Journal of Biomolecular NMR **27:** 101–113, 2003.

Perspective

New methods for fast multidimensional NMR

Ray Freeman[a,*] & Eriks Kupče[b]
[a]*Jesus College, Cambridge, U.K.;* [b]*Varian, Inc, Eynsham, Oxford, U.K.*

Received 29 April 2003; Accepted 5 May 2003

Key words: filter diagonalization method, GFT-NMR, Hadamard spectroscopy, multidimensional, projection-reconstruction, single-scan two-dimensional NMR

Abstract

Considerable excitement has been aroused by recent new methods for speeding up multidimensional NMR experiments by radically modifying the normal time-domain sampling protocols. These new schemes include the filter diagonalization method, GFT-NMR, the single-scan two-dimensional technique, Hadamard spectroscopy, and a proposal based on projection-reconstruction of three-dimensional spectra. All these methods deliver appreciable improvements in the speed of data acquisition and show promise for speeding up multidimensional NMR of proteins. This perspective aims to describe these important new procedures in simple terms and to comment on their advantages and possible limitations.

Introduction

When two-dimensional spectroscopy was first introduced (Jeener, 1971; Aue et al., 1976) it was shown to have virtually the same sensitivity as one-dimensional measurements performed in the same total time, since all the NMR signals are gathered in each and every scan. The same considerations also apply to multidimensional spectroscopy. Recently attention has focused on the rather less obvious point that such spectra do not enjoy optimum efficiency in terms of speed. Systematic exploration of the indirect evolution dimensions one step at a time has the advantage that no signal frequencies are overlooked, but this is not the fastest mode of data acquisition for relatively sparse spectra. Time is actually wasted examining regions of frequency space where there are no signals, only noise. This realization has excited a recent flurry of activity to see whether multidimensional experiments could be speeded up, provided that the inherent sensitivity is already high enough. Since it is clear that multidimensional NMR spectroscopy of proteins can be seriously handicapped by the speed factor, these are important developments.

To fix ideas, imagine a three-dimensional array of time-domain data. The traditional approach would be to examine each and every element of this array in a systematic fashion, by following the evolution of the NMR signal as a function of t_1, t_2 and t_3 independently. By contrast, all the new methods seek to simplify this process by cutting down the number of individual measurements, either in the time domain or in the frequency domain. Sensitivity is reduced, but for many samples in high-field spectrometers, sensitivity is already more than adequate; it is speed that is the critical factor. These new methods go far beyond the usual compromises adopted in multidimensional spectroscopy – less-than-optimum resolution, the reduction or elimination of phase cycling, linear prediction of the tails in the evolution dimensions, and the acceptance of aliasing where this can be tolerated. Such short-cuts do help, but normally do not go far enough in reducing the duration of high-dimensional NMR spectroscopy; an expensive spectrometer can still be tied up for days on a single investigation.

The new methods reviewed here are all quite distinct, but because these ideas have been 'in the air'

*To whom correspondence should be addressed. E-mail: rf110@cus.cam.ac.uk

for some time, one mode of attack on the speed problem tends to suggest alternatives. All have advantages and drawbacks, and it is not yet apparent that a clear winner has emerged. It is hoped that this Perspective may go some way to demystifying these new ideas, without doing too much mischief to the underlying basic theory.

The filter diagonalization method

The filter diagonalization method (Mandelshtam, 2000, 2001; Hu et al., 2000; Chen et al., 2000) was proposed as an alternative to the ubiquitous Fourier transform methodology, but initially not as a method for speeding up multidimensional spectra. It is a data-processing technique related to the better-known linear prediction method (Barkhuijsen et al., 1985), and it starts from the premise that the inherent NMR sensitivity is high enough that the effects of noise can be neglected, and that the resonance lines are Lorentzian in shape. The former assumption is common to all the fast multidimensional methods, and the latter is not a really serious restriction, since, if necessary, an arbitrary lineshape could be constructed from a superposition of Lorentzians.

Linear prediction assumes that the NMR time-domain signal is a superposition of complex sinusoids that decay exponentially with time. In an ideal world, where there is only one resonance line in the spectrum and no noise, only two complex data acquisitions are enough to define the frequency of that line and its width. In real situations the presence of nearby resonances, non-Lorentzian lineshapes and the existence of random noise can render the method rather unreliable. Furthermore the number of component sinusoids (called the 'filter length') has to be assumed; if it is too large the calculation is too slow; if it is too small the fitting becomes unreliable. So linear prediction is generally used with caution, usually to extrapolate the tail of a time-domain signal to avoid truncation artifacts, or to predict the first few points of a free induction signal that might have been falsified by transmitter breakthrough. In the presence of noise, linear prediction can have quite unpredictable behaviour, thus refuting the all-too-common misconception that this processing technique can improve sensitivity.

Filter diagonalization shares with linear prediction the invaluable property that linear algebra is used, and it attains the same goal, but it does so by fitting the experimental time-domain data *as if* in the frequency domain. That is to say, it takes into account the undeniable fact that lines that are reasonably well separated in frequency show little or no interference effects. Filter diagonalization uses a matrix representation of an operator, in which interference effects between different components are represented by off-diagonal elements. It aims to derive a representation of the NMR spectrum by a procedure that does not involve Fourier transformation at all.

Our concern here is what the filter diagonalization method can accomplish, not primarily with the mathematical details. The key theoretical paper (Wall and Neuhauser, 1995) is entitled 'Extraction, through filter diagonalization, of general quantum eigenvalues or classical normal mode frequencies from a small number of residues or short-time segments of a signal'. An operator $\hat{U}$ represents a time autocorrelation function that has the effect of computing the next time-domain data point each time it is applied. By repeated application of this operator, a model for the NMR time-domain signal can be created which may then be fitted to the actual experimental signal. If $\hat{U}$ can be diagonalized, its powers can be written in terms of powers of the eigenvalues. The eigenvalues of $\hat{U}$ then give the line frequencies and widths, while the eigenvectors give amplitudes and phases. If the fitting process addressed all the time-domain data at once, the diagonalization problem would be immense but the decomposition into small segments renders it tractable.

The crucial innovation is in the choice of basis – a set of 'frequency-like' basis functions that are roughly localized in frequency space, for example, in a set of segments uniformly distributed across the NMR spectrum. This breaks down a potentially enormous generalized eigenvalue problem into a set of much smaller problems where unimportant matrix elements lying far from the diagonal can be safely neglected. Only a 'local' fitting procedure is then required, neglecting basis functions outside the particular segments under consideration, leaving much smaller 'filtered' matrices for diagonalization. Each frequency-domain segment is multiplied by a shaped window function, usually a cosine-squared envelope with 50% overlap between adjacent windows. Intense or very broad lines outside the window may require special treatment, and spectral windows containing a strong solvent peak are normally excluded from the process. The fitting procedure is not bedevilled by the usual problems of 'false' minima because linear algebra is employed, but even a perfect fit to the experimental NMR signal is

no absolute guarantee that the derived spectrum is reliable. This is particularly serious if the basis is too small, so it is necessary to fit the segments several times, using varying basis sizes. The segments are then assembled to give the full NMR spectrum.

In the extension to two frequency dimensions the basis is much larger, being determined by the product of the number of points in each dimension. Herein lies the great strength of the filter diagonalization method applied to two-dimensional spectroscopy – the resolving power depends essentially on the *area* in the two-dimensional time domain rather than on the individual lengths in the two time dimensions taken separately. Thus a very large number of experimental data points recorded in the t_2 dimension (acquired without a significant time penalty) can compensate for a much smaller number of data points in the t_1 dimension (where time is of the essence). The minimum requirement is that there should be substantially more basis functions in the two-dimensional window than there are NMR peaks in that area of the spectrum. This offers a considerable economy in instrument time compared with the conventional sampling regime. The only change to the traditional two-dimensional experiment is the reduction in the number of samples taken in the evolution dimension; all the rest is a data processing operation. Analogous considerations apply to higher-dimensional spectra.

Unfortunately the application to multidimensional spectroscopy introduces some serious new problems, some spectroscopic and some numerical. These difficulties include the phase-twist line shape (which poses more fitting problems than the Lorentzian shape assumed for one-dimensional spectra), the presence of t_1 noise, phase shifts induced by low-level instrumental noise, an inherent lack of correlation between the F_1 and F_2 co-ordinates of a given spectral feature, large changes in the computed spectrum generated by quite small perturbations of the input data, and complications when there are exactly degenerate frequencies, as in COSY multiplets. These problems are partly addressed by the introduction of a 'regularization' procedure (Chen et al., 2000), broadly related to the method proposed by Tikhonov (1963). A regularization parameter q^2 is introduced and increased from zero until the calculation becomes stable. One of the consequences of regularization is a broadening of noise components (and small signals comparable with the noise) to the point where they disappear from the computed spectrum. Some spectroscopists would regard this kind of non-linearity as unsatisfactory, even dangerous, but it has the useful side effect of improving the presentation of the contour map.

The filter diagonalization method has been applied to the two-dimensional constant-time (Powers et al., 1991) heteronuclear single-quantum correlation (HSQC) spectrum of ubiquitin (Chen et al., 2003). The principal aim of these measurements was to demonstrate an improvement in spectral resolution compared with conventional Fourier transform experiments under the same conditions. However it also permits a significant reduction in the constant-time parameter, which in the general case would improve sensitivity by reducing relaxation losses. It further speeds up data gathering by reducing the number of samples in the evolution dimension by a factor of approximately six. Figure 1 compares the traditional Fourier transform HSQC results on one part of the ubiquitin spectrum obtained with a constant-time parameter of 26.4 milliseconds with those from the filter diagonalization method recorded with this parameter reduced to 4.25 milliseconds. The latter spectrum, obtained in less than 4 min, corresponds very closely to the conventional Fourier transform spectrum except in the very crowded region. Ubiquitin has uncharacteristically slow spin-spin relaxation; for large proteins in general it may not be feasible to follow the free induction signals for long enough to obtain fine digitization in the direct frequency dimension.

The filter diagonalization scheme is a method of spectral estimation that offers an alternative to the standard Fourier transform methodology. It can speed up the data-gathering stage of multi-dimensional spectroscopy by concentrating the sampling operations in the direct detection dimension, thereby compensating for deliberately sparse sampling in the indirect evolution dimensions. The fitting process is not always completely reliable and may require some trial-and-error adjustment of the number of basis functions and of the 'regularization parameter', so it cannot be said to be a 'push-button' procedure at the present time. The time needed for the computation is generally much longer than that of the conventional Fourier transform protocol, although it can be reduced by confining the calculations to signal-bearing regions. Parallel processing is expected to reduce the computational time to a more acceptable level.

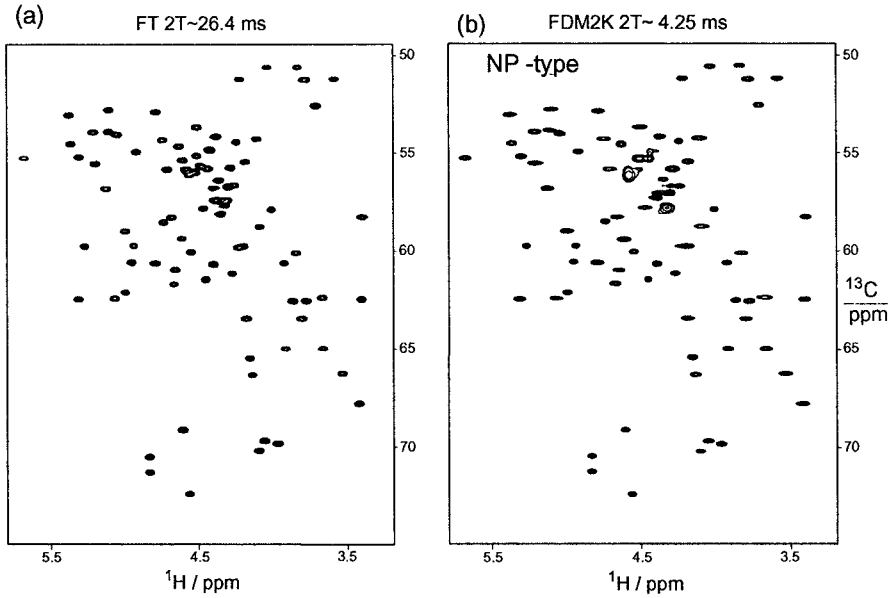

Figure 1. Part of the 500 MHz constant-time HSQC spectra of a 1 mM D_2O solution of isotopically labelled ubiquitin, showing the carbon-13-proton correlations. (a) The conventional Fourier transform spectrum with the constant time parameter CT = 26.4 ms. (b) The spectrum obtained by the filter diagonalization method with CT = 4.25 ms, and joint processing of N and P type data sets. Spectrum (b) required less than 4 min of data acquisition, but note the slight loss of resolution in the very crowded region. Spectra courtesy of A.J. Shaka.

GFT- spectroscopy

This technique attacks the speed problem of multidimensional NMR by incrementing several evolution periods jointly, either exactly in step with one another or with suitable scaling factors between them. This is clearly much quicker than incrementing these time dimensions independently. The 'reduced dimensionality' idea has evolved over quite a long period (Szyperski et al., 1993a, b, 1995, 1998, 2002; Brutscher et al., 1994, 1995; Simorre et al., 1994; Löhr and Rüterjans 1995; Ding and Gronenborn, 2002). Locking several evolutions together is the easy part; there remains the problem of disentangling the different chemical shift frequencies that evolve in the various 'locked' dimensions. The chemical shifts appear in the experimental time-domain data set as linear combinations, and they are extracted by a procedure based on a G-matrix (Kim and Szyperski, 2003).

The extraction technique is most easily understood in the frequency domain. Consider the simple example of an N-dimensional correlation experiment where the aim is to reduce the investigation to $N-2$ dimensions by incrementation of the first three evolution periods jointly and in step. In the simplest case where no scaling is used, this involves setting $t_1 = t_2 = t_3 = t$. At first sight this appears to be counter-productive, conflating signals from three different dimensions, but in fact the exploration of these dimensions can be rendered independent again by the GFT procedure. Let a typical chemical shift frequency in the first dimension be Ω_1 so that the normalized signal intensity transferred to the second evolution period is given by $\cos(\Omega_1 t)$, neglecting relaxation losses for simplicity. After evolving at the new chemical shift frequency Ω_2, it passes on a signal proportional to $\cos(\Omega_1 t)\cos(\Omega_2 t)$ to the third evolution period where it evolves at another new frequency Ω_3. At the end of the third evolution period the transferred signal is given by $\cos(\Omega_1 t)\cos(\Omega_2 t)\cos(\Omega_3 t)$. Further evolution stages, t_4, t_5, ..., etc., are monitored by the conventional method where the various time developments are followed independently.

Now comes the crucial trick. The entire sequence is repeated with a 90° phase shift of the radiofrequency pulse that initiates the second evolution period, so that the corresponding transfer function becomes $\sin(\Omega_2 t)$ rather than $\cos(\Omega_2 t)$. This stratagem is repeated in a succession of four 'scans' in which the radiofrequency phase shifts are systematically changed, keeping a fixed phase in the first evolution period:

scan (1) 0° 0° 0°,
scan (2) 0° 90° 0°,
scan (3) 0° 0° 90°,
scan (4) 0° 90° 90°.

This introduces cosine or sine modulation of the four transfer functions according to the products:

$\cos \Omega_1 t \ \cos\Omega_2 t \ \cos\Omega_3 t$,
$\cos \Omega_1 t \ \sin\Omega_2 t \ \cos\Omega_3 t$,
$\cos \Omega_1 t \ \cos\Omega_2 t \ \sin\Omega_3 t$,
$\cos \Omega_1 t \ \sin\Omega_2 t \ \sin\Omega_3 t$.

These four equations describe the time development for a typical chemically-shifted site; analogous equations can be written for all the other sites.

Fourier transformation of these time-domain signals with respect to t gives the corresponding frequency-domain representation. In order to make all the responses absorptive, the receiver phase is shifted by n times 90° where n is the number of sine modulations. Standard trigonometric functions for sums and differences of angles indicate that cosine modulation transforms as an in-phase doublet whereas a sine modulation transforms as an antiphase doublet. Schematically the frequency-domain patterns from this single site can be represented as stick spectra (Figure 2).

Without prejudice to the argument, it can be assumed that Ω_3 is larger than Ω_2, so that the 'sum and difference' frequencies A through D are given by:

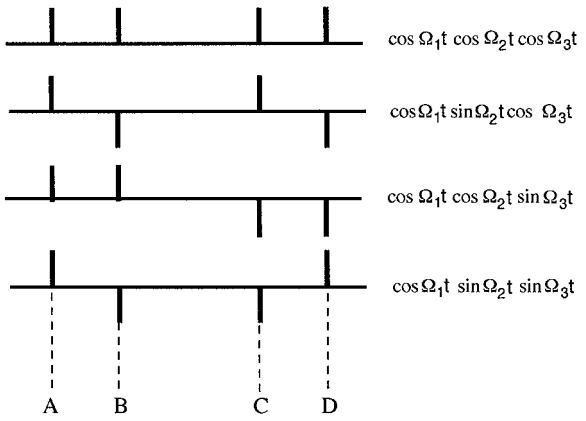

Figure 2. Schematic spectra for the GFT method from a single chemical site. The different cosine and sine modulations in four successive scans create these different patterns in the frequency domain. The individual frequencies can be extracted by adding or subtracting these spectra with the appropriate sign combinations.

$A = +\Omega_1 + \Omega_2 + \Omega_3$,
$B = +\Omega_1 - \Omega_2 + \Omega_3$,
$C = +\Omega_1 + \Omega_2 - \Omega_3$,
$D = +\Omega_1 - \Omega_2 - \Omega_3$.

It is clear that the three frequencies can be separated by combining the columns with the appropriate plus or minus signs (compare section on Hadamard spectroscopy). In this way, although the three frequency dimensions have been explored jointly in 'lock-step', the requisite chemical shift information can be disentangled.

These four simultaneous equations would be sufficient to obtain all three chemical shifts in an ideal situation where there is no accidental overlap of any of the responses; indeed the problem is slightly overdetermined. But in the real world, and particularly in large molecules, this is unlikely, and further steps must be taken to resolve ambiguities caused by degeneracies in the frequencies. This entails further measurements in which the evolution in certain dimensions is inhibited. Additional scans are performed with $t_3 = 0$, giving the transfer functions:

$\cos \Omega_1 t \ \cos \Omega_2 t$,
$\cos \Omega_1 t \ \sin \Omega_2 t$.

After Fourier transformation, the frequency-domain spectra consist of in-phase and antiphase doublets respectively, with frequencies given by

$E = +\Omega_1 + \Omega_2$,
$F = +\Omega_1 - \Omega_2$.

Finally, if both t_2 and t_3 are set to zero there is just a single peak at the 'centre of gravity' of the previous spectra, $G = \Omega_1$. This last piece of information can be very useful for discriminating between the subspectral responses from different chemical sites.

Synchronization of more than three evolution periods leads to correspondingly more complicated expressions, with larger G-matrices and larger numbers of repeated 'scans' but the principle remains the same. For example, joint evolution in four dimensions requires 15 scans, with correspondingly more complex G-matrices. The reduction in measurement time improves exponentially as the number of locked dimensions increases, despite the increase in number of mandatory scans. However, the technique sacrifices sensitivity for each stage of dimensionality reduction, and higher-dimensional experiments necessarily incur increased signal loss through relaxation, so eventually a point is reached where the sensitivity, rather than the duration, becomes the limiting factor.

The GFT method has been applied to the 600 MHz spectra of the aptly-named 76-residue protein ubiquitin (Kim and Szyperski, 2003). The sample was doubly-labelled in carbon-13 and nitrogen-15, and consisted of a 2 mM solution in 90%H_2O/10%D_2O. The conventional five-dimensional HACACONHN experiment was reduced to only two effective dimensions by incrementing the H(α), C-13(α), C-13′ and N-15 evolution dimensions together and in step. This entailed 2 h 18 min of data gathering. After assignment of the frequencies, the chemical shifts were extracted and measured by a least-squares fitting procedure. It was calculated that the equivalent conventional experiment would have required an unacceptably long duration of almost 6 days of spectrometer time.

The strength of the GFT method lies in the possibility of reducing the dimensionality of high-dimensional spectra – those that cannot be contemplated by conventional Fourier transform methods because of impossibly long data-gathering times. For example, pulse sequences for five-dimensional NMR can be readily devised but are rarely implemented by the traditional methodology. The limitations of the GFT method might be expected to become evident in cases of severe overlap and in situations where sensitivity is also a problem.

Single-scan multidimensional spectroscopy

Frydman et al. (2002, 2003) have demonstrated an ingenious 'single-scan' scheme that collapses evolution and acquisition into a single real-time dimension. The new method involves the application of a selective radiofrequency pulse in the presence of a magnetic field gradient, so that at any one instant only one thin slice of the sample is excited. Typically a z-gradient is used, acting along the long axis of the standard cylindrical sample, but in principle gradients in all three directions could be employed. The selectivity in the frequency domain is low enough in terms of chemical shift dispersion that the entire spectral width is excited by each pulse, but nevertheless sufficiently selective that the slice thickness is small in comparison with the active length of the sample. The gradients must be very intense. Typically Gaussian-shaped selective pulses are employed. This excitation cycle is repeated N_1 times with a suitable increments in the frequency of the selective pulse so that N_1 discrete slices are excited. In a typical example N_1 might be 40–60.

The conventional scan-by-scan exploration of t_1 that has proved so profligate with spectrometer time is now replaced by a procedure that maps out the entire spin evolution within a single scan. A slow serial measurement has been converted into a parallel measurement completed in a much shorter time. This spatial encoding procedure can be thought of as creating N_1 little 'boxes' in which the signals from different spin packets are temporarily stored, rather like mail sorted into different pigeon-holes. In practice these 'boxes' are slices through the sample normal to the z axis. These N_1 spatially-resolved sub-ensembles of nuclear spins evolve independently. Since the excitation of each slice is initiated at a slightly later time, the evolution time t_1 is automatically decremented at the same rate (Figure 3). This is the key feature – we can associate a specific t_1 value with each excited slice. Free precession due to a chemical shift is separated from the dispersal in the applied gradient by the application of a second equal gradient of opposite polarity; the chemical shift precession persists throughout. The repeated matched gradient pairs have no cumulative effect, and the undesirable influence of molecular diffusion is minimized. In practice the number of active slices (and the number of evolution increments N_1) is limited by the available gradient strength compared with the spectral width under investigation, suggesting that the method is best suited to proton spectroscopy.

For simplicity we consider the well-known TOCSY sequence, but a COSY experiment would follow similar lines. At this point the operation deviates from conventional methodology and becomes quite subtle. Consider the evolution of spins in one representative slice with NMR chemical shifts Ω_A, Ω_B and Ω_C. This slice (a) is associated with its characteristic evolution time $(t_1)_a$, so at the end of the excitation stage this particular spin packet has accumulated precession angles that depend on $\Omega_A(t_1)_a$, $\Omega_B(t_1)_a$ and $\Omega_C(t_1)_a$. At this point there is no applied gradient, and all the slices experience the same magnetic field, being distinguished only by their different t_1 values.

Next there is an isotropic mixing stage which is not spatially selective. Now consider the interference effects between spin magnetization vectors from *all* the active slices. These vectors have been twisted into three distinct helices with pitches in the ratio Ω_A, Ω_B and Ω_C (Figure 4). The first stage of the acquisition process employs a positive magnetic field gradient pulse that progressively unwinds the three helices, bringing them to a focus *one at a time*, as the three

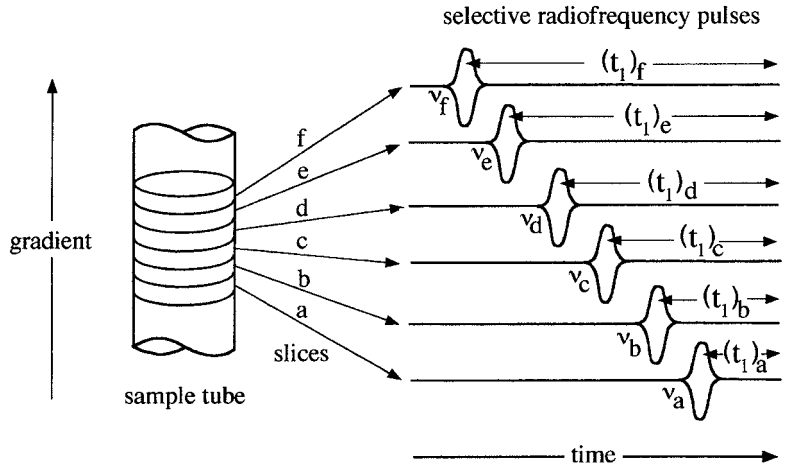

Figure 3. Schematic representation of the single-scan two-dimensional experiment. In the applied field gradient, selective radiofrequency pulses with linearly incremented frequencies (ν_a, ν_b, ν_c, etc.) define successive slices (a, b, c, etc.) in the z direction. The pulses are strong enough to excite the entire spectrum, and because they are applied with decreasing time delays, they excite spectra with increasing evolution times $(t_1)_a$, $(t_1)_b$, $(t_1)_c$, etc, thus exploring the entire evolution dimension in a single scan. Each slice is associated with a specific t_1 value.

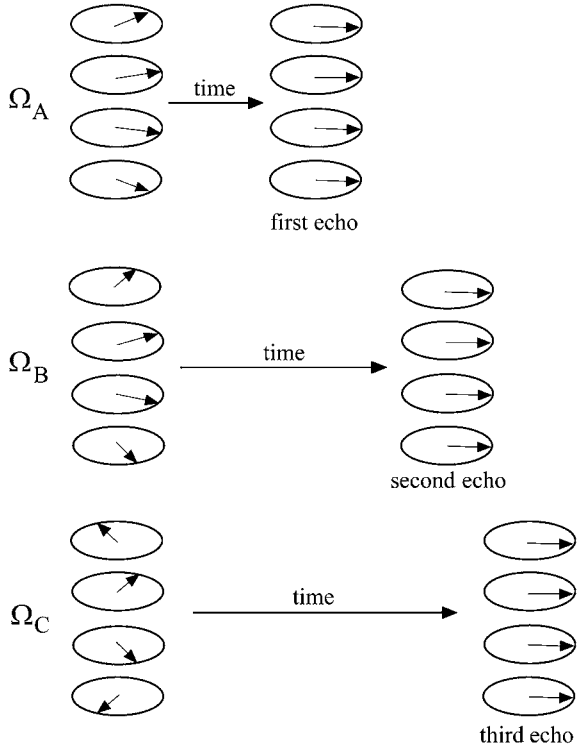

Figure 4. In the single-scan method for two-dimensional spectroscopy the excitation stage spreads magnetization vectors from different slices of the sample into helices, the pitch increasing in proportion to the chemical shifts Ω_A, Ω_B and Ω_C. During the acquisition stage the 'read' gradient unwinds these three helices, but because the initial pitches are different, three successive echoes are formed. They represent the NMR spectrum in the F_1 dimension.

different 'handicaps' are compensated one after the other. This crucial refocusing effect relies on the linear relationship between the initial t_1 evolution and the z-ordinates, as each slice experiences just the right local magnetic field to compensate for the evolution that occurred during the excitation stage. It *refocuses* (rather

than twisting the helices more tightly) because an applied magnetic field rotates magnetization vectors in a direction opposite to the sense of nuclear precession (in the same way that a linearly-increasing field sweep traces out a spectrum in the order of *decreasing* NMR frequency). This is a rather special kind of gradient-recalled echo where the initial dispersal due to free evolution is refocused by an actual field gradient. Regular sampling during this interval detects three successive echoes, representing the three peaks of the NMR spectrum in the F_1 dimension. Paradoxically, the high intensity of the field gradient is a favorable feature, ensuring that the three echoes are very sharp. Note that no Fourier transformation is involved in deriving this 'spectrum'.

The positive gradient pulse is followed by a matched gradient pulse of opposite polarity. This has the effect of refocusing the three echoes once again, generating an F_1 spectrum that is reversed in direction. To improve sensitivity these signals are processed separately and later reversed and combined with the first set. The module consisting of a pair of opposed gradient pulses is then repeated and data collected at regular intervals as before. A sequence of N_2 modules is used, where N_2 is normally considerably larger than N_1. As a result of the progressively increased delay between mixing and acquisition, the intensities in the later versions of the F_1 spectrum are modified, as coherence is transferred between sites through the TOCSY effect. This imposes a t_2 dependence on the successive F_1 spectra. Eventually the signal decays through the usual instrumental and relaxation processes. In this manner each F_1 spectrum is associated with its correlated time-domain response $S(t_2)$, and after a *single* stage of Fourier transformation with respect to t_2, the two-dimensional spectrum $f(F_1, F_2)$ is obtained.

For those of us accustomed to the normal rules of two-dimensional spectroscopy where two stages of Fourier transformation are mandatory, this is a quite surprising feature. Only the transformation with respect to t_2 is needed; the F_1 spectra arise automatically through sequential echo refocusing. The F_1 'evolution' dimension is confined to quite short time intervals (typically 480 microseconds) embedded within a much longer sequence (typically 128 milliseconds) where signals are acquired as a function of t_2. The entire measurement is completed in a fraction of a second.

It would be easy to fall into the trap of thinking that the sensitivity of this single-scan technique is compromised because, at any one instant, slice selection limits the number of nuclear spins being excited. But remember that each step of this repetitive cycle gathers information about the entire spectrum, and signals from all the active slices combine in the acquisition stage. Slice selection is simply acting as a way of temporarily introducing a new 'storage' dimension; all the available nuclear spins contribute to the final spectrum. Although the signal strength is comparable with that of a conventional free induction decay of the same duration, the noise level is higher. Compared with conventional two-dimensional spectroscopy the sensitivity is reduced because the total duration of data gathering is so much shorter; there is always this trade-off between speed and sensitivity. Furthermore, the dwell time has to be short, typically 2 microseconds, and the receiver bandwidth has to be broad, thus letting in more noise. The receiver duty cycle is not optimum, and there are certain non-idealities in the selective pulses and artifacts generated by the rapidly-switched field gradients. There is however one distinct advantage of any single-scan technique – it avoids the partial saturation that occurs in most multi-scan methods. As in the other fast multidimensional methods outlined here, it is assumed that sensitivity is not the main consideration. The limitations of the single-scan technique are more likely to lie in difficulties associated with the generation of suitably intense, reproducible, and rapidly-switched gradient pulses – technical features that may well be remedied in due course. Whatever the eventual outcome, one has to admire the ingenuity behind this remarkable experiment.

The method has been illustrated (Frydman et al., 2003) for the single-scan two-dimensional 500 MHz proton TOCSY spectrum of the dipeptide AspAla, using a 1.5 s water presaturation pulse and a 120 millisecond DIPSI sequence (Shaka et al., 1988) for isotropic mixing. There were $N_1 = 58$ selective pulses applied in a gradient strength of 65 G per cm. The detection stage involved $N_2 = 256$ repeated gradient pairs. The resulting two-dimensional TOCSY spectra show the expected correlations (Figure 5). Note that, neglecting water presaturation, the spectrum was acquired in only 0.28 s, far faster than any of the competing 'fast' two-dimensional methodologies described here. The duration is so short that the extension to three- or even four-dimensional experiments could be contemplated even if the additional evolution periods are explored by conventional Fourier transform methods.

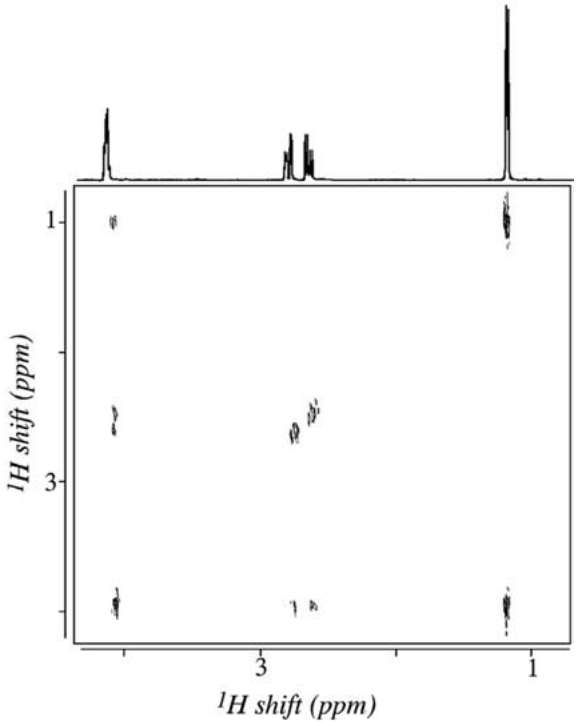

Figure 5. The single-scan two-dimensional method applied to the 500 MHz phase-sensitive TOCSY spectrum of protons in the dipeptide AspAla (the two H^{α} resonances have the same chemical shift). Apart from the water presaturation period, the measurement was completed in only 0.28 s. Spectrum courtesy of Lucio Frydman.

Hadamard spectroscopy

This approach to the speed problem exploits the sparse nature of most multidimensional NMR spectra by attacking the frequency domain directly, thereby avoiding the usual requirement for uniform sampling in the evolution dimensions (Kupče and Freeman, 2003a, b). It is predicated on the assumption that the relevant NMR frequencies can be measured from exploratory one-dimensional experiments carried out in a rather short time. This information is used to set the frequencies of an array of simultaneous selective radiofrequency pulses, one for each distinct chemical site. One multiplex array is used for the initial excitation stage of a multidimensional experiment and further arrays are employed for the subsequent coherence transfer stages. Acquisition of the free induction decays is conventional and employs Fourier transformation to derive the spectra.

The key to this method is an encoding scheme based on Hadamard matrices (Hadamard, 1893) made up of plus and minus signs. To fix ideas, consider the simple example of 15 chemical sites irradiated with 15 different radiofrequencies simultaneously. A Hadamard matrix of order 16 would be employed, modulating the individual radiofrequencies according to the signs in each row of the matrix. Each column corresponds to one of the irradiation channels.

```
+ + + + + + + + + + + + + + + +
+ − + − + − + − + − + − + − + −
+ + − − + + − − + + − − + + − −
+ − − + + − − + + − − + + − − +
+ + + + − − − − + + + + − − − −
+ − + − − + − + + − + − − + − +
+ + − − − − + + + + − − − − + +
+ − − + − + + − + − − + − + + −
+ + + + + + + + − − − − − − − −
+ − + − + − + − − + − + − + − +
+ + − − + + − − − − + + − − + +
+ − − + + − − + − + + − − + + −
+ + + + − − − − − − − − + + + +
+ − + − − + − + − + − + + − + −
+ + − − − − + + − − + + + + − −
+ − − + − + + − − + + − + − − +
```

For practical instrumental reasons the first 'all plus' column would not normally be used. For the initial excitation stage, the plus and minus signs define $+90°$ and $-90°$ radiofrequency pulses, whereas for the transfer stages they designate 'on' and 'off' 180° pulses; in either case the modulation reverses the sign of the final NMR signal. The measurements are repeated until 16 scans have been carried out, the modulation code being changed in each successive scan according to the relevant row of the H-16 matrix. All 16 scans must be completed. The free induction decay recorded at the end of each scan is a composite of the responses from all 15 chemical sites, some positive and the others negative in sense.

This encoding procedure serves as a means for disentangling the individual responses, rather like the operation of a telephone scrambling device. For example, subtraction of the sum of the free induction signals from the even-numbered scans from the sum of those from odd-numbered scans gives the 'pure' response from the first site, the one that employed column two of the matrix; all other responses cancel, being made up of equal numbers of positive and negative contributions. Formally this operation can be represented as a Hadamard transformation.

There is of course a slight mismatch between the number of chemical sites to be irradiated and the order of the nearest Hadamard matrix, but the effect is small, as matrices exist for any order $N = 4n$ where n is an integer. Common values for this type of experiment would be $N = 8$, 12, 16, 20, 24, 28, 32, etc. The number of 'wasted' columns of the encoding matrix is therefore quite small, and it is useful to be able to omit the 'all plus' column on the grounds that it could prove rather more sensitive to residual spectrometer imperfections.

Two or more encoding operations can be cascaded, allowing the Hadamard method to be applied in successive frequency dimensions. Note that the introduction of each new dimension multiplies the required numbers of scans. Thus a typical three-dimensional experiment might employ (for example) Hadamard matrices of order 16 in the first dimension and 12 in the second dimension, with the usual acquisition of the free induction decays in the third dimension. There would be 192 scans. At first sight the number of mandatory scans might seem to constitute a real limitation on the application of the Hadamard method. However, it replaces a traditional multidimensional experiment that systematically explores the corresponding evolution dimensions with a much larger total number of scans. A conventional three-dimensional experiment, equivalent to the example outlined above, might involve 128 increments in the t_1 dimension and 64 increments in t_2 to cover the desired frequency ranges while maintaining only marginal resolution in these dimensions, hence 8192 scans. This is a speed advantage of almost 43 in favour of the Hadamard mode. Furthermore, direct multiplex irradiation in the frequency domain defines the radiofrequencies unequivocally; there is no sign ambiguity and no question of aliasing, so there is no need for quadrature phase detection, saving another factor of two. Certain types of spectrometer imperfections are reduced as a result of the Hadamard modulation scheme, thus diminishing the necessity for phase cycling. When all these factors are taken into account there is a speed advantage of between two and three orders of magnitude for high-dimensional Hadamard spectroscopy.

Clearly the Hadamard experiment is completed more quickly the fewer the number of chemically distinct sites chosen for irradiation. This can be turned to good advantage in applications to protein chemistry, where non-specific enrichment in carbon-13 and nitrogen-15 is a popular stratagem. In multidimensional Hadamard spectroscopy the operator has the option of selecting only the most interesting sites for irradiation, essentially converting a globally-enriched sample into one that has been specifically labelled in the important carbon and nitrogen locations. This ability to generate a subspectrum from specific residues is illustrated for the 700 MHz three-dimensional HNCO spectrum of 0.3 mM agitoxin, a 4 kDa protein globally enriched in carbon-13 and nitrogen-15 (Kupče and Freeman, 2003c). Figure 5a shows the projection onto the proton-carbon-13 plane derived from the full three-dimensional conventional spectrum with two scans per increment, recorded in 20 h 43 min. This projection does not discriminate with respect to nitrogen-15 frequencies. Seven selected responses, indicated by arrows, were then recorded by the Hadamard technique (Figure 5b) with a time saving of more than 200.

The distinguishing feature of direct-irradiation Hadamard spectroscopy is the need for prior measurements to determine the NMR frequencies. For small and medium-sized molecules these preparatory experiments are readily implemented and cost little in spectrometer time. However, some new problems can arise with large proteins because of the very complexity of the spectra, poor resolution and low inherent sensitivity, making it difficult to attack an 'unknown' protein entirely from scratch. Unless assignment information is available from earlier work, it may be difficult to select frequencies in the second and subsequent dimensions that correlate with those chosen in the first dimension. The real power of the Hadamard method for biological macromolecules emerges in studies of protein binding or folding, where a full multidimensional spectrum has already been acquired and where the interest lies in the time development of the spectra. The ability to focus attention on residues close to the active site is a distinct advantage in this application.

Projection-reconstruction

The internal structure of a three-dimensional object can be reconstructed from projections of the X-ray absorption measured in several different directions, an idea that has been successfully exploited in X-ray tomography (Hounsfield, 1973). An analogous reconstruction of a three-dimensional NMR spectrum from its plane projections should present a rather easier problem because the absorption sites are discrete and normally well-resolved. So the question naturally arises, could all the three-dimensional NMR

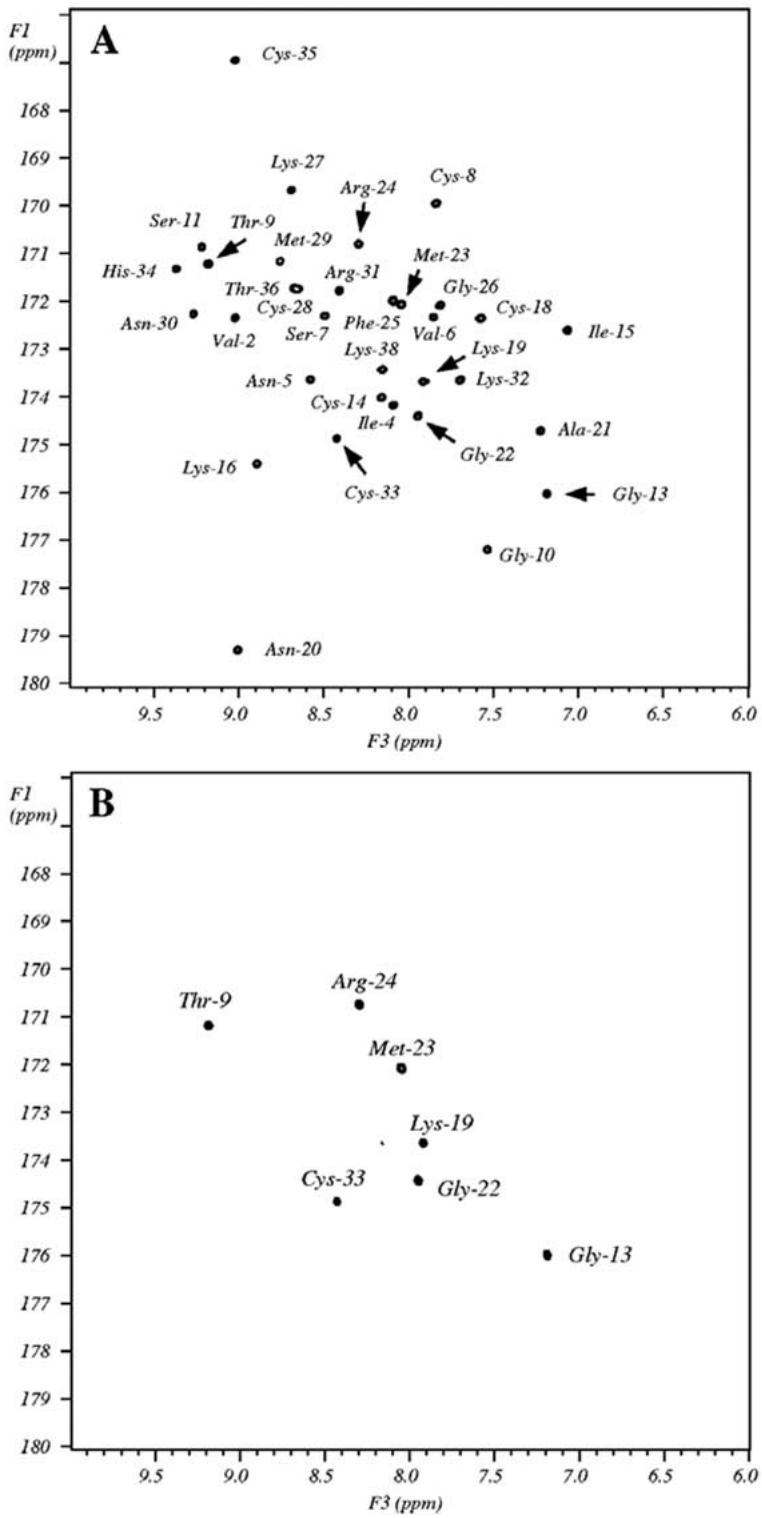

Figure 6. The 700 MHz three-dimensional HNCO spectrum of agitoxin projected onto the proton-carbon-13 plane. (A) The conventional Fourier transform spectrum with seven selected residues indicated by arrows. (B) The Hadamard subspectrum where only these seven responses were excited. The speed advantage of the Hadamard mode was more than 200.

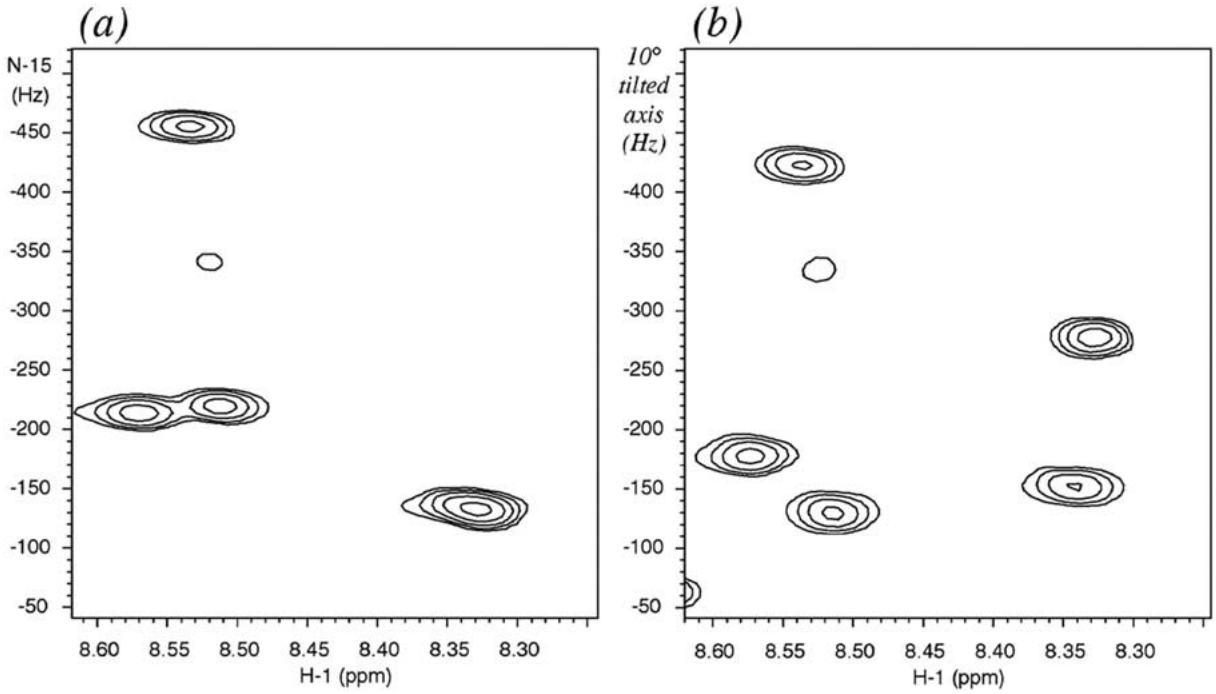

Figure 7. A small selected region of the 500 MHz proton HNCO spectrum of ubiquitin, labelled with carbon-13 and nitrogen-15. (a) Projection onto the proton-nitrogen-15 plane, obtained by setting $t_1 = 0$. Two responses at the bottom right overlap. (b) Projection onto a plane that subtends an angle of 10° with respect to the proton-nitrogen-15 plane. This displaces all the responses in a vertical direction by different amounts, separating the two overlapping peaks. In the tilted projection the vertical scale cannot be expressed in ppm.

information be extracted from a small number of two-dimensional projections measured in a relatively short time? If F_1 and F_2 are the evolution dimensions and F_3 the direct acquisition dimension, then the projections onto the orthogonal F_1F_3 and F_2F_3 planes are readily recorded by setting $t_2 = 0$ and $t_1 = 0$, respectively. These are sometimes referred to as the 'first planes' of the three-dimensional spectrum and can be measured quite quickly. In an ideal situation where not too many peaks contribute to the intensity of a single response in the projection, and where the measured NMR intensities in both projections are reliable, the entire three-dimensional spectrum can be extracted from information on these two planes. In actual practice, ambiguities are to be expected and must be resolved if the full three-dimensional spectrum is to be recovered.

A projection is made up of intensities summed along the various parallel 'rays', and if a given ray intersects several peaks, there is a potential ambiguity in the relative intensities of these peaks, which may or may not be resolved by consideration of intensities in the orthogonal projection. One solution is to measure a skew projection intermediate between the F_1F_3 and F_2F_3 planes by incrementing t_1 and t_2 jointly, with some suitable scaling factor. Judicious choice of this new projection direction should separate any overlapping peaks. In extreme cases further skew projections could be envisaged. Even if tilted projections were taken at every 10° interval between 0° and 90°, the measurement would still be completed much faster than the full three-dimensional spectrum. Skew projections are obtained by replacing the normal evolution-time increments Δt_1 and Δt_2 with $\Delta t_1 \sin\alpha$ and $\Delta t_2 \cos\alpha$ where α is the desired projection angle. The linewidths in the skew projection are related to those in the two orthogonal projections and to the projection angle α. In a favorable case, as few as three plane projections measured in a short time can map out the entire three-dimensional spectrum. The final result might well be a complete set of three-dimensional coordinates along with the associated intensities. Since the two orthogonal 'first-plane' projections are less susceptible to relaxation losses, the method is well suited to studies of large proteins, but being faster than the full three-dimensional protocol, it sacrifices some of the multiplex sensitivity advantage.

Resolution of ambiguity is demonstrated for 500 MHz proton spectra taken from a small selected region of the HNCO spectrum of ubiquitin (labelled in carbon-13 and nitrogen-15). The projection onto the proton-nitrogen-15 plane shows two overlapping responses (Figure 7a). A skew projection onto a plane tilted through 10° with respect to this plane displaces all these responses in the vertical dimension by varying amounts, neatly separating the two peaks in question (Figure 7b). Note that this tilted projection method assigns a *trajectory* to each peak.

Conclusions

The time factor stands out as the most serious curb to the application of multidimensional spectroscopy to proteins and other macromolecules of biochemical interest, and it is clear that the new methodologies outlined above go a long way to addressing this problem. Such gains naturally involve some penalties. In most cases the inherent NMR sensitivity is assumed to be so high that the effects of instrumental noise can be neglected, but as the performance of commercial spectrometers continues to improve, this may become a less demanding consideration. The key advantage of the filter diagonalization method is that it requires no operational changes at all, but the sophisticated data processing can prove time-consuming; it is not yet a 'turn-key' operation. The strength of GFT spectroscopy is that it can tackle very-high-dimensional spectra without sacrificing resolution, and the G-matrix transform is fast, but it may run into overlap problems in very crowded spectra. Frydman's single-scan technique is by far the fastest of all, acquiring a two-dimensional data set in less than a second, but it relies on high-intensity rapidly-switched magnetic field gradients, a technology that may need to be improved for the method to gain wide acceptance. Direct irradiation Hadamard spectroscopy is predicated on the prior acquisition of the chemical shift frequencies, and some correlations may also be required to set up an experiment on a large protein. However, it offers the opportunity to focus attention on the resonances of prime interest, for example those close to the active site of a protein. The proposal for reconstructing a three-dimensional spectrum from a small number of plane projections promises similar improvements in speed. It is conceivable that hybrid combinations of these techniques may serve to combine some of the respective advantages, or diminish any drawbacks.

Acknowledgements

The authors are indebted to A.J. Shaka, Thomas Szyperski, and Lucio Frydman for providing manuscripts prior to publication and for helpful advice about their respective techniques, and to A.J. Shaka and Lucio Frydman for permission to reproduce figures from published work.

References

Aue, W.P., Bartholdi, E. and Ernst, R.R. (1976) *J. Chem. Phys.*, **64**, 2229.

Barkhuijsen, H, DeBeer, R, Bovée, W.M.M.J. and Van Ormondt, D. (1985) *J. Magn. Reson.*, **61**, 465–481.

Brutscher, B., Cordier, F., Simorre, J.P. Caffrey, M.S. and Marion, D. (1995) *J. Biomol. NMR*, **5**, 202–206.

Brutscher, B., Simorre, J.P. Caffrey, M.S. and Marion, D. (1994) *J. Magn. Reson.*, **B105**, 77–82.

Chen, J., De Angelis, A.A., Mandelshtam, V.A. and Shaka, A.J. (2003) *J. Magn. Reson.*, **161**, 74–89.

Chen, J., Mandelshtam, V.A. and Shaka, A.J. (2000) *J. Magn. Reson.*, **146**, 363–368.

Ding, K. and Gronenborn, A.M. (2002) *J. Magn. Reson.*, **156**, 262–268.

Frydman, L., Lupulescu, A. and Scherf, T. (2003) *J. Amer. Chem. Soc.* (in press).

Frydman, L., Scherf, T. and Lupulescu, A. (2002) *Proc. Natl. Acad. Sci. USA*, **99**, 15859–15862.

Hadamard, J. (1893) *Bull. Sci. Math.*, **17**, 240–248.

Hounsfield, G.N. (1973) *Brit. J. Radiol.*, **46**, 1016.

Hu, H., De Angelis, A.A., Mandelshtam, V.A. and Shaka, A.J. (2000) *J. Magn. Reson.*, **144**, 357–366.

Jeener, J. (1971) Ampère International Summer School, Basko Polje, Yugoslavia.

Kim, S. and Szyperski, T. (2003) *J. Am. Chem. Soc.*, **125**, 1385.

Kupče, E. and Freeman, R. (2003a) *J. Magn. Reson.* **162**, 300–310.

Kupče, E. and Freeman, R. (2003b) *J. Magn. Reson.* (in press).

Kupče, E. and Freeman, R. (2003c) *J. Biomol. NMR*, **25**, 349–354.

Löhr, F. and Rüterjans, H. (1995) *J. Biomol. NMR*, **6**, 189–197.

Mandelshtam, V.A. (2000) *J. Magn. Reson.*, **144**, 343–356.

Mandelshtam, V.A. (2001) *Progr. NMR Spectrosc.*, **38**, 159–196.

Powers R., Gronenborn, A.M., Clore, G.M. and Bax, A. (1991) *J. Magn. Reson.*, **94**, 209–213.

Shaka, A.J., Lee, C.J. and Pines, A. (1988) *J. Magn. Reson.*, **77**, 274.

Simorre, J.P., Brutscher, B., Caffrey, M.S. and Marion, D. (1994) *J. Biomol. NMR*, **4**, 325–334.

Szyperski, T., Banecki, B., Braun, D. and Glaser, R.W. (1998) *J. Biomol. NMR*, **11**, 387–405.

Szyperski, T., Braun, D., Bartels, C. and Wüthrich, K. (1995) *J. Magn. Reson.*, **B108**, 197–203.

Szyperski, T., Wider, G., Bushweller, J.H. and Wüthrich, K. (1993a) *J. Biomol. NMR*, **3**, 127–132.

Szyperski, T., Wider, G., Bushweller, J.H. and Wüthrich, K. (1993b) *J. Am. Chem. Soc.*, **115**, 9307–9308.

Szyperski, T., Yeh, D.C., Sukumaran, D.K., Moseley, H.N.B. and Montelione, G.T. (2002) *Proc. Natl. Acad. Sci. USA*, **99**, 8009–8014.

Tikhonov, A.N. (1963) *Soviet Math. Dokl.*, **4**, 1035–1038.

Wall, M.R. and Neuhauser, D. (1995) *J. Chem. Phys.*, **102**, 8011–8022.