ARTICLE

On projection-reconstruction NMR

Clark D. Ridge · Vladimir A. Mandelshtam

Abstract Three most simple Projection-Reconstruction algorithms, namely, the Lowest-Value, Additive Back-Projection and Hybrid Back-Projection/Lowest-Value algorithms, are analyzed. A new, also simple, algorithm that reconstructs the spectrum by utilizing the amplitude histogram at each reconstruction point, is explored. The algorithms are tested using simulated spectra. While all the algorithms considered can potentially result in substantial reduction of the amount of data needed for reconstruction, they can suffer from a number of drawbacks. In particular, they often fail when the spectra are noisy and/or contain overlapping peaks. When compared to the existing algorithms, the new, histogram-based algorithm has the potential advantage of being able to deal with spectra containing peaks of opposite phase.

Projection-reconstruction (PR) NMR has been proposed as a way of reconstructing a multi-dimensional spectrum using reduced-dimensional spectra, which are referred to as projections. One reason which makes the PR methodology very appealing to the NMR community is a potentially enormous reduction of the amount of data needed to produce well-resolved spectra. Another reason is the possibility to use very simple processing algorithms (e.g., compared to those required for processing of general non-uniformly sampled data sets). Even though substantial data size reductions using PR methodology applied to real NMR data have been reported in a number of publications, unfortunately, for truly realistic (i.e., noisy and/or crowded) NMR spectra, this potential advantage of the PR

methodology often turns out to be illusory. However, due to the advances in improving the experimental sensitivity, there is always a hope, especially for certain specifically designed NMR experiments, that these methods may still be very useful.

The simplest PR case corresponds to reconstruction of a 3D spectrum from 2D plane projections. Because the directly detected dimension is treated independently, this case can be broken down into a series of 2D slices being reconstructed from 1D projections. Detailed discussions of the methodology of collecting the data are given in refs. (Kupče and Freeman 2003, Kupče and Freeman 2004a, 2004b). Here we present an analysis of three existing algorithms for the reconstruction of NMR spectra from projections and propose a fourth. For simplicity the analysis focuses on the reconstruction of 2D planes from 1D projections. All higher dimensional cases can be understood as an extension of this case. The algorithms discussed here are all simple and deterministic, and capable of reconstructing any desired point that falls in the range of the projections. The first algorithm discussed is the Lowest-Value (LV) reconstruction, followed by the Additive Back-Projection (BP) algorithm (Kupče and Freeman 2003, 2004a). A hybrid of the previous two algorithms called the Hybrid Back-Projection/Lower-Value (HBLV) algorithm (Venters et al. 2005) is then discussed. Finally, a new algorithm is proposed and its characteristics compared with the existing algorithms.

The projection–reconstruction problem

Given L-dimensional time-domain signal $c(\vec{t}) \equiv c(t_1, \ldots, t_L)$ available at some discrete subset of points $\vec{t} \equiv$

C. D. Ridge (✉) · V. A. Mandelshtam
Chemistry Department, University of California at Irvine, Irvine, CA 92697, USA
e-mail: mandelsh@uci.edu

$(t_1, \ldots, t_L)$, in principle, one is interested in estimating the full-dimensional spectrum $S(\vec{\omega})$ (with $\vec{\omega} \equiv (\omega_1, \ldots, \omega_L)$) by solving the inverse problem

$$\int d^L \vec{t} e^{i t \vec{\omega}} S(\vec{\omega}) = c(\vec{t}). \tag{1}$$

A number of approaches have been proposed in the past to solve this problem. Mathematically rigorous spectral inversion techniques are usually complicated with perhaps one exception of the discrete Fourier transform (DFT) of the truncated time-domain data.

In order to reconstruct the full spectrum, the simple PR algorithms utilize the information contained in the radially-sampled data. The latter consists of 1D spectral projections corresponding to 1D DFT's evaluated along certain directions defined by unit vectors $\vec{s}_i = (\cos \alpha_{i,1}, \ldots, \cos \alpha_{i,L})$ ($\|\vec{s}_i\| = 1$) :

$$P_i(\omega) = \int dt e^{-it\omega} c(t\vec{s}_i). \tag{2}$$

We note here that for a sufficiently dense set of data points, any non-uniformly sampled time-domain signal $c(\vec{t})$, as well as a radially-sampled one (as in Eq. 2), can be inverted by directly evaluating the inverse Fourier transform using the quadrature resulting from the non-uniform grid provided. However, for a sparsely sampled data such a direct inversion is very inaccurate, while an accurate inversion (if possible at all) should at least be very non-trivial. The simple PR algorithms are intended to invert sparse radially-sampled data. None of such algorithms are rigorous: they are rather intuitive, and are not designed to solve Eq. 1 correctly. Moreover, they may easily fail for crowded and/or noisy spectra. However, for spectra with relatively high SNR and sharp well separated peaks, where the lineshapes and amplitudes are not very important, they may be useful and efficient in terms of the total amount of data needed to obtain high spectral resolution.

Considering a 2D reconstruction problem, effectively corresponding to a 3D NMR experiment, a single line (in the $\omega_1 - \omega_2$ plane) is characterized by two pairs of frequencies and widths, (v_1, v_2) and (γ_1, γ_2). The position of the line in an α-projection, defined by angles $\alpha_{i,1} = \alpha$ and $\alpha_{i,2} = \pi/2 - \alpha$ will then be given by

$$v_\alpha = v_1 \cos \alpha + v_2 \sin \alpha, \tag{3}$$

and the linewidth will satisfy a similar formula,

$$\gamma_\alpha = \gamma_1 |\cos \alpha| + \gamma_2 |\sin \alpha|, \tag{4}$$

(i.e., the width is always positive and varies within some finite interval). The latter expression implies that the width of the projected line is α-dependent. Since the integral weight of the line,

$$\int d\omega P_i(\omega), \tag{5}$$

is preserved (i.e., does not depend on α), the height of the line in the projection will be inversely proportional to its width. This inconsistency in the line heights for projections at different angles makes it difficult to treat them on equal footing. However, by convoluting the spectra in all the projections with sufficiently broad Gaussians of the same width, one can make the projections of the same line at different angles more uniform (albeit for the price of reducing the resolution). In our further analysis we assume that such a broadening is always implemented. Yet, one still has to keep in mind that none of the algorithms discussed below are able to faithfully reproduce the shape of a line, as these algorithms all take a similar approach to the problem; namely, they all try to reconstruct the spectrum at each point independently of the neighboring points.

The Lowest-Value algorithm

The LV algorithm is the most aggressive in terms of the artifact suppression, yet the simplest approach to reconstruction, in which at each point $\vec{\omega}$ the spectral amplitude is estimated by selecting the lowest-value out of the set of the n available projections:

$$S^{\mathrm{LV}}(\vec{\omega}) = \mathrm{sgn}[P_i(\omega)] \min_{i=1}^{n} |P_i(\omega)|, \quad \omega = \vec{\omega} \cdot \vec{s}_i, \tag{6}$$

where $\mathrm{sgn}(x) = x/|x|$.

The additive back-projection algorithm

The second algorithm is similar to the first in simplicity, but in an effort to gain in signal to noise ratio (SNR), all the values at each point are summed:

$$S^{\mathrm{BP}}(\vec{\omega}) = \frac{1}{n} \sum_{i=1}^{n} P_i(\omega), \quad \omega = \vec{\omega} \cdot \vec{s}_i. \tag{7}$$

The drawbacks of both reconstruction formulas, 6 and 7, are well documented (Kupče and Freeman 2004a, 2004b, Venters et al. 2005, Yoon et al. 2006) While Eq. 6 discriminates well against false-positive peaks, it generally results in a poor SNR, as picking the smallest value (out of n available values) cannot take advantage of signal accumulation by combining the information from all the n available data sets. Moreover, due to the finite SNR, the spectrum gets worse as more projections are used.

For demonstration purposes, we made a model spectrum of four peaks, well separated in two dimensions and arranged around the center (see Fig. 1). The parameters of
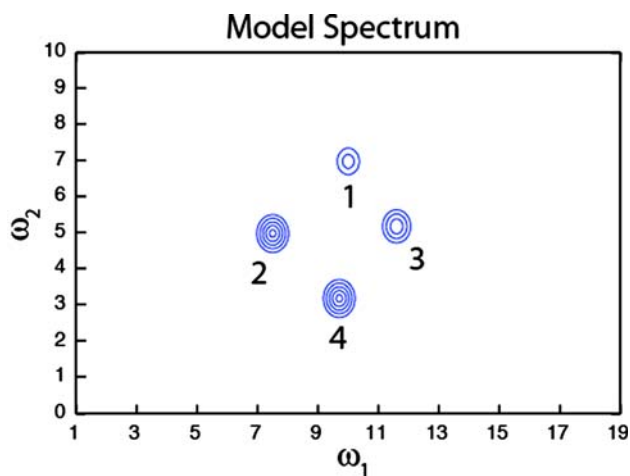
Fig. 1 The model spectrum constructed using the peak parameters from Table 1

Table 1 The parameters for the peak positions (v_1, v_2), widths (γ_1, γ_2) and amplitudes (d) in the model 2D spectrum shown in Fig. 1

peak	v_1	v_2	γ_1	γ_2	d
1	10.0	7.0	0.2	0.2	0.2
2	7.5	5.0	0.2	0.2	1
3	11.5	5.2	0.2	0.2	0.5
4	9.7	3.2	0.2	0.2	1

the peaks are given in Table 1. For simplicity, the widths of all the peaks are the same. They are also the same in both dimensions.

In this and all other numerical experiments the projection angles were always evenly dispersed in the interval [$-\pi/2;\pi/2$], with angles $\alpha = 0$ and $\alpha = \pi/2$ included.

The projections all contained 1,000 points so that simple interpolation could be used in the reconstruction process. Gaussian noise was added so that twice the standard deviation of the noise was 10% of the largest non-overlapping peak height in the $\alpha = 0$ and $\alpha = \pi/2$ projections. All four signals had greater than 1 signal to noise ratios, the smallest one being SNR = 2. Gaussian noise was always included into the simulated projections unless noted otherwise.

Reconstructed spectra are usually not smooth, which makes them hard to contour. Therefore, before making contourplots the spectra were convoluted with Gaussians. All the spectra were reconstructed to a grid with 300 points in each dimension. Contours in all of the spectra started at a level above significant noise, but low enough to see the peaks of interest where possible.

Figure 2 shows spectra reconstructed by the LV algorithm using different number, n, of projections. The LV-reconstructed spectra are well-resolved for small number ($n = 4$) of projections. That is, for sparse spectra consisting

of narrow and well separated peaks and relatively high SNR this reconstruction technique is very efficient. However, when n is increased the smallest peak breaks into pieces, and the overall spectral resolution worsens. On average, the spectral magnitude reconstructed by Eq. 6 is reduced with increasing n. At the same time, the noise level is also reduced. Depending on the noise statistics and the actual peak height, the apparent SNR for each peak may or may not improve with increasing n; at the same time, due to the noise fluctuations the peak lineshapes become rough. Even though this roughness can partially be removed by convoluting the reconstructed spectrum with a Gaussian, sufficiently small peaks would still break into pieces even after smoothing.

Unlike the LV reconstruction, Eq. 7 does accumulate the signals from all the n available projections, albeit for the price of producing false cross-peaks formed by the BP ridges. This can be seen in the lower right panel of Fig. 3.

The Hybrid Back-Projection/Lowest-Value algorithm

In the HBLV algorithm as presented in refs. (Venters et al. 2005, Coggins and Zhou 2006) rather than taking the lowest-value or summing the values, a number $k < n$ is selected. Every combination of k projections is summed. Of all these combinations, the lowest at each point is selected:

$$S^{\mathrm{HBLV}}(\vec{\omega}) = \frac{1}{k}\,\mathrm{sgn}\left[\sum_{P_i \in A_j}^{n} P_i(\omega)\right] \times \min_{j=1}^{{}_nC_k}\left|\sum_{P_i \in A_j}^{n} P_i(\omega)\right|,$$
$$\omega = \vec{\omega}\cdot\vec{s}_i,$$
(8)

where A_j ($j = 1, …, {}_nC_k$) define all the possible choices of k projections $P_i(\omega)$ out of n available ones. As argued in refs. (Venters et al. 2005, Coggins and Zhou 2006), due to the sum over k projections, expression (8) does take some advantage of signal accumulation and thus has better SNR properties than that of the LV reconstruction, while the ridge and cross-peak artifacts are still removed by the minimization step. Figure 4 demonstrates the differences between the above three algorithms in how the reconstruction at some particular point is made by selecting a value out of the n available values.

The striking drawback of Eq. 8 is that numerically it is by about a factor of ${}_nC_k = n!/(n-k)!k!$ more expensive than the other two expressions, (6) and (7). For example, one possibility suggested by Coggins and Zhou (2006) corresponds to $n = 30$ and $k = 8$, in which case reconstruction by Eq. 8 at each spectral point requires to sample as many as ${}_{30}C_8 \approx 6 \times 10^6$ terms. In the case of 4D spectral reconstruction, this can make a computer cluster busy for several days. Although the authors of PR-CALC
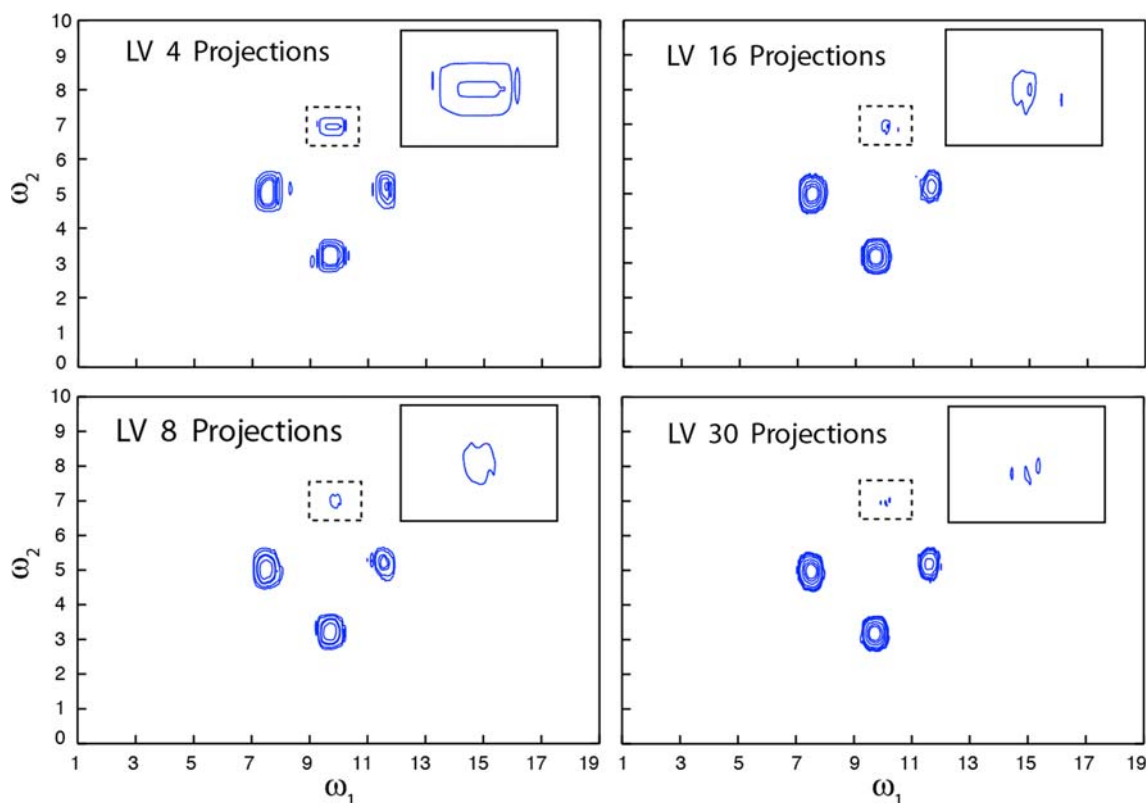
Fig. 2 Spectra reconstructed by the LV algorithm from radially-sampled data with SNR = 10 using 4, 8, 16 and 30 projections. The exact noiseless spectrum is shown in Fig. 1. All the spectra are smoothed as described in the text and the contour-levels are adjusted to adequately represent all the peaks. An enlarged view (300%) of the area enclosed in the dashed line is shown inset. As more projections are used, both the amplitudes of the reconstructed peaks and the noise level decrease with the same rate. However, as the number of projections grows the smaller peaks break into pieces due to the noise fluctuations

(Coggins and Zhou 2006) did everything to carefully optimize their code for best performance, they find that because of the $_nC_k$ factor the use of more than about 30 projections may become prohibitive.

At this point, we distinguish the following two cases. In the first case, for which the PR-NMR techniques are perhaps most suitable, the spectrum $S(\vec{\omega})$ is assumed to be positive, except for negative but small noise spikes. In such a case the negative spikes in the projections $P_i(\omega)$ can be removed before further processing.

Without loss of generality we assume that the absolute values of the amplitudes $P_i(\omega)$ have been sorted in the ascending order.

$$|P_1(\omega)| \leq |P_2(\omega)| \leq \ldots \leq |P_n(\omega)|. \tag{9}$$

Clearly, for all positive projections, the minimum arising in Eq. 8 is given by the sum of the k lowest $P_i(\omega)$ values:

$$\min_{j=1}^{_nC_k} \sum_{P_i \in A_j}^{n} P_i(\omega) = \sum_{i=1}^{k} P_i(\omega). \tag{10}$$

So, for spectra with no negative peaks the HBLV algorithm boils down to a trivial calculation, which can be accomplished by many orders of magnitude faster than that

suggested by Eq. 8. In words, for each frequency grid point of interest the value of the reconstructed spectrum is set to be equal to the arithmetic mean of the k smallest (out of n available) projection values. The modified (i.e., using Eq. 10) HBLV spectral reconstruction algorithm is demonstrated in Fig. 3 for $k = 8$, 16, 24, 30 and $n = 30$, with the $k = 30$ case formally corresponding to the BP algorithm. The best spectra are obtained using $k = 8$ and $k = 16$, while at $k = 24$ the ridge artifacts start to show above the level of the lowest peak.

The second case corresponds to spectra that may have genuine negative peaks. In this case, the minimum in Eq. 8 is not necessarily given by the k projections with smallest magnitudes $|P_i(\omega)|$, because of possible cancellations of negative and positive contributions. However, these cancellations are also the reason for Eq. 8 being not a meaningful reconstruction formula, when negative peaks are encountered. Let us demonstrate this statement using a simple example shown in Fig. 5. In this demonstration the true spectrum consists of only four peaks: a negative peak at point $(\vec{\omega}_1)$ with amplitude $S_1 = -1$ units, and three positive peaks with amplitudes $S_2 = 7$, $S_3 = 8$, $S_4 = 9$ units, located at three other positions. We are interested in
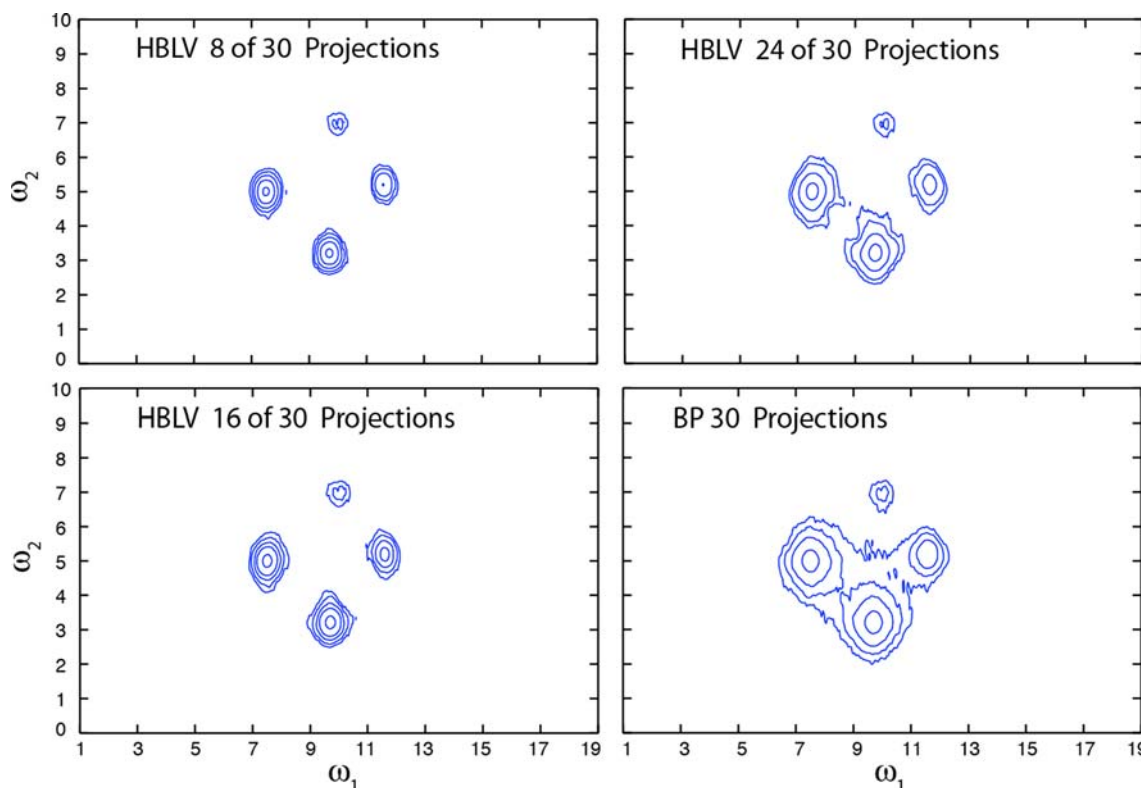
Fig. 3 The modified HBLV algorithm (cf. Eq. 10) is used here to reconstruct the model spectrum with $k = 8, 16, 24, 30$ out of $n = 30$ projections. For the $(k = 24, n = 30)$ case the artificial ridges start to show above the level of the lowest peak. The $(k = 30, n = 30)$ case (the lower right spectrum) is formally equivalent to using the BP algorithm
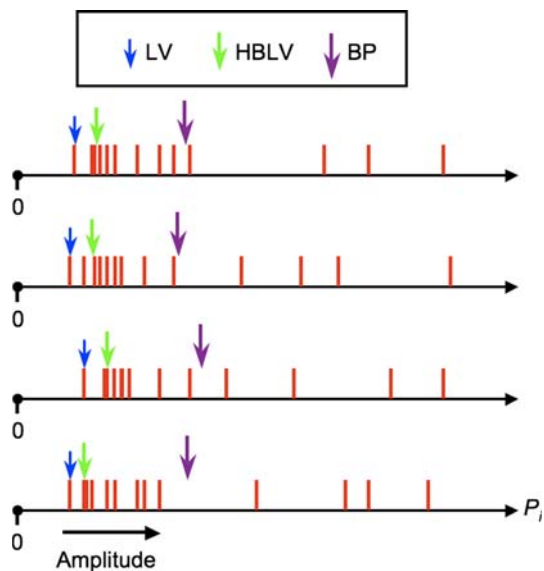

Fig. 4 For the LV algorithm the lowest-value at each point is selected; in BP the average of the values is selected; and in HBLV the lowest partial average of the projections is selected as indicated by the arrows

recovering the spectrum at $(\vec{\omega}_1)$ assuming infinite SNR in all the projections. There are three projection angles at which the negative peak overlaps with one of the positive peaks. This may result in the following set of projection amplitudes:

$$P_i(\omega) \in \{6, 7, 8, -1, -1, -1, \ldots\}, \quad \omega = \vec{\omega}_1 \cdot \vec{s}_i,$$

where we assumed that all projections, except the first three, are due to the single negative peak. Following the suggestion by Coggins and Zhou (2006) we set $k = 8$ and apply Eq. 8 to reconstruct the spectrum at point $(\vec{\omega}_1)$, where the negative peak is situated:

$$S^{\text{HBLV}}(\vec{\omega}_1) = P_2(\omega) + \sum_{j=4}^{10} P_i(\omega) = 7 + 7 \cdot (-1) = 0.$$

It is not hard to see that the HBLV reconstruction using $k = 7$ or $k = 9$ will also result in $S^{\text{HBLV}}(\vec{\omega}_1) = 0$.

At first glance our example may seem dishonest as it was carefully designed (or "cooked-up") to make the method fail. However, we argue that in practice the situation is even worse due to the presence of a large number of positive peaks with different amplitudes overlapping with small negative peaks and thus resulting in many possibilities for similar-type cancellations. Moreover, increasing the number of projections n will only increase the chances that there will be a combination of k negative and positive terms summing to some small value below the noise level.
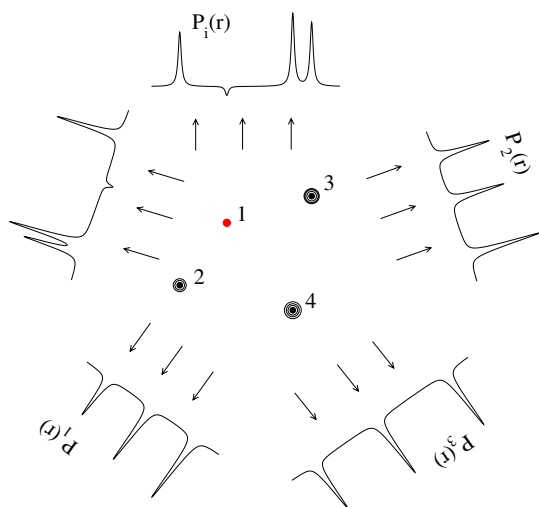
Fig. 5 Demonstration of a failure of the HBLV algorithm to reconstruct peaks with negative amplitudes (see text). The spectrum consists of one negative peak with amplitude $S_1 = -1$ and three positive peaks with amplitudes $S_2 = 7$, $S_3 = 8$ and $S_4 = 9$. The negative peak overlaps with one of the positive peaks in projections $P_i(\omega)$ ($i = 1,2,3$)

Once we are convinced that Eq. 8 does not present advantages for mixed spectra, i.e., having amplitudes of opposite sign, it becomes apparent that the problem of combinatorial search for the smallest sum is superficial. For a mixed case, the modified procedure (Eqs. 9 and 10) results in the spectrum, which is an upper-bound of that given by Eq. 8. The modified HBLV is still a valid PR algorithm, but its status is somewhat similar to that of Eq. 6: in the $n \to \infty$ limit with fixed value of k the HBLV spectrum will still suffer from the loss of sensitivity. However, depending on the k/n ratio, the SNR of the HBLV spectrum will be better than that of a single LV spectrum.

We also note ref. (Yoon et al. 2006) discussing a variety of PR methods. For example, one of the proposed extensions of the LV algorithm (similar to that using Eqs. 9 and 10 is to divide the set of n projections into k groups and apply Eq. 6 to each group independently. The resulting k LV estimates are then averaged to obtain an estimate that has better SNR properties than a single LV spectrum.

The histogram method

Inspired by the previous example, here we propose yet another deterministic PR algorithm that is simple, numerically inexpensive, but applicable to the case of negative amplitudes. Namely, for a given reconstruction point $(\vec{\omega},)$ consider the set of the corresponding projection values $P_i(\vec{\omega} \cdot \vec{s}_i)$ ($i = 1, , n$) and the distribution function $g(P)$ estimated from their histogram. Assuming that $g(P)$ has a
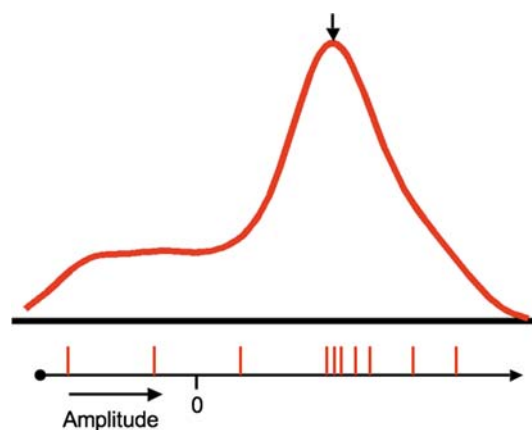

Fig. 6 Demonstration of the new histogram-based PR algorithm: instead of taking the lowest amplitude out of a set of amplitudes or averaging their values, the most likely amplitude is determined by finding the maximum of the amplitude distribution

maximum at $P = \hat{P}$, the reconstructed spectrum can be estimated by setting $S(\vec{\omega}) = \hat{P}$. One way to accomplish this idea is to use the expression:

$$g(P) = \sum_i \exp\left[-\frac{(P - P_i(\vec{\omega} \cdot \vec{s}_i))^2}{2\sigma^2}\right],\tag{11}$$

where the Gaussian width σ is an adjustable parameter. Clearly, its choice depends on the number of projections and the dynamic range of the amplitudes. σ must be large enough to make the resulting distribution $g(P)$ at each point $(\vec{\omega})$ sufficiently smooth in order to remove the ambiguities associated with the selection of its maximum $\hat{P}$ (see Fig. 6). The existence of an adjusting parameter makes the Histogram Method less deterministic and may be regarded as a disadvantage, e.g., when compared to the LV and BP algorithms. (Note that HBLV also has an additional adjusting parameter, k, whose choice is not obvious.) Although other recipes may be used, in the present work parameter σ was chosen according to

$$\sigma = \frac{2}{n}(P_{max} - P_{min}),\tag{12}$$

where P_{max} and P_{min} are, respectively, the maximum and minimum values of the projection amplitudes at a given reconstruction point. This choice worked well for the examples considered.

The value from the sum of the amplitude distributions does not depend on adding the amplitudes themselves, and therefore, there is no way for the amplitude values to cancel and cause severely attenuated or missing peaks. The sensitivity of the Histogram Method improves when more projections are used, because in this case one has better amplitude statistics at each reconstruction point and is able to produce a more accurate distribution function from the amplitude histogram. The method should generally not

produce artificial ridges that usually accompany the BP algorithm. However, as is the case with the other three simple PR algorithms, the Histogram Method does not preserve the lineshapes and it does not address some of the other issues with PR such as dealing with crowded spectra or assessing the quality of the reconstruction.

The Histogram Method is first demonstrated and compared to the other methods using the model spectrum described in Table 1, except that the bottom and right peaks (numbered 3 and 4) are now negative. The results of the reconstruction using $n = 30$ simulated projections are shown in Fig. 7. The appearance of the spectra depends on

the choice of the contour-levels: the presence of low-level contours would show artifacts (false-positive peaks), while the desire to hide the artifacts below the lowest contour-level results in the false-negative peaks. In Fig. 7 we have chosen the latter option. While the new approach can easily resolve both negative and positive peaks, the other three methods (LV, BP, and HBLV) have difficulties, particularly in resolving the smallest peak (1), so it would stand above the artifact/noise level. The BP method is barely able to reproduce peak 1, but the ridge artifacts make the overall BP spectrum inferior to that of the Histogram Method.
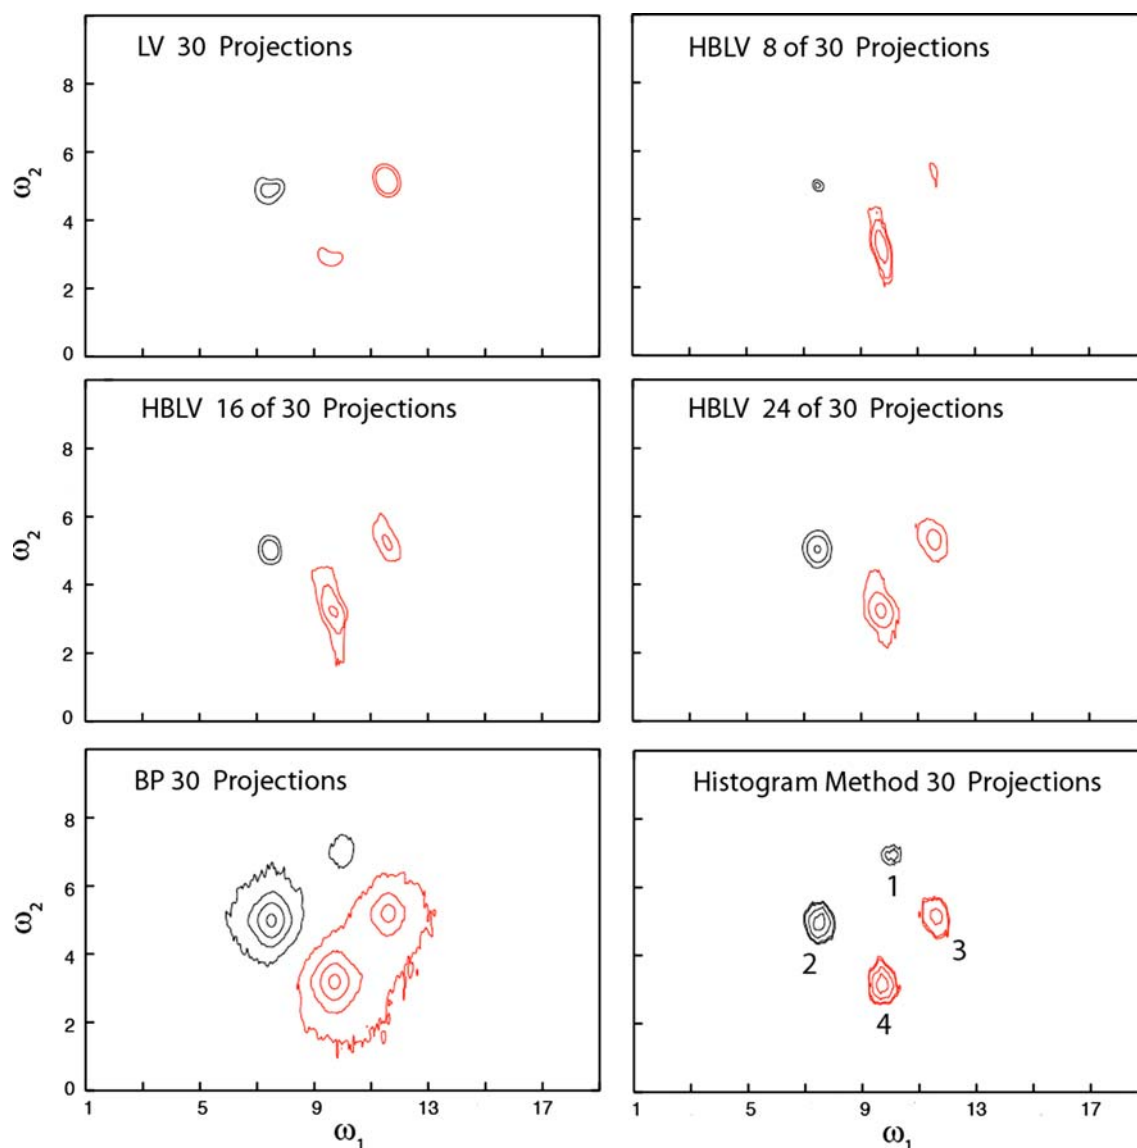

Fig. 7 Here the LV, modified HBLV, BP and the Histogram algorithms are used to reconstruct a model spectrum that is identical to that in Table 1 and Fig. 1 except that the bottom and right peaks (numbered 3 and 4) are now negative. The LV spectrum cannot repoduce the smallest of the peaks and requires double the gaussian smoothing to make the peaks contour. The modified HBLV algorithm, (cf. Eq. 10) has better results with larger k. The BP suffers from the same ridges as before and the peak intensities are attenuated down to the level of the artificial ridges
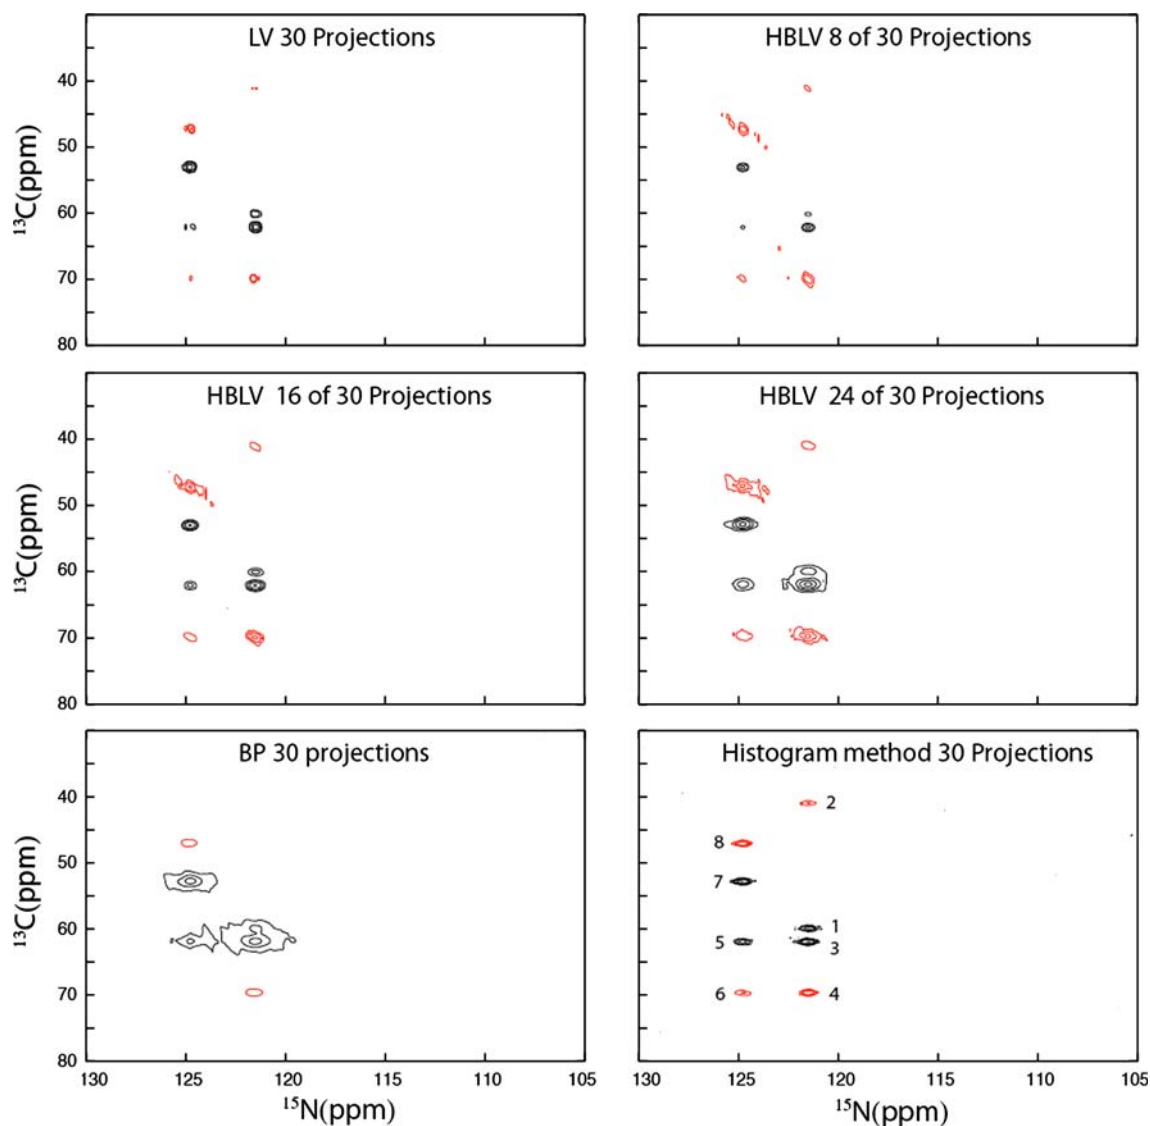
Fig. 8 Reconstruction of a simulated slice of an HNCACB spectrum of ubiquitin at 8.71 ppm in the proton dimension using the same algorithms and number of projections as in Fig. 7. The peak parameters are given in Table 2

As another demonstration of the Histogram Method, in Fig. 8 we consider a slice of simulated HNCACB spectrum of ubiquitin taken at 8.71 ppm in the proton dimension. The HNCACB experiment is one in which there are peaks of opposite phase (Wittekind et al. 1993, Sattler et al. 1998). The peak parameters used to simulate 30 projections are given in Table 2. As before, the projections were generated by adding 10% Gaussian noise. As the figure shows, the Histogram Method using 30 projections produces a reconstruction of high quality, even when the peaks are of opposite phase.

Conclusions

The algorithms tested here are all able to reconstruct model signals, but with varying degrees of success. The result can

Table 2 The peak parameters used in the simulated slice of the HNCACB spectrum of ubiquitin (Fig. 8)

peak	v_C	v_N	γ_C	γ_N	d
1	59.89	121.52	0.15	0.15	1.3
2	40.95	121.52	0.15	0.15	−0.5
3	61.94	121.52	0.15	0.15	4
4	69.65	121.52	0.15	0.15	−1.7
5	61.94	124.80	0.15	0.15	1.3
6	69.94	124.80	0.15	0.15	−0.5
7	52.83	124.80	0.15	0.15	4
8	47.07	124.80	0.15	0.15	−1.7

The chemical shifts were obtained from ref. (Ulrich et al. 2007) and other parameters were chosen using ref. (Sattler et al. 1998)

depend heavily on the noise level, number of peaks, number of projections, and other factors. The appearance of the reconstructed spectra can be much improved by the use of a Gaussian smoothing. This improvement comes with a small loss in resolution and a small attenuation in peak height. The HBLV algorithm can be processed in about the same time as all the LV or BP algorithms, when formulated correctly (i.e., using Eq. 10 rather than Eq. 8). A new histogram-based algorithm is shown to be capable of reconstructing spectra with improved or similar quality when compared to the existing algorithms, especially in the case of spectra with peaks of opposite phase.

Implementation of Eq. 11 followed by a maximization procedure within an existing PR software is straight-forward and requires only a few lines of a code that could replace a few lines of the code corresponding to, e.g., the LV algorithm. Reconstruction by the Histogram Method will become available through the Varian software in the near future.

The model examples used in this article are not sufficient to fully assess the advantages and disadvantages of the Histogram Method. Once the algorithm is implemented within the NMR software, it will be tested on real NMR data. Furthermore, the Histogram Method as presented here is simple, but perhaps not the most efficient spectral reconstruction algorithm based on the statistical analysis of the projected amplitudes. We anticipate that more sophisticated methods, which, in particular, are not restricted to the local statistical analysis, may be more efficient.

References

Coggins BE, Zhou P (2006) PR-CALC: a program for the reconstruction of NMR spectra from projections. J Biomol NMR 34:179–195

Kupče E, Freeman R (2003) Reconstruction of the three-dimensional NMR spectrum of a protein from a set of plane projections. J Biomol NMR 27:383–387

Kupče E, Freeman R (2004) The Radon transform: a new scheme for fast multidimensional NMR. Concept Magn Reson 22A:4–11

Kupče E, Freeman R (2004) Projection–reconstruction technique for speeding up multidimensional NMR spectroscopy. J Am Chem Soc 126: 6429–6440

Sattler M, Schleucher J, Griesinger C (1998) Heteronuclear multidimensional NMR experiments for the structure determination of proteins in solution employing pulsed field gradients. Prog NMR Spectrosc 34:93–158

Ulrich EL et al (2007) BioMagResBank. Nucl Acids Res. doi: 10.1093/nar/gkm957

Venters RA, Coggins BE, Kojetin D, Cavanagh J, Zhou P (2005) (4,2)D projection-reconstruction experiments for protein backbone assignment: application to human carbonic anhydrase II and calbindin D-28K. J Am Chem Soc 127:8785–8795

Wittekind M, Müller L (1993) HNCACB, a high sensitivity 3D NMR experiment to correlate amide-proton and nitrogen resonances with the alpha and beta carbon resonances in proteins. Magn Reson B 101:201–205

Yoon JW, Godsill S, Kupče E Freeman R (2006) Deterministic and statistical methods for reconstructing multidimensional NMR spectra. Magn Reson Chem 44:197–209