Published in final edited form as:

J Biomol NMR. 2004 August ; 29(4): 491–504. doi:10.1023/B:JNMR.0000034351.37982.9e.

Automated evaluation of chemical shift perturbation spectra: New approaches to quantitative analysis of receptor-ligand interaction NMR spectra

Chen Peng^a, Stephen W. Unger^a, Fabian V. Filipp^b, Michael Sattler^b, and Sándor Szalma^{c,*}

^aAccelrys Inc., 9685 Scranton Road, San Diego, CA 92121, U.S.A.

^bEMBL Heidelberg, Meyerhofstr. 1, Heidelberg D-69117, Germany

^cMeTa Informatics, 12987 Caminito Bautizo, San Diego, CA 92130, U.S.A.

Abstract

This paper presents new methods designed for quantitative analysis of chemical shift perturbation NMR spectra. The methods automatically trace the displacements of cross peaks between a perturbed test spectrum and the reference spectrum (or among a series of titration spectra), and measure the changes of chemical shifts, heights, and widths of the altered peaks. The methods are primarily aimed at the ¹H-¹⁵N HSQC spectra of relatively small proteins (<15 kDa) assuming fast exchange between free and ligand-bound states on the chemical shift time scale, or for comparing spectra of free and fully bound states in the slow exchange situation. Using the ¹H-¹⁵N HSQC spectra from a titration experiment of the 74-residue Pex13p SH3 domain with a Pex14p peptide ligand (14 residues, $K_d = \sim 40\mu\text{M}$), we demonstrate the scope and limits of our automatic peak tracing (APET) algorithm for efficient scoring of high-throughput SAR by NMR type HSQC spectra, and progressive peak tracing (PROPET) algorithm for detailed analysis of ligand titration spectra. Simulated spectra with low signal-to-noise ratios (S/N ranged from 20 to 1) were used to demonstrate the reliability and reproducibility of the results when dealing with poor quality spectra. These algorithms have been implemented in a new software module, FELIX-Autoscreen, for streamlined processing, analysis and visualization of SAR by NMR and other high-throughput receptor/ligand interaction experiments.

Keywords

chemical shift perturbation spectra; dissociation constant K_d ; FELIX-Autoscreen program; ¹H-¹⁵N HSQC spectroscopy; high-throughput screening; ligand titration experiment; peak picking; peak shape; receptor/ligand interaction; SAR by NMR

Introduction

Chemical shift perturbation spectra are very commonly used to probe the structural and functional changes of a molecule after a certain chemical or physical event, such as chemical reactions, solvation, complexation, mutation, or binding to another molecule. The

*To whom correspondence should be addressed. szasa@metainfoconsulting.com.

underlying rationale is that if the system under study has been changed chemically or physically, the NMR peaks may change their location, line shape, or disappear due to changed local chemical environment (see reviews by Otting, 1993; Pellecchia et al., 2000; Zuiderweg, 2002 and references therein). Among the many possible techniques, 2D HSQC of ^{15}N -labeled proteins (Bodenhausen et al., 1980; Bax et al., 1990) has long been used in the studies of protein-ligand binding events because of its high sensitivity to ligand interaction over a wide affinity range, and its ability to provide residue-specific information that helps map the binding site(s) on the protein. This idea has been elaborated as the SAR by NMR method for screening and designing high-affinity ligands (Shuker et al., 1996; Hajduk et al., 1997). Lately, chemical shift perturbation has also been used in the study of protein-protein docking (McCoy and Wyss, 2000; Morelli et al., 2001, Fahmy and Wagner, 2002; Dominguez et al., 2003; Clore and Schwieters, 2003).

The analysis of perturbation spectra is usually focused on the difference between the perturbed (called test) and the non-perturbed (called reference) spectrum. Theoretically, the difference can be obtained by directly subtracting the two spectra and then reading the residual peaks. This is a common practice for 1D spectra though rarely used for 2D. In a more sophisticated data analysis approach (Ross et al., 2000) to evaluate the 2D ^1H - ^{15}N HSQC spectra for high-throughput screening (HTS), the cross peaks were picked from the reference spectrum and used to define the integration areas for the test spectra. For each test spectrum, the integrations of data points from these areas are compared with the corresponding ones from the reference spectrum to calculate the similarity (called 'correlation coefficient') between the test and the reference spectrum. The changes of the integrals of a batch of test spectra were next subject to statistical analysis to cluster the experiments into groups that correspond to different binding sites or corrupted experiments due to protein aggregations or changes in sample conditions. Billeter et al. have recently introduced a mathematical method, three-way decomposition, in the analysis of a set of ^1H - ^{15}N HSQC spectra (Orekhov et al., 2001; Damberg et al., 2002). In their application, the enumeration of the set of 2D spectra constitutes the third dimension in addition to the ^1H and ^{15}N dimensions of the ^1H - ^{15}N HSQC spectra. It is demonstrated that the unchanged and changed peaks can be identified as 1D peak shapes along the ^1H and ^{15}N axes, and the change of the shapes along the 3rd dimension can be used to identify which of the spectra have been significantly affected. Obviously, the result from either method is a rough measure of how many peaks have shifted away or disappeared from the original reference peak locations, but none of them distinguishes between shifted or disappeared peaks, or, more importantly, measures how far peaks have shifted from their original locations.

In our opinion, it is more intuitive and chemically meaningful to map the changed peaks to the original reference peaks, and quantitatively measure the changes in terms of chemical shifts (and peak volumes or peak widths if they are useful). Obviously, such results are the necessary starting point for many of the studies summarized above, including the mapping of the binding site(s), determination of K_d , and the study of protein-protein docking. In the case of HTS, we also believe that such results constitute a better input for further statistical analysis and hence lead to more reliable clustering results. However, unless the test spectrum has been fully assigned using other NMR experiments, peak mapping is usually tedious and error-prone because of noise, peak overlap, and cross shifting. In manual analysis (e.g.,

Farmer et al., 1996, Williamson et al., 1997), the test and reference spectra are overlaid and the unchanged peaks are first identified. For each of the shifted peaks, instead of trying to find its real origin, a test peak is assigned to its nearest reference peak and the distance is measured to define the 'lower limit for chemical shift changes'. Such nearest peak-based approach was explicitly addressed and verified with assignment results (Williamson et al., 1997; Muskett et al., 1998). In our experience, however, it is not a rigorous method and may lead to error in some situations illustrated in Figure 1. One way to clarify such ambiguity is to use the titration spectra (e.g., Van Nuland et al., 1993; Chen et al., 1993). The tracing of the peak movement is more reliable since the displacement between each titration step is smaller and hence less ambiguous. It should also be noted that a human analyst can easily avoid matching peaks to noise and match the other peaks correctly in situations illustrated in Figure 1, based mainly on earlier experience with the peak shapes. The pattern recognizing capability of a computer program, usually done at subconscious level by human, is of crucial importance for matching peaks in a robust way.

Based on these considerations, we have developed a software module, Autoscreen, in the FELIX environment (Kumar et al., 1998; Accelrys, 2002) to automate the manual analysis tasks mentioned above. To overcome the common practical difficulties, new algorithms were devised to automatically optimize the peak picking parameters, and to use peak shape for robust peak matching. Our primary targets are 2D ^1H - ^{15}N HSQC spectra perturbed by protein-ligand interactions in the fast exchange regime as usually is the cases of HTS or conventional ligand titrations, or the saturated spectra in the slow exchange limit. The scope and limits of the methods are demonstrated by the analysis of the ^1H - ^{15}N HSQC spectra of the Pex14p peptide bound Pex13p SH3 domain. Simulated spectra with low signal-to-noise ratios (S/N ranged from 20 to 1) were used to demonstrate the reliability and reproducibility of the results when dealing with poor quality spectra.

Methods and algorithms

Mathematical consideration

Let the set of reference peaks picked from the reference spectrum be $P = \{p_1, p_2, \dots, p_n\}$, and the set of test peaks from the test spectrum be $T = \{t_1, t_2, \dots, t_m\}$. In terms of graph theory (Gross et al., 1998), the relationship between P and T can be represented by a bipartite graph $G = (P + T, E)$, where the vertex set $P + T$ is the union of the two peak sets, and the edge set $E = \{e_1, e_2, \dots, e_q\}$ includes the mapping between the reference and test peaks (Figure 2). G is bipartite because the two incident vertices of each e are from the separate sets P and T , not from a common one. Each edge $e = (p, t)$ is associated with a distance weight d , which is a measure of the change of chemical shifts and peak shape from p to t . For an unchanged peak, $d = 0$. For a peak that either shifted its location or changed its line shape, $d > 0$. In practice, it is not uncommon for some peaks to have no match in the other peak set. Such unmatched peaks usually include disappeared reference peaks or newly emerged test peaks. Although there is no edge associated with such peaks, a weight d_r or d_t can be assigned to an unmatched reference or test peak, respectively, to count for their contributions to the overall changes of the spectra.

For chemical shift perturbation spectra, we define the distance d between p and t to be the sum of the weighted absolute chemical shift differences, H and N , along the ^1H and ^{15}N dimension, respectively, augmented by the change of peak shape:

$$d = (\alpha \Delta H + \beta \Delta N) / S, \quad (1)$$

where α and β are the weights of the ^1H and ^{15}N chemical shifts respectively, used to compensate the different chemical shift dispersions of ^1H and ^{15}N nuclei. Normally $\alpha = 1$ and $\beta = 0.2$. S is the similarity of peak shapes of p and t , and has a value of 1.0 if they are identical in shape or less than 1.0 if they are different. In this way a chemical shift displacement is enlarged if the peak shape is changed. S is defined as the product of the similarities of their peak widths and peak intensities:

$$S = S_{\lambda_H} \times S_{\lambda_N} \times S_I, \quad (2)$$

where S_{λ_H} , the similarity of peak width along ^1H dimension, is calculated based on the peak width of p , λ_H^p , and that of t , λ_H^t , along the ^1H dimension:

$$S_{\lambda_H} = \begin{cases} \lambda_H^p / \lambda_H^t & \text{if } \lambda_H^p \leq \lambda_H^t, \\ \lambda_H^t / \lambda_H^p & \text{otherwise.} \end{cases} \quad (3)$$

Analogously, S_{λ_N} , the similarity of peak width along the ^{15}N dimension, is calculated based on the peak width of p , λ_N^p , and that of t , λ_N^t , along the ^{15}N dimension:

$$S_{\lambda_N} = \begin{cases} \lambda_N^p / \lambda_N^t & \text{if } \lambda_N^p \leq \lambda_N^t, \\ \lambda_N^t / \lambda_N^p & \text{otherwise.} \end{cases} \quad (4)$$

S_I , the similarity of peak intensity, is calculated based on the peak heights of p , I^p , and that of t , I^t :

$$S_I = \begin{cases} I^p / I^t & \text{if } I^p \leq I^t, \\ I^t / I^p & \text{otherwise.} \end{cases} \quad (5)$$

Once the peak mapping is available, the difference between the test and reference spectra can be obtained by summing up the contributions of the matched peaks and unmatched peaks as the following:

$$D = \sum d + \sum d_r + \sum d_t. \quad (6)$$

D is commonly used for scoring HTS experiments where such ‘scores’ of the test spectra are used to cluster the experiments into groups that correspond to non-changed, possibly binding, and corrupted experiments. For a more detailed analysis, it is often instructive to examine the displacements of the individual peaks contributing to D .

Now the key is to obtain the mapping between the test and reference peaks. Since normally the peak displacements are small compared to the chemical shift dispersion area, and the majority of the shifted peaks have relatively small changes to their peak shape, we can limit the possible matchings to only those pairs that fulfill the following criteria:

$$\Delta H \leq \max \Delta H, \quad (7)$$

$$\Delta N \leq \max \Delta N, \text{ and} \quad (8)$$

$$S \geq \min S, \quad (9)$$

where $\max H$ is the upper limit of the absolute displacement of ^1H chemical shift between p and t , $\max N$ the upper limit of the absolute displacement of ^{15}N chemical shift, and $\min S$ the lower limit of shape similarity.

The conditions in Equations 7–9 usually do not guarantee a unique peak mapping. Instead, we have to enumerate all the possible mappings and select the one with the minimum D as the optimal peak mapping, which constitutes the lower limit of the global change of the whole spectrum. This is a typical combinatorial optimization problem and can be solved using the systematic tree search methods or stochastic methods (Lawler et al., 1985; Press et al., 1992).

Intelligent peak picking in the test spectra

We assume that a well-refined, usually assigned, peak set from the reference spectrum is available before analyzing the HTS or titration experiment. The peak picking in the test spectra, however, should be done on-the-fly during scoring. Hence the quality of the test peaks is crucial to the subsequent analysis. We have devised an algorithm that automatically adjusts the pick area and threshold based on the quantity of peaks and the quality of peak matching.

First a peak picking area is determined based on the chemical shift dispersion of the reference peaks. The area is the minimum rectangle that includes all the reference peaks plus margins of $\max H$ and $\max N$ in D1 and D2 dimensions, respectively. Starting from an estimated initial threshold (calculated based on some randomly picked data points), the test peaks are picked. Next the peaks are examined and threshold is adjusted according to the following rules:

1. If the number of test peaks is less than 1.1 times of that of the reference peaks, lower the threshold. This ensures that at least 10% more test peaks are picked in order not to miss weak valid test peaks.
2. If the number of test peaks is greater than two times of that of reference peaks, raise the threshold. This ensures not too many peaks are picked.
3. Otherwise, for each reference peak, find its candidate matches based on Equations 7–9. Note that the peak intensities are not used at this stage when calculating the peak shape similarity using Equation 2, i.e., $S_I \equiv 1$, because the test spectrum may have different intensity level from that of the reference. If over 25% of the reference peaks do not have a candidate match, lower the threshold. Otherwise if each reference peak has more than three candidates on the average, raise the threshold. This aims to tune the peak picking based on the quality of peak mapping.

If the test peaks pass all the examinations, they are accepted and the peak picking is completed. Otherwise the threshold is adjusted accordingly. The maximum number of iterations is set to seven.

Mapping test peaks to reference peaks

For each reference peak p , the nearby test peaks that satisfy Equations 7–9 are considered possible candidates of its match. The distance d between p and such a candidate test peak, t , is calculated based on Equations 1–4, with the peak shape similarity considered. If there are multiple candidates, they are sorted by their d values and only the first four with the smallest d are retained. Again, at this stage, the peak intensities S_I are not used in Equation 2 since the test spectrum may have a different intensity level from the reference spectrum.

In order to calibrate the test peak intensities, an intensity ratio R_I is calculated as follows:

$$R_I = \sum I^p / \sum I^t, \quad (10)$$

where $\sum I^p$ is the sum of peak intensities of all reference peaks that have at least one candidate and $\sum I^t$ is the sum of peak intensities of the first candidate test peak of those reference peaks.

Next, the peak intensities of all test peaks are calibrated by multiplying them by R_I , and the list of candidate matches is updated by including the peak intensity in Equation 2. To guarantee that each reference peak will have a match in the final result, a dummy candidate with a distance d_t is added to the candidates for each reference peak. The remaining task is to choose one candidate for each reference peak so that D in Equation 6 is minimal. Note that a test peak can be assigned to only one reference peak in the resulting peak mapping, although some efforts are used to resolve overlapping test peaks as shown in the later section.

As the first option, the classical depth-first tree search algorithm (Gross et al., 1998) has been adopted to enumerate all the possible mappings. The search tree takes the first candidate of the first reference peak as the root, and takes the first, non-matched candidate of the second reference peaks as the next level of the search tree, and so on, until the n th (n is the total number of reference peaks) reference peak is reached. The selection of the matches constitutes one possible peak mapping and the score D is calculated based on Equation 6. Next it backtracks to the $(n - 1)$ th reference peak and continues the search so that the next candidate of the n th reference peak is selected as the match. If the i th ($1 < i < n$) test peak is exhausted, i.e., all its candidates have been selected, it backtracks to a higher level, $(i - 1)$ th reference peak, and select the next available candidate, and go on searching the new subtree. Each time a new mapping is obtained, the new score D' is compared with the current D . If $D' < D$, then the new mapping is retained and D is replaced by D' ; otherwise the new mapping is discarded and D remains unchanged. The process is complete when it backtracks to the last candidate of the first reference peak. The retained mapping and its corresponding score D are returned as the result.

Two heuristics are devised to reduce the search space and hence enhance the efficiency of search. In the first heuristic, the search tree is divided into as many independent subtrees as possible. This is done by first excluding those reference peaks with unambiguous matching test peaks, and then separating the remaining reference peaks into groups so that any two reference peaks from different groups do not share a candidate, and hence the search can be done independently for each group. The second heuristic takes advantage of the fact that the search space is partially ordered because the candidates of each reference peak have been sorted based on their distance d . Since normally the majority of the displaced peaks remain the nearest or the second nearest to their original locations, a small 'average search breadth' can be used to further limit the search space. (See Peng et al., 1995 for details of a similar method devised to quickly generate organic molecules from NMR-derived structural fragments).

The increased spectral complexity of bigger proteins may undermine the effectiveness of the above heuristics. To achieve a near-optimal matching in affordable computational time for complicated spectra, we also adapted the simulated annealing algorithm to the peak matching problem. In terms of simulated annealing (Kirkpatrick et al., 1983; Press et al., 1992), each possible peak mapping is a configuration with an energy D calculated according to Equation 6. In our annealing schedule, the initial temperature T is set to $10d_r$, which is considerably larger than the largest change of D when exchanging a pair of matches. The annealing proceeds downward in multiplicative steps each amounting to a 10 percent decrease in T . Each new value of T is held constant for $100n$ reconfigurations, or for $10n$ successful reconfigurations, whichever comes first. If T has been decreased for 100 times, or if no successful reduction of D is accomplished for a certain T value, the process is stopped and the current configuration is used as the resultant mapping. Such an algorithm is normally much faster than the tree search algorithm hence in this case the two heuristics devised for the tree search are not needed.

Handling the unmatched reference peaks

In the resultant peak mapping, an unmatched reference peak is represented as one matched to a dummy test peak with a distance of d_t . There are many possible reasons for a reference peak not to be matched. For example, the peak vanishes in the test spectrum because of considerable change of relaxation time, or the peak is displaced greater than $\max H$ or $\max N$ and is hence not matched. More often two well-separated peaks (or overlapping ones but identified from other experiments) in the reference spectrum become (or remain) partially overlapped in the test spectrum and hence only one test peak is picked, as illustrated in Figure 3.

In order to handle a case like Figure 3, we use the local optimization method of the FELIX program (Accelrys, 2002) to unravel the hidden test peaks. First, the unmatched reference peaks, together with matched ones that lie closer than four times the peak width to them, are fitted to the corresponding portion of the test spectrum. Then, the optimized peaks are matched to the corresponding reference peaks if they meet the criteria in Equations 7–9. If this is successful, their contributions are calculated according to Equation 1. In this way the originally matched test peak may also be refined (e.g., the matching of peak **2** to **a'** instead of **a** in Figure 3). Otherwise, such a reference peak remains unmatched and contributes d_t to the score D in Equation 6.

Note the value of d_t must be set significantly larger than the value d of a normal peak so that a dummy candidate is only selected when other candidates are not available. However, the value of d_t should not be too big in order to prevent noise peaks from being matched to the reference peaks. Empirically we set the value of d_t to the upper boundary of a detectable peak displacement based on the values of $\max H$ or $\max N$, i.e., $d_t = \max H + \max N \times 0.2$. In the subsequent application, 0.18 and 1.1 are used for $\max H$ or $\max N$, respectively, and 0.4 is hence used for d_t . Internally the program uses $1.5d_t$, i.e. 0.6 in this case, for peak mapping.

Handling unmatched test peaks

Since we always pick over 10% more test peaks than reference peaks, many test peaks remain unmatched. In our experience, most of the unmatched test peaks are noise, but they can also be real peaks including the newly emerged ones due to change of sample conditions, the ones corresponding to the reference peaks that one fails to observe in the reference spectrum because of overlap or other reasons, or the ones that the program fails to match because their displacements exceed the search range defined in Equations 7 and 8.

We use the peak shape to distinguish the genuine unmatched test peaks from noise. With the assumption that all matched test peaks are real peaks, statistical analysis is performed on their peak widths and intensities. The average values ($\bar{\lambda}_H$, $\bar{\lambda}_N$, and $\bar{I}$) and standard deviations (σ_H , σ_N , and σ_I) of the peak widths along the ^1H dimension, the peak widths along the ^{15}N dimension, and peak heights, respectively, are obtained. An unmatched test peak is accepted as a real peak if its peak widths λ_H and λ_N and peak height I meet the following conditions:

$$\bar{\lambda}_H - x\sigma_H \leq \lambda_H \leq \bar{\lambda}_H + x\sigma_H, \quad (11)$$

$$\bar{\lambda}_N - x\sigma_N \leq \lambda_N \leq \bar{\lambda}_N + x\sigma_N, \text{ and} \quad (12)$$

$$\bar{I} - x\sigma_I \leq I \leq \bar{I} + x\sigma_I, \quad (13)$$

where x is a user-defined coefficient (e.g., 1.2 to 2.0).

Each of the identified unmatched test peaks contributes a value d_t to the score of the experiment according to Equation 6. Normally we use $d_t = 0.2$, which is smaller than the contribution of an unmatched reference peak (e.g., $d_r = 0.4$) based on the fact that reference peaks are usually well-refined during resonance assignment while the test peaks are automatically picked and are usually less reliable.

PROPET: Progressive peak tracing in titration spectra

The automatic peak tracing (APET) algorithms discussed above are especially suited for automatically analyzing the HTS experimental data, where each ligand (or a groups of ligands) has only one test spectrum. On the other hand, if multi-point titration data is available, all the test spectra with different ligand concentrations should be evaluated against the reference spectrum to follow the titration curve. Moreover, the relatively small peak displacements between the individual test spectra can usually be used to resolve ambiguous peak mapping. For that purpose, the APET algorithms are extended to deal with titration spectra by progressive peak tracing (PROPET).

Instead of directly evaluating each titration spectra against the reference, PROPET evaluates each titration spectrum against its immediate predecessor in the titration series. It starts with the reference spectrum, and analyzes the series of test spectra along the titration curve sequentially. If it is the first test spectrum, the procedure is essentially the same as APET, i.e., the test peaks are automatically picked and then mapped to the reference peaks as described in the above sections. This establishes a mapping $M(T_1, P)$ between the test peaks T_1 and the reference peaks P , and the displacement between each matched peak pair, if shifted, is measured. Next the second test spectrum is evaluated against the first test spectrum by taking the first test spectrum as a reference, and a mapping between the test peaks T_2 (from the second test spectrum) and T_1 is established as $M(T_2, T_1)$. By replacing T_1 with their matched reference peaks in P , a mapping between the second test spectrum and the reference spectrum is established as $M(T_2, P)$, and the displacements of the test peaks referring to the reference peaks are measured again. The same procedure goes on until the last test spectrum is evaluated. The results are tabulated as the chemical shift changes for each reference peak along the ligand concentration dimension.

Experimental

FELIX-Autoscreen

The APET and PROPET algorithms are coded in C and C++ and implemented as a new module, Autoscreen, in the FELIX program (Accelrys, 2002). To use Autoscreen, the reference and test spectra files, either in time-domain or frequency-domain can be entered to Autoscreen and constitute a project. If in time-domain, the reference spectrum is first processed interactively. The parameters used for processing the reference spectrum are remembered by Autoscreen and can be automatically applied to the test spectra later. If APET is chosen to analyze a batch of HTS spectra, all the test spectra are processed, peak picked, and evaluated against the reference peaks automatically without user interaction. If PROPET is chosen to analyze titration spectra, the procedure is carried out in a more interactive way so that, for each point in the titration series, the user can choose the proceeding point to use, verify the peak mapping result and correct wrong matches if necessary. In either case, the results can be presented as histograms or tables of scores for all experiments, or chemical shift changes for an individual spectrum. A test spectrum can be overlaid on the reference spectrum as contours of different colors, and arrows be displayed to indicate the significantly shifted peaks. The arrows can be removed or changed interactively if corrections are needed. The multiple spectra overlay allows up to 12 different spectra be overlaid at the same time.

Other routines are provided in FELIX to cluster a batch of HTS experiments into different groups based on the resulting peaks displacements, to export the peak displacement results for K_d calculation and coloring the perturbed residues on the 3D structure of the protein.

NMR Experiments

The full titration curve for a Pex14p peptide ligand (14 amino acids) to the Pex13p SH3 domain (amino acids 303–373 plus amino acids –3 - –1 from proteolytic cleavage) (Douangamath et al., 2002) is used to demonstrate the effectiveness of our algorithms. Backbone resonance assignments of the SH3 domain were obtained using standard triple resonance experiments (Sattler et al., 1999) recorded on a ^{15}N , ^{13}C -labeled Pex13 SH3 domain. Chemical shift changes during the NMR titrations were monitored in 2D sensitivity enhanced ^1H - ^{15}N HSQC experiments with 512 and 128 complex points in t_2 and t_1 , respectively, recorded with 16 scans on ^{15}N -labeled SH3 domain at 600 MHz. The SH3 domain had an initial concentration of 0.5 mM in a volume of 600 μl . The peptide ligand Pex14p (30 mM) was added to a final concentration of 1.0 mM, such that the total change of the sample volume was less than 5%. The full titration curve consisted of seven ^1H - ^{15}N HSQC spectra recorded at different ligand concentrations ranging from 0.1 to 1.0 mM, as shown in Table 1.

Results and discussion

Spectral data processing and simulation

The ^1H - ^{15}N HSQC spectra were processed on a DELL personal computer (CPU: AMD Athlon™ 1G Hz, RAM: 256 Mbyte) using the FELIX 2002 software. The reference

spectrum together with the seven test spectra were organized as an Autoscreen project. The reference spectrum was first zero-filled and processed to a data matrix of 1024×512 points. The matrix was automatically phase corrected using the PAMPAS method (Dzakula, 2000), baseline corrected using the Facelift method (Chylla and Markley, 1993), and reversed in D2 dimension. The spectrum was referenced to match the previous assignment results by setting the cross peak (G357 in Figure 4) with the highest ^{15}N frequency as (8.15, 107.47) ppm. The 65 assigned amide peaks were imported and displayed on the reference spectrum. Thirteen unassigned peaks were manually picked. This led to a total of 78 peaks, including two negative peaks. Peaks were fitted to the reference spectrum with peak center optimized first, and next with peak center, width and volume optimized for two times. A few peaks lying in the overlapping areas were given abnormal peak widths and were then manually corrected. The total 78 optimized peaks were used as the reference peak set for the subsequent analysis.

The seven test spectra were processed automatically using the same parameters as used to process the reference spectra. The S/N ratio of these spectra is about 120 in the fingerprint area.

For the subsequent test, a series of low S/N ratio spectra were synthesized for the Reference, Test1, and Test7 by adding random noise with various intensity levels. Starting from an experimental spectrum, the highest well-resolved peak in the fingerprint area was picked and its height, I_0 , was measured. The S/N ratio was then calculated as $2.5 \cdot I_0$ divided by the peak-to-peak noise amplitude (Martin et al., 1980). For a target S/N ratio random noise was added to the entire spectrum to lower the S/N to the desired value. For each of the three original spectra, 20 spectra were synthesized with a S/N ratio ranging from 20 to 1.

APET analysis of the saturated test spectrum Test7

The last point, Test7, of the titration curve was evaluated directly against the reference spectrum to mimic a HTS case where usually only one test spectrum is available for each ligand. Based on the location of the reference peaks, the program limited the peak picking area to the fingerprint area. The peak picking threshold was automatically adjusted twice and 117 peaks were picked from Test7. All test peaks falling closer than ($\max H = 0.18$, $\max N = 1.10$) ppm to a reference peak, as well as sharing a similarity of peak shape $S > 0.01$, were taken as the candidate matches of the reference peak. The peak width and peak heights (calibrated by a ratio of $R_1 = 1.91$, see Equation 10) were used to compute the peak shape similarity and to augment the chemical shift distance between two matched peaks (Equations 1–5). For an unmatched reference peak, a ‘distance’ of $d_r = 0.4$ was assigned. Unmatched test peaks were filtered based on Equations 11–13 with $x = 1.8$ and each contributed $d_t = 0.2$ to D . The tree search method was used to find the peak mapping with the lowest score D (Equation 6).

The analysis took less than a second to complete. The results of peak mapping and displacement identification are illustrated in Figure 4 and some crowded areas are detailed in Figure 5. Among the 78 reference peaks, 70 were matched to a picked test peak. Among the 70 matches 47 pairs had $d > 0.1$ and are connected by green arrows. The remaining 8 reference peaks were identified as unmatched peaks (marked by green crosses). Moreover, 5

test peaks were identified as unmatched test peaks (marked by green boxes). Since the chemical shift changes are our main interest, we chose to use the peak shape only during peak mapping but not to include the change of peak shape while calculating the score D (i.e., D is recalculated with $S \equiv 1$ in Equation 1 after the peak mapping). The peak shape similarities of the matched peaks are listed separately (See Table 1). The sum of the peak displacements of matched peaks was 5.53, and unmatched reference peaks and test peaks contributed 0.4 or 0.2 each, respectively, giving to a total score of 9.73.

It is noteworthy that none of the reference peaks are matched to any noise peak, although around 30 more peaks (mostly noise) were picked. Remarkably, those spurious noise peaks were successfully filtered from the unmatched test peaks, leaving only the 5 legitimate ones (marked in green boxes). This shows that the use of the peak shape similarities makes the peak matching more robust. As listed in the last column of Table 1, the shape similarities are mostly bigger than 0.5. On the other hand, a noise peak, which usually has very different peak intensity and peak widths, typically has a similarity much smaller than 0.01 and hence has a very small chance of being chosen as the final match to a reference peak.

It is also noted that the fitting of unmatched references to the test spectrum worked well for the overlapped peaks D346 and E309 (Figure 5A). Without this function D346 would be claimed as unmatched reference peak since their corresponding test peaks are not resolved in the test spectrum either, leaving only one test peak picked for them by the program. This function, however, did not work successfully for R345 (Figure 5A) because the program did not properly fit the unmatched reference peak R345 to the shoulder peak (pointed to by the red arrow). The peak matches in the crowded area (Figure 5A) show that the global optimization worked well in tracing the cross shifted peaks such as I336 and W349, M-1 and K329. The displacements of these peaks are analogous to the situation illustrated in Figure 1A and would be hard to resolve using the nearest peak-based method. Although some of the matches are further corrected in the following PROPET analysis, we believe these results are close to what a human analyst could achieve based solely on the Test7 and reference spectra.

Reproducibility study of APET results

In real-world HTS, the spectral quality may vary from sample to sample, and the interactive verification and correction of individual peak mapping is not possible. Hence the reproducibility of the resulting D values is of essential importance for reliably profiling experiments. There are two main sources that may lead to the fluctuation of the D value. The first is the inherent stochastic nature of our methods – the estimation of the initial peak picking threshold is of random nature, and the optional simulated annealing method for peak mapping is a stochastic method. The second is the impact of poor spectral quality on the results.

To test the variation of the D values due to our stochastic algorithms, we repeated the APET analysis of Test7 for 10 times with the tree search or the simulated annealing option.

Although the number of picked test peaks varied from 96 to 153, the tree search method gave constant results (5.53 for the matched peaks, and 8 unmatched reference peaks) except that number of unmatched test peaks varied from 4 to 5, hence D ranged from 9.53 to 9.73 with an average of 9.60 and a standard deviation of 0.10. The simulated annealing, however,

resulted in more fluctuation with D values ranging from 8.97 to 10.03, with an average D value of 9.59 and standard deviation of 0.40. As expected, simulated annealing gave more variation in the results than the tree search method.

To test the reproducibility of the D values on spectra with low S/N ratio, we used the Reference, Test1, and Test7 spectra (which mimic a non-binding, partially binding, and fully binding case, respectively) with artificially lowered S/N ratio ranging from 20 to 1 and repeated the APET analysis with the identical parameters as in the previous APET analysis section, except that $x = 1.2$ was used to give a more rigorous filtering of unmatched test peaks (see Equations 11–13). As illustrated in Figure 6A, in the S/N range of 20 to as low as 5, the resulting average D values are 0.67, 5.45, and 8.51, with standard deviation of 0.48, 0.49 and 0.68, respectively. These results are separated well enough to reliably profile the three samples. A study of the peak mapping in these spectra found that, while the majority of the matched peaks remained remarkably consistent with the results of the experimental spectra, the numbers of unmatched reference and test peaks varied considerably. As illustrated in Figure 6B, if the contribution of the unmatched peaks is excluded, the resulting D scores show much smaller standard deviations of 0.20, 0.34, and 0.47, respectively, in the same range of S/N ratios. The reasons are that as the spectrum becomes more noisy, there is a greater chance that artifacts are picked as test peaks, and the peak fitting for unmatched reference peaks becomes less reliable.

PROPET analysis of the whole titration curve

The first titration point, Test1, was evaluated against the reference spectrum with the same parameters as those used for the APET analysis of Test7, except that the search range was shrunk to $\max H = 0.10$ and $\max N = 0.80$ ppm since smaller peak displacements were expected. Of the total 78 reference peaks 70 were matched to a test peak, and 8 remain unmatched. The peak matchings were inspected and verified manually by browsing through each of them and a few matches were corrected interactively.

Next the second titration point, Test2, was evaluated against the preceding point, Test1, with the same parameters, and the results were verified before proceeding to the next titration point. Similar procedures were taken sequentially along the whole titration curve until the last point, Test7, was analyzed. The peak mapping obtained by this progressive analysis is also shown in Figure 4, with only the 11 different peak matches being illustrated as red arrows. These different peak matches are detailed in Figure 5 with selected multiple spectral contours overlaid in different colors.

It can be seen that by following the stepwise peak shifts, the PROPET analysis clarified many peak matchings which were not possible in the APET analysis. Figures 5B and 5C show two such typical examples. In Figure 5B, peaks M324 and 73 were originally matched to two closer test peaks (the green arrows). By following the intermediate peaks during the PROPET analysis, their matches were switched and bigger displacements were assigned to both of them (the red arrows). In Figure 5C, peak K340 was originally identified as an unmatched peak (marked by a green cross) and the test peak close to it was matched to peak E323 (green arrow). The PROPET analysis showed that E323 was actually shifted to exactly

the same location as D316 (which did not shift). The other test peak was hence matched to K340 (red line).

By scrutinizing the multi-spectra overlay display around the unmatched peaks, the unmatched reference peaks D304 and I362 could be unambiguously matched to two of the unmatched test peak (purple arrows in Figures 4 and 5C). These matches were not found during the APET analysis because their displacements ($H = 0.47$ ppm for D304 and $N = 2.54$ ppm for I362) exceed the used search range (max $H = 0.18$ ppm, max $N = 1.10$ ppm). If the search range was extended to cover their shifts, they would be properly matched. However, the tree search algorithm would be significantly slowed down since the search space is enlarged for each peak. In that case, it would be better to adopt the simulated annealing option for the search. It should be noted, however, that missing such peak matches constitutes little impact upon the purpose of APET analysis, since all the relevant peaks are recognized as unmatched reference or test peaks and each of them contributes a reasonable contribution to the score D .

The afore-mentioned two peak pairs were not matched by PROPET either because in the first titration experiment Test1, none of them showed a recognizable test peak due to intermediate exchange. It should be noted that the PROPET analysis does not have the capability to correct the peak tracking *retrospectively*. For that reason, it is recommended to verify and correct the peak mapping at each step manually. Sometimes one even needs to go back to correct a match from the analysis of a previous titration point. An interesting example is the shift of the two partially overlapping peaks M-1 and D348 in Figure 5A. During the PROPET analysis of the first titration point, Test1, they were matched to the two red Test1 peaks **a** and **b**, respectively. However, those matches, when relayed to the relevant test peaks in the subsequent titration points, would lead to deviation from the original peak shifting direction (see the green dashed line from peak M-1, the other line not shown). On the contrary, if M-1 and D348 are switched, the displacement vectors (the two red arrows starting from M-1 and D348) fit well with the centers of the intermediate peaks. Hence we corrected the matches in Test1 and re-analyzed the other experiments so that the correct trajectory was followed.

Figure 5D shows an example where even a titration experiment may not resolve the matching unambiguously. Both APET and PROPET analyses did not find a match for the reference peak Y361. By carefully inspecting the multi-spectrum overlay we found that one of the spectra, Test3, showed a shoulder peak (the green contours marked by a black arrow). This shoulder peak appears to be an intermediate peak that supports the proposal that Y361 shifted to exactly the same place as peak T354 did (illustrated as the purple arrow). The significantly raised peak height of the resulting test peak is another piece of supporting evidence. However, more exclusive experimental evidence (e.g. by changing the sampling condition to separate the peaks, or by using other experimental techniques to get the resonance assignment of Test7) is needed to draw an unambiguous conclusion in this case.

The step-wise chemical shift displacements of some of the significantly perturbed peaks ($d > 0.17$ in Test7) are listed in Table 1. From such a list it is straightforward to map the binding

sites from the PDB structure of the protein and calculate the K_d of the ligand binding, as has been described by Douangamath et al. (2002).

Conclusion

In this paper we presented a suite of algorithms designed for quantitative evaluation of chemical shift perturbation spectra by tracing the peak displacements. Compared to other methods, this approach is more intuitive and gives more complete and detailed results such as which cross peaks have changed and how much they have changed. In order to match peaks reliably and efficiently, we developed novel algorithms that automatically optimize the peak picking area and threshold to guarantee an appropriate set of test peaks are picked, and use the similarity of peak shape to enhance the matching of displaced peaks between different spectra. Based on a rigorous definition of the measure of difference between a perturbed spectrum and the reference spectrum, we adapted the classic combinatorial optimization methods, tree search and simulated annealing, to search the peak mapping that corresponds to a global minimum of the spectral difference. We also presented methods that deal with the disappeared and emerged peaks, by fitting unmatched reference peaks to the test spectrum and filtering unmatched test peaks based on the statistics of peak shapes. It is demonstrated that, even for titration of a peptide ligand which introduces far more chemical shift perturbations than in typical SAR by NMR type data, the APET algorithms can quickly trace the changed peaks with an accuracy that is close to that of a human analyst. For protein-ligand titration spectra, it is shown that the elaborated PROPET algorithms monitor the peak displacements in a more reliable way by following the stepwise peak movements along the titration curve. Furthermore, it is worth mentioning that the tools provided by our program, such as those for overlay display of multiple spectra, interactive correction of peak mapping and displacements, and automatic bookkeeping of such results in both database and spreadsheets, are very useful for a human analyst especially when dealing with complicated spectra.

We also note that the PROPET algorithms can be further improved by tracking the peak movements not only prospectively but also retrospectively. A relatively straightforward improvement is to check whether the displacements corresponding to the same peak have roughly the same directions. A more sophisticated improvement is to choose the displacements by checking if their distances can be fit to the expected mathematical curve along the titration trajectory (e.g., a hyperbolic function when a binary binding event is expected). Moreover, a more reliable algorithm for peak fitting is needed to match a reference peak to a partially overlapping test peak such as peak R345 in Figure 5A. Finally, the use of peak shape information not only makes the peak mapping more robust to noise and signal overlap, but also provides an opportunity to generalize our program for evaluating other types of spectral changes. An example is the transferred NOE spectroscopy (Moore, 1999), where changes of peak intensities instead of peak locations are monitored.

Acknowledgments

The authors acknowledge Drs Jens Quant and Lesley MacLachlan for stimulating discussions and suggestions. We thank Drs Mary Donlan, Sunil Patel and Zeljko Dzakula for their valuable suggestions. We also thank the anonymous referees for their suggestions to include the analysis of the effect of noise on the reproducibility of D

values and to use the fit of the displacements to an expected mathematic function as a possible future improvement of the PROPET algorithm. M.S. acknowledges support by the Deutsche Forschungsgemeinschaft.

Abbreviations

APET	automated peak tracing
<i>D</i>	the score (or overall changes) of a test spectrum
<i>d</i>	chemical shift displacement of a matched test peak
<i>d_r</i> and <i>d_t</i>	the contribution of an unmatched reference or test peak to <i>D</i> , respectively
HSQC	heteronuclear single quantum correlation spectroscopy
HTS	high-throughput screening
<i>K_d</i>	dissociation constant
max <i>H</i> and max <i>N</i>	the upper limit of ¹ H or ¹⁵ N chemical shift displacement, respectively
PROPET	progressive peak tracing
<i>R_I</i>	the ratio of intensity level of the test spectrum vs. that of the reference spectrum
<i>S</i>	similarity of peak shape
SAR by NMR	structure-activity relationships by nuclear magnetic resonance

References

- Accelrys Inc. FELIX User Guide, Version 2002. San Diego: 2002.
- Bax A, Ikura M, Kay LE, Torchia DA, Tschudin R. J. Magn. Reson. 1990; 86:304–318.
- Bodenhausen G, Buben DJ. Chem. Phys. Lett. 1980; 69:185–189.
- Chen Y, Reizer J, Saier MH Jr, Fairbrother WJ, Wright PE. Biochemistry. 1993; 32:32–37. [PubMed: 8418852]
- Chylla RA, Markley JL. J. Magn. Reson. Ser. B. 1993; 102:148–154.
- Clore GM, Schwieters CD. J. Am. Chem. Soc. 2003; 125:2902–2912. [PubMed: 12617657]
- Damberg CS, Orekhov VY, Billeter M. J. Med. Chem. 2002; 45:5649–5654. [PubMed: 12477348]
- Dominguez C, Boelens R, Bonvin AMJJ. J. Am. Chem. Soc. 2003; 125:1731–1737. [PubMed: 12580598]
- Douangamath A, Filipp F, Klein ATJ, Barnett P, Zou P, Voorn-Brouwer T, Vega M-C, Mayans OM, Sattler M, Distel B, Wilmanns M. Mol. Cell. 2002; 10:1007–1017. [PubMed: 12453410]
- Dzakula Z. J. Magn. Reson. 2000; 146:20–32. [PubMed: 10968954]
- Fahmy A, Wagner G. J. Am. Chem. Soc. 2002; 124:1241–1250. [PubMed: 11841293]
- Farmer BT II, Constantien KL, Goldfarb V, Friedrichs MS, Wittekind M, Yanchunas J Jr, Robertson JG, Mueller L. Nat. Struct. Biol. 1996; 3:995–997. [PubMed: 8946851]
- Gross, J.; Yellen, J. Graph Theory and its Applications. Boca Raton: CRC Press; 1998.

- Hajduk PJ, Dinges J, Miknis GF, Merlock M, Middleton T, Kempf DJ, Egan DA, Walter KA, Robins TS, Shuker SB, Holzman TF, Fesik SW. *J. Med. Chem.* 1997; 40:3144–3150. [PubMed: 9379433]
- Kirkpatrick S, Gelatt CD Jr, Vecchi MP. *Science*. 1983; 220:671–680. [PubMed: 17813860]
- Kumar, RA.; Bhakta, K.; Szalma, S.; Donlan, M.; Carter, B. An Integrated High Throughput Solution for SAR by NMR, 39th Experimental Nuclear Magnetic Resonance Conference; March 22–27, 1998; Asilomar (USA). 1998. p. 258Proceedings
- Lawler, EL.; Lenstra, JK.; Rinnooy Kan, AHG.; Shmoys, DB. *The Traveling Salesman Problem*. Chichester: John Wiley & Sons Ltd; 1985.
- Martin, ML.; Martin, GJ.; Delpuech, J-J. *Practical NMR Spectroscopy*. London: Heydon & Son Ltd; 1980. p. 45–47.
- McCoy MA, Wyss DF. *J. Biol. NMR*. 2000; 18:189–198.
- Moore JM. *Biopolymers*. 1999; 51:221–243. [PubMed: 10516573]
- Morelli X, Palma PN, Guerlesquin F, Rigby AC. *Protein Sci.* 2001; 10:2131–2137. [PubMed: 11567104]
- Muskett F, Frenkiel TA, Feeney J, Freedman RB, Carr MD, Williamson RA. *J. Biol. Chem.* 1998; 273:21736–21743. [PubMed: 9705310]
- Orehhov VY, Ibraghimov IV, Billeter M. *J. Biomol.NMR*. 2001; 20:49–60. [PubMed: 11430755]
- Otting G. *Curr. Opin. Struct. Biol.* 1993; 3:760–768.
- Pellecchia M, Stevens SY, Vander Kooi CW, Montgomery DH, Feng EH, Gierasch LM, Zuiderweg ERP. *Nat. Struct. Biol.* 2000; 7:298–303. [PubMed: 10742174]
- Peng C, Yuan S, Zheng C, Shi Z, Wu H. *J. Chem. Info. Comput. Sci.* 1995; 35:539–546.
- Press, WH.; Teukolsky, SA.; Vetterling, WT.; Flannery, BP. *Numerical Recipes in Fortran 77: The Art of Scientific Computing*. 2nd. Cambridge: Cambridge University Press; 1992. p. 436–443.
- Ross A, Schlotterbeck G, Klaus W, Senn H. *J. Biomol. NMR*. 2000; 16:139–146. [PubMed: 10723993]
- Sattler M, Schleucher J, Griesinger C. *Prog. NMR Spectrosc.* 1999; 34:93–158.
- Shuker SB, Hajduk PJ, Meadows RP, Fesik SW. *Science*. 1996; 274:1531–1534. [PubMed: 8929414]
- Van Nuland NAJ, Kroon GJA, Dijkstra K, Wolters GK, Scheek RM, Robillard GT. *FEBS*. 1993; 315:11–15.
- Williamson RA, Carr MD, Frenkiel TA, Feeney J, Freedman RB. *Biochemistry*. 1997; 36:13882–13889. [PubMed: 9374866]
- Zuiderweg ERP. *Biochemistry*. 2002; 41:1–7. [PubMed: 11771996]

Figure 1.

Common situations where the nearest peak-based approach may fail to find the correct matching between the reference peaks (in solid contours) and test peaks (in dotted contours). The hypothetical reference peaks **1** and **2** are shifted to locations of the test peaks **a** and **b**, respectively, in the test spectrum as indicated by the solid arrows. (A) If the peak mapping starts from peak **2**, the closest test peak **a** will be matched to it, and hence **b** will be matched to **1**. This leads to an erroneous, sequence-dependent peak mapping shown as the dashed arrows. (B) If a noise peak **c** appears closer to peak **1** than to **a**, the noise will be matched to **1** and hence **a** to **2**, leaving **b** unmatched.

Figure 2. Schematic illustration of the peak mapping problem for chemical shift perturbation spectra. The solid circles represent the reference peak set P , the dotted circles the test peak set T , and the edges E the mapping between the matched reference and test peaks. The reference peak marked by a cross represents an unmatched reference peak. The test peak marked by a square represents an unmatched test peak.

Figure 3.

Resolving overlapped test peaks by fitting control peaks to the test spectrum. The well separated control peaks **1** and **2** become overlapped in the test spectrum, hence only the peak top **a** is picked and matched to **2**, while **1** remains unmatched. After fitting **1** and **2** to the envelope, the resolved peaks (dotted lines) **b** and **a'** are matched to **1** and **2**, respectively.

Figure 4.

Overlay view of Test7 (violet and cyan contours for positive and negative peaks, respectively) and Reference (black and yellow for positive and negative peaks, respectively) ^1H - ^{15}N HSQC spectra of Pex13p SH3 domain. The green arrows show the peak displacements with $d > 0.15$ (Equation 1, whereas $S \equiv 1$) resulting from the APET analysis of Test7 vs. Reference. The peaks marked by the green crosses or green boxes are unmatched reference peaks or unmatched test peaks, respectively, from the APET analysis. The red arrows show the peak displacements that were corrected by the PROPET analysis of

all titration experiments. The purple arrows show the three pairs of peaks that the program failed to match. Peaks are marked by their assignments or, if not assigned by their peak IDs. Areas in the red dashed boxes are detailed in Figure 5.

Figure 5. Multiple overlay view of some significantly displaced areas in the titration ^1H - ^{15}N HSQC experiments of Pex13p SH3 domain with Pex14p peptide. For clarity, only selected experiments are shown and up to two contours are drawn for each experiment in the designated color. The experiment names and their corresponding colors are indicated in each box. The analysis results are shown in the same way as in Figure 4. Black crosses are used to indicate the centers of the some reference peaks in crowded areas.

Figure 6.

D values of the APET analysis on selected spectra with various artificially lowered S/N ratios. (A) *D* values include the contribution from both matched peaks and unmatched peaks. (B) *D* values include only the contribution of matched peaks. Diamonds denote the Reference experiment, mimicking a non-binding experiment. Squares denote the Test1 experiment, mimicking a partially binding experiment. Triangles denote the Test7 experiment, mimicking a fully binding experiment.

Peak displacements [ppm] and shape similarities from the titration experiments of Pex13p SH3 domain (0.5 mM) with the Pex14p peptide (0.1–1.0 mM)^a

Table 1

Peak ID ^b	Assign-ment ^b	Titration experiments and [Pex14p] in mM									
		Test1 (0.1)	Test2 (0.2)	Test3 (0.4)	Test4 (0.5)	Test5 (0.6)	Test6 (0.75)	Test7 (<i>d</i> and <i>S</i> ^c) (1.0)			
2	Asp304	<i>_d</i>	<i>_d</i>	0.45	0.48	0.49	0.51	0.54	0.54	0.58	
14	Phe317	0.08	0.13	0.16	0.18	0.18	0.18	0.19	0.19	0.72	
18	Glu323	0.12	0.19	0.26	0.26	0.27	0.27	0.27	0.27	0.37	
19	Met324	0.0	0.12	0.15	0.18	0.19	0.18	0.19	0.19	0.44	
23	Leu328	0.12	0.20	0.27	0.28	0.28	0.30	0.30	0.30	0.83	
34	Lys339	0.06	0.14	0.17	0.18	0.18	0.18	0.18	0.18	0.71	
42	Asp348	0.08	0.15	0.19	0.20	0.20	0.21	0.21	0.21	0.61	
43	Trp349	0.11	0.17	0.22	0.22	0.25	0.25	0.25	0.25	0.69	
55	Ile362	<i>_d</i>	0.43	0.53	0.55	0.57	0.59	0.60	0.60	0.81	
59	Tyr366	0.11	0.21	0.27	0.28	0.28	0.31	0.31	0.31	0.59	
60	Ile367	0.14	0.23	0.30	0.31	0.33	0.34	0.34	0.34	0.77	
73	(7.06, 113.4)	0.26	0.25	0.32	0.33	0.33	0.34	0.34	0.34	0.30	

^aOnly matched peaks with displacements greater than 0.17 ppm in Test7 are shown. Displacements are calculated according to Equation 1 with $\alpha = 1$, $\beta = 0.2$ and $S \equiv 1$. Peak similarities are shown only for Test7.

^bThe peak assignments (or peak IDs for unassigned peaks) correspond to the peak labels in Figures 4–5.

^cPeak similarities *S* are calculated according to Equation 2 and are not included in the peak displacements. Note the peak heights have been calibrated using $R_1 = 1.91$ based on Equation 10.

^dPeaks only visible at very low contour levels.