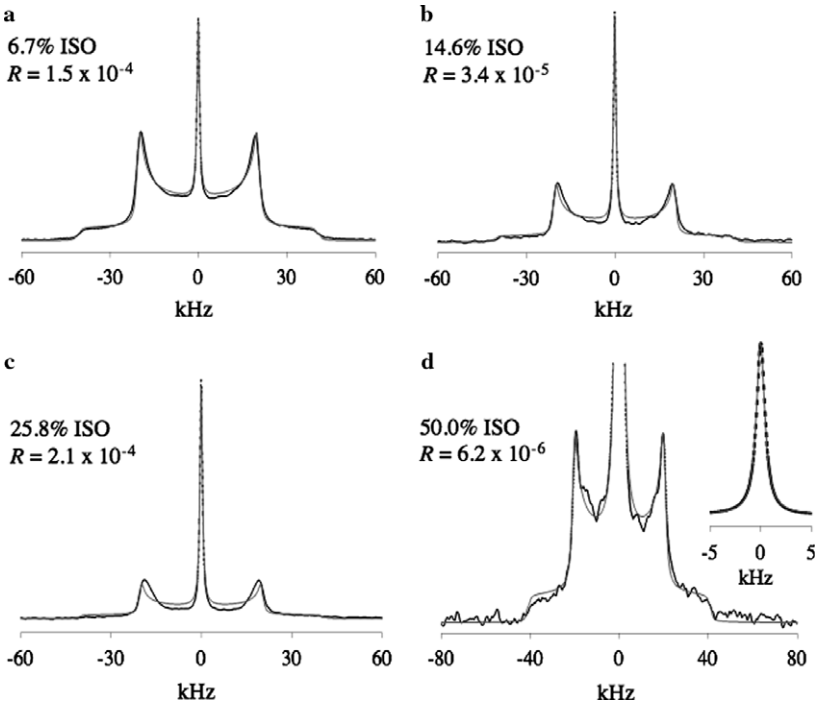

Fig. 6. Examples of experimental combined solid and aqueous alanine-d_3 spectra obtained with 40 μs τ (black lines) and overall best fitted spectra from simulated annealing (SA) fit with libraries AlaMeRot167, AlaMeRot160, and AlaISO167 (gray lines). Values of R ranged from 6.2×10^{-6} to 2.1×10^{-4}. Methyl rotation (MeRot) rates ranged from 1.8×10^{7} to 1×10^{9} s^{-1} and isotropic (ISO) rates ranged from 1.7×10^{7} to 1×10^{9} rad/s. The ratio of the contribution from AlaMeRot167 to that from AlaMeRot160 was also inconsistent between fits and ranged from 0.35 to 1.63. Spectra are of (a) 6.7% aqueous sample, (b) 14.6% aqueous sample, (c) 25.8% aqueous sample, and (d) 50.0% aqueous sample, scaled to show up to show details of the MeRot pattern and expanded to show ISO pattern only.

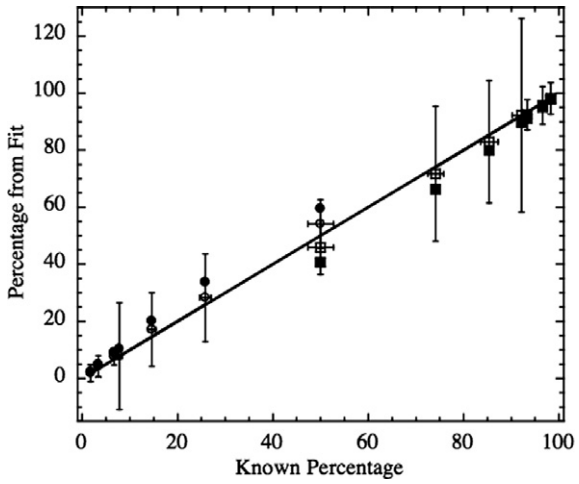

Fig. 7. Plot of percentages of alanine-d_3 isotropic (ISO) and methyl rotation (MeRot) motional models obtained from spectral fitting against the percentages known from sample composition. Square symbols represent MeRot percentage and circles represent ISO percentage. Open (closed) symbols are from fits of spectra with 40 μs (100 μs) τ. Error bars are shown only for the data with 40 μs τ for clarity. Fits were done with simulated annealing using two libraries for MeRot with χ of 160 and 167 kHz, and one library for ISO with χ of 167 kHz. Examples of fit spectra are shown in Figs. 5c and 6.

these percentages can be the basis for calculating enthalpy changes, for example if an ISO motion can be regarded as representing free ligand and another motional pattern associated with bound ligand [5]. The rate determined for ISO motion in aqueous alanine-d_3 is not necessarily correct, as the motion may be faster than can be detected through line shape analysis.

Our results suggest a situation in which our programs, like any fitting method, must be applied with particular care. Only motional spectral patterns that can be visually identified in the experimental spectrum, or can reasonably be expected to be present from other evidence, can be entered into the fit. For example, if features representing each component are not visually identifiable in a multiple-model spectrum, unlike the case in Figs. 1 and 2, DFP may have trouble characterizing a hidden component in experimental spectra. Similarly, a small amount of a component whose spectra have relatively small amplitude may cause the same difficulty. Notice that as percent of methyl rotation (MeRot) decreases in Fig. 7, the inaccuracy in percent from the fit increases. At low percent of MeRot, the percentage is expected to be very poorly determined. This may be accounted for in part by the poor quality of the fit of the MeRot pattern, but is also expected because the ISO spectral pattern is so much higher in amplitude than the MeRot pattern. Hence, as the percentage of ISO increases, the MeRot pattern becomes less distinguishable from the noise. For perfectly shaped, well-fit spectra with essentially no noise, this problem should not be seen. However, for experimental spectra with commonly obtainable noise levels only certain percentages of different

models in combination may be reliably extractable from fits.

Future program development might include incorporating the Deuterium Motional Simulation (DMS) spectral calculation program into DFP to allow calculation of spectra during the fitting process, for refining fits after fitting with a library with large increments in the rates. Any inclusion of spectral calculation at each SA step must consider the time for each spectral calculation, and so will be practical mainly for simpler models with fewer sites. Also, the usefulness of refinement will depend upon whether a single exact fit or an average is best. In cases in which averaging over a range with the attendant estimation of errors through the standard deviations of parameters is appropriate, homing to a single best fit with refinement may not be helpful. Currently, DFP is not set up conveniently for searching for the best values of the quadrupole parameters, but rather assumes that these are known. To facilitate searching for best values of these parameters, they could be included as variables in the libraries, or the capability to vary the quadrupole parameters in SA fits in which spectral calculations are done at each step could be included in DFP. The ability to vary χ in a library and average over it would have been helpful in dealing with the alanine-d_3 MeRot pattern here, where two separate libraries with different χ where employed to represent a range of χ values (Figs. 5c and 6). Enhancements to the simulated annealing algorithm will be required to fit with six or more libraries, but this probably will not be pursued since most applications will not require so many individual motional models. If many models are required because a range of motions is present, averaging over the appropriate parameters to represent the range probably makes more sense than using several distinct models. Additional code for streamlining the fitting process, such as a program to set up inputs for DFP interactively and a procedure to automatically plot spectra from DFP output files would be useful.

5. Experimental

Deuterium-labeled L-alanine-3,3,3-d_3 (99 at.% D) was purchased from CDN Isotopes, and re-crystallized from a saturated H_2O solution [10]. Solid samples were weighed in 4mm glass containers designed to hold the entire sample inside the transceiver coil. Aqueous solutions were placed in a shortened 4 mm NMR tube that allowed some of the sample to be outside the coil; however, quantities were estimated for these samples by assuming that only a 1 cm long column of liquid inside the coil could be detected in the deuterium quadrupole echo spectrum. The ability to take spectra of a combination of solid alanine-d_3 and aqueous solutions of alanine-d_3 held in separate tubes in the NMR transceiver coil provides a means to acquire spectra representing two motional models, for which the quantities of deuterium undergoing each motion are quantitatively known independently of the NMR experiment. The percentages of methyl rotation and isotropic models were calculated from the mass of alanine-d_3 in solid samples, and the mass of alanine-d_3 and volume of the liquid samples. Deuterium-labeled L-phenylalanine-ring-d_5 (98%) was purchased from Cambridge Isotope Laboratories.

Solid-state deuterium quadrupole echo NMR spectra were acquired at 46.204–46.2065 MHz on a Chemagnetics CMX-II 300 MHz solid-state NMR spectrometer, using a static (non-sample-spinning) ^2H probe equipped with 5 mm and 10 mm transceiver coils. The 90° pulse lengths were 1.8 or 1.9 μs. Alanine-d_3 spectra were collected with delays between quadrupole echo pulses, τ, of 20, 40, or 100 μs, delay between scans of 1.5 s, spectral width of 500 kHz, 1024 or 8192 points, and 512 to 4096 scans. Spectra of samples including aqueous alanine-d_3 where corrected by subtracting from the echo signal before Fourier transformation a signal acquired in an identical fashion for H_2O, to subtract out a very small D_2O signal. Noise spectra with no sample, glass, or Teflon tape in the coil were obtained similarly to alanine-d_3 spectra, but with differing numbers of scans. A spectrum of phenylalanine-d_5 was obtained with τ of 40 μs, delay between scans of 0.1 s [23], 20,000 scans, 512 points and dwell time of 0.5 μs. This data was zero-filled to 2048 points for transformation, and then cut down to a spectrum of 512 points with spectral width 500 kHz for fitting. For all data collection the probe temperature was controlled to 25 ± 0.1 °C with a Chemagnetics temperature controller using compressed air as the cooling gas.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. 0210839. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation (NSF).

References

[1] H.W. Spiess, Molecular dynamics of solid polymers as revealed by deuteron NMR, Colloid Polym. Sci. 261 (1983) 193–209.
[2] J.H. Ok, R.R. Vold, R.L. Vold, M.C. Etter, Deuterium nuclear magnetic resonance measurements of rotation and libration of benzene in a solid-state cyclamer, J. Phys. Chem. 93 (1989) 7618–7624.
[3] D.F. Shantz, R.F. Lobo, Guest–host interactions in zeolites as studied by NMR spectroscopy: implications in synthesis, catalysis and separations, Top. Catal. 9 (1999) 1–11.
[4] E. Gedat, A. Schreiber, J. Albrecht, Th. Emmler, I. Shenderovich, G.H. Findenegg, H.-H. Limbach, G. Buntkowsky, ^2H-solid-state NMR study of benzene-d_6 confined in mesoporous silica SBA-15, J. Phys. Chem. B 106 (2002) 1977–1984.
[5] J. Xiong, G.E. Maciel, Deuterium NMR studies of local motions of benzene adsorbed on Ca–montmorillonite, J. Phys. Chem. B 103 (1999) 5543–5549.
[6] J. Xiong, H. Lock, I.-S. Chuang, C. Keeler, G.E. Maciel, Local motions of organic pollutants in soil components, as studied by ^2H NMR, Environ. Sci. Technol. 33 (1999) 2224–2233.

[7] R.G. Griffin, K. Beshah, R. Ebelhauser, T.H. Huang, E.T. Olejniczak, D.M. Rice, D.J. Siminovitch, R.J. Wittebort, Deuterium NMR studies of dynamics in solids, in: G.J. Long, F. Grandjean (Eds.), The Time Domain in Surface and Structural Dynamics, Kluwer Academic Publishers, 1988, Chapter 7.
[8] M.S. Greenfield, A.D. Ronemus, R.L. Vold, R.R. Vold, P.D. Ellis, T.E. Raidy, Deuterium quadrupole-echo NMR spectroscopy. III. Practical aspects of lineshape calculations for multiaxis rotational processes, J. Magn. Reson. 72 (1987) 89–107.
[9] R.J. Wittebort, E.T. Olejniczak, R.G. Griffin, Analysis of deuterium nuclear magnetic resonance line shapes in anisotropic media, J. Chem. Phys. 86 (1987) 5411–5420.
[10] K. Beshah, E.T. Olejniczak, R.G. Griffin, Deuterium NMR study of methyl group dynamics in L-alanine, J. Chem. Phys. 86 (1987) 4730–4736.
[11] A.E. Aliev, K.D.M. Harris, ^{2}H NMR lineshape analysis using automated fitting procedures on local and quasi-global optimization techniques, Magn. Reson. Chem. 36 (1998) 855–868.
[12] S. Kirkpatrick, C.D. Gelatt Jr., M.P. Vecchi, Optimization by simulated annealing, Science 220 (1983) 671–680.
[13] P.J.M. van Laarhoven, E.H.L. Aarts, Simulated Annealing: Theory and Applications, Kluwer Academic Publishers, 1987.
[14] P.R. Bevington, Data Reduction and Error Analysis for the Physical Sciences, McGraw-Hill Book Company, 1969.
[15] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical Recipes in C The Art of Scientific Computing, Cambridge University Press, 1992.
[16] T.M. Barbara, M.S. Greenfield, R.L. Vold, R.R. Vold, Deuterium quadrupole echo NMR spectroscopy. I. Effects of chemical exchange during single and composite pulses, J. Magn. Reson. 69 (1986) 311–330.
[17] A.D. Ronemus, R.L. Vold, R.R. Vold, Deuterium quadrupole echo NMR spectroscopy II. Artifact suppression, J. Magn. Reson. 70 (1986) 416–426.
[18] W.J. Thompson, Computing in Applied Science, John Wiley & Sons, New York, 1984, pp. 55–61, pp. 239–249.
[19] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, D. Sorensen, LAPACK Users' Guide, Society for Industrial and Applied Mathematics, 1999.
[20] J.M. Koons, E. Hughes, H.M. Cho, P.D. Ellis, Extracting multitensor solid-state NMR parameters from lineshapes, J. Magn. Reson. A 114 (1995) 12–23.
[21] M. Bak, N.C. Nielsen, REPULSION, a novel approach to efficient powder averaging in solid-state NMR, J. Magn. Reson. 125 (1997) 132–139.
[22] H. Conroy, Molecular Schrodinger equation. VIII. A new method for the evaluation of multidimensional integrals, J. Chem. Phys. 47 (1967) 5307–5318.
[23] C.M. Gall, J.A. DiVerdi, S.J. Opella, Phenylalanine ring dynamics by solid-state ^{2}H NMR, J. Am. Chem. Soc. 103 (1981) 5039–5043.
[24] R.R. Vold, Deuterium NMR studies of dynamics in solids and liquid crystals, in: Nuclear Magnetic Resonance Probes of Molecular Dynamics, Kluwer Academic Publishers, 1994, pp. 27–112.