Journal of Magnetic Resonance 213 (2011) 1–13

Contents lists available at SciVerse ScienceDirect

Journal of Magnetic Resonance

journal homepage: www.elsevier.com/locate/jmr

Perspectives in Magnetic Resonance

Parallel receivers and sparse sampling in multidimensional NMR

Ēriks Kupče [a,*], Ray Freeman [b]

[a] Agilent Technologies, 6 Mead Road, Yarnton, Oxford OX5 1QU, UK
[b] Jesus College, Cambridge CB5 8BL, UK

ARTICLE INFO

Article history:
Received 31 July 2011
Revised 19 August 2011
Available online 30 August 2011

Dedicated to Professor Lewis Kay on the occasion of his 50th birthday

Keywords:
Sparse sampling
Multiple receivers
Hadamard NMR
Multidimensional NMR
Projection reconstruction
Controlled aliasing
Random sampling
Hyperdimensional NMR

ABSTRACT

The recent introduction of NMR spectrometers with multiple receivers permits spectra from several different nuclear species to be recorded in parallel, and several standard pulse sequences to be combined into a single entity. It is shown how these improvements in the flow and quality of spectral information can be significantly augmented by compressive sensing techniques – controlled aliasing, Hadamard spectroscopy, single-point evaluation of evolution space (SPEED), random sampling, projection-reconstruction, and hyperdimensional NMR. Future developments of these techniques are confidently expected to mitigate one of the most serious limitations in multidimensional NMR – the excessive duration of the measurements.

© 2011 Elsevier Inc. All rights reserved.

1. Introduction

The recent introduction of NMR spectrometers with several receivers operating in parallel [1–3] offers an important new way to increase the flow of spectral information, while at the same time virtually eliminating the differential effects of slow changes in the environmental conditions or chemical degradation of the sample. More importantly, several complementary pulse techniques can be combined into a single entity [4–6], so that all the requisite structural information can be gathered in a single pass, bringing NMR into line with the usual practice in X-ray, infrared and mass spectroscopy. There is no longer any need to revisit the NMR machine at a later date to gather supplementary spectroscopic information. The timesaving is appreciable, and the method is compatible with other schemes for speeding up multidimensional NMR. This Parallel Acquisition NMR Spectroscopy (PANSY) protocol also lends itself to the processing of a sequence of different samples introduced by sample-changing equipment. Here we focus on experiments involving parallel receivers and explore the potential benefits of combining such methods with various sparse sampling techniques.

In today's world where speed and efficiency are constant concerns, economy of data sampling assumes paramount importance. This is particularly true of multidimensional NMR measurements, where the overall duration of the experiment is often determined by the sampling density in evolution space. The gold standard for time-domain sampling is of course the Nyquist condition, which requires that the sampling rate be equal to (or exceed) twice the highest frequency component in the experimental signals. Sparse sampling (also known as compressive sampling, compressive sensing or compressed sensing) seeks to circumvent this requirement, often by taking into account some prior knowledge or assumption about the form and complexity of the raw spectral data. In this manner the time penalty associated with the Nyquist criterion can be significantly mitigated. As in all NMR experiments the sensitivity increases as the square root of the number of accumulated scans. Consequently the advantage of sparse sampling methods is predicated on the assumption of adequate sensitivity – those experimental situations where extensive signal accumulation is *not* required. This has been called the 'sampling-limited' case as opposed to the 'sensitivity-limited' regime. Fortunately many multidimensional experiments fall into this category.

Many different practical schemes have been suggested for recording multidimensional spectra in the shortest possible time. Note that for any sampling interval Δ, there is an associated frequency f_N, called Nyquist frequency:

$$f_N = \frac{1}{2\Delta} \tag{1}$$

The sampling theorem states that if a function $x(t)$ sampled at regular intervals Δ contains no frequencies higher than f_N then it is completely determined by its samples, x_n. Compressive sampling

* Corresponding author.
E-mail address: eriks.kupce@agilent.com (Ē. Kupče).

1090-7807/$ - see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmr.2011.08.027

of evolution space involves deliberate flouting of the Nyquist condition of regular sampling and tacit acceptance of the consequent complications – in particular the introduction of artefacts into the recovered spectrum. This is the principal limitation of most of these data compression protocols.

There are two very simple and widely used sparse sampling methods. One employs early truncation of the sampling operation, reintroducing values for the missing evolution data by the process of linear prediction [7]. This saves a factor of about two in each independent evolution dimension. The other is undersampling – deliberately allowing aliasing of the spectrum by digitizing at a rate appreciably below the Nyquist frequency [8]. In the general case this introduces ambiguity in the interpretation of the aliased spectrum. However in the special case of INADEQUATE spectra [9], multiple aliasing in the double-quantum dimension (F_1) is of little consequence for it is only used to separate the F_2 traces, and even where there is direct overlap the interpretation is still straightforward.

An important precursor for compressive sampling is the knowledge that the data is indeed compressible; in the NMR context this implies plenty of space between the peaks. In many cases this prior information can be easily derived from one- or two-dimensional measurements, and can usually be obtained with high digitization. It is a relatively fast operation compared with the planned multidimensional investigation. It would be unwise to ignore this information in subsequent data acquisition and processing steps. Accurate values for one set of chemical shifts are obtained with little penalty in instrument time. One implementation uses this prior knowledge to replace the conventional evolution period with direct selective excitation at each chemical site [10], relying on the relatively small number of such sites. If all sites are irradiated simultaneously and the excitation radiofrequency fields are coded by reference to a Hadamard matrix [11–14], the signal from each individual site can be recovered in turn by decoding according to the same matrix. The sensitivity per unit time is not impaired, while the experimental duration is reduced in the approximate ratio $2N/C$ where C is the number of distinct chemical sites, and N is the number of samples in the evolution domain that would be required by the Nyquist criterion.

Another proven data compression technique employs limited radial sampling of the evolution dimensions instead of the complete exploration of all points on a regular Cartesian grid. It relies on a Fourier transform theorem formulated by Bracewell in the context of radioastronomy [15] which relates a section through a time-domain function at an angle α to the projection of the corresponding frequency domain spectrum at the same angle α. There is a program available that predicts those values of α that avoid eclipsed peaks. A small number of such projections can be processed to reconstruct the full multidimensional spectrum. This important method is described in more detail below (Section 2). Compared with systematic conventional Cartesian sampling of evolution space, projection-reconstruction enjoys a speed advantage of an order of magnitude for each new frequency dimension beyond the second. This is no trivial achievement – an expensive NMR spectrometer would be tied up for several days performing a four-dimensional experiment where each of the three evolution dimensions is systematically explored with 64 time increments, and with one second allowed for acquisition and recovery through spin–lattice relaxation. Even with faster-than-normal pulsing and Ernst-angle excitation [16], or polarization sharing with adjacent nuclei [17], this long experimental duration is only modestly reduced. Note that a special case of radial sampling ($\alpha = 45°$) has been exploited for speeding up multidimensional spectroscopy using an entirely different data-processing protocol known as G-matrix Fourier transformation [18–20].

The most impressive speed advantage of all is achieved by spatially-selective single-scan NMR [21] where evolution information is coded into a spatial dimension using selective excitation in an applied magnetic field gradient. Information corresponding to each evolution 'increment' is stored as a small volume element within the gradient, and is then read out by a repeated succession of matched gradient pairs. This ingenious technique poses several instrumental problems but it has nevertheless demonstrated that an entire two-dimensional spectrum can be recovered from an experiment lasting less than a second. Through the imposition of additional applied field gradients the scheme can be extended into further evolution dimensions.

Many other data compression protocols have been suggested. One approach is to replace the usual *regular* sampling of the evolution dimension by a scheme where the sampling rate reduces with time according to an exponential function [22], so that although later data points are more and more sparsely measured, the resolution requirements are still satisfied. With a view to limiting the amplitudes of sampling artefacts, restricted *random* sampling has been employed [23]. In this context the question that must be faced is whether a large number of weak artefacts is preferable to a small number of stronger artefacts, particularly when the location of the latter can often be predicted. Bestowing a noise-like character on sampling artefacts is a dangerous game.

Sparse sampling methods that follow a spiral trajectory [24] or concentric circles [25] in evolution space also have their advocates. Covariance spectroscopy [26] and multi-way decomposition [27] have also been employed to address the speed problem. An entirely different approach is the filter diagonalization method [28] which compensates for sparse sampling of the evolution dimension by very fine digitization of the acquisition domain.

However it must always be remembered that violation of the full Nyquist criterion inevitably creates sampling artefacts in the reconstructed multidimensional spectrum. Various remedies have been proposed for addressing this question, for example the 'CLEAN' algorithm first devised to correct problems that arise in radio astronomy [29,30]. This is a general iterative program designed to separate a particular physical response from its associated artefacts provided that the form of both are known *a priori*. In radio astronomy it is employed to suppress distortion of the image caused by known deficiencies in the antenna array; in NMR spectroscopy it can remove the characteristic dispersion-mode components from the phase-twist lineshape [31]. The CLEAN algorithm was then reinvented in multidimensional NMR [32], where limited radial sampling creates well-defined ridge-like artefacts. It searches for the tallest peak in the reconstructed spectrum and then subtracts that response along with its associated ridges, storing the intensity and frequency co-ordinates in a table. The next stage attacks the next tallest peak and removes it from the spectrum. Each new iteration subtracts an additional set of the offending ridges, thus reducing the number of residual false cross-peaks ('false positives'). The algorithm continues until the search routine has found all responses that lie just above the baseplane noise. The CLEAN program lists a set of recovered intensity and frequency co-ordinates; they can be used to derive the structural information directly, or to reconstruct a cleaner version of the multidimensional spectrum. The same algorithm can be exploited to reduce artefacts generated by other sparse sampling experiments.

2. Projection-reconstruction

Of all the sparse sampling protocols, radial sampling appears to show the most promise because the resulting artefacts are well defined and readily removed at the data processing stage [33–38]. The three-dimensional time-domain signal matrix $s(t_1, t_2, t_3)$ is sampled by incrementing the evolution parameters t_1 and t_2 jointly along a skew diagonal where the new variable $t = t_1/\cos\alpha = t_2/\sin\alpha$, and $\tan\alpha = \Delta t_2/\Delta t_1$. The Fourier transform of this signal is a

projection of the frequency-domain spectrum $S(F_1, F_2, F_3)$ onto a pair of planes subtending this same angle $\pm\alpha$ with respect to the F_1F_3 plane. In the evolution space $s(t_1, t_2)$ this is given by a simple Fourier transform [37]:

$$F_{ij}\{^4s_\alpha(t)\} = \int_{-\infty}^{+\infty 2} s_{+\alpha}(t)\cdot e^{-i\omega_{+\alpha}t}dt + \int_{-\infty}^{+\infty 2} s_{-\alpha}(t)\cdot e^{-j\omega_{-\alpha}t}dt = {}^2S(\omega_{+\alpha}) + {}^2S(\omega_{-\alpha}) = {}^4S(\omega_{\pm\alpha}) \quad (2)$$

where ω denotes the frequency domain. The full three-dimensional spectrum can then be reconstructed from information contained in a very small number of such projections, using algorithms related to those employed in X-ray tomography and magnetic resonance imaging.

Where the F_1 and F_2 axes represent different nuclear species (say, ^{1}H and ^{15}N) there may be reasons to select a projection angle that favors the higher chemical shift dispersion and narrower lines. There is an advantage in choosing projection angles that avoid occultation – where two or more resonances are caught in *enfilade* along the projection ray.

2.1. Predictive projection spectroscopy

An adaptive procedure is available to choose a projection angle that avoids this eclipsed case [37]. Consider a representative F_1F_2 plane of a three-dimensional spectrum constructed by convolution of the spectra recorded along the F_1 and F_2 axes. This represents a *provisional* lattice of cross-peaks, some genuine and some false. A program computes the projection of this lattice, retaining *not* the usual integral along the projection beam, but the highest value (this has been called the 'skyline' projection). As the projection angle α is varied between ±90°, a count is made of the number of peaks M detected in each projection. Whenever occultation occurs, the graph of M against α shows a dip, indicating a projection angle to be avoided. Owing to the presence of false cross-peaks in the F_1F_2 plane, this test overestimates the incidence of *genuine* occultations, but it nevertheless serves to select a suitable tilt angle for the next stage of data gathering. The number of required projections depends on the complexity of the spectrum, and it is possible to estimate the number of expected peaks so that the experiment can be terminated when there is sufficient information for a satisfactory reconstruction. This point is achieved in a significantly shorter time than with conventional sampling of the full time-domain grid. This gain in speed can be an order of magnitude for each new evolution dimension.

2.2. Reconstruction from projections

The inverse procedure – reconstruction – poses more serious problems. Consider the reconstruction of a three-dimensional spectrum $S(F_1, F_2, F_3)$ one F_1F_2 plane at a time, followed by the assembly of several such planes as a function of the third frequency F_3. Suppose that three projections have been recorded. For reasons of sensitivity, two of these may often be projections onto the axes F_1 and F_2, obtained by setting $t_2 = 0$ and $t_1 = 0$ respectively. The third would be sampled along a skew time axis where t_1 and t_2 are incremented in concert. Reconstruction involves the process of 'back-projection' of the signals from each of the three projection traces. This generates three sets of parallel ridges running across the F_1F_2 plane. The point where ridges from all three projections intersect defines the co-ordinates of a genuine correlation peak, but an intersection of only two ridges represent a 'false positive' response.

Ridges and false correlation peaks need to be severely attenuated or suppressed completely if the reconstruction is to be reliable. Consider now the general case where there are P projections. A simple and effective solution is to examine the contributions to the intensity at any given pixel in the F_1F_2 plane and retain only the lowest value [34,39,40]. Ridges and false correlation peaks are then replaced by baseplane noise, but true correlation peaks survive the operation intact. This is the recommended solution whenever the intrinsic sensitivity of the experiment supports it, but it has the disadvantage that the signal-to-noise is determined by the noisiest projection, and hence does not improve as more projections are used (indeed it may even degrade). The alternative is to *sum* all the contributions to the pixel intensity, thereby increasing the signal-to-noise ratio in proportion to $\sqrt{P}$, but sacrificing the ability to suppress the back-projection artefacts. Nevertheless, as P is increased, these artefacts become less obtrusive and may eventually sink below the baseplane noise.

Consequently the lowest-value algorithm is preferred for situations where P is small and the intrinsic sensitivity is high, whereas the additive algorithm is reserved for instances where P is large and sensitivity is the prime concern. Clearly some intermediate 'hybrid' solution has much to offer – balancing the competing demands of artefact suppression and signal-to-noise. This is achieved [41] by comparing all the contributions to the intensity at a given pixel and accepting, not merely the lowest, but the K lowest values. The ratio K:P determines the balance between the conflicting aims of artefact suppression and the desire to improve the signal-to-noise ratio. The choice $K = 1$ corresponds to pure artefact suppression, while $K = P$ represents the situation where sensitivity is paramount.

Spectra of high complexity may require a different approach. For example, multidimensional experiments on proteins generate extremely large data sets, so some data compression is advisable. The simple solution is to sacrifice all information about intensities and linewidths, and retain only the center frequencies of resonance lines in the projections. This invokes a peak-picking algorithm with the attendant danger that some lines may be missed if they are below the search threshold or if they lie on the shoulder of a stronger resonance. Within these caveats, the problem has a simple 'algebraic' solution [36,42]. Each correlation is defined by a set of P simultaneous equations defining the intersection of P back-projected straight lines, one from each projection. In practice some slight latitude is permitted, so that an intersection is accepted as valid if these straight lines meet within a small area of uncertainty. Because of the inherent unreliability of the peak-picking algorithm, the distinction between true and false correlations is now less clear-cut, but iterative programs can be written [43,44] to address this complication.

A quite different approach to reconstruction relies on statistical methods. A provisional model of the spectrum $S(F_1, F_2, F_3)$ is tested for its compatibility with the set of experimental projections. The model is then refined in an iterative program until a satisfactory convergence is achieved [45]. The algorithms used may include maximum entropy, maximum likelihood or Bayesian inference, and the speed of convergence may be slow unless some prior knowledge is involved, obtained, for example, by the deterministic methods outlined above. Although the final result may appear as a clean, noise-free spectrum, this 'improvement' is illusory because the non-linear processing algorithm suppresses noise (along with weak signals comparable with that noise) and may even promote artefacts to the status of genuine resonances, by making prior assumptions about the shape of the responses, for example that they are two-dimensional Gaussians.

2.3. A practical application: the HNCO spectrum of ubiquitin

The need for speed is most evident in the application of multidimensional NMR to proteins, for some experiments can tie up an

expensive spectrometer for several days at a time. A simple illustrative example is provided by the three-dimensional HNCO spectrum of 1 mM aqueous ubiquitin (10% D_2O) recorded on a 500 MHz spectrometer [46]. Orthogonal projections onto the F_1 (^{13}C) and F_2 (^{15}N) axes were supplemented by a tilted projection subtending an angle of 30° with respect to the F_1F_3 plane. This data gathering stage required 29 min and 11 s to complete. The 'lowest-value' reconstruction algorithm was used because sensitivity in this case was not critical. Finally the reconstructed parallel F_1F_2 planes were assembled according to the appropriate F_3 (^1H) values to give the full three-dimensional spectrum. A comparison was then made with the conventional method which systematically sampled every point on the three-dimensional time-domain grid, and which required 18 h and 54 min of spectrometer time, almost 39 times longer. In fact the projection-reconstruction results were more finely digitized than those from the conventional experiment; apart from this difference the comparison of selected F_1F_2 planes is excellent (Fig. 1). Normally ^{13}C–^{15}N spectra are little studied in this direct correlation mode because of poor sensitivity.

3. Hadamard spectroscopy

The Hadamard transform is an integral transform that is often described as a sub-class of the Fourier transform. The basis functions of the Hadamard transform can assume only two values, usually written +1 or −1. For many applications, such as digital electronics and digital logic, these 'square wave' functions are considerably better suited than smooth functions, such as the cosines and sines that form the basis of the Fourier transform.

In fact, the Hadamard transform has been used in NMR since the invention of phase cycling [47]. For example, the classical four-step phase cycle, Φ_4 that is used in the heteronuclear multiple-quantum correlation (HMQC) sequence [48] and many other experiments, is a simple derivative of the Hadamard $\mathbf{H_4}$ matrix:

$$\Phi_4 = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 2 \\ 0 & 0 & 2 & 2 \\ 0 & 2 & 2 & 0 \end{bmatrix} \quad \mathbf{H_4} = \begin{bmatrix} + & + & + & + \\ + & - & + & - \\ + & + & - & - \\ + & - & - & + \end{bmatrix}$$

The in-phase/antiphase (IPAP) method [49] and the so-called 'sensitivity-enhanced' and 'time-share' experiments are also examples of the Hadamard transform, involving the ($\mathbf{H_2}$) matrix. There are many other techniques that are intrinsically based on Hadamard matrices – magnetic resonance imaging, stochastic NMR, and quantum computing.

Hadamard encoding is probably one of the simplest ways to reduce the experiment time in multi-dimensional NMR spectroscopy. Given one-dimensional spectra corresponding to two different species of nuclei A and B, a two-dimensional correlation experiment establishes the connectivities between the nuclei A and B. In the case of single-bond correlations, each cross peak represents a direct connection (a chemical bond) between the individual atoms A and B. Often in sparse spectra of small molecules there are significant empty spectral regions and it is easy to imagine an experiment that avoids encoding these empty spaces and thus reduces the total measurement time. This can be achieved using site selective excitation, based on prior knowledge of the chemical shifts, in combination with the Hadamard encoding (see Fig. 2). Since the frequency selection is achieved using a set of simultaneous, multiply-selective (shaped) radiofrequency pulses, the Hadamard transform has the effect of replacing conventional time domain (evolution) data with frequency domain excitation. Following the Hadamard transform, the data matrix is reassigned to the appropriate frequency grid, so that the final format is equivalent to a spectrum recorded using the conventional methodology.

The only difference from the conventional experiment is that the spins evolve in the presence of multiple radiofrequency fields, instead of evolving freely. Because all encoding pulses are of equal duration, the experiment bears a strong resemblance to the constant time experiments (also with respect to the relaxation losses). In experiments that use the native Hadamard encoding the all '+' column is often discarded because it tends to accumulate interfering signals ('residuals') that do not respond to the phase variations.

Since the Hadamard matrix is real there is no need to record complex (phase sensitive) data in the indirectly detected dimension F_1. Thus the experiment time is reduced by another factor of two. There is also a $\sqrt{2}$ advantage in sensitivity compared to the conventional experiments due to the use of square-wave (rather than sinusoidal) modulation. Extension to higher-dimensional experiments [12] involves one or more Hadamard encoded

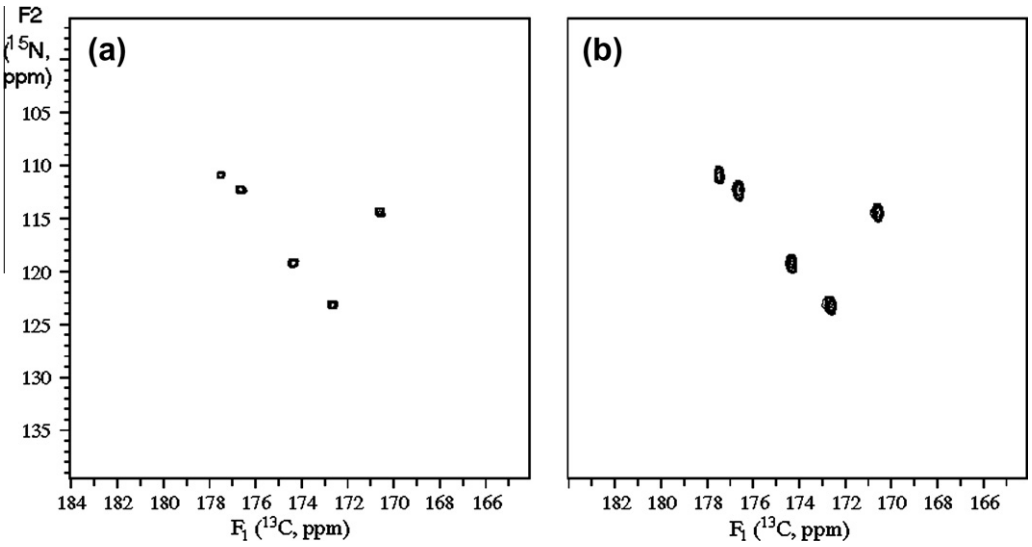

Fig. 1. The F_1F_2 plane extracted at proton frequency of 8.77 ppm from the three-dimensional 500 MHz HNCO spectrum of ubiquitin and comparing the projection-reconstruction results (left) with the conventional method (right).

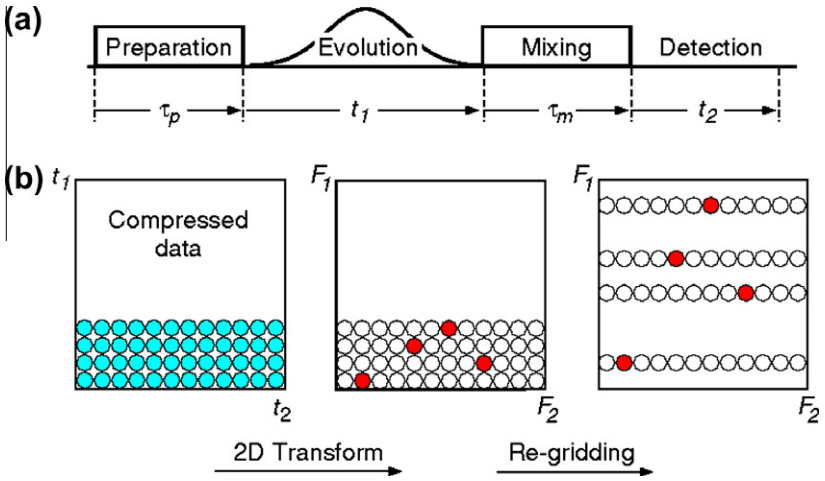

Fig. 2. The Hadamard and phase-encoding experiments (a) depicted in terms of the classical scheme for a two-dimensional experiment, and (b) showing the data processing steps.

dimensions and relies on prior knowledge of chemical shifts in these dimensions.

3.1. Phase encoding

One minor drawback of the Hadamard method is that the number of F_1 encoding sites (the number of NMR peaks) is not usually a power of two or a multiple of four. As a result one usually needs to record slightly more experiments than strictly necessary. Phase encoding [50] can be used instead. Rather than a combination of alternating phases, as in the Hadamard scheme, the phase-encoding method increments the phases at the individual sites in regular increments of $\Psi_i = 2\pi i/N$ ($i = 0 \ldots N-1$). Thus the spins evolve in the presence of the encoding radiofrequency fields as if there were no gaps between the peaks in the F_1 dimension. This requires recording a complex data set.

4. Multiple receivers

There are many different kinds of spin manipulation that involve polarization or coherence transfer between heteronuclei, but in most cases indirect detection methods are used and the actual acquisition is restricted to protons. We can now widen our horizons by exploiting multiple receivers [1,2] that operate in parallel (PANSY). An existing broadband or 'triple resonance' radiofrequency probe is connected via a separate preamplifier and mixer to an independent receiver channel comprising a local oscillator, amplifier, digitizer and controller board for each nuclear species under investigation. These hardware modifications are complemented by software that allows the different signals to be acquired either simultaneously or at different stages of the pulse sequence. One advantage is that ^{13}C or ^{15}N spectra with wide spectra bands and high resolution may be recorded directly with no significant time penalty, rather than through the usual indirect detection schemes such as HSQC or HMQC. Combinations of different pulse sequences may require compromises to be made to accommodate potential conflicts between phase cycles, field gradients, or decoupling protocols, but these complications can easily be resolved. The general result is a significant improvement in the quality and flow of spectral information.

4.1. An 'all-in-one' experiment: PANACEA

One exciting application of the new multiple-receiver technology is the PANACEA [4–6] scheme (Parallel Acquisition NMR, an All-in-one Combination of Experimental Applications). The consolidation of INADEQUATE, HSQC and HMBC in a single sequence often yields enough information to solve the structure of a small organic molecule, since the INADEQUATE feature provides reliable evidence for the carbon framework, while the HSQC sequence establishes the proton positions, and the HMBC experiment gives the longer-range interactions and indicates the links through heteroatoms such as oxygen or nitrogen. While PANACEA represents a cure for most of the ills of small molecules, it needs more serious work before it can be applied to investigate macromolecules such as proteins [51].

There remains the problem of speeding up the INADEQUATE feature, since the response requires an interaction between two directly adjacent ^{13}C spins, and in samples with natural isotopic abundance this represents only one molecule in 857. The experiment can be speeded up by multiple selective excitation with Hadamard encoding. Once the ^{13}C chemical shifts have been recorded, all the ^{13}C sites are irradiated simultaneously with a set of selective radiofrequency pulses. A one-dimensional trace showing a single ^{13}C–^{13}C correlation for each directly bound pair of carbon atoms is extracted at the decoding stage.

Applied to a solution of menthol (30% in $CDCl_3$) in a 500 MHz spectrometer, this extension of the PANACEA protocol delivers the INADEQUATE spectrum in only 56 s (Fig. 3). There are ten carbon sites. In principle a 12 by 12 Hadamard matrix would suffice for this case but in practice a 16 by 16 matrix was employed for the practical reason of better suppression of certain residual artefacts. Note that in the same PANACEA experiment the HSQC and HMBC features are also recorded simultaneously. The complete structure is obtained in less than a minute [6].

4.2. Multiple receivers in protein NMR

The parallel acquisition schemes outlined above were designed for application to small molecules. For larger biomolecules quite different pulse sequences are required and clearly the INADEQUATE experiment is no longer applicable. Furthermore, it has become the custom to employ global isotopic enrichment in ^{13}C and ^{15}N for biochemical studies, whereas the small molecule experiments dealt with natural-abundance samples that had independent coupled spin systems. It is nevertheless quite feasible to use the same basic concept to combine two or more standard protein NMR sequences into a single entity using parallel acquisition of signals from two or more nuclear species.

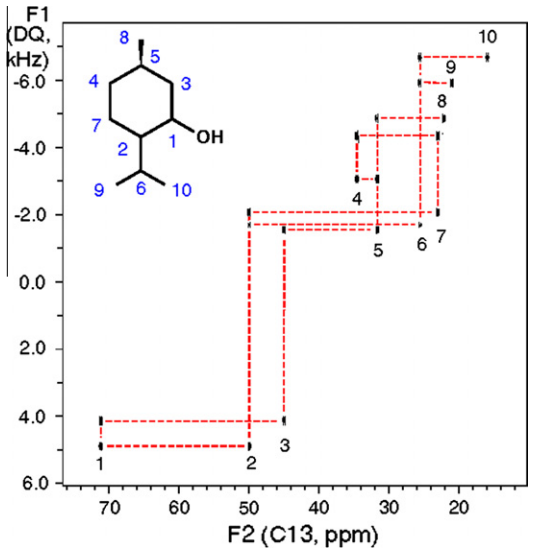

Fig. 3. The INADEQUATE spectrum of menthol (30% in $CDCl_3$) recorded in only 56 s on a 500 MHz spectrometer with a cold probe optimized for ^{13}C detection. An HSQC spectrum was also recorded in parallel in this PANACEA experiment.

One practical example will serve to indicate the usefulness of combining standard protein NMR sequences, using several receivers. Sensitivity considerations normally dictate that the majority of experiments on proteins focus on proton detection, but interest is growing on direct detection of ^{13}C or ^{15}N nuclei, partly because they are less susceptible than protons to broadening by paramagnetic species. One such ^{13}C direct-detection scheme is the two-dimensional (HA)CACO experiment where C_α shifts appear in the F_1 dimension and CO shifts in the F_2 dimension. Parallel acquisition of ^1H and ^{13}C signals permits this sequence to be consolidated with the three-dimensional (HA)CA(CO)NNH experiment with essentially no penalty in measurement duration. This 'all-in-one' pulse sequence is set out in Fig. 4. The first stage (indicated in black) detects the ^{13}C signals from the (HA)CACO sequence, followed by the (HA)CA(CO)NNH sequence (red). The two are 'tied together' by exploiting a residual ^{13}C signal from the two-dimensional experiment in the parallel three-dimensional experiment. The first part of the CO free induction decay is used in the 2D (HA)CACO experiment and the remaining weak tail of the free induction decay (afterglow) is retrieved for the 3D (HA)CA(CO)NNH experiment [51]. Although the ^{13}C afterglow is weak, the sensitivity is significantly enhanced by magnetization transfer to protons. By a judicial choice of the truncation point of the free induction decay, a compromise is struck between resolution in the direct-detection CO spectrum and the signal-to-noise in the ^1H detected spectrum.

The method was tested in a study of the 600 MHz spectrum of an aqueous solution (10% D_2O) of the 143-residue protein nuclease A inhibitor, globally enriched in ^{13}C and ^{15}N. A triple resonance HCN probe was employed, with cryogenically-cooled ^1H and ^{13}C coils and preamplifiers. The (HA)CACO and (HA)CA(CO)NNH data sets were acquired in a single 3h experiment [51]. The spectra are illustrated in Fig. 5.

To get an indication of whether this concept might be applied to even larger (more sluggish) proteins, a test was made to judge the influence of the rate of molecular tumbling in solution. The nuclease A inhibitor experiments at 25 °C were repeated at 2 °C with essentially comparable spectra, albeit with slightly reduced signal-to-noise ratios. Experiments on ^{15}N relaxation [52] had already demonstrated that the effective correlation time for molecular reorientation increases from about 9 ns at 25 °C to about 17 ns at 2 °C, suggesting that these combined NMR sequences could be considered for studies on larger proteins.

While the nuclease A inhibitor experiment required 3 h of spectrometer time, in principle this could be considerably shortened by employing the projection-reconstruction technique. In order to demonstrate this point, the three-dimensional (HA)CA(CO)NNH spectrum of a small protein GB1 was reconstructed from just three projected planes [51]. Two of these were the orthogonal 'first planes' F_1F_3 (obtained with t_2 set to zero) and F_2F_3 (with t_1 set to zero). The third was a projection tilted about the F_3 (NH) axis through an angle 69.5° from F_1 (^{13}CA) towards F_2 (^{15}N) and obtained by Fourier transformation of a radially sampled signal where t_1 and t_2 were incremented in concert. This particular tilt angle was chosen to take advantage of the narrower lines and better shift dispersion of the ^{15}N signals. Fig. 6 compares the conventional three-dimensional spectrum (sampled on a full Cartesian grid and requiring 3 h 10 min of spectrometer time) with this projection-reconstruction version recorded in only 15 min. Where sensitivity is not the primary concern, significant speed advantages should be generally achievable by related projection-reconstruction methods.

5. Single-Point Evaluation of the Evolution Dimension (SPEED)

In three-dimensional spectroscopy, both the Hadamard and phase encoding techniques are clear about the number of data points that are required in the indirect dimensions – it is equal to the number of cross-peaks (correlations) that need be detected. Note, that even if all the peaks overlap in the directly-detected dimension (F_N), both the Hadamard and the phase encoding techniques still resolve all the correlations in the indirect dimensions. However in practice, complete overlap is very unlikely, and even in very crowded spectra few peaks overlap at a single frequency in F_N. This means that a smaller number of samples can be used. For instance, in the simplest case when all peaks are resolved in F_N, a single (hypercomplex) measurement can provide the required information. Measurement of the phase of the observed peaks provides the information about the chemical site where the magnetization originates, and hence the two-dimensional correlation spectrum can be reconstructed. In fact, there is no need for selective pulses or phase encoding in such an experiment; it is only necessary to measure the 'natural' phase evolution after a period t_1. This technique is called SPEED [53]. The phase needs be measured with a high enough accuracy to distinguish between the encoded sites, and this requires an adequate signal-to-noise ratio. The frequencies can be calculated according to a simple relationship between the phase, ϕ_i and the frequency, f_i:

$$\phi_i = 2\pi f_1 \Delta t \tag{3}$$

where Δt is the duration of the phase evolution time. If the cosine and sine-modulated magnetization components are detected simultaneously, the two-dimensional correlation spectrum can be reconstructed from this single measurement. In practice, the frequency that is calculated from Eq. (1) may not be very accurate and will depend strongly on the signal-to-noise ratio and on the incidence of various instrumental imperfections (eddy currents, gradient recovery times, probe ring-down times, pulse imperfections, *et cetera*). Use of prior knowledge to measure the chemical shifts (namely the ^{13}C spectrum), will provide more accurate correlation maps. On the other hand, the accuracy of the frequency determination will also improve with increasing the number of measurements. A simple linear least-squares fit can be used to determine the frequency from several measurements. If the time is incremented in regular intervals, the conventional Fourier transform combined with one of the linear prediction techniques may provide a similar result. In fact, the filter diagonalization technique [28,54] is a more sophisticated version of such an approach.

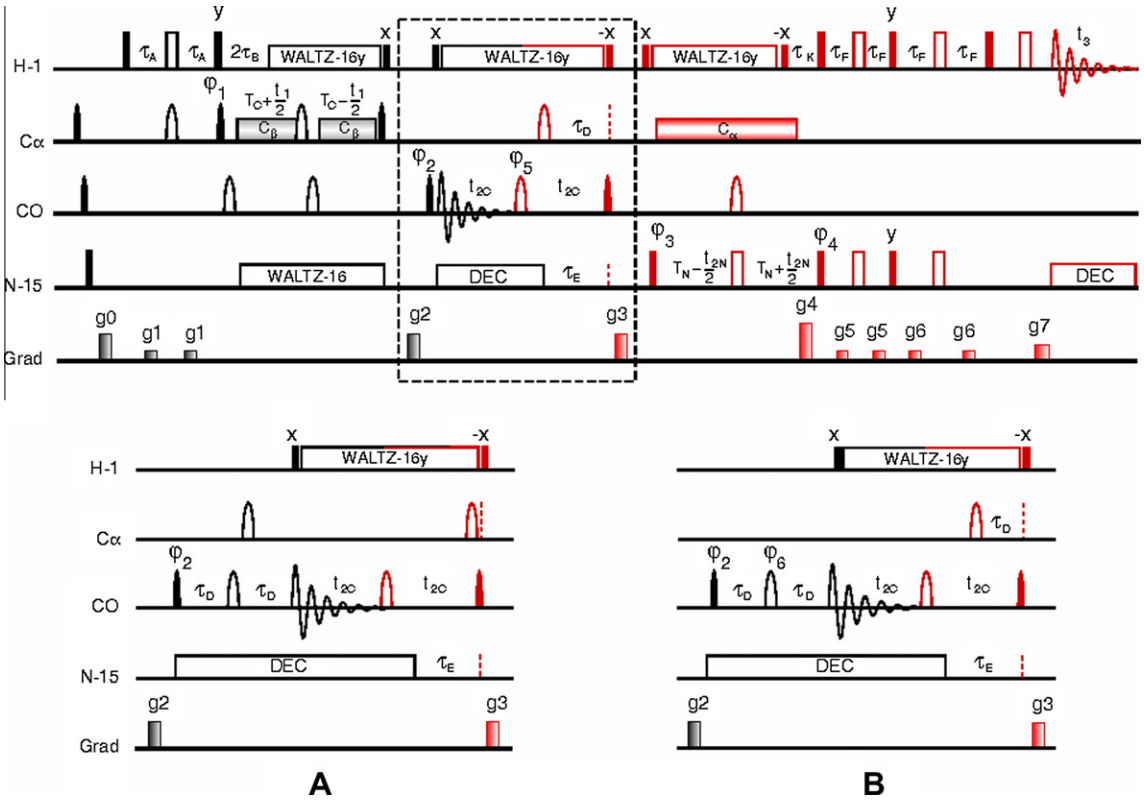

Fig. 4. The combined 2D (HA)CACO (black)/3D (HA)CA(CO)NNH (black + red) dual-receiver experiment [51]. The C′–C_a couplings are eliminated using the IPAP versions of the pulse sequence where the dashed box is replaced by panels A and B.

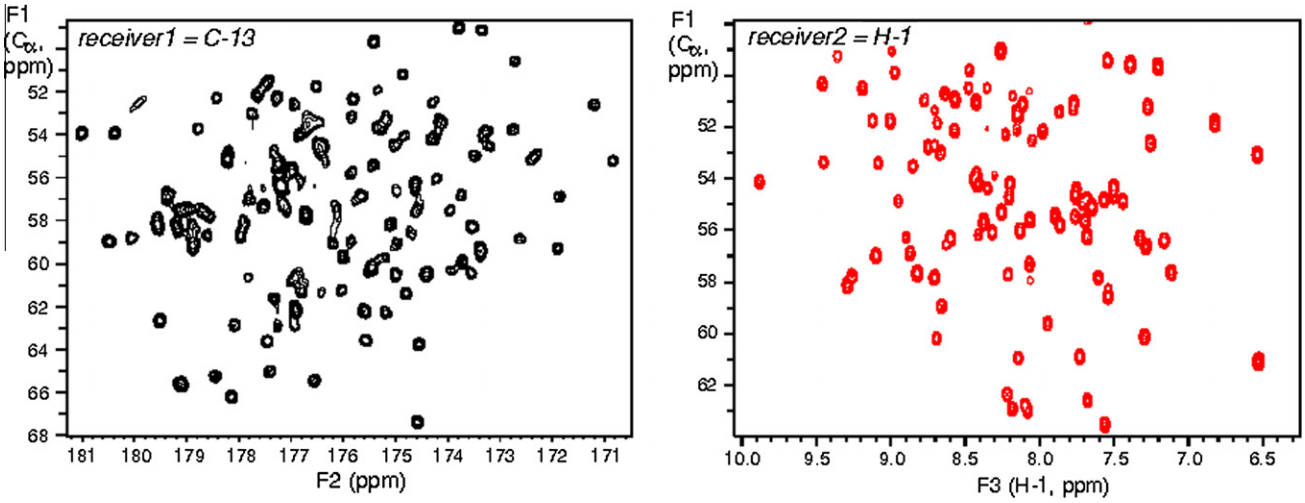

Fig. 5. 600 MHz spectra of globally enriched 143 residue protein nuclease A inhibitor recorded with the pulse sequence of Fig. 4 on a system equipped with two receivers, ^{13}C (left) and ^{1}H (right) in a single 3-h experiment.

5.1. A practical application of 'SPEED'

In a three-dimensional experiment, good sensitivity is achieved in the 'first planes' F_1F_3 and F_2F_3 recorded as the Fourier transforms of signals $S(t_1, 0)$ and $S(0, t_2)$ respectively. It is the common practice to run these measurements as part of the setting-up procedure. Consider for the purpose of illustration that we are interested in the ^{13}C–^{15}N correlations in an HNCO spectrum by direct detection of the NH proton signals [55]. In the simple case where all the NH signals are well resolved, all the correlations are obtained directly from the spectra recorded in the first planes, since each individual NH resonance is coupled to just a single ^{15}N and a single ^{13}C resonance. Complications only arise where two NH signals overlap and cannot be resolved. Suppose we then perform a 'single-point' measurement where the ^{15}N spins evolve for a fixed interval t_1^* and the ^{13}C signals for an interval t_2^*. The degenerate NH site is then associated with two evolving ^{15}N frequencies and two evolving ^{13}C frequencies, suggesting *four* possible correlation peaks. Two of these must be false. If the two overlapping NH peaks have significantly different intensities, then the ambiguity is resolved merely by inspection, since the strong cross-peak correlates with strong, and weak with weak.

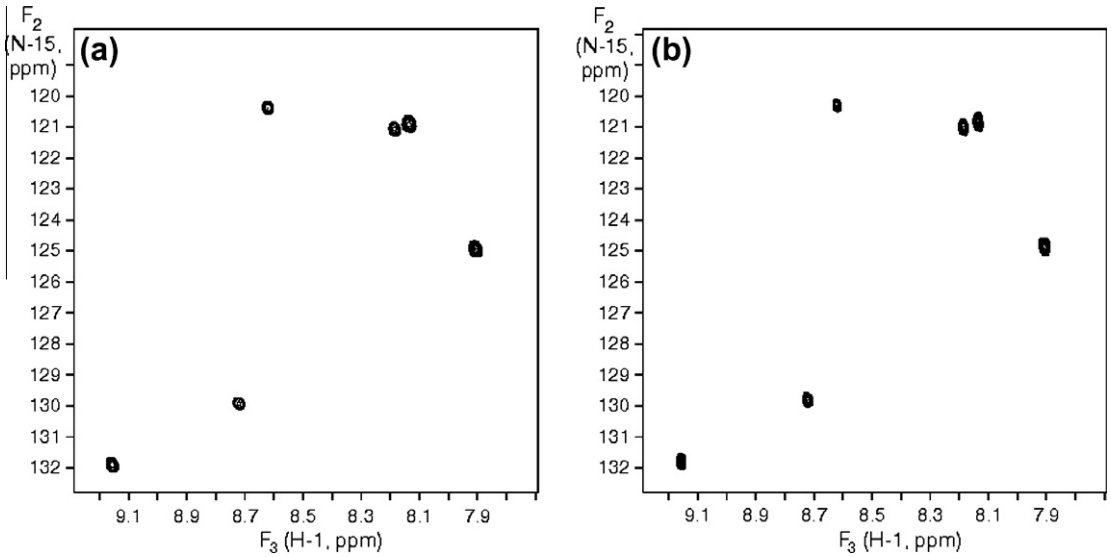

Fig. 6. Comparison of conventional (a) and projection-reconstruction (b) (HA)CA(CO)NNH spectra of a small protein GB1 recorded with the pulse sequence of Fig. 4 in 3 h (a) and 15 min (b) respectively. The spectra represent the F_2F_3 planes extracted at F_1 = 51.93 ppm. The optimum projection angle of 69.5° was used to record the spectrum (b).

The general case is more challenging. Suppose there are N degenerate NH frequencies (with intensities A_i). As an example, when $N = 4$ there would be 24 possible arrangements of the pattern of the cross-peak. Signals are recorded at the single point (t_1^*, t_2^*) with the phases of the radiofrequency pulses set (0°, 0°), (0°, 90°), (90°, 0°), (90°, 90°). The corresponding NH proton intensities are modulated by the four terms:

$$S(0^\circ,0^\circ)=\sum_{i=1}^{N}A_i\cos(\omega_C t_1^*)\cos(\omega_N t_2^*) \tag{4}$$

$$S(0^\circ,90^\circ)=\sum_{i=1}^{N}A_i\cos(\omega_C t_1^*)\sin(\omega_N t_2^*) \tag{5}$$

$$S(90^\circ,0^\circ)=\sum_{i=1}^{N}A_i\sin(\omega_C t_1^*)\cos(\omega_N t_2^*) \tag{6}$$

$$S(90^\circ,90^\circ)=\sum_{i=1}^{N}A_i\sin(\omega_C t_1^*)\sin(\omega_N t_2^*) \tag{7}$$

The key point is that these *experimentally recorded* intensities can be *calculated* on the basis of the intensities and frequencies observed in the spectra in the first planes F_1F_3 and F_2F_3. Of the several possible solutions, only one set matches the experimental intensities. All ambiguity is resolved. Furthermore, the information recorded in the spectra in the first planes can also be used to suggest the best choice for t_1^* and t_2^*. Note that this method for resolving ambiguity needs to be applied only where there are instances of NH overlap; all other correlations are straightforward. Fig. 7 shows a comparison between the conventional 500 MHz correlation spectrum of ^{13}C,^{15}N enriched agitoxin, with that reconstructed from a single-point measurement ($t_1^* = 2.30$ ms and $t_2^* = 1.93$ ms), together with the information from the first planes F_1F_3 and F_2F_3. The fact that the two spectra were run at slightly different temperatures accounts for the small differences in chemical shifts [55]. In the general case, as more and more frequency dimensions are introduced, the probability of undesirable overlap recedes and therefore the method for single-point sampling of evolution space becomes progressively more attractive.

6. Aliasing

Some years ago oversampling and digital signal processing were introduced as a new means of improving the dynamic range of NMR systems and the quality of NMR spectra. Increasing the sampling rate by a factor of four beyond the Nyquist condition is equivalent of two additional bits of resolution in the analog-to-digital converter. Undersampling has the opposite effect – data compression is achieved at the expense of reducing the dynamic range. Although this might be regarded as undesirable or counterproductive, it yields valuable dividends in speed. In general, recording reduced amount of data using non-conventional sampling techniques is known as *compressive sensing* and (in fields other than NMR) is often performed by simultaneously collecting and compressing large amounts of data. The James Bond style sampling (the 'shoot first, ask questions later' approach) can be expensive in terms of equipment and sensors that are needed, or in terms of the required time. The compressive sensing which deliberately samples at a rate lower than specified by the Nyquist condition is particularly useful in multidimensional NMR. By accepting the consequent aliasing, it offers a dramatic reduction in measurement time. Some less critical features of the spectrum are sacrificed in return for speed.

6.1. A practical application of multiple aliasing

In the general case, aliasing of spectra [8] must be used with a certain degree of caution because it leaves ambiguities about chemical shift assignments and it suffers problems associated with increased spectral overlap. However, in the particular case of the INADEQUATE experiment, undersampling in the evolving double-quantum dimension (F_1) can be exploited with far more freedom. There is no fine structure in the F_1 dimension to create entangled spin multiplets. More importantly, even when multiple aliasing causes F_2 traces to overlap exactly in the F_1 dimension, the correlation information is not usually jeopardized. As a consequence, an appreciable improvement in speed can be achieved by recording INADEQUATE spectra in a regime that violates the Nyquist condition to a significant degree.

Fig. 8 shows the PANACEA ^{13}C–^{13}C correlation spectrum of 1 M cholesterol in $CDCl_3$ recorded on a 500 MHz spectrometer [6]. Controlled aliasing in the F_1 dimension gives a fourfold reduction in the experimental duration. The assignment of ^{13}C–^{13}C connectivities is unaffected by the aliasing, permitting the PANACEA sequence to be completed in only 23 min, including the multiplicity-edited HSQC measurement (Fig. 9). The INADEQUATE scheme

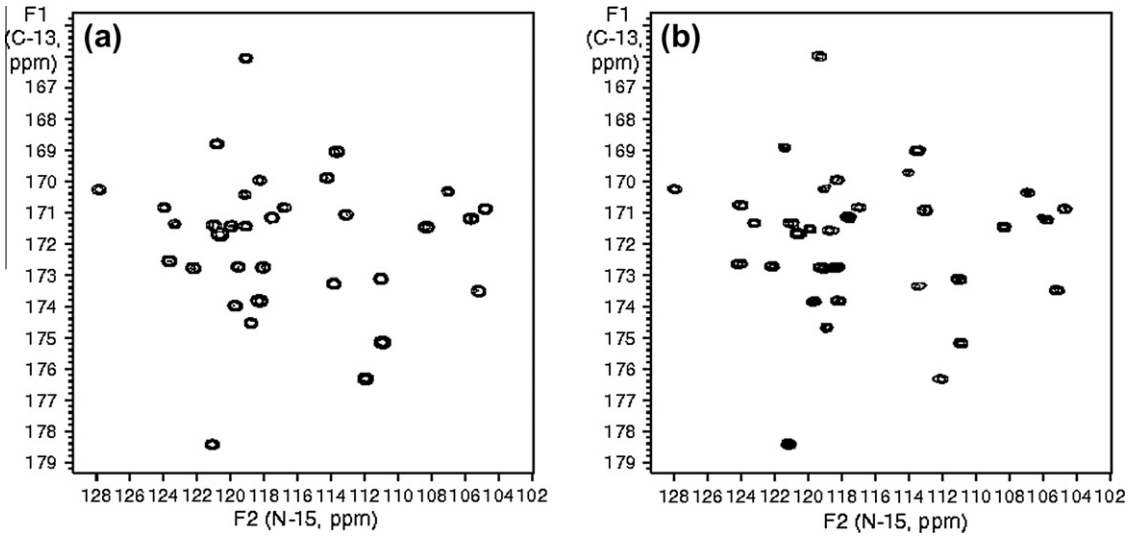

Fig. 7. Comparison between the F_1F_2 projections of (a) the conventional 500 MHz HNCO spectrum of agitoxin recorded in 11 h, with (b) that reconstructed from an 18 s single-point measurement (t_1^* = 2.30 ms and t_2^* = 1.93 ms) and the prior information from the first planes F_1F_3 and F_2F_3.

is not limited to ^{13}C–^{13}C correlations. A second PANACEA application demonstrates that ^{29}Si–O–^{29}Si correlations can be measured in pure silicone oil (Fig. 10). Magnetization transfer from the protons has been exploited to enhance the ^{29}Si sensitivity (all the ^{29}Si sites are coupled to protons). An internal frequency regulation scheme operates on a strong ^{29}Si resonance, correcting any slow drifts of the applied magnetic field [6]. Although the ^{29}Si nucleus is 4.7% abundant, this advantage is countered in the INADEQUATE experiment by the small magnitudes of the ^{29}Si–O–^{29}Si coupling constants (~2 Hz). Aliasing in the double-quantum dimension reduces the spectral width by a factor of 16. The INADEQUATE and the simultaneous HMBC measurement were complete in only 8 min 22 s.

7. Incoherent sampling

Although magnetic resonance is principally concerned with coherent processes, there are occasions where it is advantageous to impose some incoherence. For example, a radiofrequency transmitter can be deliberately rendered incoherent in order to excite spins over a wide frequency band. One case in point is the super-regenerative oscillator [56], a pulsed amplifier with a high degree of positive feedback, in which oscillations are initiated by Johnson noise in the receiver coil. Oscillations build up exponentially but are quickly damped by quench pulses. In this incoherent mode, the device is capable of exciting NMR signals anywhere within a wide frequency range, but once such an NMR response is detected,

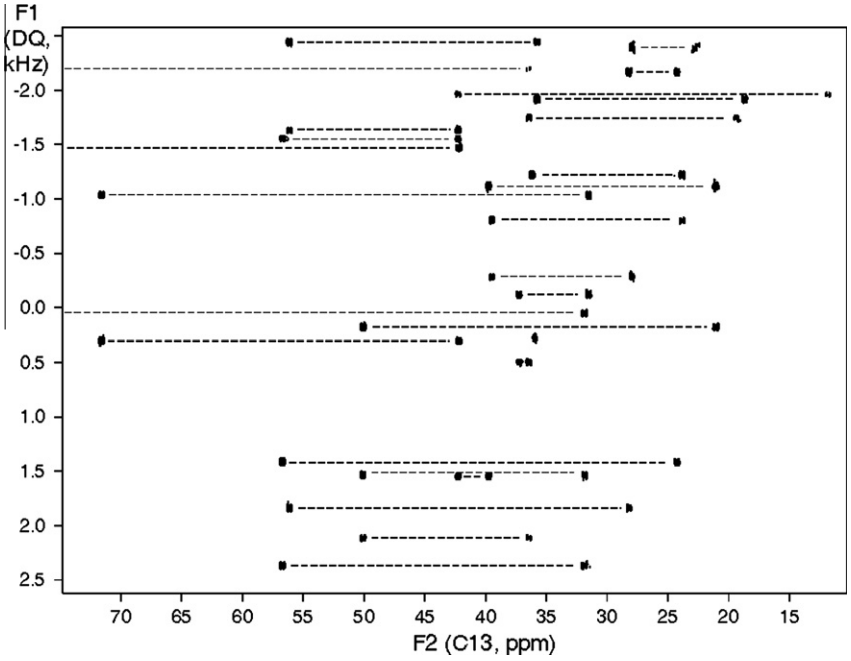

Fig. 8. The INADEQUATE spectrum of 1 M cholesterol in $CDCl_3$ recorded on a 500 MHz spectrometer, with fourfold aliasing in the double-quantum (F_1) dimension, induced by undersampling. Experimental duration: 23 min.

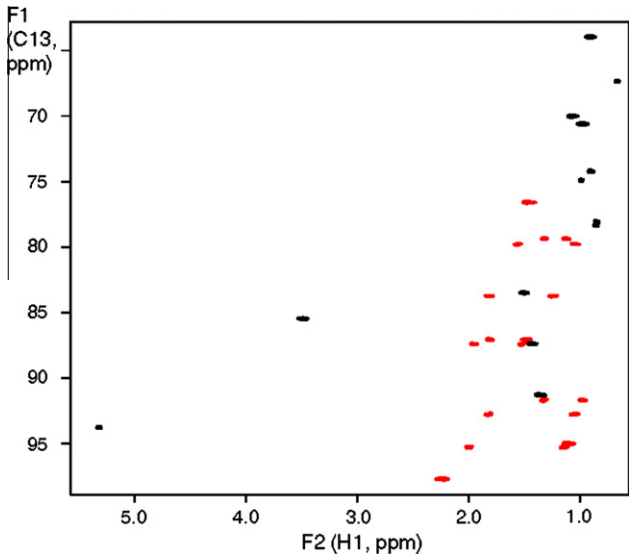

Fig. 9. The multiplicity-edited HSQC spectrum of cholesterol recorded in the same PANACEA experiment as Fig. 8.

the oscillator reverts to a coherent mode, where the phase of each oscillation is determined by the NMR signal rather than by the random noise. The useful result is that the detected NMR frequency then follows the variations in the applied magnetic field with only a very small error.

A second example is noise decoupling [57] where an incoherent radiofrequency oscillator decouples protons over an appreciable range of chemical shifts. The radiofrequency phase is modulated with a time-dependent signal that approaches a true random function, with a view to avoiding any small *coherent* contribution to the frequency-domain spectrum. In practice, because a strictly random function is difficult to implement, a deterministic pseudo-random sequence is employed. This was highlighted in the early versions of noise decoupling where 'incoherence' was approximated by a digital function derived from the telephone numbers listed in the Palo Alto telephone directory. Encryption science is bedevilled by similar problems related to the generation of a long pseudo-random function that is sufficiently complex that the code cannot be 'broken'.

There are some experiments that generate an undesirable coherent NMR response that must be suppressed. If radiofrequency pulses are applied to a spin system at intervals shorter than the spin–spin relaxation time T_2, transverse magnetization settles into a steady-state regime, but because the precession phase varies in a cyclic manner across the spectrum, undesirable phase and intensity anomalies are generated. A simple way to suppress these is to deliberately introduce incoherence by adding a slight 'jitter' on the timing of the radiofrequency pulses, thus scrambling the NMR phases [58]. A more sophisticated solution disperses the unwanted transverse signals in a sequence of applied magnetic field gradients [59], but there remains the problem of 'randomizing' the magnitudes and directions of these gradients to avoid echo effects.

We may conclude that *artificial* generation of a *truly* random function is probably impossible to achieve. Nevertheless the term 'random sampling' is widely used to denote schemes for recording NMR signals at time co-ordinates selected in a pseudo-random manner. Often there is the additional constraint that these chosen data points should fall on the usual Cartesian grid in evolution space, allowing the Fourier transform to be applied directly. If the desire for speed demands *sparse* sampling, incoherence has the effect of dispersing the sampling artefacts into 't_1-noise' like ridges spanning a wide frequency range across the NMR spectrum. The nearer the induced incoherence approaches true randomness, the more these residual artefacts resemble true 'white' t_1-noise. In applications where signal-to-noise is no longer a critical parameter, this is an acceptable strategy – there is simply a relatively moderate sacrifice in sensitivity. In the opposite case of poor signal strength, random sampling would normally be avoided, unless it provides some other overriding advantage. A smaller set of

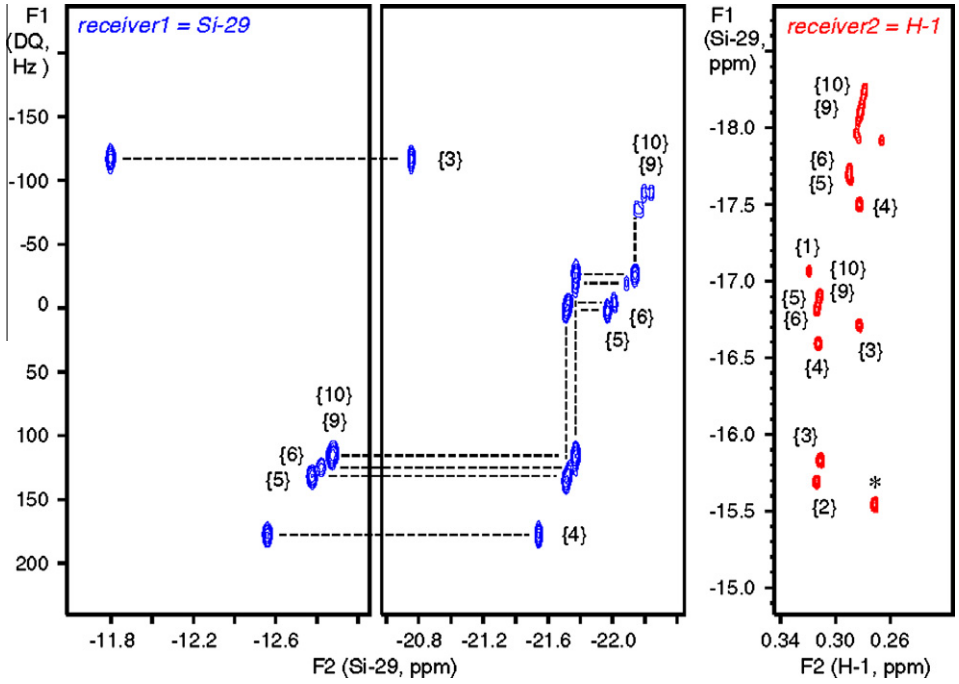

Fig. 10. INADEQUATE ^{29}Si–O–^{29}Si correlations measured in pure silicone oil. Aliasing has reduced the spectral width in the double-quantum dimension by a factor of 16. This result and the corresponding HMBC measurement were completed in only 8 min 22 s.

stronger artefacts at *predictable* locations can usually be suppressed more efficiently by means of a processing algorithm such as CLEAN [29–32].

7.1. A practical application of random sampling

In spite of these general caveats about random sampling, there are certain experiments where it offers real advantages compared with the conventional regular digitization protocol. One such application is based on the desire to circumvent the restrictions imposed by the Nyquist condition. Although multidimensional spectra provide better *effective* resolution in the sense that peaks are dispersed into additional frequency dimensions, the *actual* resolution (as determined by the line width) normally remains poor, because of the practical limits on the number of sampling points in evolution space. For example, with such coarse digitization in the evolution dimensions, a four-dimensional spectrum may have apparent linewidths of hundreds of Hz, causing undesirable overlap and complicating the assignment. Normally when the Nyquist condition is violated there are serious problems of aliasing. In contrast, random sampling transforms the aliased signals into noisy artefacts, rather evenly distributed across the spectrum. This allows the sampling rate to be greatly reduced compared with the Nyquist condition, and the maximum evolution time to be set much longer than in the conventional mode [60]. Resolution in the indirect dimensions is greatly enhanced, but at the expense of these additional spectral artefacts which compromise the sensitivity.

The key parameter in random sampling is the peak-to-artefact ratio, which is proportional to the ratio of the square root of the number of time-domain data points compared with the number used for full conventional sampling. However it is independent of the number of dimensions and the maximum evolution times [60]. Consequently it is better to increase the number of sampling points than to increase the number of repeated scans. In addition there must be sufficient time-domain sampling points for an adequate definition of the narrow lines in the frequency domain, of the order of three points per line width. This implies problems for the transformation of such a large amount of time-domain data, particularly as the number of dimensions increases. Fortunately the high dispersion of peaks in these randomly-sampled spectra allows the Fourier transformation to be focused on specific resonance frequencies (for example those of amide protons or carbonyl carbons) recorded in prior measurements in lower-dimensional spectra. This reduces a formally four-dimensional spectrum to a set of two- or three-dimensional projections, each with high digital resolution.

There are several possible schemes for deciding the detailed form of the randomization. Zawadzka–Kazimierczuk et al. [61] decided *not* to restrict the evolution-domain sampling points to the usual Cartesian grid. However for sparse sampling on a fine grid, it could be argued that the difference between *on-grid* and *off-grid* random sampling is small enough to be neglected, while on-grid sampling simplifies the transformation process. In addition they replaced a purely random array with a 'Poisson-disc' distribution, where there is a sample-free circle around each sampling point in the time domain. This reduces low-frequency noisy modulation of the NMR signals caused by random sampling, creating a low artefact region around each resonance in the frequency domain, thus improving the signal-to-artefact ratio (but only in that particular region). These ideas have been exploited to record high-resolution four-dimensional spectra of a 96-residue protein and a 341-residue maltose-binding protein on a 700 MHz spectrometer [61]. This example illustrates the general point that acceptance of sparse random sampling artefacts (with the consequent sacrifice in sensitivity) can be justified if a crucial new feature (better resolution) can be achieved in return.

8. Seeing the bigger picture: hyperdimensional NMR

The complexity of protein molecules drives NMR spectroscopy inexorably towards higher dimensions, where the sheer volume of data and the necessary spectrometer time soon become quite unrealistic. While four-dimensional NMR spectroscopy has become a rather 'routine' technique in protein studies, increasing the dimensionality of NMR experiments any further is clearly not so simple. One research group [62] has recorded a 'conventional' five-dimensional spectrum using random sampling to make the acquisition times and data matrices manageable, and a variation of the projection-reconstruction technique has been used to reach seven frequency dimensions [63]. Although a 10-dimensional investigation might seem an attractive project, a measurement made by the conventional procedure with standard parameters would tie up the spectrometer indefinitely.

Hyperdimensional NMR [64,65] describes a new way to study NMR spectra with dimensionality exceeding the threshold of four dimensions. It bypasses the unmanageable data space and experiment time constraints and provides an easy and comprehensible approach to study spectra of any dimensionality. We should really associate the concept of dimensionality with the size of the repetitive coupled spin system of the molecule under investigation, in this case the protein backbone or side-chain. The actual measurements are lower-dimensional projections of the corresponding hypermatrix; for obvious practical reasons of visualization, spectra of any dimensionality above two are studied as plane projections.

We imagine a *virtual N*-dimensional matrix, and use it as a scaffold on which two or more lower-dimensional results are assembled, so that desired new correlations can be derived. The human brain has evolved the capacity to form a three-dimensional impression of the outside world, by combining information from a pair of slightly different two-dimensional images. These images have the vertical dimension in common, but differ with respect to the horizontal dimension. An extension of this principle can be employed in multidimensional NMR. In the simplest example, two three-dimensional spectra S(F_1, F_2, F_3) and S(F_1, F_4, F_5) that have a frequency dimension (F_1) in common can be joined together to give information about correlations in the F_3F_5 plane. (We assume no exact resonance overlap in the common dimension F_1). Note that interactions between spins with resonances in the dimensions F_3 and F_5 might not be accessible from conventional experiments.

In general any pair of low-dimensional results can be combined to construct a higher-dimensional spectrum in order to highlight some more inaccessible interaction within the extended coupling network. In practice the full *N*-dimensional matrix would never be completed, for one thing it would require a completely unrealistic storage capacity. It merely acts as a conceptual framework for constructing a new matrix and then projecting the sub-spectra of interest. In this sense hyperdimensional NMR is a sparse sampling technique.

To fix ideas consider a practical example. Two adjacent amino-acid residues constitute a coupled 10-dimensional spin system, so the hyperdimensional approach lends itself well to this problem, particularly where sequencing of the residues is important. The 10 spins in question would be labeled NH(i–1), N(i–1), CA(i–1), CH(i–1), CO(i–1), N(i), NH(i), CA(i), CH(i), and CO(i). If we consider all the possible pairwise interactions, they would correspond to 45 different plane projections of the hypothetical 10-dimensional matrix. Any one of these 45 subspectra may be obtained by the procedure outlined above.

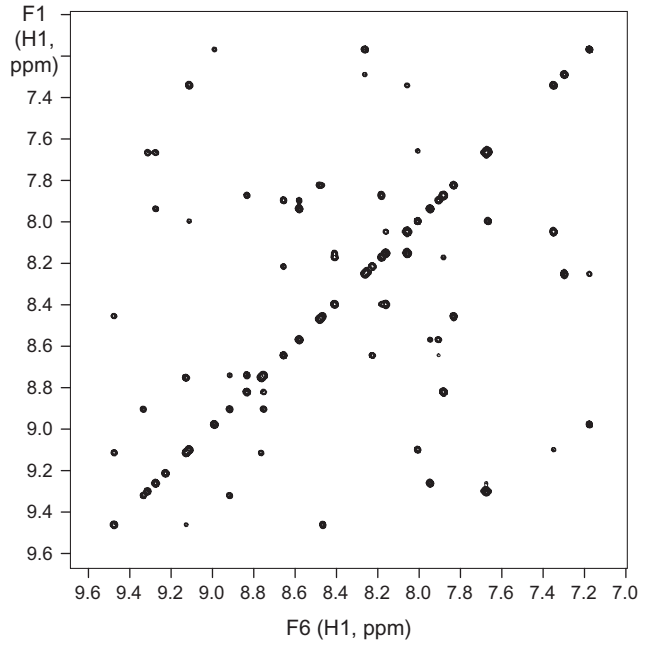

Fig. 11. Hyperdimensional NMR of agitoxin. The F_1F_6 plane showing correlation between NH groups on adjacent residues. It represents a projection of the five-dimensional spectrum obtained by combining three-dimensional HNCA and HN(CO)CA spectra on an imaginary 10-dimensional data matrix.

8.1. A practical application: agitoxin

A single application must suffice to demonstrate the method. Consider the case of a small 39-residue protein, agitoxin, isotopically enriched in ^{13}C and ^{15}N, and studied as a 0.3 mM aqueous solution (10% D_2O) on an 800 MHz spectrometer. Agitoxin contains a large number of adjacent pairs of aminoacid residues; hence any single plane projection of the virtual ten-dimensional matrix comprises many correlation peaks. Suppose we wish to establish a homonuclear correlation between NH protons on adjacent residues, corresponding to a spectrum with the characteristic line of responses along the principal diagonal with correlations defined by off-diagonal peaks. A standard three-dimensional HNCA sequence links NH(i–1) with C_α(i–1), while a standard three-dimensional HN(CO)CA sequence links NH(i) with C_α(i–1). These experiments have a common frequency axis (F_3) corresponding to the C_α(i–1) sites. The combination of the two allows us to construct a virtual five-dimensional spectrum linking NH(i–1) with NH(i). This is a five-dimensional projection of the virtual ten-dimensional matrix (Fig. 11).

The general procedure is compatible with many of the other sparse sampling schemes described above. In particular, the multiple receiver technology provides a new way to obtain the required information even faster as many of the traditionally indirectly detected dimensions become available for simultaneous direct observation.

9. Conclusions

The Achilles heel of multidimensional NMR is the long duration of the experiment, imposed by the need for systematic exploration of all of evolution space, while satisfying the Nyquist condition and the demands for high resolution. The problem can be attacked in several different ways. One increases the flow of spectral information by employing multiple receivers to detect signals from several different nuclear species in parallel (PANSY). An important extension of this procedure allows two or more standard pulse sequences to be combined into a single entity (PANACEA) with no significant increase in experimental duration. A second scheme speeds up individual elements of the pulse sequence by, for example, drastic reduction in the time allowed for recovery by spin–lattice relaxation. In situations where the intrinsic sensitivity is already adequate, a third approach employs deliberate data compression methods that challenge the standard rules for adequate sampling of evolution space, often invoking prior knowledge from earlier one- or two-dimensional measurements. These last schemes include projection-reconstruction (radial sampling), Hadamard spectroscopy (multiple selective excitation) and several non-uniform sampling protocols such as exponential or random sampling. All sparse sampling schemes generate undesirable artefacts that can be suppressed by an iterative processing algorithm (such as CLEAN); if they arise from incoherence effects (random sampling) the resulting t_1-noise-like artefacts are spread evenly across the spectrum. Often the consequent loss in the dynamic range and sensitivity can prove acceptable in return for higher resolution or faster access to the spectra. Finally the concept of hyperdimensional NMR indicates a route to spectra of very high dimensionality.

This *Perspective* has focused on applications of sparse sampling that involve the use of multiple parallel receivers. Of course we have barely covered the tip of the 'iceberg'. There is a huge ongoing effort to further increase the efficiency of the NMR experiments with important contributions from the groups of Wüthrich et al. [66], Wagner et al. [67], Mishkovsky and Frydman [68], Markley et al. [69], Brutcher et al. [70], Zhang and Bruschweiler [71], Orekhov et al. [72], Hoch et al. [73], Celik and Shaka [74], Kozminski et al. [75] and many others. Cross-fertilization of ideas with similar developments in magnetic resonance imaging [75,76] will undoubtedly produce new techniques and new applications in this rapidly developing field.

References

[1] Ē. Kupče, R. Freeman, B.K. John, Parallel acquisition of two-dimensional NMR spectra of several nuclear species, J. Am. Chem. Soc. 128 (2006) 9606–9607.
[2] Ē. Kupče, S. Cheatham, R. Freeman, Two-dimensional spectroscopy with parallel acquisition of ^{1}H–X and ^{19}F–X correlations, Magn. Reson. Chem. 45 (2007) 378–380.
[3] Ē. Kupče, B. Wrackmeyer, Multiple receiver experiments for NMR spectroscopy of organosilicon compounds, Appl. Organomet. Chem. 24 (2010) 837–841.
[4] Ē. Kupče, R. Freeman, Molecular structure from a single NMR experiment, J. Am. Chem. Soc. 130 (2008) 10788–10792.
[5] Ē. Kupče, R. Freeman, High-resolution NMR correlation experiments in a single measurement (HR-PANACEA), Magn. Reson. Chem. 48 (2010) 333–336.
[6] Ē. Kupče, R. Freeman, Molecular structure from a single NMR sequence (fast PANACEA), J. Magn. Reson. 206 (2010) 147–153.
[7] H. Gesmar, J.J. Led, Two-dimensional linear prediction NMR spectroscopy, J. Magn. Reson. 83 (1989) 53–64.
[8] D. Jeannerat, High resolution in heteronuclear ^{1}H–^{13}C NMR experiments by optimizing spectral aliasing with one-dimensional carbon data, Magn. Reson. Chem. 41 (2003) 3–17.
[9] A. Bax, R. Freeman, T.A. Frenkiel, An NMR technique for tracing out the carbon skeleton of an organic molecule, J. Am. Chem. Soc. 103 (1981) 2102.
[10] J. Friedrich, S. Davies, R. Freeman, Two-dimensional spectroscopy without an evolution period – "pseudo-COSY", J. Magn. Reson. 75 (1987) 540–545.
[11] Ē. Kupče, R. Freeman, Two-dimensional Hadamard spectroscopy, J. Magn. Reson. 162 (2003) 300–310.
[12] Ē. Kupče, R. Freeman, Fast multidimensional Hadamard spectroscopy, J. Magn. Reson. 163 (2003) 56–63.
[13] Ē. Kupče, T. Nishida, R. Freeman, Hadamard NMR spectroscopy, Prog. NMR Spectrosc. 42 (2003) 95–122.
[14] J. Hadamard, Résolution d'une Question Relative aux Déterminants, Bull. Sci. Math. 17 (1893) 240–248.
[15] R.N. Bracewell, The Fourier Transform and its Applications, McGraw-Hill Higher Education, 2000.
[16] P. Schanda, B. Brutscher, Very fast two-dimensional NMR spectroscopy for real-time investigation of dynamic events in proteins on the time scale of seconds, J. Am. Chem. Soc. 127 (2005) 8014–8015.
[17] Ē. Kupče, R. Freeman, Fast multidimensional NMR by polarization sharing, J. Magn. Reson. 45 (2007) 2–4.

[18] K. Ding, A.M. Gronenborn, Novel 2D triple-resonance NMR experiments for sequential resonance assignments of proteins, J. Magn. Reson. 156 (2002) 262–268.
[19] S. Kim, T. Szyperski, G.F.T. NMR, A new approach to rapidly obtain precise high dimensional NMR spectral information, J. Am. Chem. Soc. 125 (2003) 1385–1393.
[20] W. Koźmiński, I. Zhukov, Multiple quadrature detection in reduced dimensionality experiments, J. Biomol. NMR 26 (2003) 157–166.
[21] L. Frydman, T. Scherf, A. Lupulescu, The acquisition of multidimensional NMR spectra within a single scan, Proc. Natl. Acad. Sci. USA 99 (2002) 15858–15862.
[22] J.C.J. Barna, E.D. Laue, M.R. Mayger, J. Skilling, S.J.P. Worrall, Exponential sampling: an alternative method for sampling in two-dimensional NMR experiments, J. Magn. Reson. 73 (1987) 69–77.
[23] K. Kazimierczuk, A. Zawadzka, W. Koźmiński, I. Zhukov, Random sampling of evolution time space and Fourier transform processing, J. Biomol. NMR 36 (2006) 157–168.
[24] K. Kazimierczuk, W. Koźmiński, I. Zhukov, Two-dimensional fourier transform of arbitrarily sampled NMR data sets, J. Magn. Reson. 179 (2006) 323–328.
[25] B.E. Coggins, P. Zhou, Sampling of the NMR time domain along concentric rings, J. Magn. Reson. 184 (2007) 207–221.
[26] F.L. Zhang, R. Bruschweiler, Indirect covariance NMR spectroscopy, J. Am. Chem. Soc. 126 (2004) 13180–13181.
[27] D. Malmodin, M. Billeter, Robust and versatile interpretation of spectra with coupled evolution periods using multi-way decomposition, Magn. Reson. Chem. 44 (2006) 185–195.
[28] J. Chen, V.A. Mandelshtam, A. Shaka, Regularization of the filter diagonalization method: FDM2K, J. Magn. Reson. 146 (2000) 363–368.
[29] J.G. Ables, Astron. Astrophys. Suppl. 15 (1974) 383.
[30] J.A. Högbom, Astron. Astrophys. Suppl. 15 (1974) 417.
[31] A.J. Shaka, J. Keeler, R. Freeman, Separation of chemical shifts and spin coupling in proton NMR: elimination of dispersion signals from two-dimensional spectra, J. Magn. Reson. 56 (1984) 294–313.
[32] Ē. Kupče, R. Freeman, Fast multidimensional NMR: radial sampling of evolution space, J. Magn. Reson. 173 (2005) 317–321.
[33] R. Freeman, Ē. Kupče, New methods for fast multidimensional NMR, J. Biomol. NMR 27 (2003) 101–113.
[34] Ē. Kupče, R. Freeman, Projection-reconstruction of three-dimensional NMR spectra, J. Am. Chem. Soc. 125 (2003) 13958–13959.
[35] Ē. Kupče, R. Freeman, Fast reconstruction of four-dimensional NMR spectra from plane projections, J. Biomol. NMR 28 (2004) 391–395.
[36] Ē. Kupče, R. Freeman, The radon transform: a new scheme for fast multidimensional NMR, Concepts Magn. Reson. 22A (2004) 4–11.
[37] Ē. Kupče, R. Freeman, Projection-reconstruction technique for speeding up multidimensional NMR spectroscopy, J. Am. Chem. Soc. 126 (2004) 6429–6440.
[38] R. Freeman, Ē. Kupče, Distant echoes of the accordion: reduced dimensionality, GFT-NMR, and projection-reconstruction of multidimensional spectra, Concepts Magn. Reson. 23A (2004) 63–75.
[39] R. Baumann, G. Wider, R.R. Ernst, K. Wüthrich, Improvement of 2D NOE and 2D correlated spectra by symmetrization, J. Magn. Reson. 44 (1981) 402–406.
[40] L. McIntyre, X.-L. Wu, R. Freeman, Fine structure of cross-peaks in truncated COSY experiments, J. Magn. Reson. 87 (1990) 194–201.
[41] C.D. Ridge, V. Mandelshtam, On projection-reconstruction NMR, J. Biomol. NMR 43 (2010) 51–159.
[42] Z.P. Liang, P.C. Lauterbur, Principles of Magnetic Resonance Imaging. A Signal Processing Perspective, IEEE Press, New York, 2000.
[43] S. Hiller, F. Fiorito, K. Wüthrich, G. Wider, Automated projection spectroscopy, Proc. Natl. Acad. Sci. 102 (2005) 10876–10881.
[44] S. Hiller, R. Joss, G. Wider, Automated NMR assignment of protein side-chain resonances using automated projections spectroscopy (APSY), J. Am. Chem. Soc. 130 (2008) 12073–12079.
[45] J.W. Yoon, S. Godsill, Ē. Kupče, R. Freeman, Deterministic and statistical methods for reconstructing multidimensional NMR spectra, Magn. Reson. Chem. 44 (2006) 197–209.
[46] Ē. Kupče, R. Freeman, Reconstruction of the three-dimensional NMR spectrum of a protein from a set of plane projections, J. Biomol. NMR 27 (2003) 383–387.
[47] G. Bodenhausen, R. Freeman, D.L. Turner, Suppression of artifacts in two-dimensional J-spectroscopy, J. Magn. Reson 27 (1977) 511–514.
[48] M. Ikura, L.E. Kay, A. Bax, Three-dimensional NOESY–HMQC spectroscopy of a ^{13}C-labeled protein, J. Magn. Reson. 86 (1990) 204–209.
[49] M. Ottinger, F. Delaglio, A. Bax, Measurement of *J* and dipolar couplings from simplified two-dimensional NMR spectra, J. Magn. Reson. 131 (1998) 373–378.
[50] Ē. Kupče, R. Freeman, Multisite correlation spectroscopy with soft pulses. A New phase-encoding scheme, J. Magn. Reson. A 105 (1993) 310–315.
[51] Ē. Kupče, L.E. Kay, R. Freeman, Detecting the "afterglow" of ^{13}C NMR in proteins using multiple receivers, J. Am. Chem. Soc. 132 (2010) 18008–18011.
[52] N.A. Farrow, R. Muhandiram, A.U. Singer, S.M. Pascal, C.M. Kay, G. Gish, S.E. Shoelson, T. Pawson, J.D. Forman-Kay, L.E. Kay, Backbone dynamics of a free and a phosphopeptide-complexed Src homology 2 domain studied by ^{15}N NMR relaxation, Biochemistry 33 (1994) 5984–6003.
[53] Ē. Kupče, R. Freeman, SPEED: single-point evaluation of the evolution dimension, Magn. Reson. Chem. 45 (2007) 711–713.
[54] V.A. Mandelshtam, FDM: the filter diagonalization method for data processing in NMR experiments, Prog. NMR Spectrosc. 38 (2001) 159–196.
[55] Ē. Kupče, R. Freeman, Fast multidimensional NMR by minimal sampling, J. Magn. Reson. 191 (2008) 164–168.
[56] R.V. Pound, R. Freeman, Frequency control of an oscillator by nuclear magnetic resonance, Rev. Sci. Instrum. 31 (1960) 96–102.
[57] R.R. Ernst, "Nuclear magnetic double resonance with an incoherent radio-frequency field", J. Chem. Phys. 45 (1966) 3845.
[58] R. Freeman, H.D.W. Hill, Phase and intensity anomalies in Fourier transform NMR, J. Magn. Reson. 4 (1971) 366–383.
[59] B. Vitorge, G. Bodenhausen, P. Pelupessy, Speeding up Nuclear magnetic resonance spectroscopy by the use of SMAll recovery times – SMART NMR, J. Magn. Reson. 207 (2010) 149–152.
[60] K. Kazimierczuk, A. Zawadzka, W. Koźmiński, I. Zhukov, Lineshapes and artifacts in multidimensional Fourier transform of arbitrary sampled NMR data sets, J. Magn. Reson. 188 (2007) 344–356.
[61] A. Zawadzka-Kazimierczuk, K. Kazimierczuk, W. Koźmiński, A set of 4D experiments of enhanced resolution for easy resonance assignments in proteins, J. Magn. Reson. 202 (2010) 109–116.
[62] K. Kazimierczuk, A. Zawadzka, W. Koźmiński, Non-uniform frequency domain for optimal exploitation of non-unform sampling, J. Magn. Reson. 205 (2010) 286–292.
[63] S. Hiller, C. Wasmer, G. Wider, K. Wüthrich, Sequence-specific resonance assignment of soluble nonglobular proteins by 7D APSY-NMR spectroscopy, J. Am. Chem. Soc. 129 (2007) 10823–10828.
[64] Ē. Kupče, R. Freeman, Hyperdimensional NMR spectroscopy, J. Am. Chem. Soc. 128 (2006) 6020–6021.
[65] Ē. Kupče, R. Freeman, Hyperdimensional NMR spectroscopy, Prog. Magn. Reson. Spectrosc. 52 (2008) 2–30.
[66] S. Hiller, G. Wider, K. Wüthrich, APSY-NMR with proteins: practical aspects and backbone assignment, J. Biomol. NMR 42 (2008) 179–195.
[67] S.G. Hyberts, K. Takeuchi, G. Wagner, Poisson-gap sampling and forward maximum entropy reconstruction for enhancing the resolution and sensitivity of protein NMR data, J. Am. Chem. Soc. 132 (2010) 2145–2147.
[68] M. Mishkovsky, L. Frydman, Principles and progress in ultrafast multidimensional nuclear magnetic resonance, Ann. Rev. Phys. Chem. 60 (2009) 429–448.
[69] H. Eghbalnia, A. Bahrami, M. Tonelli, K. Halenga, J.L. Markley, High-resolution iterative frequency identification for NMR as a general strategy for multidimensional data collection, J. Am. Cem. Soc. 127 (2005) 12528–12536.
[70] E. Lescop, P. Schanda, R. Rasia, B. Brutcher, Automated spectral compression for fast multidimensional NMR and increased time resolution in real-time NMR spectroscopy, J. Am. Chem. Soc. 129 (2007) 2756–2757.
[71] F. Zhang, L. Bruschweiler-Li, R. Bruschweiler, Simultaneous de novo identification of molecules in chemical mixtures by doubly indirect covariance NMR spectroscopy, J. Am. Chem. Soc. 132 (2010) 16922–16923.
[72] V. Jaravine, A.V. Zhuravleva, P. Permi, I. Ibraghimo, V.Yu. Orekhov, Hyperdimensional NMR spectroscopy with nonlinear sampling, J. Am. Chem. Soc. 130 (2008) 3927–3936.
[73] A.S. Stern, D.L. Donoho, J.C. Hoch, NMR data processing using iterative thresholding and minimum l_1-Norm reconstruction, J. Magn. Reson. 188 (2007) 295–300.
[74] H. Celik, A.J. Shaka, Filter diagonalization using a "sensitivity-enhanced basis": improved performance for noisy NMR spectra, J. Magn. Reson. 207 (2010) 17–23.
[75] K. Kazimierczuk, A. Zawadzka, W. Kozminski, I. Zhukov, Determination of spin–spin couplings from ultrahigh resolution 3D NMR spectra obtained by optimized random sampling and multidimensional Fourier transformation, J. Am. Chem. Soc. 130 (2008) 5404–5405.
[76] M. Blaimer, F. Breuer, M. Mueller, R.M. Heidemann, M.A. Griswold, P.M. Jakob, S.M.A.S.H. SENSE, SMASH SENSE, PILS, GRAPPA: how to choose the optimal method, Topics Magn. Reson. Imag. 15 (2004) 223–236.