

석사 학위논문
Master's Thesis

스펙트럼 분포의 구조를 이용한 고해상도 다차원
NMR 분광 분석 기법

Exploitation of Structure for Enhanced Reconstruction of
Multidimensional NMR Spectra



Finn Wolfreys
바이오및뇌공학과
Department of Bio and Brain Engineering

KAIST

2012

스펙트럼 분포의 구조를 이용한 고해상도 다차원
NMR 분광 분석 기법

Exploitation of Structure for Enhanced Reconstruction of
Multidimensional NMR Spectra



Exploitation of Structure for Enhanced Reconstruction of Multidimensional NMR Spectra

Advisor : Professor Ye, Jong Chul

by

Finn Wolfreys

Department of Bio and Brain Engineering

KAIST

A thesis submitted to the faculty of KAIST in partial fulfillment of the requirements for the degree of Master of Science in Engineering in the Department of Bio and Brain Engineering. The study was conducted in accordance with Code of Research Ethics¹.

2012. 12. 12.

Approved by

Professor Ye, Jong Chul

[Advisor]

¹Declaration of Ethical Conduct in Research: I, as a graduate student of KAIST, hereby declare that I have not committed any acts that may damage the credibility of my research. These include, but are not limited to: falsification, thesis written by someone else, distortion of research findings or plagiarism. I affirm that my thesis contains honest conclusions based on my own careful research under the guidance of my thesis advisor.

스펙트럼 분포의 구조를 이용한 고해상도 다차원 NMR 분광 분석 기법

Finn Wolfreys

위 논문은 한국과학기술원 석사학위논문으로
학위논문심사위원회에서 심사 통과하였음.

The logo for KAIST (Korea Advanced Institute of Science and Technology) is displayed in a light blue color. It consists of the letters 'KAIST' in a bold, sans-serif font, with a horizontal line underneath the letters.

2012년 12월 12일

심사위원장 예 종 철 (인)

심사위원 이 관 수 (인)

심사위원 최 명 철 (인)

MBIS
20104186

Wolfreys, Finn. Exploitation of Structure for Enhanced Reconstruction of Multidimensional NMR Spectra. 스펙트럼 분포의 구조를 이용한 고해상도 다차원 NMR 분광 분석 기법. Department of Bio and Brain Engineering . 2012. 61p. Advisor Prof. Ye, Jong Chul. Text in English.

ABSTRACT

Multidimensional Nuclear Magnetic Resonance (NMR) spectroscopy is truly a multidisciplinary research area. It is born of advances in physics, engineering, signal processing, and the demands of biomolecular structural determination applications. One of the key bottlenecks for NMR is the acquisition time – taking up to a week even for relatively modest multidimensional experiments. Attempts to overcome this bottleneck through hardware, pulse sequence, and signal processing methods have been making steady advances for almost for the entire history of NMR. However, there appears no broadly applicable technique in common use. Compressed Sensing is a framework that exploits signal sparsity for reconstructing signals sampled well below the Nyquist rate, and has already seen application in NMR in the context of acquisition time reduction. Here the motivation of exploiting prior signal knowledge was taken a step further than assuming just sparsity with the exploitation of the structural relationship between cross and diagonal peaks observable in many NMR experiments. The approach taken was to use weighted iterative soft, and hard thresholding with statistical updates according to the structural relationship between cross and diagonal peaks performed using belief propagation. This reconstruction method was applied to one of the quintessential examples of both multidimensional NMR, and experiments exhibiting signal structure: COSY. Simulations performed on a COSY spectrum for nonactin and lactose demonstrated that imposing structure on the reconstruction yielded good improvement to the normalised mean square error of the reconstruction. It was found that IST proved more stable than IHT, but in general had slightly higher reconstruction error. In the case of IST imposition of structure consistently improved reconstruction with reconstruction possible for upwards of 85% data removal. Reconstruction influenced by structure also reduced the artefacts in areas of the spectrum not included in the signal while improving the strength of peaks degraded by general iterative thresholding algorithms. The results are a positive reflection on the exploitation of structure that is common to a whole host of NMR experiments and is able to be generalised to higher dimensional experiments. This promises improved reconstruction quality with reduced acquisition time for multidimensional NMR in the future.

Contents

Abstract	i
Contents	ii
List of Figures	iv
Chapter 1. Basic Multidimensional NMR and Its Signals	1
1.1 NMR Signals	1
1.2 Multidimensional NMR	2
1.3 COSY	4
1.4 Practical Implications	10
1.5 Signal Structure in the Fourier Domain	10
Chapter 2. Methods for the Reconstruction of NMR Spectra	12
2.1 Multidimensional NMR Spectra Interpretation (MUNIN)	12
2.2 G-Matrix Fourier Transform NMR	15
2.3 Projection-Reconstruction NMR	17
2.4 Compressed Sensing and NMR	19
Chapter 3. Review of Supporting Theory	22
3.1 Compressed Sensing	22
3.1.1 Reweighted ℓ_1 -Norm	25
3.1.2 Approaches to Compressed Sensing with Models	25
Chapter 4. Proposed Algorithm for the Exploitation of Structure During Reconstruction of NMR Spectra	28
4.1 Intitial Estimate and Posterior Probabilities	28
4.2 Incorporating Structure: the Belief Propagation Step	30
4.3 Incorporating Probability with Weighted Iterative Thresholding	34
Chapter 5. Methods	36
5.1 Data Preparation and Simulation	36
5.2 Thresholds and Transition Probabilities	36
5.3 Reconstruction: Comparison with Normal Iterative Thresholding	36
5.4 Reconstruction: Accessing the Effect of Belief Propagation	37

Chapter 6.	Results	38
6.1	Nonactin	38
6.1.1	Threshold and Transition Probabilities	38
6.2	Comparison of Proposed Algorithm to Previous Techniques . .	39
6.2.1	Iterative Hard Thresholding	39
6.2.2	Iterative Soft Thresholding	42
6.2.3	Reconstructed Spectrum	44
6.3	Lactose	51
Chapter 7.	Discussion and Conclusions	55
	References	59
	Summary (in Korean)	62



List of Figures

1.1	The COSY experiment	4
1.2	Representation of a spectrum showing only diagonal peaks (left), and (right) the same spectrum with all <i>possible</i> locations for cross peaks indicated by open diamonds.	11
1.3	A cross peak (shaded diamond) essentially defines the two corresponding diagonal peaks (open diamonds), as well as a peak symmetric about the body diagonal of the spectrum. Note that the same relationship can be imposed with the cross peak on the opposite side of the body diagonal.	11
4.1	A simple Bayesian network in which C is a parent and D and child of C	31
4.2	The simple parent child relationship between cross and diagonal peaks first modelled. . .	32
4.3	The primary method for obtaining coefficients in cross and diagonal peaks. Searching along the body diagonal when coefficients above a threshold are found a search is done vertically to find the corresponding cross peak.	33
4.4	The method used for matching coefficients when performing belief propagation in this work. Here the three cross peak coefficients have been paired with the central three diagonal peak coefficients. The two groups of coefficients that constitute the slices through the cross and diagonal peaks are separated by coefficients that are below the threshold (background).	33
4.5	A model for spectrum structure that links both cross peaks to a single diagonal peak. The arrows indicate the direction from parent to child.	34
4.6	The method used to match coefficients for the second structural model utilised. The original diagonal and cross peak referred to in the text are the right most diagonal and cross peaks aligned along the vertical direction. The remaining two are referred to as the second, or corresponding, diagonal and cross peak. Note the smaller of the original diagonal and cross peak defines the maximum cluster of coefficients that will ultimately be connected.	35
4.7	An overview of the proposed reconstruction process.	35
5.1	A 256×1024 sampling matrix in which approximately 80% of the data has been removed in row form.	37
6.1	COSY spectrum for nonactin used in the simulations.	38
6.2	NMSE plotted against the ratio of measurements to the total amount of data for IHT, IHT with weighting, IHT with weighting and belief propagation on a two-level parent and child node structure, and IHT with weighting and belief propagation with two parents connected with a single child. The NMSE resulting from the use of 10 different randomly selected sampling matrices at each level of measurements were averaged. The data set was nonactin COSY (magnitude mode). The method for determining the structure coefficients was simple thresholding (threshold = 5×10^6). Transition probabilities were: $P(d = 0 c = 0) = 0.1$ and $P(d = 1 c = 1) = 0.9$	40

6.3	The affect of parameter changes: NMSE plotted against the ratio of measurements to the total amount of data for IHT, IHT with weighting, IHT with weighting and belief propagation on a two-level parent and child node structure, and IHT with weighting and belief propagation with two parents connected with a single child. The NMSE resulting from the use of 5 different randomly selected sampling matrices at each level of measurements were averaged. The data set was nonactin COSY (magnitude mode). The method for determining the structure coefficients was simple thresholding (threshold = 1×10^7). Transition probabilities were: $P(d = 0 c = 0) = 0.1$ and $P(d = 1 c = 1) = 0.9$. Here the higher threshold has negatively affected the reconstruction.	41
6.4	NMSE plotted against the ratio of measurements to the total amount of data for IST, IST with weighting, IST with weighting and belief propagation along a two-level parent and child node structure, and IST with weighting and belief propagation with two parents connected with a single child. The NMSE resulting from the use of 5 different randomly selected sampling matrices at each level of measurements were averaged. The data set was nonactin COSY (magnitude mode). The method for determining the structure coefficients was simple thresholding (threshold = 5×10^6). Transition probabilities were: $P(d = 0 c = 0) = 0.1$ and $P(d = 1 c = 1) = 0.9$. Note these are the same parameters as, and the same sensing matrices used in the results shown in Figure 6.3.	42
6.5	NMSE versus fraction of measurements kept for IHT and IST demonstrating divergence in the case of IHT. The NMSE resulting from the use of 5 different randomly selected sampling matrices at each level of measurements were averaged. The data set was nonactin COSY (magnitude mode). The method for determining the structure coefficients was simple thresholding (threshold = 1×10^7). Transition probabilities were: $P(d = 0 c = 0) = 0.1$ and $P(d = 1 c = 1) = 0.9$. Note the divergence at measurement fraction 0.2. In cases of convergence IHT with structure has a consistently lower error than just IHT and weighted IHT.	43
6.6	A selected reconstruction of a nonactin COSY spectrum using 15% of all measurements with NMSE indicated in the title of each graph. The spectra are: (top) the direct Fourier Transform, (bottom) the direct Fourier transform.	45
6.7	A selected reconstruction of a nonactin COSY spectrum using 15% of all measurements with NMSE indicated in the title of each graph. The spectra are: (top) the reconstruction using probabilistic weighting with no belief propagation, (bottom) reconstruction with weighting and belief propagation in which diagonals are matched with two cross peaks where appropriate. Transition probabilities were: $P(d = 0 c = 0) = 0.1$ and $P(d = 1 c = 1) = 0.7$, and the threshold was 1×10^6 . Note the reduction in artefacts along the right side in the vertical direction in the reconstruction.	46
6.8	Reconstruction of Nonactin COSY spectrum. The spectra are: (top) the reconstruction using IHT, (bottom) reconstruction with weighting and no belief propagation. The view is a zoomed version of that previously presented in which artefact presence and absence has be highlighted.	47

6.9	Reconstruction of Nonactin COSY spectrum. The spectra are: (top) reconstruction with weighting and no belief propagation, and (bottom) reconstruction with weighting and belief propagation in which diagonals are matched with two cross peaks where appropriate. The view is a zoomed version of that previously presented in which artefact presence and absence has be highlighted.	48
6.10	Reconstruction of Nonactin COSY spectrum based on iterative soft thresholding. The spectra are: (top) the reconstruction using just IST, (bottom) reconstruction with weighting and no belief propagation. All parameters are the same as those used in the hard thresholding case.	49
6.11	Reconstruction of Nonactin COSY spectrum using IST and belief propagation.	50
6.12	A one dimensional slice through the spectrum of nonactin at $\omega_2 = 1.62\text{ppm}$. The reconstruction using structure is shown in red, with that using just iterative thresholding in blue. Note the large intensity scale, and that, in general, the “sinc-wiggles” are smaller for the structure-based approach which yields larger peaks.	50
6.13	COSY spectrum for nonactin used in the simulations.	51
6.14	A selected reconstruction of a lactose COSY spectrum using 15% of all measurements with NMSE indicated in the title of each graph. The spectra are: (top) the direct Fourier transform, (bottom) the reconstruction using just iterative hard thresholding. Transition probabilities were: $P(d = 0 c = 0) = 0.1$ and $P(d = 1 c = 1) = 0.7$, and the threshold was 1.3×10^7 . Note the artefacts along the middle and right side in the vertical direction in the reconstruction.	52
6.15	A selected reconstruction of a lactose COSY spectrum using 15% of all measurements with NMSE indicated in the title of each graph. The spectra are: (top) the reconstruction using probabilistic weighting with no belief propagation, (bottom) reconstruction with weighting and belief propagation in which diagonal coefficients are matched with cross peak coefficients in a two node chain where appropriate. Transition probabilities were: $P(d = 0 c = 0) = 0.1$ and $P(d = 1 c = 1) = 0.7$, and the threshold was 1.3×10^7 . Note the reduction in artefacts along the middle and right side in the vertical direction in the reconstruction.	53
6.16	COSY spectrum for lactose used in the simulations with contour level set to show “natural” artefacts (highlighted).	54

Chapter 1. Basic Multidimensional NMR and Its Signals

For the majority of its history NMR spectroscopy has had a close relationship with signal processing. After all, one of the most momentous developments in NMR spectroscopy, the move away from continuous wave NMR to pulsed methods, can in many ways be attributed to the application of the Fourier transform to NMR spectroscopy. In continuous wave spectroscopy the sample is excited and its signal measured simultaneously as the frequency of electromagnetic radiation is “swept” across a range and the resulting absorption signal obtained. In contrast, pulsed methods first excite the sample with a sequence of electromagnetic pulses and then measure the resulting time domain free induction decay (FID) from the excited nuclei. The pulses, taking the form of a square wave in the frequency domain, contain a range of frequencies sufficient to excite the nuclei of the sample over the range being investigated. The Fourier transform takes the time domain data into the frequency domain and thus yields absorption information in the form of intensity as a function of frequency. The increasing availability of digital computers for scientific investigation and Cooley and Tuckey’s Fast Fourier Transform (FFT) algorithm helped to make this technique practical and presented a significantly more efficient alternative to the continuous wave approach. This is but one of the influences signal processing has had on NMR spectroscopy and as will be explained in the coming chapters, it has ushered in new developments, and with them new problems, that increasingly advanced signal processing methods are directed at addressing.

As well as improving the efficiency of NMR spectroscopy, Fourier spectroscopy paved the way for multidimensional NMR. Multidimensional NMR vastly improved the range and variety of samples that NMR could be used to investigate and has brought tremendous developments in a number of fields that rely on chemical structure elucidation. This chapter provides an overview of the key points of multidimensional NMR from a signal processing perspective. It also introduces the popular *COSY* experiment and through exposition of the nature of its spectra demonstrates the frequency domain structure of NMR signals – one of the main principles that this work is based upon.

1.1 NMR Signals

In one dimension the FID signal is modelled as a sum of complex sinusoids enveloped by an exponential decay [28] [23]. in the continuous case this is given by the following expression:

$$S(t) = \sum_{j=1}^L (A_j e^{i\phi_j}) e^{-R_j t} e^{2\pi i f_j t} + \epsilon, \quad (1.1)$$

and in the discrete case:

$$S[k] = \sum_{j=1}^L (A_j e^{i\phi_j}) e^{-R_j k \Delta t} e^{2\pi i f_j k \Delta t} + \epsilon[k]. \quad (1.2)$$

The sum represents the combination of exponentially decaying sinusoids (complex exponential functions) coming from each nuclei at a given Larmor frequency excited by a radio frequency (RF) pulse. In the above equations, A is the amplitude while ϕ represents the phase of each sinusoid with f being its

frequency. R is a decay constant for the exponential decay. The final term in each equation, ϵ , represents noise due to electronic circuits, magnetic field inhomogeneity, and digital quantisation error, amongst other sources. It is frequently modelled as being zero mean random Gaussian noise, but it must be noted that this is an approximation often for computational convenience. Strictly speaking when using a digital computer to process a sampled FID, then the FID must be modelled in the discrete domain. However, in many cases it is simpler, while still adequate, to consider the signals as being in the continuous domain.

The previous models are for the FID, which is in the time domain. The next step is to obtain the frequency domain spectrum. This is done by taking the Fourier transform. The clearest way to see the resulting form of the signal in the frequency domain is to consider just one of the terms in the sum of Equation 1.1. To enforce the fact that the signal begins at $t = 0$ s the exponentially decaying sinusoid is multiplied by the unit step function $H(t)$ which is defined as:

$$H(x) = \int_{-\infty}^x \delta(t) dt.$$

Assuming for convenience that the phase of the sinusoid is zero, and making the substitution $\omega = 2\pi f$, the problem is given by the following:

$$\begin{aligned} \mathcal{FT}\{H(t)e^{-Rt}e^{2\pi i\omega_0 t}\} &= \frac{1}{2\pi} (\mathcal{FT}\{H(t)e^{-Rt}\} * \mathcal{FT}\{e^{i\omega_0 t}\}) \\ &= \frac{1}{2\pi} \left(\int_{-\infty}^{\infty} H(t)e^{-Rt}e^{i\omega t} dt * \int_{-\infty}^{\infty} e^{i\omega_0 t}e^{i\omega t} dt \right) \\ &= \frac{1}{2\pi} \left(\int_0^{\infty} e^{-Rt}e^{i\omega t} dt * \int_{-\infty}^{\infty} e^{i\omega_0 t}e^{i\omega t} dt \right) \\ &= \frac{1}{2\pi} \left(\frac{1}{R+i\omega} * 2\pi\delta(\omega-\omega_0) \right) \\ &= \frac{1}{2\pi} \left[\left(\frac{R}{R^2+\omega^2} + i\frac{\omega}{R^2+\omega^2} \right) * 2\pi\delta(\omega-\omega_0) \right] \\ &= \frac{R}{R^2+(\omega-\omega_0)^2} + i\frac{(\omega-\omega_0)}{R^2+(\omega-\omega_0)^2} \\ &= A(\omega-\omega_0) + iD(\omega-\omega_0). \end{aligned} \tag{1.3}$$

Here the first step is taken by using the property that multiplication in the time domain corresponds to convolution in the frequency domain, and convolution with a delta function is a shifting operation (represented by $*$). $A(\omega)$ represents the so-called *absorption* line shape, and $D(\omega)$ is referred to as the *dispersion* line shape. The overall function is known as a *Lorentzian*. In terms of practical spectroscopy absorption line shapes make it far easier to discern the frequency of a peak than do dispersion line shapes. These equations imply that the spectrum due to an exponentially decaying sinusoid is given by a real absorption peak centred on the sinusoid's frequency and an imaginary dispersion line shape. Taking account of the summation in the original expression for the FID, it can be clearly seen that the overall spectrum will be comprised of absorption and dispersion line shapes located about their respective frequencies.

1.2 Multidimensional NMR

The best way to encapsulate the basic principles of a multidimensional NMR experiment is perhaps with the simplest available: the two-dimensional experiment. On a structural level two-dimensional NMR experiments can be divided into four stages: *preparation*, *evolution*, *mixing*, and *detection*. In many ways

the two-dimensional experiment can be thought of as two one-dimensional experiments where one directly follows the first. The first experiment comprises the preparation and evolution periods while the second is composed of the mixing and detection periods.

Multidimensional experiments begin with equilibrium magnetisation aligned with the B_0 field along the z -direction. The preparation period involves a sequence of pulses and delays which manipulate this equilibrium magnetisation in some way. For instance, a 90° pulse in the x -direction would rotate the magnetisation into the transverse plane¹. The second stage is referred to as evolution. In the one dimensional case this would be the start of data acquisition – following the excitation by electromagnetic radiation in the preceding stage the resulting transverse magnetisation would evolve and an FID would be recorded. The difference in the two dimensional case, however, is that although the magnetisation evolves (for a period of time, t_1), it is not directly measured. Essentially, like the name suggests, the evolution period is time over which the magnetisation is permitted to evolve. It is in the time delay t_1 that the experiment begins to take on its multidimensional characteristics, as will be elaborated upon in the following paragraphs.

Following the evolution period is another application of pulses and delays referred to as the mixing period. This period permits the application of any of a vast number of sequences of pulses and delays available to the spectroscopist. Although the possible sequences vary, there are often some conceptual similarities between the manipulations of magnetisation in the mixing stage. For instance, since the mixing period directly precedes the stage in which the FIDs will actually be recorded, the mixing period is often used to transfer magnetisation that has evolved throughout the experiment in a way that is of some experimental interest, to a form that will be observable in the detection period². The fourth and final stage of a two-dimensional NMR experiment is aptly named detection. It is a period of evolution, for a time t_2 , during which the NMR signal is sampled and recorded. This period will result in bringing the information “encoded” into the magnetisation during previous steps into a form that will be processed and ultimately presented as the NMR spectrum.

As was previously mentioned, it is from the t_1 delay that the additional dimension is introduced into a two-dimensional NMR spectrum. Unlike the excite and measure scheme of one-dimensional NMR, two-dimensional experiments repeat the sequence of preparation, evolution, mixing, and detection in order to build up arrays of free induction decays. Every time the two dimensional sequence is repeated the t_1 evolution time delay is incremented by a multiple of a small amount Δt . This means that not only is the signal a function of the t_2 (the time over which it is directly measured), but it is also a function of the delay which is incremented at each repetition.

In terms of signal processing, the signal from a two-dimensional experiment is a function of two variables t_1 and t_2 : $S(t_1, t_2)$. However, for this to be useful for spectroscopy the frequency domain information of intensity as a function of frequency is required. For the case of a uniformly sampled array of FIDs resulting from a two-dimensional experiment, the signal $S(t_1, t_2)$ in its discrete digitized format can be thought of as a matrix of data points where the t_1 axis is aligned along the rows, and the t_2 axis is aligned along the columns. The problem of obtaining the frequency domain representation of this data is a simple extension of the Fourier transform used in one-dimensional spectroscopy. In the two-dimensional case the two-dimensional Fourier transform is required. Although many computer processing utilities encapsulate the two-dimensional Fourier transform into one function, in NMR spectroscopy the process

¹In the vernacular of NMR, this transverse magnetisation is described as being the result of coherence amongst the spins.

²In NMR spectroscopy only transverse (i.e. xy -plane) magnetisation is observable through measurement of the FID

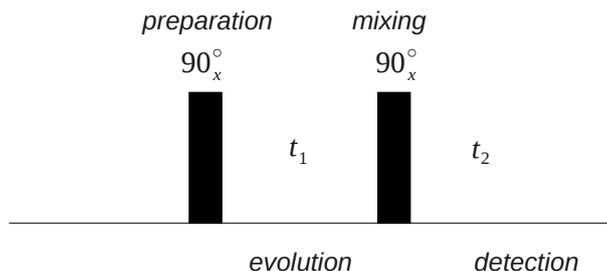


Figure 1.1: The COSY experiment

of taking two-dimensional time domain data and transforming it into frequency domain data is generally described as a two step process. The first step is the normal one-dimensional transform along the t_2 dimension corresponding to the Fourier transform of each FID in the array. This results in the form: $S(t_1, F_2)$, known in NMR spectroscopy as an *interferogram*. Following this the second step takes the Fourier transform along the rows of the matrix (the t_1 dimension) and results in the form: $S(F_1, F_2)$ i.e. the spectrum.

1.3 COSY

The *CO*rrelation *Sp*ectroscopY (COSY) family of experiments is one of the oldest and arguably most widely used NMR experiments for structure determination. Its origins are in the proposal of two-dimensional spectroscopy itself with the sequence of two pulses separated by a t_1 delay pulse sequence originally presented by Jean Jeener as the first attempt at two-dimensional NMR spectroscopy [35].

To understand more about the nature of the signals from the COSY family (and other) two-dimensional experiments a product operator analysis of the magnetisation manipulation and resulting signal will be presented. This will be followed by a higher level explanation of why the experiment is useful and the inherent signal structure introduced by the pulse sequence. Introduction to the product operator formalism is outside the purview of this section; however, an introduction can be found in Chapter 6 of [28], and more detailed explanations in [33] and [9] amongst others.

The COSY pulse sequence is shown in Figure 1.1. The preparation and mixing periods both consist of a 90° x -pulse. These are separated by the two periods of free evolution: evolution and detection. It is the evolution of magnetisation under the Hamiltonian for free evolution and J-coupling that enables the deduction of structural correlations between the spin systems. The following is a description of each stage of the COSY experiment using the product operator formalism based on that given in [28].

Preparation

Focusing on the case of two coupled spins, the first spin, I , is considered. The Hamiltonian for a pulse of strength (in radians per second) ω_1 is given by $\hat{H}_x = \omega_1 I_x$. This occurs over a time of t_p and the evolution of the equilibrium magnetisation for a 90° pulse ($\omega_1 t_p = 90^\circ$) in the x -direction results in

the following:

$$\begin{aligned}
\hat{I}_z \xrightarrow{\omega_1 t_p \hat{I}_x} & \cos(\omega_1 t_p) \hat{I}_z - \sin(\omega_1 t_p) \hat{I}_y \\
= & \cos(90^\circ) \hat{I}_z - \sin(90^\circ) \hat{I}_y \\
= & -\hat{I}_y.
\end{aligned}$$

The expression is simplified in this case as the pulse results in 90° rotation and the cosine and sine of 90° are zero and one respectively. This completes the preparation period in which the magnetisation has been rotated into the transverse plane, aligned along the $-y$ -direction.

Evolution

The transverse magnetisation generated by the first pulse then evolves for a period t_1 . Assuming I is coupled to a second spin S , the magnetisation will evolve under the offset as well as J-coupling with the Hamiltonian given by $\hat{H} = \Omega_1 \hat{I}_z + \Omega_2 \hat{S}_z + 2\pi J_{12} \hat{I}_z \hat{S}_z$. The offset, Ω_1 , is the difference between the Larmor frequency (in radians per second) and the frequency of the rotating frame. It essentially represents the so-called ‘‘chemical shift’’ in the rotating reference frame. J_{12} is the coupling constant between the two spins I and S in Hertz, and has been multiplied by a factor of 2π in order to keep the units consistently in radians per second. Since the evolution under each term of the Hamiltonian can be considered separately, first the effect of the chemical shift on the magnetisation’s evolution is considered. This gives:

$$-\hat{I}_y \xrightarrow{\Omega_1 t_1 \hat{I}_z} -\cos(\Omega_1 t_1) \hat{I}_y + \sin(\Omega_1 t_1) \hat{I}_x.$$

As the Hamiltonian for the second spin’s evolution has no effect on the first it can be skipped, and the effect of coupling for each term in the previous equation considered next giving:

$$\begin{aligned}
-\cos(\Omega_1 t_1) \hat{I}_y \xrightarrow{2\pi J_{12} t_1 \hat{I}_z \hat{S}_z} & -\cos(\pi J_{12} t_1) \cos(\Omega_1 t_1) \hat{I}_y \\
& + \sin(\pi J_{12} t_1) \cos(\Omega_1 t_1) 2\hat{I}_x \hat{S}_z, \\
\sin(\Omega_1 t_1) \hat{I}_x \xrightarrow{2\pi J_{12} t_1 \hat{I}_z \hat{S}_z} & \cos(\pi J_{12} t_1) \sin(\Omega_1 t_1) \hat{I}_x \\
& + \sin(\pi J_{12} t_1) \sin(\Omega_1 t_1) 2\hat{I}_y \hat{S}_z.
\end{aligned}$$

These equations represent the state of the magnetisation at the end of the evolution period.

Mixing

As is shown in the COSY pulse sequence diagram of Figure 1.1, the mixing period for the two-dimensional COSY sequence consists of a 90° x -pulse. The effect of this pulse on the previous magnetisation is shown by the next equations. Each term is shown separately for simplicity, and due to space constraints. Unlike the preparation period, the initial magnetisation is now on both coupled spins which

means a Hamiltonian of the form: $\hat{H} = \omega_1 \hat{I}_x + \omega_1 \hat{S}_x$ must be used. The terms are given by:

$$- \cos(\pi J_{12} t_1) \cos(\Omega_1 t_1) \hat{I}_y \xrightarrow{(90^\circ)(\hat{I}_x + \hat{S}_x)} - \cos(\pi J_{12} t_1) \cos(\Omega_1 t_1) \hat{I}_z, \quad (1.4)$$

$$\sin(\pi J_{12} t_1) \cos(\Omega_1 t_1) 2 \hat{I}_x \hat{S}_z \xrightarrow{(90^\circ)(\hat{I}_x + \hat{S}_x)} - \sin(\pi J_{12} t_1) \cos(\Omega_1 t_1) 2 \hat{I}_x \hat{S}_y, \quad (1.5)$$

$$\cos(\pi J_{12} t_1) \sin(\Omega_1 t_1) \hat{I}_x \xrightarrow{(90^\circ)(\hat{I}_x + \hat{S}_x)} \cos(\pi J_{12} t_1) \sin(\Omega_1 t_1) \hat{I}_x, \quad (1.6)$$

$$\sin(\pi J_{12} t_1) \sin(\Omega_1 t_1) 2 \hat{I}_y \hat{S}_z \xrightarrow{(90^\circ)(\hat{I}_x + \hat{S}_x)} - \sin(\pi J_{12} t_1) \sin(\Omega_1 t_1) 2 \hat{I}_z \hat{S}_y. \quad (1.7)$$

The previous expressions represent the state of the magnetisation before the final period of free evolution that will give rise to the FID. Of the four components that make up the expression for the magnetisation, Equation 1.5, containing the term $\hat{I}_x \hat{S}_y$, represents double quantum coherence, which does not result in detectable magnetisation and will not contribute to the final signal. Also, Equation 1.4, the term containing \hat{I}_z , which is not aligned in the transverse direction, will not be detectable in the final signal. Therefore, it is acceptable to consider only the remaining two terms given by Equations 1.6 and 1.7.

Detection

The magnetisation that will ultimately be detected during this period comes from the evolution under J-coupling and the chemical shift of \hat{I}_x and $\hat{I}_z \hat{S}_y$. These are already modulated by sine and cosine terms from previous periods of evolution. In addition to this the evolution of \hat{I}_x will result in further modulation. This evolution can be deduced by considering the effect of the offset and J-coupling. The evolution due the offset is given by:

$$\hat{I}_x \xrightarrow{\Omega_1 t_2 \hat{I}_z} \cos(\Omega_1 t_2) \hat{I}_x + \sin(\Omega_1 t_2) \hat{I}_y.$$

Considering the J-coupling, the magnetisation of each term in the previous equation evolves further as:

$$\begin{aligned} \cos(\Omega_1 t_2) \hat{I}_x \xrightarrow{2\pi J_{12} t_2 \hat{I}_z \hat{S}_z} & \cos(\pi J_{12} t_2) \cos(\Omega_1 t_2) \hat{I}_x \\ & + \sin(\pi J_{12} t_2) \cos(\Omega_1 t_2) 2 \hat{I}_y \hat{S}_z, \end{aligned}$$

$$\begin{aligned} \sin(\Omega_1 t_2) \hat{I}_y \xrightarrow{2\pi J_{12} t_2 \hat{I}_z \hat{S}_z} & \cos(\pi J_{12} t_2) \sin(\Omega_1 t_2) \hat{I}_y \\ & + \sin(\pi J_{12} t_2) \sin(\Omega_1 t_2) 2 \hat{I}_x \hat{S}_z. \end{aligned}$$

From this point no more evolution will result in detectable magnetisation. This means that in terms of magnetisation that will give rise to a detectable signal, only the transverse terms \hat{I}_x and \hat{I}_y need to be considered. In order for there to be sign discrimination in the final spectrum NMR signals are detected in quadrature. This means that the detector measures the signal $x(t)$ directly, and a signal $y(t)$ measured at an offset of 90° . The signal is represented as a complex number with $x(t)$ comprising the real part and $y(t)$ the imaginary. Although the signal that is actually being measured is the current induced in a coil due to the precession of the spins, the form of the signal can be gauged by considering the magnetisation in the xy -planes deduced from the above equations.

$$\begin{aligned} x \text{ magnetisation } (M_x): & \quad \cos(\pi J_{12} t_2) \cos(\Omega_1 t_2) \\ y \text{ magnetisation } (M_y): & \quad \cos(\pi J_{12} t_2) \sin(\Omega_1 t_2) \end{aligned}$$

$$\begin{aligned}
S(t_2) &\propto M_x + iM_y \\
&= [\cos(\pi J_{12}t_2) \cos(\Omega_1 t_2) + i \cos(\pi J_{12}t_2) \sin(\Omega_1 t_2)] e^{-R_2 t_2} \\
&= \cos(\pi J_{12}t_2) [\cos(\Omega_1 t_2) + i \sin(\Omega_1 t_2)] e^{-R_2 t_2} \\
&= \cos(\pi J_{12}t_2) e^{i\Omega_1 t_2} e^{-R_2 t_2} \\
&= \frac{(e^{i\pi J_{12}t_2} + e^{-i\pi J_{12}t_2})}{2} e^{i\Omega_1 t_2} e^{-R_2 t_2} \\
&= \frac{e^{i(\Omega_1 + \pi J_{12})t_2} + e^{i(\Omega_1 - \pi J_{12})t_2}}{2} e^{-R_2 t_2}
\end{aligned} \tag{1.8}$$

To set up the expression for the NMR signal the x component is simply added to the imaginary number i multiplied by y . In order to take into account the relaxation phenomena that manifest in an exponential decay of the FID, the term $e^{-R_2 t_2}$ is appended to the expression. This signal represents purely that which is a result of evolution during the t_2 period. To consider the full multidimensional signal the modulation from the t_1 evolution period must be considered also. From Equation 1.6 this is given by

$$\cos(\pi J_{12}t_1) \sin(\Omega_1 t_1).$$

Once again the decay of this signal must be accounted for through the attachment of an exponential decay in t_1 . This results in

$$\cos(\pi J_{12}t_1) \sin(\Omega_1 t_1) e^{-R_1 t_1}. \tag{1.9}$$

The full two-dimensional time domain signal can be expressed by multiplying Equations 1.9 and 1.8:

$$S(t_1, t_2) \propto \cos(\pi J_{12}t_1) \sin(\Omega_1 t_1) e^{-R_1 t_1} \frac{e^{i(\Omega_1 + \pi J_{12})t_2} + e^{i(\Omega_1 - \pi J_{12})t_2}}{2} e^{-R_2 t_2}. \tag{1.10}$$

The next step is to transform this signal into the frequency domain using the Fourier transform. Since the signal is a function of both t_1 and t_2 a two-dimensional Fourier transform is required. Noticing that to obtain Equation 1.10, two functions of two different variables have been multiplied, the two dimensional time domain signal can be expressed in separated form as

$$S(t_1, t_2) = S(t_1)S(t_2).$$

This means that the multidimensional Fourier transform becomes the multiplication of Fourier transforms over a single variable:

$$\mathcal{FT}\{S(t_1, t_2)\} = \mathcal{FT}\{S(t_1)\}\mathcal{FT}\{S(t_2)\}. \tag{1.11}$$

The first transform is calculated by using the *cosine Fourier transform* which yields the following expression for the spectrum:

$$\frac{1}{2}[(A_1(\Omega_1 + \pi J_{12}) + A_1(\Omega_1 - \pi J_{12}))]. \tag{1.12}$$

The subscript on the absorption peak is introduced to differentiate it from the absorption peaks arising from evolution in the detection period. The form of this signal is very similar to that used to

obtain Equation 1.3. Applying similar analysis yields:

$$\begin{aligned}
\mathcal{FT}\{S(t_2)\} &\propto \frac{1}{2\pi} \mathcal{FT} \left\{ \frac{e^{i(\Omega_1 + \pi J_{12})t_2}}{2} \right\} * \mathcal{FT} \{e^{-R_2 t_2}\} \\
&\quad + \frac{1}{2\pi} \mathcal{FT} \left\{ \frac{e^{i(\Omega_1 - \pi J_{12})t_2}}{2} \right\} * \mathcal{FT} \{e^{-R_2 t_2}\} \\
&= \frac{1}{2} [A_2(\Omega_1 + \pi J_{12}) + iD_2(\Omega_1 + \pi J_{12})] \\
&\quad + \frac{1}{2} [A_2(\Omega_1 - \pi J_{12}) + iD_2(\Omega_1 - \pi J_{12})] \tag{1.13}
\end{aligned}$$

To obtain the full frequency domain representation of the signal arising from the \hat{I}_x term in Equation 1.6, following Equation 1.11, Equations 1.12 and 1.13 must be multiplied. Since the absorption peaks are of the most practical interest, only the real part of the spectrum is considered.

$$\begin{aligned}
\Re(S(\omega_1, \omega_2)) &\propto \frac{1}{2} [(A_1(\Omega_1 + \pi J_{12}) + A_1(\Omega_1 - \pi J_{12})) \\
&\quad \times \frac{1}{2} (A_2(\Omega_1 + \pi J_{12}) + A_2(\Omega_1 - \pi J_{12}))] \\
&= \frac{1}{4} A_1(\Omega_1 + \pi J_{12}) A_2(\Omega_1 + \pi J_{12}) \\
&\quad + \frac{1}{4} A_1(\Omega_1 + \pi J_{12}) A_2(\Omega_1 - \pi J_{12}) \\
&\quad + \frac{1}{4} A_1(\Omega_1 - \pi J_{12}) A_2(\Omega_1 + \pi J_{12}) \\
&\quad + \frac{1}{4} A_1(\Omega_1 - \pi J_{12}) A_2(\Omega_1 - \pi J_{12}) \tag{1.14}
\end{aligned}$$

Equation 1.14 is composed of four absorption peaks. Looking at the arguments of both A_1 and A_2 it is clear that in both frequency dimensions, ω_1 and ω_2 , the absorption peaks are positioned around Ω_1 . Geometrically speaking, this corresponds to a group of four peaks located along the body diagonal of the spectrum. Consequently this collection of peaks, or *multiplet*, is referred to as a *diagonal peak*.

The previous analysis deduces the signal from the magnetisation present in Equation 1.6. To account for the full magnetisation, the other term that gives rise to observable magnetisation, Equation 1.7, must also be considered. The steps are the same for finding the form of the signal for the diagonal peaks. First, the evolution of the magnetisation must be considered. In this case the transverse magnetisation is on the second of the two spins, so in terms of offset, the only Hamiltonian term that has any effect is $\Omega_2 t_2 \hat{I}_z$. This results in:

$$2\hat{I}_z \hat{S}_y \xrightarrow{\Omega_2 t_2 \hat{I}_z} \cos(\Omega_2 t_2) 2\hat{I}_z \hat{S}_y - \sin(\Omega_2 t_2) 2\hat{I}_z \hat{S}_x.$$

J-coupling between the two spins is considered; yielding:

$$\begin{aligned}
&\cos(\Omega_2 t_2) 2\hat{I}_z \hat{S}_y \xrightarrow{2\pi J_{12} t_2 \hat{I}_z \hat{S}_z} \cos(\pi J_{12} t_2) \cos(\Omega_2 t_2) 2\hat{I}_z \hat{S}_y \\
&\quad - \sin(\pi J_{12} t_2) \cos(\Omega_2 t_2) \hat{S}_x, \\
&-\sin(\Omega_2 t_2) 2\hat{I}_z \hat{S}_x \xrightarrow{2\pi J_{12} t_2 \hat{I}_z \hat{S}_z} -\cos(\pi J_{12} t_2) \sin(\Omega_2 t_2) 2\hat{I}_z \hat{S}_x \\
&\quad - \sin(\pi J_{12} t_2) \sin(\Omega_2 t_2) \hat{S}_y
\end{aligned}$$

Next the signal is determined using the previous method of combining the x and y magnetisation in quadrature.

$$\begin{aligned}
x \text{ magnetisation } (M_x): &\quad -\sin(\pi J_{12} t_2) \cos(\Omega_2 t_2) \\
y \text{ magnetisation } (M_y): &\quad -\sin(\pi J_{12} t_2) \sin(\Omega_2 t_2) \hat{S}_y
\end{aligned}$$

$$\begin{aligned}
S(t_2) &\propto M_x + iM_y \\
&= [-\sin(\pi J_{12}t_2) \cos(\Omega_2 t_2) - i \sin(\pi J_{12}t_2) \sin(\Omega_2 t_2)] e^{-R_2 t_2} \\
&= -\sin(\pi J_{12}t_2) [\cos(\Omega_2 t_2) + i \sin(\Omega_2 t_2)] e^{-R_2 t_2} \\
&= -\sin(\pi J_{12}t_2) e^{i\Omega_2 t_2} e^{-R_2 t_2} \\
&= -\frac{(e^{i\pi J_{12}t_2} - e^{-i\pi J_{12}t_2})}{2i} e^{i\Omega_2 t_2} e^{-R_2 t_2} \\
&= i \frac{e^{i(\Omega_2 + \pi J_{12})t_2} - e^{i(\Omega_2 - \pi J_{12})t_2}}{2} e^{-R_2 t_2}.
\end{aligned} \tag{1.15}$$

The modulation term in t_1 is given by

$$-\sin(\pi J_{12}t_1) \sin(\Omega_1 t_1) e^{-R_1 t_1}. \tag{1.16}$$

The modulation comes from Equation 1.7 and the exponential term has been added to reflect the relaxation of the magnetisation. Combining Equation 1.16 and Equation 1.15, gives the full time domain signal

$$S(t_1, t_2) \propto -\sin(\pi J_{12}t_1) \sin(\Omega_1 t_1) e^{-R_1 t_1} i \frac{e^{i(\Omega_2 + \pi J_{12})t_2} - e^{i(\Omega_2 - \pi J_{12})t_2}}{2} e^{-R_2 t_2}. \tag{1.17}$$

Again, the signal can be written as the product of two single variable functions as in Equation 1.11. Dealing with $S(t_1)$ first, the *sine Fourier transform* can be used to obtain:

$$\frac{1}{2}[A_1(\Omega_1 + \pi J_{12}) - A_1(\Omega_1 - \pi J_{12})] \tag{1.18}$$

The Fourier transform of $S(t_2)$ differ slightly to the result for the diagonal peaks since the signal is multiplied by the imaginary number i . Bringing this constant outside the Fourier transform gives a similar expression to that used to obtain Equation 1.3. The result of the Fourier transform is deduced as:

$$\begin{aligned}
\mathcal{FT}\{S(t_2)\} &\propto \frac{i}{2\pi} \mathcal{FT} \left\{ \frac{e^{i(\Omega_2 + \pi J_{12})t_2}}{2} \right\} * \mathcal{FT} \{e^{-R_2 t_2}\} \\
&\quad + \frac{i}{2\pi} \mathcal{FT} \left\{ \frac{-e^{i(\Omega_2 - \pi J_{12})t_2}}{2} \right\} * \mathcal{FT} \{e^{-R_2 t_2}\} \\
&= \frac{i}{2} [A_2(\Omega_2 + \pi J_{12}) + iD_2(\Omega_2 + \pi J_{12})] \\
&\quad - \frac{i}{2} [A_2(\Omega_2 - \pi J_{12}) + iD_2(\Omega_2 - \pi J_{12})].
\end{aligned} \tag{1.19}$$

Equation 1.19 differs from the case of Equation 1.13 essentially in its imaginary coefficient. This means that the desirable absorption peaks will be in the imaginary, rather than real part of the spectrum. As explained in [28] this can be remedied by applying a phase shift. In this case it is desired to remove the imaginary coefficient, and this can be done using a -90° phase shift. This implies multiplication by $e^{-i\pi/2} = \cos(\pi/2) - i \sin(\pi/2) = -i$. This multiplication gives:

$$\begin{aligned}
S(\omega_2) &\propto \frac{1}{2} [A_2(\Omega_2 + \pi J_{12}) + iD_2(\Omega_2 + \pi J_{12})] \\
&\quad - \frac{1}{2} [A_2(\Omega_2 - \pi J_{12}) + iD_2(\Omega_2 - \pi J_{12})].
\end{aligned} \tag{1.20}$$

Combining Equations 1.18 and 1.20 yields the final expression for the signal in the Fourier domain.

$$\begin{aligned}
\Re(S(\omega_1, \omega_2)) &\propto \frac{1}{2}[(A_1(\Omega_1 + \pi J_{12}) - A_1(\Omega_1 - \pi J_{12})) \\
&\quad \times \frac{1}{2}(A_2(\Omega_2 + \pi J_{12})) - A_2(\Omega_2 - \pi J_{12}))] \\
&= \frac{1}{4}A_1(\Omega_1 + \pi J_{12})A_2(\Omega_2 + \pi J_{12}) \\
&\quad - \frac{1}{4}A_1(\Omega_1 + \pi J_{12})A_2(\Omega_2 - \pi J_{12}) \\
&\quad + \frac{1}{4}A_1(\Omega_1 - \pi J_{12})A_2(\Omega_2 + \pi J_{12}) \\
&\quad - \frac{1}{4}A_1(\Omega_1 - \pi J_{12})A_2(\Omega_2 - \pi J_{12}). \tag{1.21}
\end{aligned}$$

Comparing Equation 1.14 and Equation 1.21, there are some interesting differences. Firstly, the peaks are not all positive. This demonstrates what can be thought of as anti-phase magnetisation in two dimensions. Secondly, the A_2 absorption peaks are centred around Ω_2 , while the A_1 absorption peaks are centred around Ω_1 . The implication graphically for the spectrum is that the group of four absorption peaks will be concentrated at the coordinates (Ω_1, Ω_2) . This results in what is referred to as a *cross peak*. The implications of cross and diagonal peaks are explored in the next section.

1.4 Practical Implications

The analysis of the previous section demonstrates that there are two types of peaks comprising the spectrum from a COSY experiment: cross peaks, and diagonal peaks. Practically speaking, the diagonal peaks correspond to chemical shifts analogous to a simple one-dimensional system. These peaks represent the presence and chemical environment of a *spin system*. Cross peaks, on the other hand, represent coupling between spin systems defined by the diagonal peaks. The cross peaks arise from magnetisation that has evolved on one spin, and due to the pulse applied during the mixing period, had this magnetisation “passed” to the spin system it is coupled to. Due to evolution under the effect of the Hamiltonian of the second spin system this magnetisation essentially “picks up” modulation with respect to the offset of the second spin system. This means that coupled spin systems, represented by diagonal peaks at the appropriate locations, will have cross peaks between them due to the magnetisation transfer. This is useful in the laboratory as it reveals the spin systems present in a particular sample, and the connection between them – hence the name *Correlation Spectroscopy*.

1.5 Signal Structure in the Fourier Domain

The diagonal and cross peak composition of COSY spectra is one that can be found, in one form or another, in many different types of NMR spectra. It also means that in the Fourier domain, where the spectrum “lives”, there is structure. For instance, for a cross peak to exist it must be between two diagonal peaks, otherwise there would not have been transfer of magnetisation (and subsequent modulated evolution). Thus diagonal peaks essentially define the possible locations of cross peaks as demonstrated in Figure 1.2. Although there may not be a cross peak between two diagonal peaks, if there is, it can only occur in a limited number of places – defined by the diagonal peaks. Going the other way the relationship is much stronger. If a cross peak exists, then it completely characterises the location of the two diagonal peaks it is between. Furthermore, COSY spectra are symmetric about the diagonal.

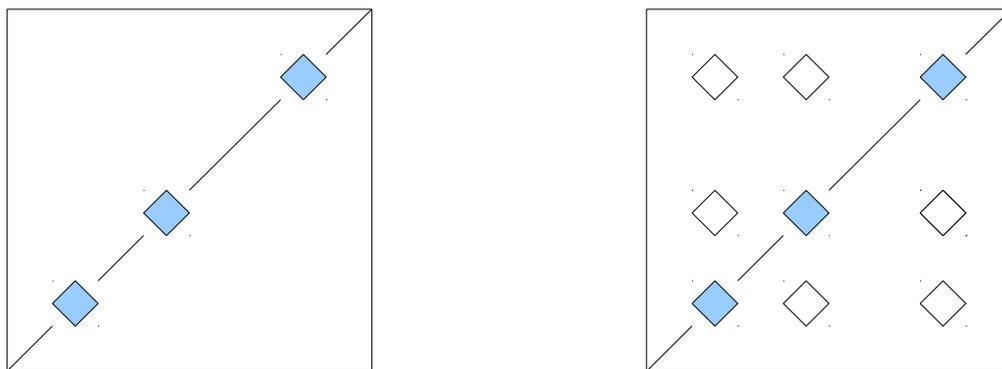


Figure 1.2: Representation of a spectrum showing only diagonal peaks (left), and (right) the same spectrum with all *possible* locations for cross peaks indicated by open diamonds.

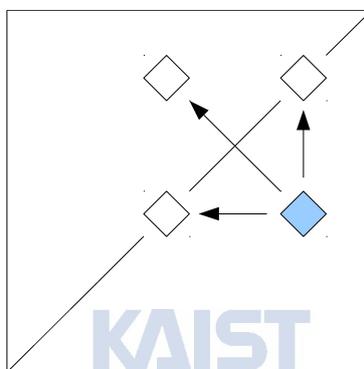


Figure 1.3: A cross peak (shaded diamond) essentially defines the two corresponding diagonal peaks (open diamonds), as well as a peak symmetric about the body diagonal of the spectrum. Note that the same relationship can be imposed with the cross peak on the opposite side of the body diagonal.

This means that knowledge of a cross peak not only gives the location of the two diagonal peaks, but also the location of the opposing cross peak as shows in Figure 1.3. Here the focus is on the COSY experiment as it provides a representative structure observable in many other types of spectra. The ideas presented here for instance, are easily extendable to experiments like *Total Correlation Spectroscopy* (TOCSY), and *Nuclear Overhauser Effect Spectroscopy* (NOESY), which have similar structures. The exploitation of this structure through the framework provided by Compressed Sensing will form the essence of the approaches taken here to enhance reconstruction of NMR spectra.

Chapter 2. Methods for the Reconstruction of NMR Spectra

Multidimensional NMR spectroscopy, and NMR spectroscopy in general, have always had a close relationship with signal processing. During the history of NMR a multitude of signal processing techniques have found application in NMR. To this day many signal processing techniques find themselves called upon routinely by spectroscopists. Zero-filling FIDs to increase resolution, apodization, and first-point correction name but a few of these techniques. Other methods such as digital filtering have also been proposed [36], and can be seen implemented many *Bruker* spectrometers. More advanced methods such as Time Proportional Phase Increments (TPPI) for indirect dimension quadrature detection in multidimensional NMR are examples of the development of techniques based on signal processing theory specifically for NMR. These are all techniques that serve to either enhance spectrum quality or increase the ability to deduce information from a spectrum.

In slight contrast to the previous algorithms, a major group of signal processing methods for NMR are concerned with one of the key bottlenecks for multidimensional NMR: acquisition time. The high time demands of multidimensional NMR essentially stem from the uniform sampling requirement of the Fourier transform. In order to obtain a spectrum of adequate resolution it is necessary to uniformly sample vast numbers of points. For instance an N -dimensional spectrum requires a total number of data points equal to $n_1 \times n_2 \dots \times n_N$ (where n_i is the number of points in each dimension). This severely limits the number of points in the indirect dimension that can be sampled and can negatively affect the resolution. Most signal processing methods intended to address this problem are directed at overcoming in some way the uniform sampling requirement of the Fourier transform. Although the number of algorithms proposed to overcome the acquisition requirements of multidimensional NMR spectroscopy prevents an all encompassing discussion, a few of notable and recent developments will be presented in this chapter. The first, *MUNIN*, represents approaches towards overcoming the uniformity requirement of the Fourier transform. The next two, *GFT-NMR* and *Projection-Reconstruction NMR* are members of a class known as “reduced dimensionality approaches”. Finally, the application of techniques from the relatively new field of Compressed Sensing to NMR spectroscopy will be discussed.

2.1 Multidimensional NMR Spectra Interpretation (MUNIN)

Structural determination by NMR spectroscopy is a process that contains several avenues through which considerable time costs can be introduced. The primary focus of many NMR signal processing algorithms is to decrease the acquisition time in costly high resolution multidimensional NMR experiments. One of the other major areas that demands not only time, but also a large degree of human input and supervision, is analysing the resulting spectra obtained from these experiments. This delay in the structural determination process has spurred interest in the development of automated spectroscopy. The ultimate goal of this area is complete automation of the structural determination process and thus removal of the laborious human analysis steps. This, however, is no easy task, as the discrimination of actual peaks amongst the artefacts and noise that invade any practical spectrum is often done through

reasoning and experience on the part of the spectroscopist analysing the spectrum. Furthermore, complications with resolution such as peak overlap add additional layers of complication to the task. Despite these difficulties, several methods have been proposed for the automation of NMR spectroscopy.

Multidimensional NMR Spectra Interpretation (MUNIN) [39] is an approach based on a mathematical technique known as *three-way decomposition*. Although it was initially introduced within the context of automated spectroscopy, specifically automated interpretation, it also has some interesting reconstruction applications. The basic idea behind MUNIN is that a multidimensional NMR spectrum can be decomposed into a sum of one-dimensional components. As stated in [4], mathematically this is expressed as:

$$S(x_1, x_2, \dots, x_N) \approx \sum_{k=1}^K a_k \cdot [\mathbf{F}_1^k(x_1) \otimes \mathbf{F}_2^k(x_2) \dots \otimes \mathbf{F}_N^k(x_N)]. \quad (2.1)$$

This equation describes a number of concepts which need to be clarified. $S(x_1, x_2, \dots, x_N)$, which represents the (N-dimensional) data, may be in either the time or frequency domain. Since in the case of NMR the concern is discrete data, each argument x_n specifies an index along a dimension of the signal. Thus it can be seen that $S(x_1, x_2, \dots, x_N)$ is an N-dimensional matrix containing the data which is to be decomposed. This expression differs slightly from that presented in [39] in that it represents approximate equality. The reason for this is that the expression presented here does not account for noise or artefacts.

As previously touched upon, the essence of MUNIN is decomposition of data. The summation is over what is referred to in MUNIN literature as *components*, of which there are K . Going a step further it can be seen that these components are made up of a direct product between N vectors, \mathbf{F}^k . These vectors are one-dimensional and referred to as *shapes*. Finally, a_k represents an amplitude term for each component. The decomposition (and compression) properties of the result MUNIN seeks to obtain is to break the data into components that can be summed and decompose the multidimensional components into one-dimensional vectors thus doing away with the need for a considerable amount of data.

From a mathematical point of view the decomposition required by MUNIN can be accomplished through the solution to a minimisation problem. This problem, as given in [4], is:

$$\min \left[\left\| S(x_1, x_2, \dots, x_N) - \sum_{k=1}^K a_k \cdot (\mathbf{F}_1^k(x_1) \otimes \mathbf{F}_2^k(x_2) \dots \otimes \mathbf{F}_N^k(x_N)) \right\|_2 + \lambda \sum_{k=1}^K a_k^2 \right] \quad (2.2)$$

This is essentially a minimisation of the error between the data and the decomposition with an additional penalisation term for large amplitudes discussed further in [4]. It turns out that a unique solution to the minimisation problem can be found for data of three or greater dimensions i.e. $N \geq 3$.

The applications of MUNIN within multidimensional NMR are two-fold. From an interpretation point of view, the interest is in decomposing a full data set and using the components to deduce some additional or more easily interpreted information. From the perspective of reducing acquisition time, the interest is in computing components from a partial data set, and then using these components as part of Equation 2.1 to generate a more complete data set. Thus the interpretation applications is interested in decomposition; whereas reduced acquisition time applications are interested in going in the opposite direction to reconstruct the signal. To examine this in a little more detail, first, the application to interpretation is described.

In many practical cases the components computed have a close correspondence with actual spin systems present in the sample. This means that the information about peak location and couplings is

decomposed into a single one-dimensional component. This introduces many advantages to the human labour intensive analysis task in NMR spectroscopy. Peak picking and automated analysis is a much easier task in one dimension. Since the full multidimensional spectrum is decomposed into one-dimensional components that, pragmatically speaking, can be considered to encapsulate the presence and interactions of a spin system, the analysis of complicated multidimensional spectra can be broken down into the analysis of more simple one-dimensional representations. A common example in MUNIN literature is the use of NMR spectroscopy to monitor ligand binding in drug discovery experiments. In this case attention is focused on the presence or absence of a particular resonance indicating ligand binding and thus structural and/or conformational change. Assuming the resonance associated with the ligand binding is decomposed as a single component, the analysis task becomes simply a matter of detecting the presence or absence of a peak in one dimension.

Turning to the desire to reduce the number of sampled FIDs and thus acquisition time, additional applications of MUNIN to nonuniform sampling can be seen. For the sparse sampling case, the original representation of Equation 2.1 can be recast as:

$$\mathbf{S}(t_1, t_2, \dots, t_N) \approx \mathbf{G}(t_1, t_2, \dots, t_N) \bullet \sum_{k=1}^K a_k \cdot [\mathbf{F}_1^k(t_1) \otimes \mathbf{F}_2^k(t_2) \dots \otimes \mathbf{F}_N^k(t_N)].$$

Here the matrix $\mathbf{G}(t_1, t_2, \dots, t_N)$ is a sampling matrix, where recorded data is marked with a one at the appropriate index in \mathbf{G} and marked zero otherwise. The matrix is multiplied directly with the component sum making all entries that have not been measured zero (\bullet indicates element-wise multiplication). In a practical NMR context, whole FIDs are recorded at one time. For this reason, whole FIDs are not included when collecting the downsampled data that will form the measurements from which the components are generated. The indexes have been changed from x_n to t_n to emphasize that the downsampling occurs in the time domain (since this is where the sampling occurs). As long as a sufficient number and type of measurements are obtained, the minimisation problem of Equation 2.2 (slightly modified to incorporate the matrix \mathbf{G}) can be solved and a full time domain data set obtained. From the full time domain dataset the usual Fourier transform steps can be applied and finally produce a full spectrum. In this case the techniques of MUNIN have been applied to obtain a full spectrum from non-uniformly sampled time domain data.

The nonuniform sampling possibilities introduced by MUNIN also provide the necessary framework for concurrent recording and analysis. This is generally prohibited in Fourier NMR experiments due to the requirement for uniform sampling. MUNIN however offers the ability to compute components and generate a full data set from a small number of measurements. It is therefore conceivable that while an experiment is progressing, the data obtained up until that point could be used to provide a rough estimate of how the spectrum will look [4]. One way this is useful is when the approximate number of resonances is known prior to performing the experiment. In this case the spectrum can be checked as the experiment (utilising random sampling) progresses, and once the expected number of resonances are observed, the experiment can be stopped. The result of this is that only the minimum amount of data required to produce the expected resonances is recorded and the total experiment time is significantly reduced.

It must be noted that although MUNIN offers a number of promising applications for NMR spectroscopy it can be computationally very expensive. One of the goals of MUNIN is to make recording high dimension NMR spectra realistic. The share amount of data associated with experiments with

dimension larger than four mean that, due to time constraints, they are out of reach for most investigators. MUNIN offers the possibility for nonuniform and thus reduced sampling times, but the large amount of data associated with these experiments can make the task of determining the components that compose a spectrum expensive in terms of CPU processing time and memory demands. However, this is rather slight in comparison to the demands of a fully sampled high-dimension NMR experiment. Also drawbacks of MUNIN are that convergence is not guaranteed, and there is no rigorous analysis for the minimum amount of data required for accurate decomposition. From a pragmatic point of view this does not necessarily introduce a lot of problems because even though convergence is not guaranteed, if there is a good working knowledge of the minimum data required and parameters for the model, decomposition is generally successful. Despite a few shortcomings, MUNIN offers considerable benefits in its application to spectra interpretation, nonuniform sampling, and related applications. It also demonstrates, through the often one-to-one relationship between components and spin systems, a nice correspondence between mathematical theory and the underlying science of NMR spectroscopy.

2.2 G-Matrix Fourier Transform NMR

As has been previously mentioned, the requirement of the Fourier transform for uniform sampling introduces major time demands for multidimensional NMR. G-Matrix Fourier Transform NMR (GFT-NMR) is aimed at two major drawbacks of multidimensional NMR that stem from the uniform sampling requirement. As explained in [29], firstly, high-dimension NMR experiments require a lot of sampling in indirect dimensions. In the end however the amount of data gleaned from the total samples that is actually used is small. Furthermore, if through some method this “useful” information could be directly obtained, the actual number of points that would be required to obtain the data with adequate signal-to-noise ratio would be comparatively small. It seems as if measurements are ultimately wasted. Secondly, the practical limits on an NMR experiment mean that often the number of points sampled in the indirect dimension must be limited. This has a direct affect on the resolution, decreasing it in the indirect dimensions and ultimately making it more difficult to determine chemical shift which differ by very small amounts.

GFT-NMR employs an idea that has been brewing in NMR for quite some time: joint evolution [19] [20]. The basic idea is instead of varying the evolution time in only one dimension at a time, vary them together. This is referred to as *reduced dimensionality*. The idea is best furnished with an example. Taking the case of a three dimensional experiment, with indirect evolution periods t_1 and t_2 , and directly measured dimension t_3 . If t_1 and t_2 vary together then effectively the number of indirect time variables has been reduced as they can both be expressed as a function of a single variable t . Here two evolution periods have become one, and following the cant and notation of GFT-NMR it is said that $K = 1$ i.e. one evolution dimension has been reduced. It is still however unclear as to how this method actually reduces the acquisition time. This will be explained next when considering the effect joint evolution has on the chemical shift.

The approach described previously explains in practice how the data can be obtained, but the real essence of GFT-NMR comes in the processing stage. Removing time dimensions is by no means “a free lunch” and with the decrease in acquisition time comes the problem of *intermodulation* of the chemical shifts. In conventional multidimensional NMR, evolution periods often serve to modulate the amplitude of the signal by $\cos(\Omega_i t_i)$ or $\sin(\Omega_i t_i)$ where Ω_i represents and chemical shift and t_i an indirect dimension. In simple terms, further evolution periods can be considered to simply appending another modulation

term e.g. after t_1 and t_2 the modulation term is $\cos(\Omega_1 t_1) \cos(\Omega_2 t_2)$ (assuming that the phase is set appropriately for cosine rather than sine modulation and the evolution occurs with a different chemical shift in each dimension). The Fourier transform effectively separates the frequencies when taken with respect to each variable as demonstrated previously. In reduced dimension experiments, however, the evolution is occurring jointly, so instead of having t_1 and t_2 , there is simply one variable. This means that the modulation will be of the form:

$$\cos(\Omega_1 t) \cos(\Omega_2 t).$$

It becomes apparent that the evolution is now linked and cannot be separated in the usual fashion. The answer of GFT-NMR is to use phase shifts on all but one of the indirect dimensions, resulting in sine modulation. The calculation performed in GFT-NMR can be done in either the time or frequency domain; however, those performed in the frequency domain are somewhat clearer. For the purposes of exposition of the technique the frequency domain approach is taken here. The example used here is based on that used in [19] and the supplementary information provided with [29]. The result of phase shift is best seen in the case where three indirect dimensions have been combined into one. Here the expression for the intermodulation is given as follows.

Modulation term	Phase		
$\cos(\Omega_1 t) \cos(\Omega_2 t) \cos(\Omega_3 t)$	0°	0°	0°
$\cos(\Omega_1 t) \sin(\Omega_2 t) \cos(\Omega_3 t)$	0°	90°	0°
$\cos(\Omega_1 t) \cos(\Omega_2 t) \sin(\Omega_3 t)$	0°	0°	90°
$\cos(\Omega_1 t) \sin(\Omega_2 t) \sin(\Omega_3 t)$	0°	90°	90°

After performing phase correction to ensure that all absorption mode lines are real, the Fourier transform of the previous expressions yields one-dimensional intermodulated spectra. Since cosine modulation results in *in phase* peaks, and sine modulation results in *anti-phase* peaks [28], linear combinations of the one dimensional spectra can be used to cancel all but one of each peak (see Figure (2) of [29]). In GFT-NMR the linear combinations of each spectrum are taken care of through matrix multiplication. Following the notation used in [29] *supplementary information*, in this case the expression is given by the following equations:

$$\begin{aligned} \hat{\mathbf{B}}(K) &= \hat{\mathbf{F}}(K) \cdot \hat{\mathbf{A}}(K) \\ \hat{\mathbf{B}}(2) &= \hat{\mathbf{F}}(2) \cdot \hat{\mathbf{A}}(2) \\ \begin{bmatrix} B1 \\ B2 \\ B3 \\ B4 \end{bmatrix} &= \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix} \cdot \begin{bmatrix} A1 \\ A2 \\ A3 \\ A4 \end{bmatrix}, \end{aligned}$$

where $A1 \dots A4$ represent the spectrum obtained by Fourier transformation of the differently modulated signals from the first part of the experiment. $B1 \dots B4$ are spectra containing only one of the peaks present in the “full” A spectra. With this information, the actual frequencies corresponding to

each peak can be expressed in terms of the others.

$$B1 = \Omega_1 + \Omega_2 + \Omega_3$$

$$B2 = \Omega_1 - \Omega_2 + \Omega_3$$

$$B3 = \Omega_1 + \Omega_2 - \Omega_3$$

$$B4 = \Omega_1 - \Omega_2 - \Omega_3$$

In matrix form this is:

$$\begin{bmatrix} B1 \\ B2 \\ B3 \\ B4 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 \\ 1 & -1 & 1 \\ 1 & 1 & -1 \\ 1 & -1 & -1 \end{bmatrix} \cdot \begin{bmatrix} \Omega_1 \\ \Omega_2 \\ \Omega_3 \end{bmatrix}.$$

This is an overdetermined system of linear equations and the approach taken to solve this in GFT-NMR is generally least squares [29]. Theoretically this could be the end of the experiment, since all chemical shifts can be determined from the solution to this system of equations. However, in reality, peaks may overlap the multiplet being observed and make it difficult to interpret the spectra, and thus obtain chemical shifts. To overcome this GFT-NMR employs further reduced dimensionality experiments, but this time removes one of the evolution dimensions. The effect this has is to produce a system of equations similar to that shown in the previous analysis, but the number of frequencies will be reduced by one. Further reduction to zero yields a single peak. These steps help to discern peaks in the frequently encountered case that they overlap. Finally, it must be restated that the analysis performed here has been in the frequency domain. GFT-NMR allows analysis to be performed in either the time or frequency domain. In the case of the time domain the process is very similar, with the F-matrix being replaced by a so-called G-matrix (where GFT-NMR gets its name).

2.3 Projection-Reconstruction NMR

Projection-reconstruction NMR [31] is a method designed to speed up multidimensional NMR spectroscopy by primarily exploiting: 1. the projection-slice theorem, and 2. the discrete nature of NMR spectra. The projection slice theorem [6] has its origins in radio astronomy, but now finds itself used throughout signal processing, and particularly in medical imaging. On a conceptual level, it can be described for the two-dimensional case as follows, paraphrasing [31]. Consider a two dimensional time domain signal, such as an FID, represented in the time domain as $S(t_1, t_2)$. Taking the data along a line at an angle α from the t_1 axis gives the so-called “slice”. It turns out, that the Fourier transform of this slice is equivalent to the projection of the signal in the frequency domain, $S(F_1, F_2)$, onto a line inclined at α in the frequency domain. This theorem is readily applied into higher dimensions. In three dimensions, for instance, the projection is onto a plane rather than a line.

It is from the discrete nature of NMR spectra that the potential for faster multidimensional NMR experiments can be seen. In Computed Tomography (CT) imaging, a very large number of measurements is needed for the projections to accurately capture the continuous nature of an object’s density. However, NMR spectra can be thought of as discrete in that the peaks occupy discrete positions within an often sparse frequency domain. This means that theoretically it should require only a small number of projections to be able to reconstruct this information. Considering the simplest case of reconstruction a single point from projections, it is clear only two projections will be needed to completely determine its position (the intersection of the lines in the “X” shape the projections form will be the “correct location”).

In order to implement the non-uniform radial sampling required by projection-reconstruction NMR, the time delays in the indirect dimensions are manipulated. For instance, in the case of a three-dimensional experiment the two indirect dimension time delays (t_1 and t_2) are set according to the following rules:¹

$$t_1 = t \cos(\alpha)$$

$$t_2 = t \sin(\alpha)$$

Here, once again α represents the inclination with respect to the t_1 axis. The above equations represent a deviation away from uniform sampling method, which varies one time delay at one time before recording the resulting FID. In this case, for each repetition of the pulse sequence both t_1 and t_2 are varied simultaneously. Since the measurements in t_3 are made as they would be for any FID, the resulting data form a plane inclined at an angle α , and with the use of the projection-slice theorem, will yield a projection inclined at the same angle in the frequency domain. Obtaining additional slices is simply a matter of repeating a series of pulse sequences where the α parameter is varied so as to give a slice at a different angle. This can essentially be thought of as radial sampling in three-dimensions.

Having obtained the data using the method in the preceding text, then comes the task of reconstructing the spectrum from the projections. It must be noted that the mathematically rigorous method for performing this reconstruction (namely: the Radon transform) would require a very large number of measurements, and would result in no gain for reconstruction over the traditional uniform sampling method. One of the tenets of projection reconstruction NMR is that it exploits the fact that spectra are relatively discrete signals and intuitively should not require full sampling to obtain some form of reconstruction. To this end, a number of methods have been developed that tackle this problem from a somewhat heuristic standpoint. Here, two of these methods, elaborated upon in [31] will be presented; further discussion can be found in [43] and references therein.

The *lowest value algorithm* uses a set of projections obtained from a projection-reconstruction NMR experiment. It examines each point and finds the lowest value across all projections that contain that point. Conceptually this algorithm can be justified by the desire to avoid false peaks. One of the dangers of using projections of NMR spectra for reconstruction is that some projections may provide an ambiguous assessment of the exact peak locations – overestimating the total number of peaks. A property of a genuine peak is that it will appear as a peak in all projections. False peaks on the other hand will only be implied by some projections. Hence if the lowest value for each point in the spectrum from all projections is taken, genuine peaks should still remain as peaks since they appear large in all projections. Meanwhile, ambiguous peaks will be eliminated by virtue of not appearing in all projections. It is clear that this algorithm assumes good measurement sensitivity, as a weak peak whose presence is masked by noise in as few as one of the projections will be discarded as false and thus result in an inaccurate reconstruction. It must also be noted that the algorithms used for projection reconstruction NMR do not preserve the shape of each peak. In general, according to the reconstruction method used in the projection-reconstruction scheme, the peak shapes will be different.

Additive back projection is another algorithm aimed at reconstructing NMR spectra from projections that, as opposed to the lowest value algorithm, is intended to handle low sensitivity measurements. Once again a set of projections is used, but instead of selecting the lowest value at each point, all values are summed. This summation is intended to improve the signal-to-noise ratio (SNR) of the resulting

¹Note that in t_3 fully sampled as a normal FID.

spectrum as the noise, owing to its random nature, upon summation will increase much more slowly than the actual signal. It also means that if a peak is masked by noise in one of the projections the final spectrum will be more likely to represent the position as being one of a peak rather than making the all or nothing decision that it is not. One of the downsides of using the additive back projection algorithm is that as a product of back projection there are ridges across the spectrum. As long as enough projections are made, and the measurements are of reasonable quality, true peaks will be of much greater intensity than these artifacts and easily discernible in the final reconstruction. Due to this, often a simple remedy is setting the contour level such that the much smaller ridges are not visible; this yields a spectrum with the expected peaks as demonstrated in [32].

It can clearly be seen that Projection-reconstruction NMR and GFT-NMR share a number of similarities in their approach to data acquisition. The two methods are members of a larger collection known as “reduced dimensionality” approaches. This refers to Projection-reconstruction NMR and GFT-NMR both varying multiple evolution periods at the same time (e.g. in the projection reconstruction technique for three dimensional spectra t_1 and t_2 vary as a function of a single variable t). There appears to be some contention over which method is a generalisation of which, with claims made to the effect that GFT-NMR is the generalisation in [47] and PR-NMR is the generalisation in [31]. Furthermore, the removal of PR-CALC [10] from public access amid GFT-NMR patent infringement allegations [11] highlights this contention.

A perhaps more reasoned historical perspective of the techniques can be found in [20]. Here it is demonstrated that they are clearly members of the same family. This family of techniques samples a much smaller subset of the data required by normal Fourier transform NMR and thus offers significant savings in terms of experiment time. The major difference between the two techniques is in their reconstruction. GFT-NMR attempts to obtain numerical values for the chemical shifts by solving an overdetermined system of equations from spectra that are the result of joint evolution in the indirect dimensions. On the other hand Projection-reconstruction NMR seeks to reconstruct a full multidimensional spectrum from what amounts to (radially) downsampled data.

2.4 Compressed Sensing and NMR

The motivation and ideas that have gone into some of the previous techniques developed for NMR fit closely with those of Compressed Sensing. For instance the decomposition operation of MUNIN closely parallels the Compressed Sensing approach of representing signals in a sparse basis. Furthermore, in [29] one of the drawbacks of conventional NMR that is listed as a motivation for development of GFT-NMR is that “[sampling in indirect dimensions] may by far exceed the measurement time required to achieve sufficient signal-to-noise ratios”. This is aligned with one of the practical motivations of Compressed Sensing, namely sampling as much as is needed, rather than full sampling and then compression.

The Compressed Sensing approach is to reconstruct the signal (or at least a sparse approximation to the signal in some basis) from downsampled measurements. This means that the sampling scheme is nonuniform. It also means that the full spectrum, complete with background and full dimensionality, is reconstructed. This is similar to the reconstruction provided through MUNIN, and different to that produced by GFT-NMR (chemical shifts; not full dimensions) and projection-reconstruction NMR (only projections; full dimensions). With Compressed Sensing also comes a solid, and steadily growing body of theoretical support that helps to provide certainty to the reconstruction of non-uniformly sampled NMR spectra. This is something that is perhaps somewhat lacking in the case of techniques such as MUNIN.

Compressed sensing is also a nice fit for NMR spectra due to their inherent sparsity in the Fourier basis. Even in the most crowded spectra, the useful information provided by the peaks represents only a small fraction of the total amount of data. Since the background is not useful to the spectroscopist, NMR spectra are well approximated by sparse representations.

Two of the earliest direct applications of Compressed Sensing to NMR were that of [17] and [46]. The work [46] considered a truncated one dimensional FID. Truncation does not necessarily violate the Shannon-Nyquist sampling theorem since the samples are still being taken at the right rate; some have just been removed from the end. The effect of this is the introduction of *ringing artefacts*, or “sinc-wiggles”, around the spectrum peaks. The algorithm developed and applied sought to reduce these from the final spectrum. The method used was based on *iterative soft thresholding* (IST) and is conceptually relatively easy to follow.

Firstly, since the uniformity condition of the Fourier transform has not been violated, the truncated FID can be transformed into the Fourier domain (albeit with significant artefacts). Next, a threshold is determined below which the spectrum can be considered artefacts and noise, this generally corresponds to just above the level of the artefacts. This threshold is then used as part of an iterative soft threshold operation. This operation entails setting all values below the threshold to zero, and reducing those above by the threshold. The inverse Fourier transform is then used to transform the thresholded data back into the time domain. The unique point about this technique over generic IST applications is that the measured part of the FID is kept, and the truncated part replaced by the corresponding part of the time domain representation of the thresholded data. This process is repeated iteratively until the resulting spectrum has converged sufficiently. Theoretical analysis in this work shows the equivalence of IST and ℓ_1 -norm minimisation, and the similarities between IST and maximum entropy reconstruction. This work provides a practical implementation of Compressed Sensing for NMR, but it must be noted that the primary concern is on one-dimensional NMR spectra. In general one-dimensional spectra take only on the order of a few seconds to obtain, and so the decreased acquisition time is not particularly useful.

Around a similar time as [46], Compressed Sensing was also being applied to multidimensional NMR as evidenced by [17]. In this work, IST was applied, but this time to the more time consuming acquisition of multidimensional NMR and without the data replacement operation used previously. While [46] used the Fourier basis to represent the signal in a sparse manner, in [17] the wavelet transform was applied to the Fourier domain spectrum to obtain a sparse representation. IST, as well as iterative soft thresholding by least squares were the algorithms of choice – both, again, effectively means to a minimum ℓ_1 -norm solution. Since the work focused on two-dimensional spectra, uniform radial line, random radial line, and random horizontal line sampling patterns were used, with random radial line sampling found to be the most effective in terms of final reconstruction quality and artefact suppression. This work demonstrated through simulations that iterative soft thresholding algorithms could be effectively applied to the problem of reconstructing undersampled NMR data, as well as introducing the wavelet transform to make NMR spectra sparse for reconstruction purposes.

Turning to practical NMR experiments, there has also been work towards applying Compressed Sensing in an experimental data context. In [27] the same approach as [17] was used, utilising IST and the iterative reweighted least squares algorithm. This was applied to NOESY data from the protein azurin and human ubiquitin. Although not explicitly stated, it appears that the difference between this work and that of [17] is that only the Fourier transform was used rather than the wavelet transform as well. The results obtained in [27] were positive indicating accurate reconstruction was possible for highly downsampled measurements (84% of measurements removed). At the same time [24] also applied

Compressed Sensing techniques to the reconstruction of NMR spectra. This work looked at the reconstruction of a two-dimensional HSQC data set for human ubiquitin and three-dimensional HNCA and HN(CO)CA data sets for a transmembrane protein. The results demonstrated a significant decrease in acquisition time thanks to the use of Compressed Sensing. Also examined was the robustness of the reconstruction to low SNR. As mentioned previously many techniques, such as Projection-reconstruction NMR using the lowest value algorithm struggle to produce good reconstructions for data with low SNR. However, it was demonstrated in [24] that the Compressed Sensing based reconstruction yielded spectra that were still useful, even when the SNR had been reduced by an order of magnitude. The works [27] and [24], essentially replicate the results of [17] for real experimental data, demonstrating the applicability of Compressed Sensing in a pragmatic NMR spectroscopy context.

Finally, a recent paper [42] examined the choice of basis for obtaining sparse spectra and the implications for their reconstruction. This work utilised simulations to reconstruct downsampled NMR spectra. It differs from previous works in that it minimises the $\ell_{0.5}$ -norm. Theoretical analysis, and simulations showed that in terms of reducing reconstruction error using a Fourier, rather than wavelet basis to represent the spectrum in a sparse manner, and minimising the $\ell_{0.5}$ -norm, was advantageous.



Chapter 3. Review of Supporting Theory

3.1 Compressed Sensing

Compressed Sensing [15] is a framework that provides the tools necessary to reconstruct *sparse* signals from a total number of measurements substantially less than the signal dimension. Sparsity is the key constraint, with a signal said to be *s*-sparse taken to mean that the signal has at most *s* non-zero entries.

The motivation of compressed sensing is perhaps most clearly described through its mathematical formulations. Specifically, the signal and measurements have a linear relationship given by:

$$y = Ax \tag{3.1}$$

Here the measurements are represented by a vector y of length m , and the underlying signal, x , is represented by a vector of length n . It is generally assumed that $n > m$ i.e. the number of measurements is smaller than the number of elements the signal is composed of. A is an $m \times n$ matrix through which the measurements and signal are related. For instance in the case of a one-dimensional NMR spectroscopy experiment, y is a free induction decay (FID), x is the spectrum, and A is the one-dimensional discrete (sampled) Fourier transform (DFT) matrix. In the case of two, and higher dimensional experiments, y and x are “vectorised” with each column of the matrix representing the signal and measurements concatenated onto the end of the next.

At first it may seem strange to attempt a solution to a linear equation of such a form as it is underdetermined and there could be infinitely many solutions. This is where sparsity comes into play. Essentially the goal is to find the sparsest solution to Equation 3.1 (in some basis). A sparse basis is taken to mean that x can be represented as $x = \Psi\theta$, where Ψ is a matrix in which each column is a vector of the basis the signal is represented in and θ is sparse vector of coefficients. This idea is probably more familiar in the context of image compression. For example, the JPEG-2000 image standard [48] represents images in a wavelet basis, removes small coefficients, and then transforms the result back into the image domain. The result is an image that is sparse in the wavelet domain, looks almost the same in the image domain, and has had its file size significantly reduced. However, this process is a little strange as an image is fully sampled, often with many millions of pixels, then compressed and most of the information discarded. On a very conceptual level, compressed sensing seeks to recover this sparse representation with the minimal number of measurements necessary by combining measurement and compression [7].

Before turning to the actual process of reconstruction, some additional constraints that compressed sensing places on the sensing matrix, A , must be mentioned. Compressed sensing requires that sensing matrices be as *incoherent* as possible. This can be explained by first expressing the full measurement model:

$$y = \Phi\Psi\theta.$$

Here a sampling matrix, Φ , has been introduced. The dimensions of this matrix are $m \times n$, and so it can be seen that when multiplying $\Psi\theta$ with the signal vector of dimension n in some basis Ψ , it will result in measurements y of length m . Coherence considers the vectors, ϕ_k , that are the columns of the

sampling matrix, and vectors, ψ_j , of the basis that the signal is represented in. Coherence is defined by [7]:

$$\mu(\Phi, \Psi) = \sqrt{n} \max_{1 \leq k, j \leq n} |\langle \phi_k, \psi_j \rangle|.$$

Essentially it is the maximum dot product between any two columns of the sampling matrix and signal basis, or maximum correlation. Compressed Sensing reconstruction requires that coherence be small for reconstruction via the minimisation of the ℓ_1 -norm. Interestingly, a way to ensure high incoherence is to randomise the sampling matrix – randomly sampling the signal [7].

With the basic motivation of Compressed Sensing established, attention is now turned to how to recover the signal. The most direct method would be reconstruction in which the “ ℓ_0 -norm” is minimised, represented as:

$$\min \|x\|_0 \quad \text{subject to} \quad y = Ax. \quad (3.2)$$

This represents the total number of elements in the vector x that are nonzero. Strictly speaking $\|x\|_0$ is not a norm, but it is generally referred to as such. As explained in [17] this problem is NP and so solving it is not practical. Instead, the general approach taken in compressed sensing is to minimise the ℓ_1 -norm, defined as:

$$\|x\|_1 = \sum_{i=1}^N |x_i|.$$

Thus the minimisation problem becomes:

$$\min \|x\|_1 \quad \text{subject to} \quad y = Ax. \quad (3.3)$$

This, of course, will not yield the same result as Equation 3.2 since the ℓ_0 -norm measures the number of nonzero coefficients, and the ℓ_1 -norm measures magnitude. One of the breakthroughs for compressed sensing was the demonstration that if the sensing matrix satisfies the *restricted isometry property* (RIP) sufficiently, then the solution of the ℓ_1 -norm minimisation is the same as that of the ℓ_0 -norm problem [7]. Practically speaking, this turns the NP-hard problem of Equation 3.2 into one that can be solved using linear programming. The RIP as defined in [7] is:

$$(1 - \delta_s)\|x\|_2^2 \leq \|Ax\|_2^2 \leq (1 + \delta_s)\|x\|_2^2,$$

where δ_s is referred to as the *isometry constant*, and is the smallest number such that the RIP holds for all s -sparse vectors x . Paraphrasing [7]: from this expression it can be seen that the RIP requires multiplication of an s -sparse vector x by a sensing matrix A to not alter the Euclidian length more than an amount specified by the bound $(1 \pm \delta_s)\|x\|_2^2$. As explained in [13] the RIP is connected to a property known as the *null space property*. This property is motivated by the reconstruction of s -sparse vectors necessitating that the null space of A may not contain s -sparse vectors. If the null space was to contain them it would mean that there would be no way to discern two different s -sparse vectors, and reconstruction would be impossible.

The specific requirement for Compressed Sensing is that a sensing matrix must satisfy the RIP, with isometry constant $\delta_{2s} < \sqrt{2} - 1$. Roughly speaking, the sensing matrix must preserve the Euclidian length of all $2s$ -sparse vectors to a fraction less than $\sqrt{2} - 1$. If this is satisfied, then the solution to the ℓ_1 -norm minimisation problem will be the same as that of the ℓ_0 -norm minimisation problem. However, the cost is that more measurements are required for reconstruction in the ℓ_1 -norm case.

The minimisation problems stemming from Equation 3.1 are for the case in which the measurements are not perturbed by noise. In the generally encountered case, the expression for the measurements

becomes:

$$y = Ax + w, \quad (3.4)$$

where w represents additive noise. Assuming the noise is bounded, this turns the minimisation into:

$$\min \|x\|_1 \quad \text{subject to} \quad \|y - Ax\|_2 \leq \epsilon. \quad (3.5)$$

Again, as mentioned in [17] Equation 3.5 can be solved using linear programming. These techniques, however, can be come computationally very expensive as the amount of data increases. This is especially true for multidimensional NMR where the data size increases as the product of the number of points sampled in each dimension. Because of this special algorithms are introduced that are a little more practical. Some of these algorithms are already familiar in Compressed Sensing NMR such as *Iterative Soft Thresholding* (IST) (analysed in [12]) and *Iterative Hard Thresholding* (IHT) (analysed in [5]).

Algorithm 1 The basic scheme for iterative thresholding algorithms

Input: matrix A , measurement vector y

Output: $\min \|x\|_1$ subject to $\|y - Ax\| \leq \epsilon$

while $\|y - Ax_i\| \geq \epsilon\|y\|$ (or haven't reached maximum iterations) **do**

$$x_{i+1} \leftarrow x_i + \delta_{t_i}(x_i + A^T(y - Ax_i))$$

$$t_{i+1} \leftarrow \mu t_i$$

end while

Iterative thresholding algorithms continuously refine their estimates until convergence, or until the total number of iterations has been reached. The general form of iterative thresholding algorithms is that given in Algorithm 1. The algorithm itself exists in many forms but can generally be classed according to the type of thresholding function they employ. Like the name implies, IST employs a soft thresholding function defined by the following piecewise function.

$$\delta_t(y) = \begin{cases} \text{sgn}(y)(|y| - t) & |y| > t \\ 0 & |y| \leq t \end{cases}, \quad (3.6)$$

where $\text{sgn}()$ is the sign function and t is a threshold. Application of Equation 3.6 results in all values below the threshold being set to zero, and those above being scaled by the threshold. *Iterative Hard Thresholding* (IHT) employs a simpler thresholding function given by

$$\delta_t(y) = \begin{cases} y & |y| > t \\ 0 & |y| \leq t \end{cases}. \quad (3.7)$$

This is similar to Equation 3.6, but without scaling of coefficients above the threshold. Iterative thresholding algorithms provide an approximate solution to Equation 3.5 that satisfies many of the constraints laid down by compressed sensing. Other greedy algorithms, such as *orthogonal matching pursuit* (OMP) also exist and are frequently used in practical Compressed Sensing applications. Iterative thresholding algorithms have been the algorithms of choice for the application of Compressed Sensing to NMR. Iterative thresholding algorithms, are not only easy to implement, but can also be made fast. Although Algorithm 1 specifies matrix multiplication by what corresponds to the DFT and inverse DFT matrix in NMR, the use of the *Fast Fourier Transform* FFT algorithm can significantly speed up this process. The implications of this are that the reconstruction process can be efficient and practical for multidimensional NMR.

3.1.1 Reweighted ℓ_1 -Norm

Due to the computational intractability associated with minimising the ℓ_0 norm, generally the approach taken in Compressed Sensing is to minimise the ℓ_1 -norm. This means additional sampling is required in the case of ℓ_1 -norm minimisation. A means to remedy this disparity to some degree was presented in [8]. In this work the weighted norm approach was demonstrated. This approach uses an initial estimate to reweight the coefficients and obtain a minimisation closer to the ℓ_0 -norm.

Specifically, the weighted minimisation problem is considered [8]. This is given by:

$$\min_x \sum_{i=1} w_i |x|_i \quad \text{subject to} \quad y = Ax$$

The weights, represented by w_i , can be made the diagonal entries of a matrix W and the minimisation written as:

$$\min_x \|Wx\|_1 \quad \text{subject to} \quad y = Ax. \quad (3.8)$$

Paraphrasing [8], when the solution to the equation of the same form as Equation 3.8, but minimising the ℓ_0 -norm is unique, then it is the same as the solution to Equation 3.8. However, the solution to the a problem of the form of Equation 3.3 is different to Equation 3.8. This means that provided the choice of weights is appropriate, the solution the weighted problem could be better than that of simply minimising the ℓ_1 -norm.

Using weights introduces another parameter to the optimisation – something which can increase the difficulty of obtaining an optimal solution. The approach taken in [8] is directly motivated by the desire to make the ℓ_1 -norm minimisation more like that of ℓ_0 -norm. Thus the weights are chosen according to:

$$w_i = \frac{1}{|x_i| + \epsilon}.$$

Here the estimate of each coefficient x_i comes from an earlier iteration in which minimisation is performed without weighting. The ϵ in the denominator is a regularisation parameter that is included to prevent the denominator being zero. It can be seen from this expression that for small coefficients the weighting will be very large. When it comes to minimising this weighted problem small coefficients, due to their large weights they will be penalised. This means that the solution will then be comprised of coefficients of smaller weights. This work also follows the reweighted approach, but aspires to incorporate probabilistic information through the weights into the minimisation procedure.

3.1.2 Approaches to Compressed Sensing with Models

In some respects the term Compressed Sensing with models is an oxymoron. Compressed Sensing techniques exploit the prior knowledge that a signal is sparse and thus already assumes a signal model. The term compressed sensing with models refers to the use of prior knowledge above and beyond that of just sparsity. This is an emerging area in Compressed Sensing, and so far most of the work has considered the only the structure inherent in the wavelet transforms of natural images. This structure, known as *persistence across scales* means that coefficients at one level of the wavelet transform have children in the next. This implies that it is likely that if a coefficient is large in one scale its children will be large, and if it is small its children also small.

Thus far, there appear to have been two major classes of approaches: deterministic, and probabilistic. The work [3] presented theoretical analysis of the deterministic model-based approach, and its implications for reconstruction in general. The analysis is based on the subspaces in which the signal

can lie. Logically, imposing structure should significantly reduce the size of the subspace in which the signal exists. Properties of the reconstruction of signals assuming structure were presented with the introduction of the *restricted amplification property*, analogous to the usual *restricted isometry property* of Compressed Sensing. It was shown that signal structure constraints can easily be incorporated into the existing algorithms used in Compressed Sensing such as *Compressive Sampling Matching Pursuit* (CoSaMP) [38] and IHT with the inclusion of a “pruning” step according to the signal structure. The structures demonstrated were block sparsity – where large coefficients are clustered in blocks, and the wavelet persistence across scales tree structure.

Other methods to incorporate structure into Compressed Sensing have primarily been based on Bayesian statistical models. They focus on the persistence across scales structural properties of the wavelet transform. This structure can be captured well by using a hidden Markov tree model. As mentioned in [44] it appears the first attempt at incorporating this kind of model into structured compressed sensing was in [18]. The approach uses the reweighted ℓ_1 -norm approach similar to that of [8]. The difference here though is that the weights are based on the probability that a coefficient is large. This is represented as:

$$W_n^i = \frac{1}{[p(s_n = 1|\theta^{i-1}, \Theta) + \delta]^q}.$$

In this expression θ_{i-1} represents the data estimate from the previous iteration, and Θ are the parameters of the statistical model. δ plays the role of a regularisation parameter, and q appears to act to smooth the range of weights. Combining this, $p(s_n = 1|\theta^{i-1}, \Theta)$, represents the probability that a given coefficient is large (hidden state $s_n = 1$), conditioned on the data value and model parameters. To obtain these weights and initial signal estimate is produced based on the measurements. Then, the expectation maximisation algorithm is used to fit model parameters to the estimated data. The model is one of mixed Gaussians given by:

$$\begin{aligned} f(\theta_n, S_n = 0) &= \mathcal{N}(0, \sigma_{0,n}^2) \\ f(\theta_n, S_n = 1) &= \mathcal{N}(0, \sigma_{1,n}^2). \end{aligned}$$

There are two distributions that coefficients can belong to. The small valued coefficients are assumed to belong to a Gaussian distribution with small variance – implying they are concentrated around around the mean value of the distribution: zero. The large values, on the other hand, belong to a Gaussian distribution with large variance, allowing the coefficients to drift away from the mean value of zero. It appears the posterior probability values of the coefficients in the first scale are learned from the mixture probability attached to each Gaussian. The *Viterbi* algorithm is then used with transition probabilities from expectation maximisation. These probabilities to compute the weights, and run subsequent iterations of weighted ℓ_1 -norm minimisation.

Other, more strictly Bayesian approaches, have also been employed. The two approaches used in this area are Markov-chain Monte-Carlo methods [21], and Variational Bayes [22]. As their names suggest, these employ Bayesian techniques to augment the Compressed Sensing reconstruction of undersampled signals. Again the persistence across scales structure of the wavelet transform is modelled. It has been noted [44] that despite the the accuracy of Monte-Carlo methods, due to their inherent reliance on random methods they are very time consuming.

To finish this discussion of methods in which prior structural knowledge has been incorporated into Compressed Sensing reconstruction, the *TurboAMP* algorithm is discussed. In [44] and [45] the *Approximate Message Passing* (AMP) algorithm [16] was extended through the sum-product algorithm

on a factor graph. Since this factor graph contains loops, a method similar to *Turbo Decoding* [41] in communications was used. The prior distribution for the coefficients assumed in the derivation of the algorithm is a Bernoulli-Gaussian distribution given by:

$$p(\theta_n | s_n) = s_n \mathcal{N}(\theta_n; 0, \sigma_n^2) + (1 - s_n) \delta(\theta_n).$$

Here, once again, s_n is the hidden Markov state, and θ the signal estimate in the wavelet basis. This model differs slightly from the mixture model commonly used to model wavelet transform structure in that large coefficients are distributed according to a Gaussian distribution and small coefficients concentrated at zero by the delta function. In structure the algorithm is similar to [18] in that it first estimates the signal using AMP¹ and then using a log-likelihood ratio computes a posterior probability that each coefficient is large. Other model parameters are estimated through their hyperprior distributions and the signal estimation. This probability becomes the prior for a hidden Markov tree, on which belief propagation [40] is performed and the probabilities refined according to the signal model. These probabilities are then passed back to AMP where the signal is re-estimated with the prior knowledge from the belief propagation step. The techniques used in the works eloquently tie together signal estimates and probabilistic models for signal structure. The results demonstrated improvements over the Bayesian techniques mentioned in the preceding text, and significant improvements over the deterministic approach described in this section.

There is a clear structural similarity between algorithms developed to exploit signal structure using Compressed Sensing. Both deterministic approaches such as [3] and probabilistic approaches such as [44] begin with an signal estimation phase. This is then used to refine the estimation, or parameters or a model according to prior structural knowledge. This refinement then influences the subsequent signal estimation, and the process of model-based refinement and estimation is repeated.

¹AMP-type algorithms are similar to iterative thresholding algorithms, but differ in their inclusion of term similar to the Onsager reaction term used in statistical physics [16]. In this case the derived algorithm employed a non-linear threshold based on the prior probability that a coefficient is large.

Chapter 4. Proposed Algorithm for the Exploitation of Structure During Reconstruction of NMR Spectra

In this section an algorithm for the exploitation of structure in the reconstruction of NMR spectra using Compressed Sensing developed in this work will be presented. As was already mentioned in Section 1.5, in the Fourier basis NMR spectra from a large subset of different NMR experiments exhibit a large degree of structure. The motivation was to utilise this structural prior knowledge to enhance the reconstruction of NMR spectra that have been downsampled, thus reducing acquisition time and increasing quality over methods that just assume sparsity.

Initial approaches were based on modifying the TurboAMP method for NMR signal structure, however due to the parameter estimation requirements and horizontal line sampling pattern demanded by NMR spectroscopy experiments, it was found that this led to unstable reconstruction. Iterative thresholding algorithms have already been applied in Compressed Sensing reconstruction of NMR spectra [46] [17], and have proved themselves to be robust. As mentioned in the previous chapter, Compressed Sensing algorithms that are intended to exploit signal structure are often similar in their approach of: initial estimation, model refinement, and estimation influenced by the structure. The task of initial signal estimation is easily accomplished using iterative thresholding algorithms. Also, performing statistical updates based on signal structure can be accomplished using techniques such as belief propagation [40]. What remains, is estimation of prior probabilities from the signal estimates for use in updates of a probabilistic model, and a method for enabling this statistical information to have some effect on the subsequent signal estimation steps. Here, the reweighted norm approach of [18] was utilised as it appears the most direct approach for incorporating probabilistic information into signal reconstruction. In terms of initial posterior probability estimation, a Rician distribution was assumed, and a log-likelihood method similar to that used in *Detection Theory* [26] in signal processing applied. In this chapter, these steps will be described in further detail.

4.1 Initial Estimate and Posterior Probabilities

The initial signal estimate is done using iterative or hard thresholding (shown in Algorithm 1). This provides the estimate of the signal, \hat{x} , needed to estimate the posterior probabilities of the data given the measurements. Following the signal structure in the Fourier domain explained in section 1.5, and the work [42] which indicated that the Fourier basis may provide a sparser representation than the wavelet basis, the NMR signal is represented in the Fourier basis.

NMR data, due to quadrature detection, is complex. Owing to the large amount of data points the large and small coefficients that comprise the signal in the Fourier domain are assumed to belong to complex Gaussian distributions. In this case the absolute value of this data will be distributed according to the Rician distribution [26]. Following [26], assuming the real and imaginary parts of the signal to be distributed according to the Gaussian distributions $\mathcal{N}(\mu_1, \sigma^2)$ and $\mathcal{N}(\mu_2, \sigma^2)$ respectively, the definition

of this distribution is given as:

$$\text{Rice}(x; \alpha, \sigma) = \begin{cases} \frac{x}{\sigma^2} \exp\left[-\frac{1}{2\sigma^2}(x^2 + \alpha^2)\right] I_0\left(\frac{\alpha x}{\sigma^2}\right) & x > 0 \\ 0 & x < 0 \end{cases}, \quad (4.1)$$

where $\alpha^2 = \mu_1^2 + \mu_2^2$, and $I_0(u)$ is the *modified Bessel function of the first kind* with order $r = 0$ given by:

$$I_0(u) = \int_0^{2\pi} \exp(u \cos \theta) \frac{d\theta}{2\pi}.$$

Coefficients are assumed to belong to either of two Rician distributions: one for large coefficients, $\text{Rice}(x; \alpha_1, \sigma)$, and one for small, $\text{Rice}(x; \alpha_0, \sigma)$. Conceptually this means that a base level of complex noise is assumed, the absolute value of which is distributed by the Rice distribution with parameters α_0 , and σ . The effect of a coefficient being large is to perturb this base level of noise. This alters the mean values of the Gaussian distribution which the real and imaginary components are distributed by. Since the parameter α is equal to $\mu_1^2 + \mu_2^2$, the affect of a peak being present in the spectrum is to essentially change the α parameter of the Rician distribution describing distribution of coefficients that comprise the peak. Here the goal is to determine the probability that a coefficient is large or small i.e. the relative probability of a coefficient belonging to one distribution over the other. This is similar to the problem of detection in statistical signal processing, specifically: the *Neyman-Pearson* theorem [26].

The objective is to determine a response amongst noise using a threshold and likelihood ratio. Mathematically a decision is to be made with respect to two hypotheses: 1. the signal is large and distributed by the Rician distribution for large coefficients (\mathcal{H}_1) 2. the coefficient is small and distributed by the Rician distribution for small coefficients (\mathcal{H}_2). Here the ratio of the two probabilities is of interest:

$$L = \frac{p(x|\mathcal{H}_1)}{p(x|\mathcal{H}_0)}. \quad (4.2)$$

This is generally used in conjunction with a threshold τ , and if the value of this likelihood is above or below a threshold a decision is made (with some error level depending on the threshold) as to whether or not the null hypothesis (\mathcal{H}_0) holds. This threshold represents how much the two distribution functions overlap. Analysing the likelihood ratio after substituting a given data point into the each distribution gives a measure of how far into either distribution the data point is and thus whether it can be considered to belong to a peak or noise. Here the interest is in turning this into a probability, which can be thought of like a posterior probability for a coefficient being large given the data, and thus get refined according to a statistical model. Speaking in terms of implementation, the first task is to determine the parameters of the distribution. In order to do this the data is thresholded above the noise and artefact level. Data below the threshold is used to estimate parameters for the Rician distribution for small coefficients, and data above the threshold used to estimate parameters for large coefficients.

The next step is to obtain a probability from the likelihood ratio. This was done by computing a log-likelihood ratio, given as:

$$\text{LLR} = \ln \left[\frac{p(x|\mathcal{H}_1)}{p(x|\mathcal{H}_0)} \right].$$

In the case of the natural logarithm function, when the ratio of the numerator and denominator is one, then the result is zero. When the numerator is larger, the equation becomes the natural logarithm of a number greater than one, which is positive. When the denominator is bigger, the equation will be the natural logarithm of a positive number less than one, which is negative. Thus the sign of the log-likelihood ratio serves to indicate which distribution function for a given coefficient is larger or smaller. The next

step taken here is to turn this information into a usable probability. If the threshold for deciding whether a given coefficient is large or small using the likelihood ratio is one (i.e. magnitude of both distributions is the same), then the actual ratio of the two distribution values when data is substituted will in effect measure the distance from this threshold. Then, assuming a logistic regression, the log-likelihood can be turned into a probability. This is done using the following expression:

$$p(s_n = 1|y) = \frac{1}{1 + \exp(-\text{LLR})}.$$

Although this is in some respects an artificial measure of the posterior probability, practically speaking, the motivation stems from the nature of the NMR spectrum itself. The step in the algorithm being performed here is the first link between estimation, and structural prior knowledge: conversion of estimated values to probabilities. NMR spectra differ from most popular reconstruction targets like natural images in that the magnitude of the coefficients varies widely across the spectrum. Large areas of relatively low value background are interrupted by small clusters of peaks several orders of magnitude larger than the background. Other peaks, although large, can differ in size by orders of magnitude. This makes simple logistic-regression models based on data values almost impossible, as they become saturated and admit only the largest of peaks with probability greater than zero. The method used here was a means to yield reasonable probabilities that could ultimately be refined through later steps in the algorithm. Also, the relative nature of the weights, as applied in iterative thresholding, mean that so long as the probabilities are consistent relative to each other, generally the reconstruction result is unaffected. The next step in the algorithm takes this posterior probability and assumes it to be a prior probability that a coefficient is large, meaning it can be incorporated into, and influence the result of belief propagation.

4.2 Incorporating Structure: the Belief Propagation Step

In section Section 1.5 the structure present in the COSY experiment is explained. This experiment was chosen as it is representative of the signal structure seen in other experiments such as TOCSY and NOESY. Unfortunately, in the case of NMR the signal structure is not as geometrically consistent as in the wavelet case. The size and positions of the scales that comprise a wavelet transform are easily determined. From here, the coefficients can be built into a hidden Markov tree that models their persistence across each scale. For NMR, although the coefficients form a structure that in a statistical sense is predictable, the large coefficients that make up this structure first must be found in the signal. For consistency, the threshold used to separate coefficients as being large or small is re-used in this step. All coefficients above this threshold are taken as potential members of the signal structure. To update the belief in each node, and thus impose structure, belief propagation was used. To perform this calculation this algorithm utilises two constructs: *messages*, and *transitions matrices*. To illustrate, the equations for the nodes shown in Figure 4.1 are given in the following expressions [30] [40] [25]:

$$\mu_{C \rightarrow D}(d) = \sum_{\sim\{d\}} \left(P(d|c) \prod_{I \in n(C) \setminus D} \mu_{I \rightarrow C}(i) \right).$$

The notation here is a little unique, and so requires some explanation. The term $\mu_{C \rightarrow D}(d)$ represents the “message” sent from C to D ¹. Technically speaking this message is the probability of C given all the

¹Note that the convention of using capital letters to denote variables and lower case letters to denote values has been adopted here.

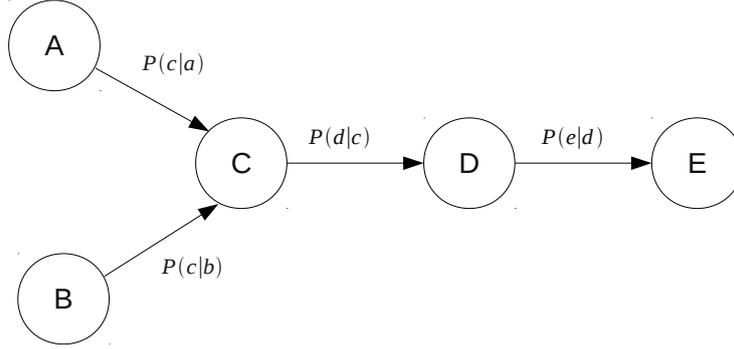


Figure 4.1: A simple Bayesian network in which C is a parent and D and child of C .

nodes that are connected to D through C [40]. The summation is over $\sim \{d\}$, which is shorthand notation for “all variables except for d ”. This is a marginalisation operation as it is, in a sense, “removing” all variables except for d by summing over all their possible states so that the final function will be one of only one variable: d . This function is multiplied by the product of all messages connected to C except for D ; this is represented with set notation where $n(C)$ is the neighbourhood (all directly connected nodes) of C . This rule applies for the messages sent from any node to another. For instance, the message sent from A to C is given by:

$$\mu_{A \rightarrow C}(c) = \sum_{\sim \{c\}} \left(P(c|a) \prod_{I \in n(A) \setminus C} \mu_{I \rightarrow A}(i) \right).$$

In the case of messages opposite the direction of causality (e.g. D to C or C to A), the result is a little different. Here, for the case of a message from D to C , the expression is still of a similar form:

$$\mu_{D \rightarrow C}(c) = \sum_{\sim \{c\}} \left(P(d|c) \prod_{I \in n(D) \setminus C} \mu_{I \rightarrow D}(i) \right).$$

Due to the order of the conditioning ($P(d|c)$) this message is in the form of a likelihood of the information contained in the nodes separated from C by D conditioned on c [40].

To compute the so-called “belief” in a variable, all incoming messages are multiplied [40]. In the case of the node C for example, the belief is given by:

$$\text{BEL}(C) = \alpha \prod_{I \in n(C)} \mu_{I \rightarrow C}(i),$$

where α is a normalising constant. As demonstrated in [40], this can be efficiently implemented using matrix multiplication. Defining a transition matrix, containing information about the probability of each state gives (assuming without loss of generality only two possible states for each variable):

$$\mathbf{T} = P(D = d|C = c) = \begin{bmatrix} P(d_1|c_1) & P(d_2|c_1) \\ P(d_1|c_2) & P(d_2|c_2) \end{bmatrix}.$$

This turns the belief propagation equations from a “parent” to a “child” (as in C to D) into:

$$\mu_{C \rightarrow D}(d) = \left(\prod_{I \in n(C) \setminus D} \mu_{I \rightarrow C}(i) \right) \cdot \mathbf{T},$$

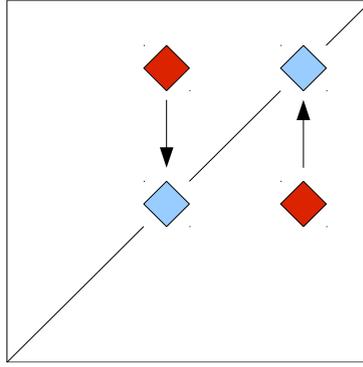


Figure 4.2: The simple parent child relationship between cross and diagonal peaks first modelled.

and from “child” to parent (as in D to C) into:

$$\mu_{D \rightarrow C}(c) = \mathbf{T} \cdot \left(\prod_{I \in n(D) \setminus C} \mu_{I \rightarrow D}(i) \right),$$

where the product of incoming messages is assumed to be a row vector when multiplying the transition matrix and a column vector when being multiplied by the transition matrix.

In the case of the algorithm developed here, the interest is in modelling the structure present in NMR spectra (explained in section Section 1.5). The basic structure for NMR signals is the cross peak and diagonal peak relationship. The motivation is to model this structure using a graphical model, as it is easily understood, and allows efficient computation of beliefs via the belief propagation algorithm. The basic structure represents a spin system and its interactions with two diagonal peaks representing two spin systems and a cross peak either side of the body diagonal representing a correlation between them. Perhaps the first point about the structure for NMR spectra is the that the relationship between two cross peaks and two diagonal peaks forms a loop. This is problematic for belief propagation which has practically been used in “loopy”-form in applications like Turbo-Coding [41] [34], and has been analysed as a means for approximate inference [37], but is still not guaranteed to converge. For this reason the initial structure modelled focused on the relationship between one diagonal peak and a single cross peak as in Figure 4.2. Since cross peaks define diagonal peaks, but diagonal peaks don’t necessarily define cross peaks, the causal direction is taken to be from cross peak to diagonal peak (following methods in [25] and [40]).

In the implementation, the relationship is between individual pixels (coefficients) that comprise the spectrum. These pixels are chosen by searching along the diagonal of a spectrum which has been thresholded so that only coefficients above the noise level are considered. When a diagonal coefficient above the threshold is found, the column it is in is examined. A search along this column is then performed starting from the middle of the diagonal. The length of the “slice” of the diagonal peak is determined by the points where the series of coefficients it is composed of drops below the threshold. Searching along this column upwards and downwards, further series of coefficients above and below the the threshold indicate slices through cross peaks as in Figure 4.3. There are often different sizes so the smallest number of coefficients are matched as in Figure 4.4. After this, belief propagation on a simple two-level parent child structure is performed. For illustrative purposes, the form of the equation representing the message

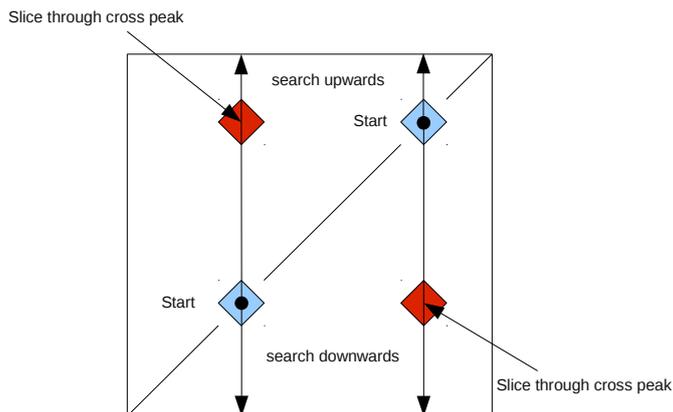


Figure 4.3: The primary method for obtaining coefficients in cross and diagonal peaks. Searching along the body diagonal when coefficients above a threshold are found a search is done vertically to find the corresponding cross peak.

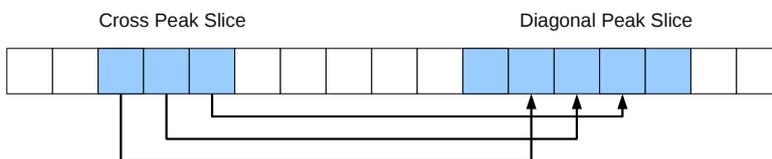


Figure 4.4: The method used for matching coefficients when performing belief propagation in this work. Here the three cross peak coefficients have been paired with the central three diagonal peak coefficients. The two groups of coefficients that constitute the slices through the cross and diagonal peaks are separated by coefficients that are below the threshold (background).

passed from the parent cross peak to child diagonal peak is given in the following:

$$\mu_{C \rightarrow D}(d) = \begin{bmatrix} P(c=0) & P(c=1) \end{bmatrix} \cdot \begin{bmatrix} P(d=0|c=0) & P(d=1|c=0) \\ P(d=0|c=1) & P(d=1|c=1) \end{bmatrix},$$

here D represents the diagonal peak coefficient, and C the cross peak coefficient. The term $\mu_{C \rightarrow D}(d)$ is a 2×1 vector. From this message the belief in D can be calculated:

$$\text{BEL}(D) = \alpha \begin{bmatrix} \mu_{C \rightarrow D}(d=0) \\ \mu_{C \rightarrow D}(d=1) \end{bmatrix} \bullet \begin{bmatrix} P(d=0) \\ P(d=1) \end{bmatrix},$$

where \bullet represents element-wise multiplication. The same algorithm is used to update C , using the previously described formula for a message being passed from a child to a parent. This belief is the probability that will be used as a weight in the next part of the algorithm, which uses weighted reconstruction to re-estimate the signal.

Although the previous method for belief propagation succeeds in imposing prior structural information on the estimates of the coefficients, it does not incorporate some of the stronger structural relationships. Stopping short of loopy belief propagation, the next method matches a diagonal peak with both corresponding cross peaks. Figure 4.5 provides a pictorial representation of this relationship. One of the motivations for this technique is that if one cross peak is not being reconstructed as it is too small, but the other is, then through the influence they exert on each other by belief propagation, the chance it

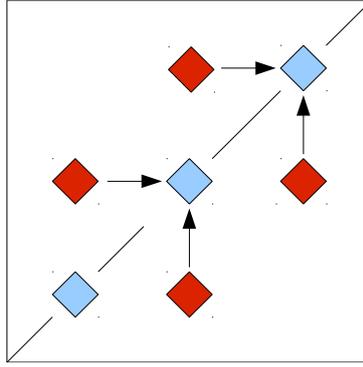


Figure 4.5: A model for spectrum structure that links both cross peaks to a single diagonal peak. The arrows indicate the direction from parent to child.

will be reconstructed will increase. However, using this method introduces some degree of complication in finding the location of the peaks.

The method used to find the coefficients comprising the structure begins in the same manner as that described previously with a search along the body diagonal of the spectrum. After locating a large coefficient on the diagonal, a column slice is taken, and the coefficients belonging to the cross peaks and diagonal peaks identified. Then, for the smaller of the cross and diagonal peak, the length in that column slice, and the width for the peak at coefficient determined. Now, the corresponding diagonal and cross peaks must be found. First the diagonal peak is found by finding the midpoint of the cross peak and projecting a horizontal line until it intersects the body diagonal. This is assumed to be the approximate location of the middle of the second diagonal peak. An area equal to the size of the smaller of the original diagonal peaks and cross peaks is selected. The second cross peak will be along the same columns as the second diagonal peak, and the same rows as the first diagonal peak. Determining the offset from the center of the smaller peak in the original diagonal and cross peak slice, means that same offset can be applied to the second diagonal peak and cross peak slice. The coefficients can then be matched in the previous manner, with a coefficient in each cross peak acting as a separate parent for a coefficient in the first diagonal peak. Figure 4.6 demonstrates the geometry of the pairing method. To finish, belief propagation is performed using the previously described rules this time for the simple structure of two parents and one child. The resulting beliefs then become weights for the weighted reconstruction step.

4.3 Incorporating Probability with Weighted Iterative Thresholding

The final step in closing the reconstruction loop is incorporating the beliefs from the previous step into the signal estimation. This is done through weighted ℓ_1 -norm minimisation given by the following optimisation problem:

$$\min \|q\|_1 \quad \text{subject to} \quad y = Ax, \quad (4.3)$$

where:

$$x = Wq \Rightarrow y = Ax = AWq.$$

This is equivalent to Equation 3.8, but with weights computed based on the probability for each of

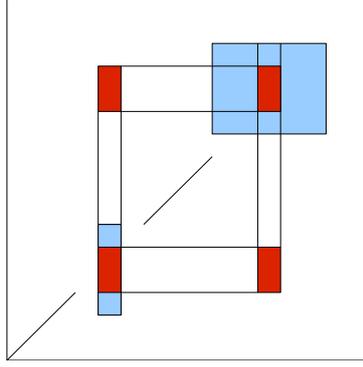


Figure 4.6: The method used to match coefficients for the second structural model utilised. The original diagonal and cross peak referred to in the text are the right most diagonal and cross peaks aligned along the vertical direction. The remaining two are referred to as the second, or corresponding, diagonal and cross peak. Note the smaller of the original diagonal and cross peak defines the maximum cluster of coefficients that will ultimately be connected.

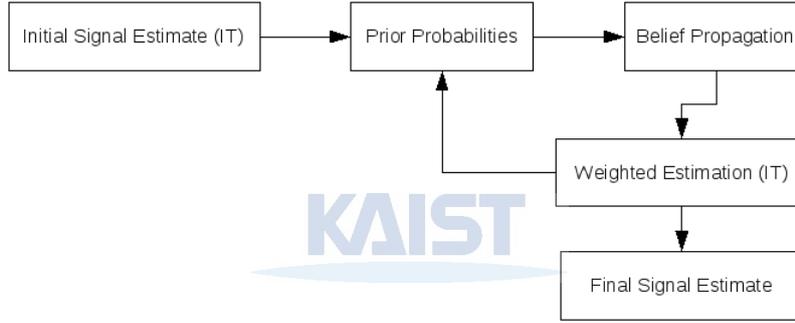


Figure 4.7: An overview of the proposed reconstruction process.

the components is large. In the form of Equation 4.3 the weights are

$$W_n = p(s_n = 1|\hat{x}).$$

These follow, in slightly different form, the weighting method used in [18], except in this case no need was found for regularisation or smoothing parameters. The approximate solution to the ℓ_1 -norm minimisation of Equation 3.8 is easily obtained using iterative soft thresholding or iterative hard thresholding. From here the algorithm repeats. The estimates formed from the weighted estimation form the estimates from which posterior probabilities will be estimated using the likelihood-ratio, and then refined using belief propagation; finally being used for estimation again. This process is repeated so as to optimise the quality of the reconstruction. Figure 4.7 provides a pictorial representation of the reconstruction process.

Chapter 5. Methods

5.1 Data Preparation and Simulation

The data used in these simulations was two-dimensional and obtained from the *Biological Magnetic Resonance Data Bank* [49]. Specific measures taken to process each particular data set are given in the next section. The general process followed: correction for Bruker digital filtering, first point correction, zero-filling, apodization, and Fourier transformation. This formed the fully sampled data. To simulate measured FIDs from this data the inverse Fourier transform was performed, and FIDs removed according to the downsampling required. Two-dimensional NMR spectroscopy requires whole FIDs to be removed at a time, and so the closest number of rows removing enough data points to match the downsampling rate were removed. An example sampling matrix can be seen in Figure 5.1.

5.2 Thresholds and Transition Probabilities

From the previous descriptions of the algorithm developed and applied in this work it is clear that, in terms of the statistical structural model, two parameters are important. The first is the threshold that separates large coefficients from small, and the second is the transition probability matrix used in the model to refine the initial probability estimates. To find a threshold, the method employed here was to observe the spectrum and choose a threshold such that spectrum artefacts are lower than the chosen value. This is somewhat similar in motivation to the method for choosing a threshold used in [46]. This threshold was then used for both estimating the prior probabilities using the Rician distribution and determining which will form the structure used in belief propagation.

Determining the appropriate transition probabilities was a slightly more complex task. The approach taken here was based analysing the reconstruction result stemming from different combinations of probabilities. In steps of 0.1 each probability combination between 0.1 and 0.9 for $p(d = 1|c = 1)$ and $p(d = 0|c = 0)$ was used in belief propagation for a particular sampling matrix. The resulting reconstruction error was then compared to the reconstruction error when no belief propagation was used (but weighting according to the Rician distribution occurred), and only thresholding. The measure for reconstruction error was the *normalised mean square error* (NMSE), which is defined (in the decibel measure) as:

$$\text{NMSE}(\text{dB}) = 10 \log \left(\frac{\|x_0 - \hat{x}\|_2^2}{\|x_0\|_2^2} \right).$$

Belief propagation was implemented in C++ and interfaced with the rest of the code (in Matlab[®] *m-files*) through *mexfiles*. All simulations were performed on a PC running *Debian GNU/Linux*.

5.3 Reconstruction: Comparison with Normal Iterative Thresholding

One of the central motivations of introducing prior structural knowledge into the reconstruction of NMR spectra is to improve upon the reconstruction results achieved by existing methods. This was

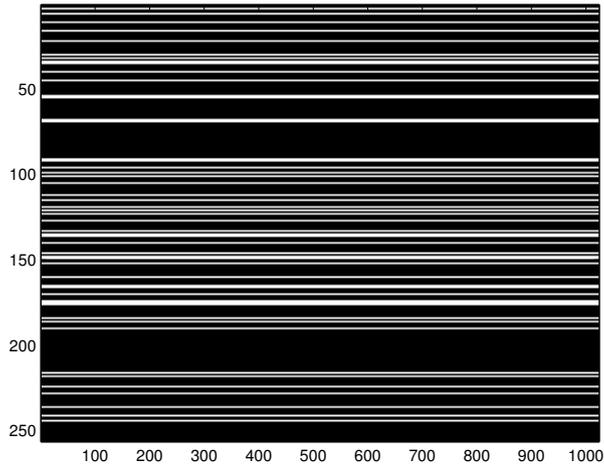


Figure 5.1: A 256×1024 sampling matrix in which approximately 80% of the data has been removed in row form.

accessed by comparing the NMSE of the reconstruction of the same data set, with the same sampling matrices, for a series of downsampling rates. Performing multiple trials, and averaging the results meant that the reconstruction results obtained through the previous, and developed method could be easily compared. Both IST and IHT were used as structure can be incorporated into both through the methods previously outlined.



5.4 Reconstruction: Accessing the Effect of Belief Propagation

The primary point of difference in the algorithm developed when compared to previous techniques is the incorporation of structure into the reconstruction. It was clear early on that weighting the data to be reconstructed had a positive effect on the reconstruction result. What was not clear, however, was what improvement was due to simply utilising weights based on the probability obtained from the Rician distribution, and those refined using belief propagation. Essentially, the effect, in quantitative terms, of belief propagation (which is based on the assumed prior structure of the data) on the reconstruction quality was desired. This was done by the same method used to compare the new algorithm to previous techniques. The NMSE obtained from the reconstruction using just weighting, and that using weighting and belief propagation was computed for a series of downsampling rates and number of trials. Graphing the average reconstruction error of each method enabled efficient evaluation of the affect of belief propagation on reconstruction.

Chapter 6. Results

6.1 Nonactin

The first real spectrum used to test the developed algorithm was a COSY spectrum for nonactin obtained from the *Biological Magnetic Resonance Data Bank* [49] (specifically [2]). Nonactin is a cyclic *ionophore* produced by some bacteria and it was chosen here because of the public availability of its spectra and the clean COSY spectrum with good number of diagonal and cross peaks it produces. The spectrum was processed using *nmrPipe* [14], and *matNMR* [50], in addition to Matlab[®] scripts developed in the course of this work for processing NMR spectra. Sine apodization was performed, as well as zero filling and correction for Bruker digital filtering (following the parameter *GRPDLY*). The spectrum for this data set can be seen in Figure 6.1.

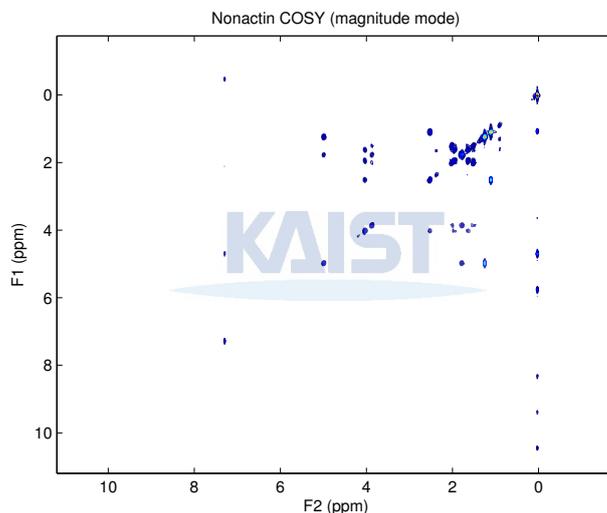


Figure 6.1: COSY spectrum for nonactin used in the simulations.

6.1.1 Threshold and Transition Probabilities

Exhaustive simulation of different thresholds and transition probabilities, and the resulting NMSE, was used to determine the effect on reconstruction error. This first step was necessary because the rest of the algorithm relies on wise selection of these two parameters.

Although it was found that threshold does have an effect on the quality of reconstruction; in general, so long as a threshold was chosen such that it eclipsed the level of artefacts and preserved the peaks, the reconstruction results were positive. Thresholds that are too low, will include artefacts and noise into the belief propagation step, resulting in incorrect inference. On the other hand, thresholds that are too high will not include enough coefficients in the belief propagation step meaning that it will yield little benefit for the reconstruction.

In the case of a threshold that is higher than the noise and artefact level, but one that is still able to capture enough coefficients for reasonable inference, probabilities that favour persistence of coefficient

state between cross and diagonal peaks yielded the best reconstruction results. This implies relatively high values for the probability a diagonal coefficient will be large given the corresponding cross peak is small, $P(d = 1|c = 0)$, and high probability that a diagonal peak coefficient is large given the corresponding cross peak coefficients are large, $P(d = 1|c = 1)$.¹ It must be noted, however, that probabilities that were too high (above 0.9) proved too strong, and were a detriment to reconstruction. The range of good probabilities that yielded consistently low reconstruction errors was approximately 0.6 to 0.9. For the case of small cross peak coefficients transitioning to large diagonal peak coefficients, the probabilities could be somewhat stronger (around 0.9) and reconstruction quality still good. This indicates that coefficients that may have been incorrectly estimated in the cross peaks need a strong probability of transition to ensure that they do not decrease the probability of the corresponding diagonal peak coefficients being large.

6.2 Comparison of Proposed Algorithm to Previous Techniques

6.2.1 Iterative Hard Thresholding

The simulations performed in this section were motivated by accessing the gains to reconstruction of the proposed method versus those previously demonstrated methods which used only thresholding [17] [46] [24] [27]. Iterative hard thresholding, with and without statistical updates was the algorithm chosen here. The average NMSE from each technique over a range of downsampling rates for two level parent to child belief propagation, and two parent belief propagation was determined (Figure 6.2).

They demonstrate that the reconstruction error for IHT with belief propagation is lower. Belief propagation using only two level inference, rather than a two parent structure appears to have a lower error. It is suspected that the reason for this is that the two parent network only matches the smallest group of coefficients in the collection of two cross peaks, and two diagonal peaks. On the other hand, the single chain method matches coefficients vertically for each cross and diagonal peak separately. This means that it captures, and updates more coefficients. Nevertheless, the reconstruction error is improved using both techniques, and the difference between the two rather slight (around 0.2dB). For the lowest downsampling case, the NMSE for IHT based reconstruction was better than that using statistical updates. However, as will be illustrated in the next section with a COSY spectrum for lactose, the NMSE may not gauge the reconstruction quality completely.

One observation made in the tests of IHT is its sensitivity to model parameters. Altering the threshold or transition probabilities could have a negative effect on the reconstruction results. Figure 6.3 demonstrate the effect of a lower threshold on reconstruction in a similar manner to that already presented. It appears the problem here is that the threshold, being set too high, is not updating enough coefficients to have a significant effect on the reconstruction error level.

¹Note that when referring to high probabilities it is meant probabilities above the non-informative value 0.5.

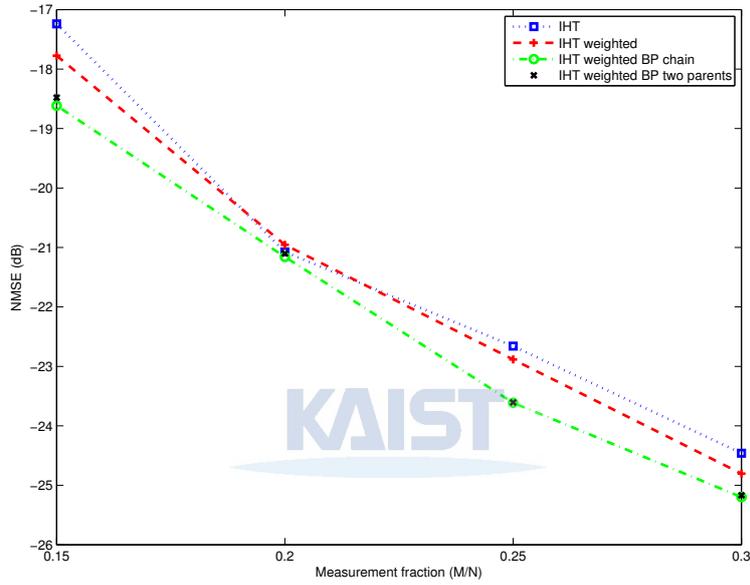


Figure 6.2: NMSE plotted against the ratio of measurements to the total amount of data for IHT, IHT with weighting, IHT with weighting and belief propagation on a two-level parent and child node structure, and IHT with weighting and belief propagation with two parents connected with a single child. The NMSE resulting from the use of 10 different randomly selected sampling matrices at each level of measurements were averaged. The data set was nonactin COSY (magnitude mode). The method for determining the structure coefficients was simple thresholding (threshold = 5×10^6). Transition probabilities were: $P(d = 0|c = 0) = 0.1$ and $P(d = 1|c = 1) = 0.9$.

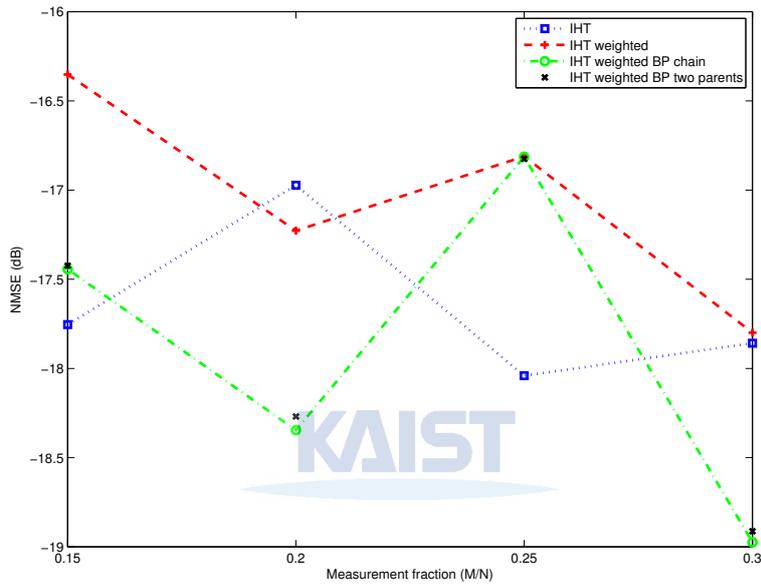


Figure 6.3: The affect of parameter changes: NMSE plotted against the ratio of measurements to the total amount of data for IHT, IHT with weighting, IHT with weighting and belief propagation on a two-level parent and child node structure, and IHT with weighting and belief propagation with two parents connected with a single child. The NMSE resulting from the use of 5 different randomly selected sampling matrices at each level of measurements were averaged. The data set was nonactin COSY (magnitude mode). The method for determining the structure coefficients was simple thresholding (threshold = 1×10^7). Transition probabilities were: $P(d = 0|c = 0) = 0.1$ and $P(d = 1|c = 1) = 0.9$. Here the higher threshold has negatively affected the reconstruction.

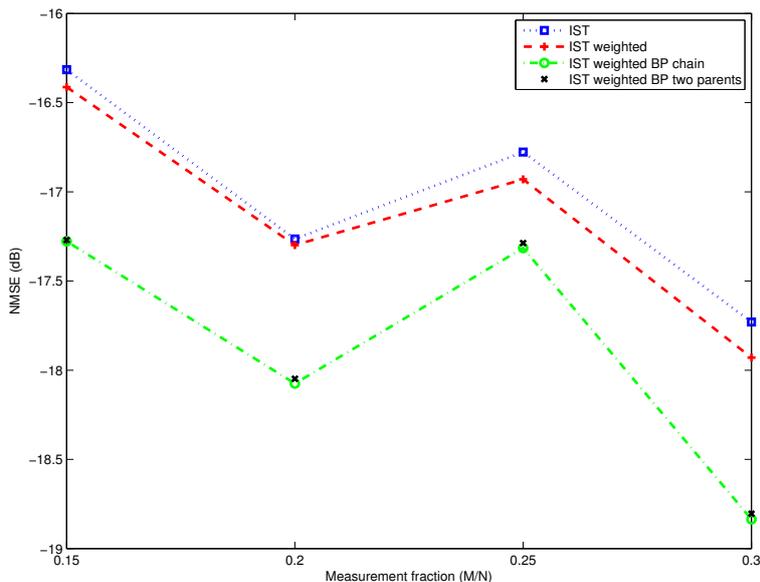


Figure 6.4: NMSE plotted against the ratio of measurements to the total amount of data for IST, IST with weighting, IST with weighting and belief propagation along a two-level parent and child node structure, and IST with weighting and belief propagation with two parents connected with a single child. The NMSE resulting from the use of 5 different randomly selected sampling matrices at each level of measurements were averaged. The data set was nonactive COSY (magnitude mode). The method for determining the structure coefficients was simple thresholding (threshold = 5×10^6). Transition probabilities were: $P(d = 0|c = 0) = 0.1$ and $P(d = 1|c = 1) = 0.9$. Note these are the same parameters as, and the same sensing matrices used in the results shown in Figure 6.3.

6.2.2 Iterative Soft Thresholding

The same simulations as conducted for IHT were also used to test the same technique, but this time utilising IST. Figure 6.4 demonstrates the results for IST in the same format as previously utilised using the same parameters, data, and sensing matrices as those shown in Figure 6.3 to yield poor results. They demonstrate consistent improvement of the NMSE over the previous technique (IST) used in [17] [46]. In comparison with IHT it was found that in general IST is more robust to high downsampling – observed to seldom diverge even for high downsampling. On the other hand, IHT, below around 80% to 85% data removal, diverges often and is unpredictable. This makes it unsuitable around this range. Figure 6.5 Shows an NMSE comparison of reconstructions from IHT and IST demonstrating IST improved stability.

In addition to this, IST tolerates a wider range of threshold levels and permits stronger transition probabilities (in excess of 0.9, and below 0.1). The trade-off here is that, despite its sensitivity, IHT is capable of, in general, yielding lower reconstruction error.

Results for both IHT and IST also demonstrate an important quality: belief propagation positively affects reconstruction error. As mentioned in the previous chapter, weighting can have a positive effect on reconstruction even without belief propagation. The simulations for both IHT and IST with and without belief propagation clearly show that belief propagation consistently improves upon the reconstruction error of weighted iterative thresholding without belief propagation. This is a demonstration of the power imposing signal structure can have on the reconstruction.

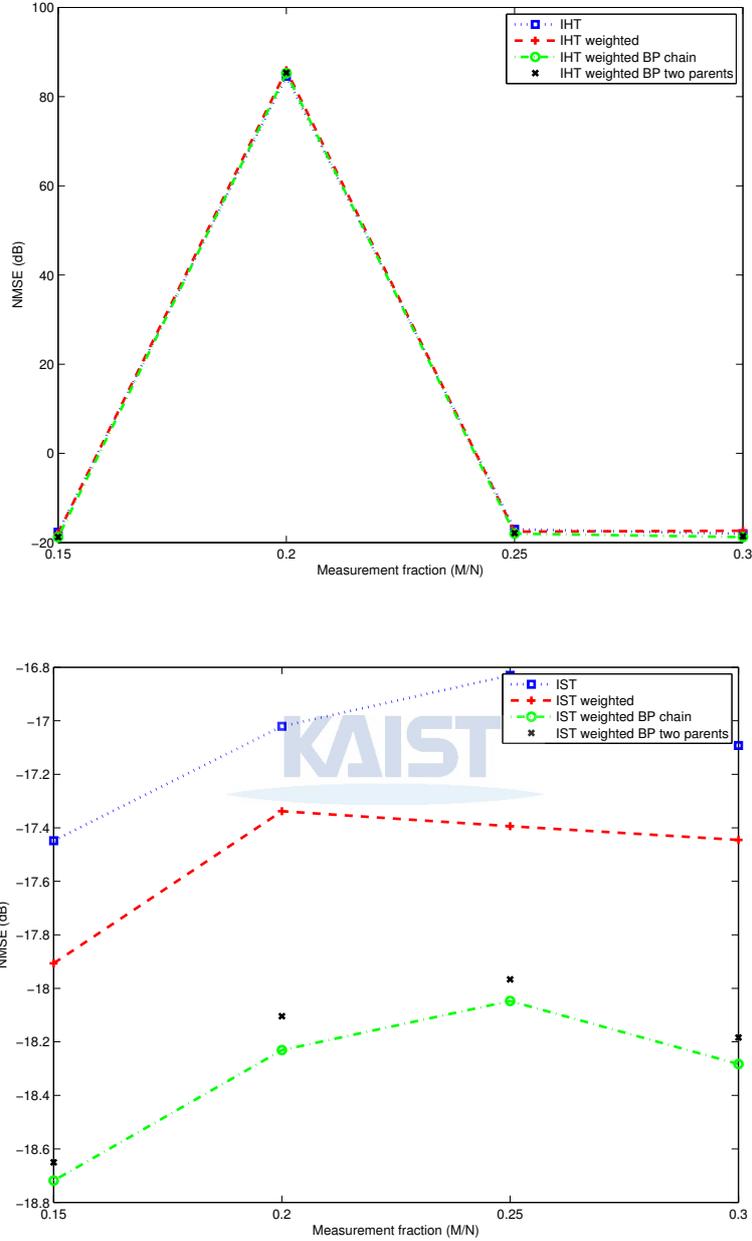


Figure 6.5: NMSE versus fraction of measurements kept for IHT and IST demonstrating divergence in the case of IHT. The NMSE resulting from the use of 5 different randomly selected sampling matrices at each level of measurements were averaged. The data set was nonactin COSY (magnitude mode). The method for determining the structure coefficients was simple thresholding (threshold = 1×10^7). Transition probabilities were: $P(d = 0|c = 0) = 0.1$ and $P(d = 1|c = 1) = 0.9$. Note the divergence at measurement fraction 0.2. In cases of convergence IHT with structure has a consistently lower error than just IHT and weighted IHT.

6.2.3 Reconstructed Spectrum

From a practical spectroscopy point of view, the NMSE is not the whole story. It was noted in some cases the the NMSE for reconstruction utilising structure was in fact lower than simple IHT. However, the spectrum had fewer artefacts in the case of IHT with structure. Therefore, example reconstructed spectra are shown in Figures 6.6 and 6.7. These figures indicate that in the case of the incorporation of structure in the signal model the artefacts (which appear dotted lines, but are in fact the contour representation of the tops of “sinc-wiggles”) are reduced (Figure 6.8 and 6.9). From a practical point of view this highlights the benefit of incorporation of structure which produces a cleaner spectrum with less artefacts. The artefact level in the spectrum with belief propagation, versus that without, is also lower demonstrating the influence structure has had on the reconstruction. It is interesting to note that in the case demonstrated in Figures 6.6 and 6.7, the NMSE for the reconstruction that is weighted, but does not refine the probabilities using belief propagation, the NMSE is lower for than for both simple IHT and IHT with weighting and belief propagation. Also shown in Figures 6.10 and 6.11 are the spectrum for reconstruction using IST. To further demonstrate the effect of belief propagation Figure 6.12 shows a one-dimensional slice through the spectrum reconstructed using IST.



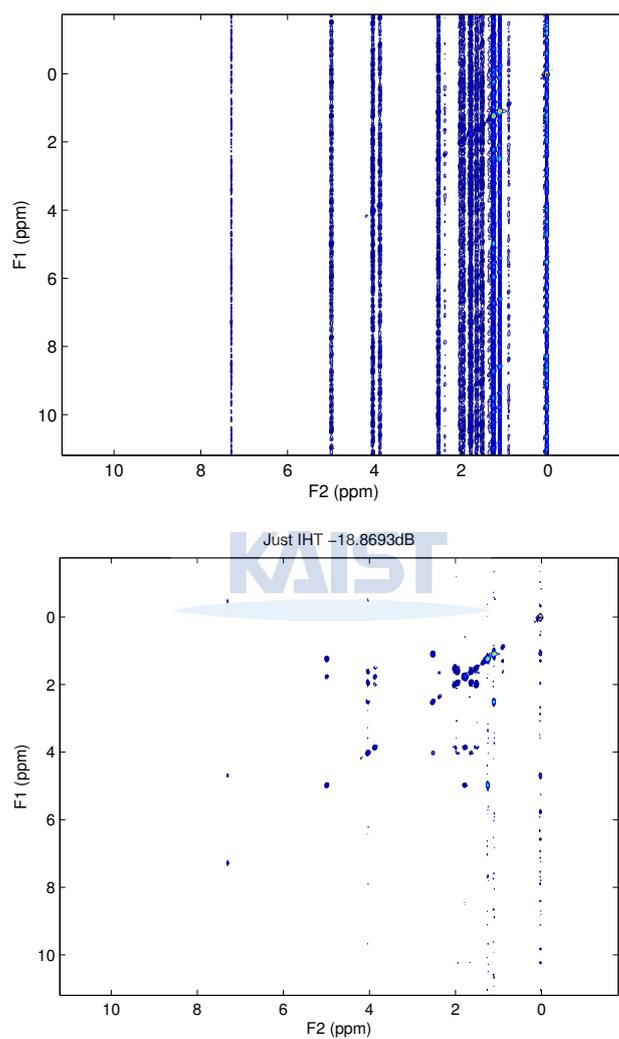


Figure 6.6: A selected reconstruction of a nonactin COSY spectrum using 15% of all measurements with NMSE indicated in the title of each graph. The spectra are: (top) the direct Fourier Transform, (bottom) the direct Fourier transform.

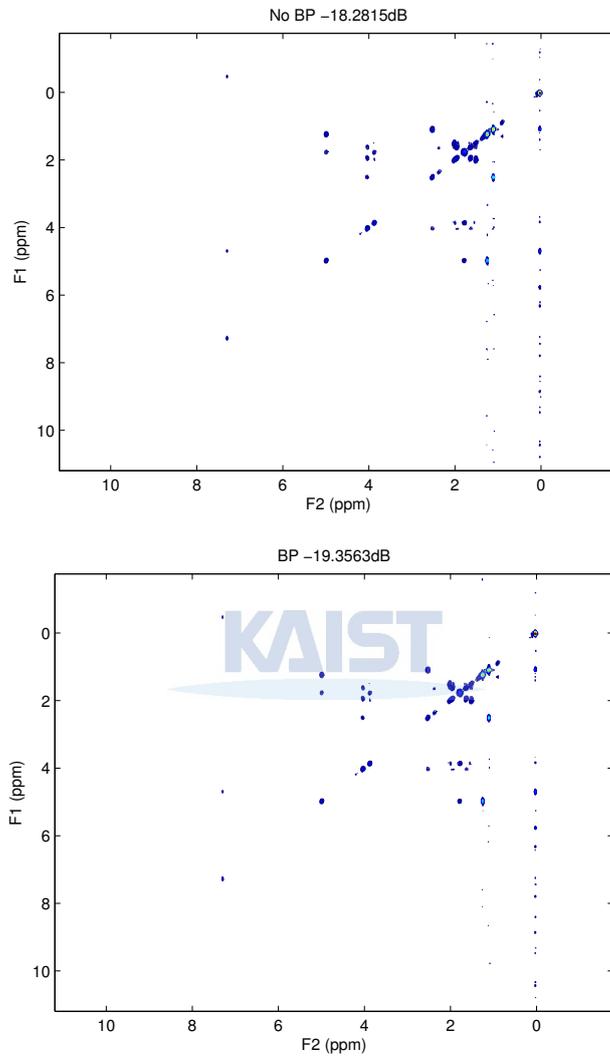


Figure 6.7: A selected reconstruction of a nonactin COSY spectrum using 15% of all measurements with NMSE indicated in the title of each graph. The spectra are: (top) the reconstruction using probabilistic weighting with no belief propagation, (bottom) reconstruction with weighting and belief propagation in which diagonals are matched with two cross peaks where appropriate. Transition probabilities were: $P(d = 0|c = 0) = 0.1$ and $P(d = 1|c = 1) = 0.7$, and the threshold was 1×10^6 . Note the reduction in artefacts along the right side in the vertical direction in the reconstruction.

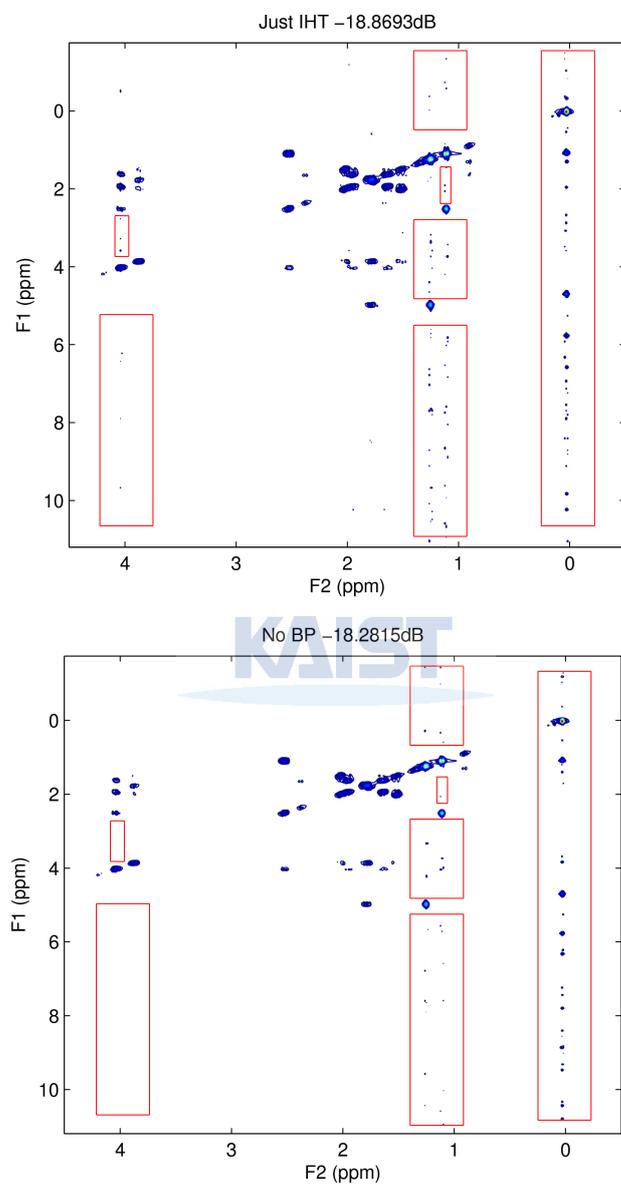


Figure 6.8: Reconstruction of Nonactin COSY spectrum. The spectra are: (top) the reconstruction using IHT, (bottom) reconstruction with weighting and no belief propagation. The view is a zoomed version of that previously presented in which artefact presence and absence has be highlighted.

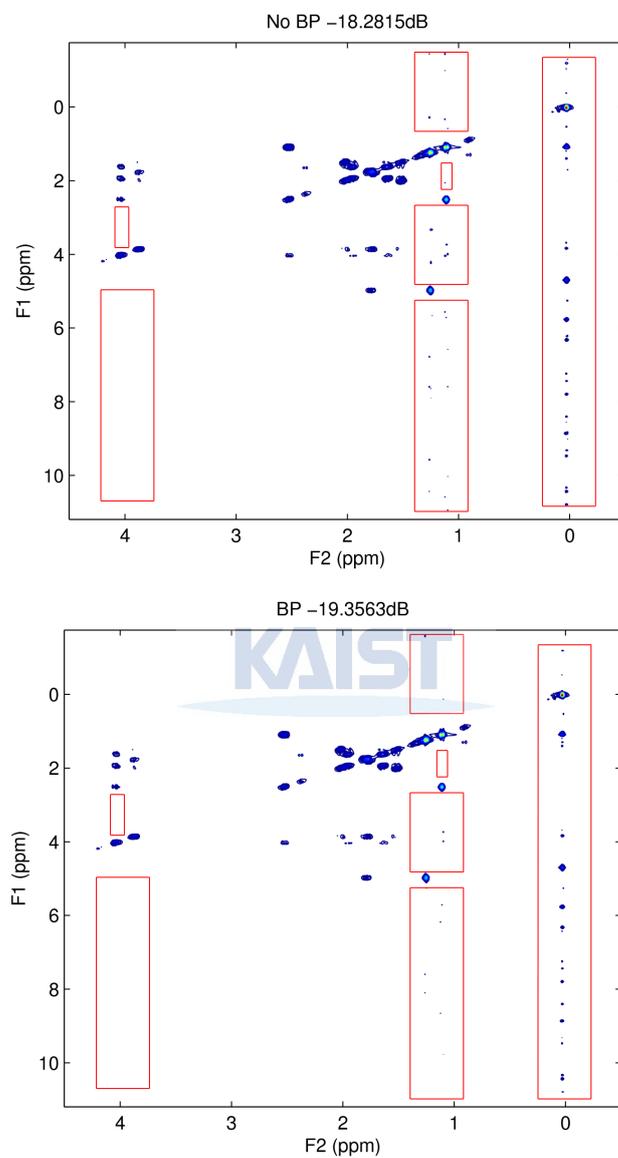


Figure 6.9: Reconstruction of Nonactin COSY spectrum. The spectra are: (top) reconstruction with weighting and no belief propagation, and (bottom) reconstruction with weighting and belief propagation in which diagonals are matched with two cross peaks where appropriate. The view is a zoomed version of that previously presented in which artefact presence and absence has be highlighted.

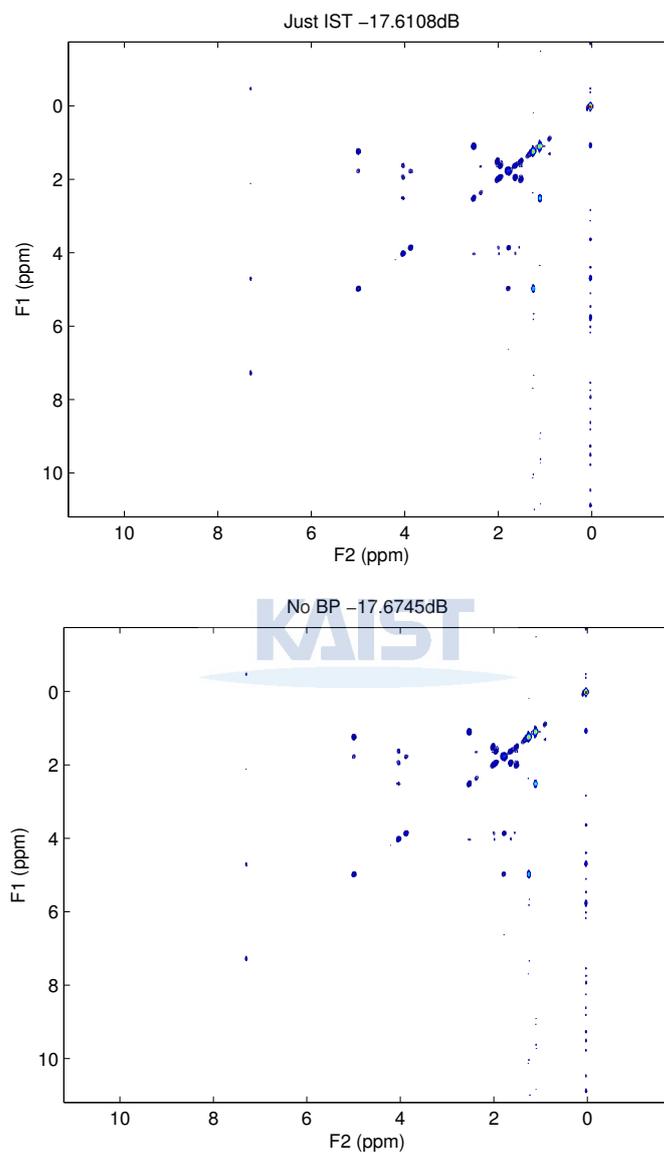


Figure 6.10: Reconstruction of Nonactin COSY spectrum based on iterative soft thresholding. The spectra are: (top) the reconstruction using just IST, (bottom) reconstruction with weighting and no belief propagation. All parameters are the same as those used in the hard thresholding case.

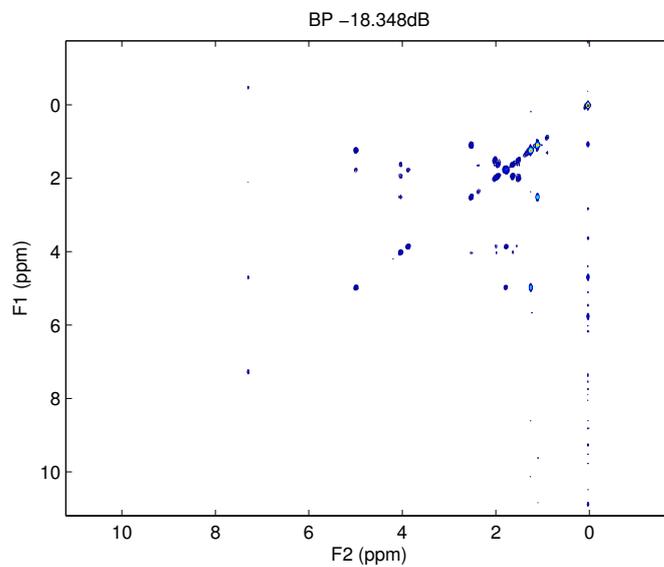


Figure 6.11: Reconstruction of Nonactin COSY spectrum using IST and belief propagation.

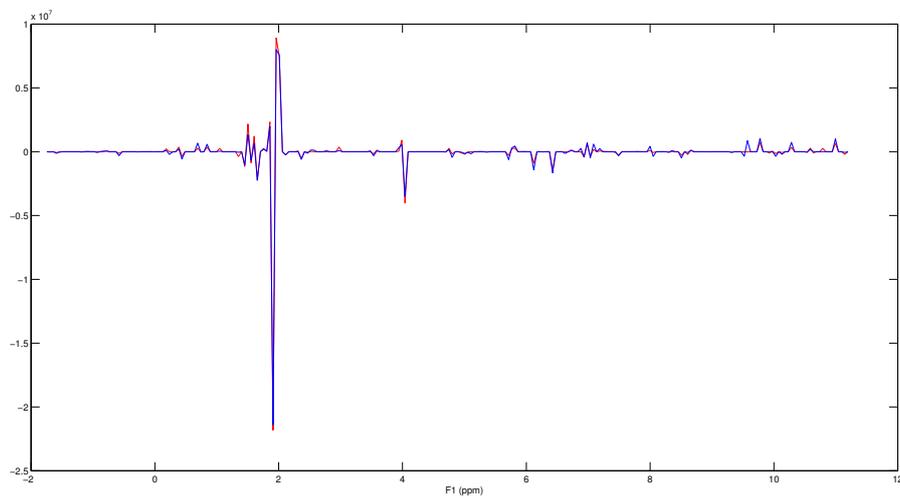


Figure 6.12: A one dimensional slice through the spectrum of nonactin at $\omega_2 = 1.62\text{ppm}$. The reconstruction using structure is shown in red, with that using just iterative thresholding in blue. Note the large intensity scale, and that, in general, the “sinc-wiggles” are smaller for the structure-based approach which yields larger peaks.

6.3 Lactose

In order to demonstrate the general applicability of the proposed algorithm, it was applied to the reconstruction of another downsampled spectrum. The spectrum used was from a COSY experiment for the disaccharide lactose. The data was from the *Biological Magnetic Resonance Data Bank* [49] (specifically [1]). Again, the spectrum was processed using *nmrPipe* [14], and *matNMR* [50], in addition to Matlab[®] scripts developed in the course of this work for processing NMR spectra. Sine apodization was performed, in addition to zero filling and correction for Bruker digital filtering (following the parameter *GRPDLY*). The spectrum for this data set can be seen in Figure 6.13.

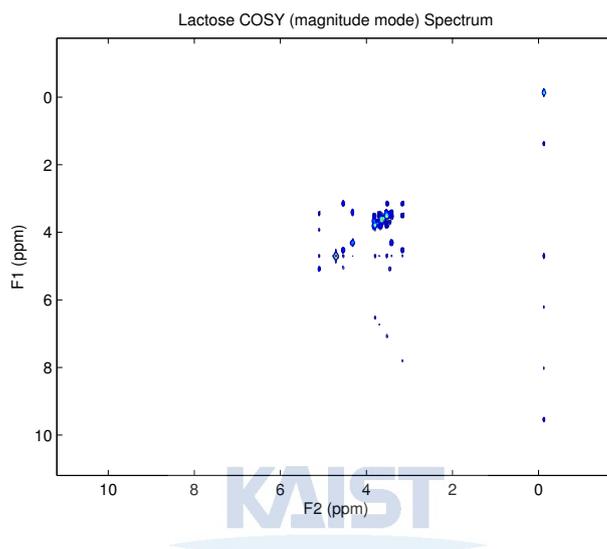


Figure 6.13: COSY spectrum for nonactin used in the simulations.

The results of the reconstruction are shown in Figures 6.14 and 6.15. they demonstrate that once again, visually the reconstruction with belief propagation is superior to that of simple IHT and weighted IHT without belief propagation. Interestingly though, the NMSE for reconstruction using just IHT is lower than that for IHT with belief propagation. Clearly the spectrum is better in the case were belief propagation has been incorporated in the reconstruction, so the validity of using the NMSE to measure reconstruction quality in these cases must be noted. Looking closely at the spectrum used for the simulations it was clear that, but setting the contour level appropriately, there were artefacts even in the original spectrum (Figure 6.16). These artefacts are reduced in the reconstruction using belief propagation and weighting. It appears that the NMSE, which compares the error between each point of the signal, and reconstruction reflects the fact that the reconstruction has improved on the original signal and reduced these “natural” artefacts. It is clear the weighted reconstruction has similar numbers of artefacts, if not slightly more than the reconstruction using belief propagation. The difference in NMSE in this case appears to come from the difference in the major peaks, which is the main improvement belief propagation brings about over just weighting the coefficients. Therefore, the benefits of the belief propagation step are two-fold. First the artefact level is reduced, secondly, the quality of the major peaks when compared with the simple weighting approach is enhanced.

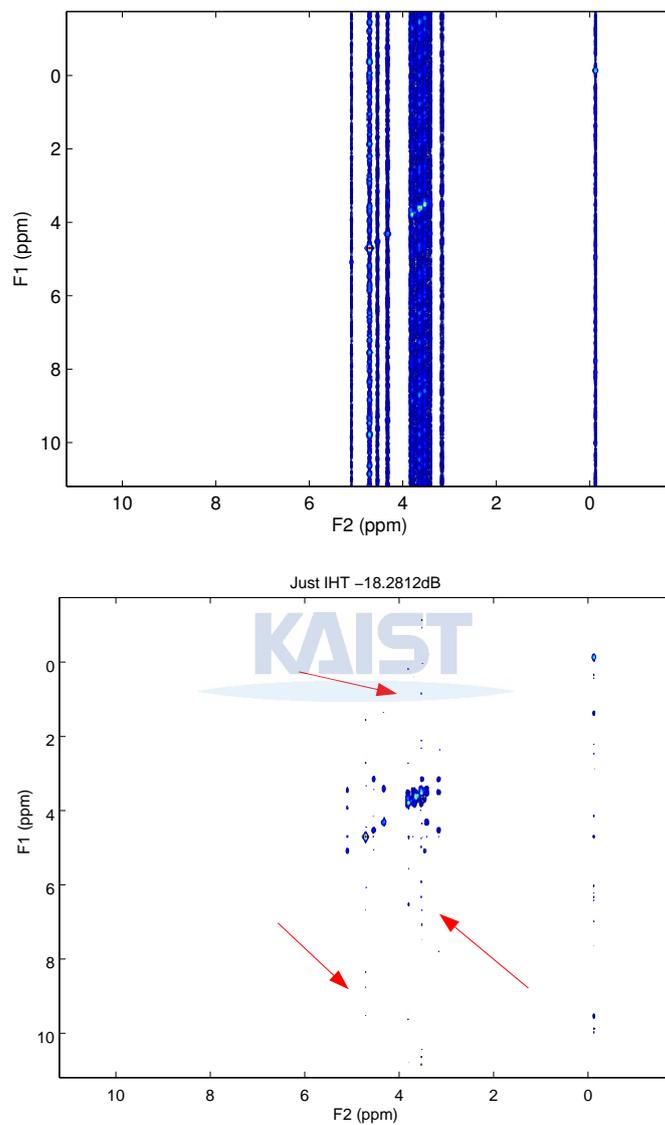


Figure 6.14: A selected reconstruction of a lactose COSY spectrum using 15% of all measurements with NMSE indicated in the title of each graph. The spectra are: (top) the direct Fourier transform, (bottom) the reconstruction using just iterative hard thresholding. Transition probabilities were: $P(d = 0|c = 0) = 0.1$ and $P(d = 1|c = 1) = 0.7$, and the threshold was 1.3×10^7 . Note the artefacts along the middle and right side in the vertical direction in the reconstruction.

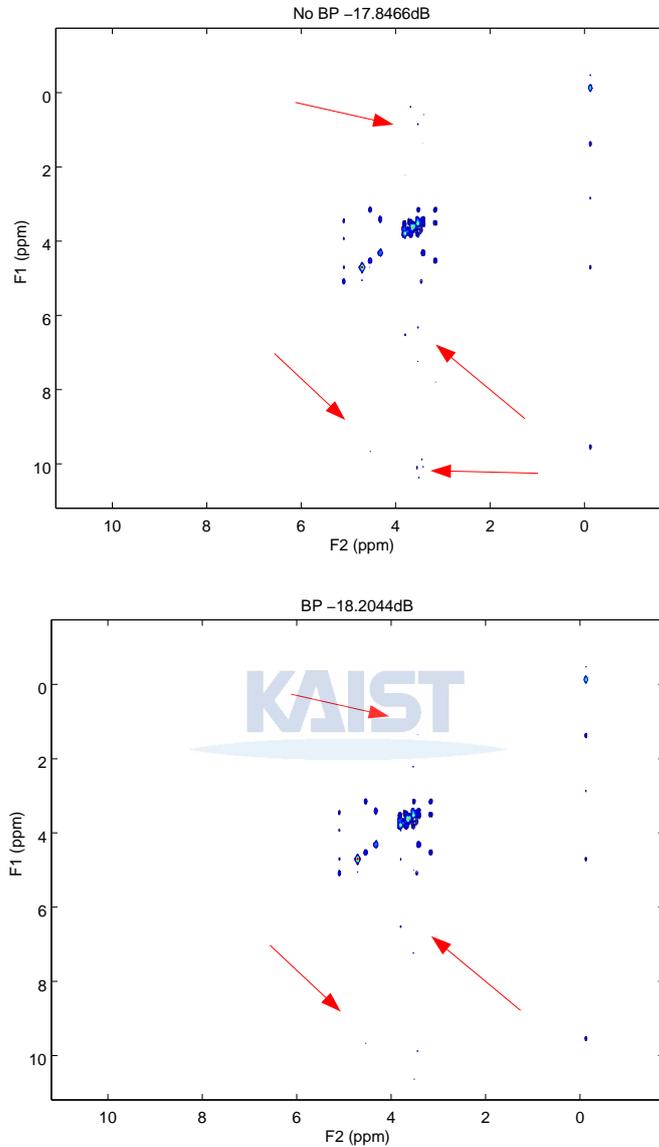


Figure 6.15: A selected reconstruction of a lactose COSY spectrum using 15% of all measurements with NMSE indicated in the title of each graph. The spectra are: (top) the reconstruction using probabilistic weighting with no belief propagation, (bottom) reconstruction with weighting and belief propagation in which diagonal coefficients are matched with cross peak coefficients in a two node chain where appropriate. Transition probabilities were: $P(d = 0|c = 0) = 0.1$ and $P(d = 1|c = 1) = 0.7$, and the threshold was 1.3×10^7 . Note the reduction in artefacts along the middle and right side in the vertical direction in the reconstruction.

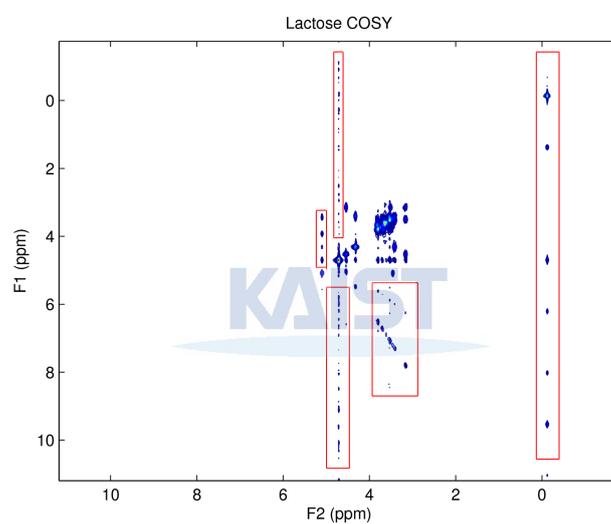


Figure 6.16: COSY spectrum for lactose used in the simulations with contour level set to show “natural” artefacts (highlighted).

Chapter 7. Discussion and Conclusions

Multidimensional NMR spectroscopy has now become established as a standard, and powerful tool for biomolecular structure determination. With this has come an increasingly strong push towards applying NMR spectroscopy to ever more complex problems – an application often demanding high resolution. The general method to increase resolution in the context of multidimensional NMR (more sampling, or more dimensions) is hindered by one of the largest bottlenecks in the structure determination process: acquisition time. The Fourier transform requires uniform sampling, which due to the large acquisition time associated with indirect dimensions, causes multidimensional experiments to be limited to around four-dimensions. There have been many attempts, both from a signal processing and experiment design perspective to decrease the amount of sampling required, and in this work an approach based on the relatively new framework of Compressed Sensing that exploits the NMR signal structure was presented.

The key advance here is that rather than simply applying the prior assumption of sparsity to the reconstruction of the NMR spectra, their inherent structure in the Fourier domain was also exploited. The example of choice was that of COSY. This experiment produces spectra in which there is a strong relationship between the geometric location within a spectrum of cross peaks and their corresponding diagonal peaks. Reconstruction using an iterative hard and thresholding to minimise the ℓ_1 -norm of data weighted corresponding to probabilities refined using belief propagation according to the structural relationships between cross and diagonal peaks was the technique used. This is similar to the more direct approaches used in the past for the application of Compressed Sensing to NMR spectroscopy, but differs in its inclusion of a step that imposes statistical structure. Two methods: a simple two-level parent to child network between diagonal and cross peaks, and two cross peaks matched to a child diagonal peak, were utilised and belief propagation used to refine the probabilities of each coefficient comprising the signal structure. Analysing the resulting reconstructed COSY spectra of nonactin and lactose reconstructed using previous techniques, and the proposed structured technique indicated that the incorporation of weighting, and probability refinement according to structure through belief propagation meant that the reconstruction error attained by the proposed algorithm as well as artefact level was much lower. This reflects positively on the techniques developed and the imposition of signal structure on the reconstruction of NMR spectra from COSY experiments. In addition to this it also suggests positive results for the same approach applied to other types of experiments such as TOCSY and NOESY, which exhibit similar structural relationships in the Fourier domain.

It is clear from experience in one-dimensional NMR that the major artefact introduced into spectra due to insufficient data is that of so-called “sinc-wiggles”. This can be modelled as multiplication in the time domain as an FID multiplied by a rectangular function. In the frequency domain this corresponds to convolution with a sinc function and thus introduces the characteristic wiggle-like pattern observed in the spectrum. This is normally overcome by apodization functions, which weight certain portions of the data more heavily than others, creating a more continuous FID. The result for downsampled spectra is similar in the indirect dimension as portions of the magnetisation’s time domain evolution along this dimension have been removed. This introduces artefacts along this dimension which have the potential to distort peaks and hide relationships, not to mention making determining the chemical shift more challenging and ambiguous. There is significant motivation to create signal processing algorithms

that overcome the uniform sampling requirement without compromising spectrum quality as it means reduced acquisition time and paves the way for high dimension NMR experiments.

The results obtained in this work are interesting in that they demonstrate that statistical weighting influenced by structure has the potential to decrease the artefact level. The weighting updates are between groups of coefficients comprising cross and diagonal peaks that can be assumed to have a statistical dependence. However, curiously, this reduced the artefacts level for areas of the spectrum not included in the model. It appears that combining statistical updates with weighting, and then minimisation means that the significant peaks that compose the useful information in the spectrum have their coefficients pushed to a higher reconstruction priority. This happens because their weighting (proportional to their probability) forces the solution to the minimisation to include these coefficients. This is because a sparse solution is being sort, and so roughly speaking there will be a limited number of large coefficients that can comprise the final solution. Since these coefficients have been prioritised in the sparse solution (by virtue of their weight), the large coefficients in the solution have been “intelligently” concentrated to where they would realistically be expected to reside. This means that there is little opportunity for coefficients that would usually comprise artefacts in any great amount since a sparse solution has been encouraged to contain peaks for which there is good evidence stemming from the signal structure over those for which there is little evidence. Essentially, increasing the probability of true peaks, decreases the chance of artefacts through the enforcement of a model and sparsity.

In addition to reducing the artefact level, it appears the belief propagation step acts to, in some respects, “appodize” the signal in the frequency domain. It was shown that for the COSY spectrum of lactose, the reconstruction using the proposed method resulted in a spectrum that was visually better than that using previous methods, but had a similar NMSE. The reason appears to be that the original spectrum from which the simulation data was obtained, contained natural artefacts. Using belief propagation reduced these artefacts, which visually improved the spectrum, while causing the spectrum to differ from the original in a way that resulted in the NMSE being higher. Furthermore, the NMSE of reconstruction using belief propagation, versus that with just weighting, shows similar artefact levels, but much better NMSE in the case of belief propagation. This suggests that the artefacts are reduced in both cases, and the peaks strengthened in the case of belief propagation. This, in form at least, is similar to an appodization operation (albeit in the frequency domain), since through weighting, it enhances the peaks, and reduces the artefact level given the original data.

An understandable objection to the methods used in this work is that the setting the threshold may be difficult when the difference between artifact and peak is ambiguous. While this can pose a problem for spectra of poor quality, the Compressed Sensing approach, through minimisation, is generally able to overcome this. It could be argued that falsely identifying a peak or artefact would mean that belief propagation would be performed on an incorrect model. This would challenge the artefact reducing properties of the proposed algorithm previously discussed. However, mathematically speaking, artefacts and peaks have very different origins. Iterative thresholding if used conservatively, with high, but steadily decreasing threshold, will in general suppress the noise as it is refining the signal estimation in such a way as to maximise the sparsity (which generally means less noise and artefacts). The probabilistic weighting done here is not an all or none decision, and so works within this minimisation framework (which for iterative thresholding is non-linear) to find a better sparse solution – essentially prioritising peaks that fit the assumed structure of the signal. Ultimately if there is evidence for a peak in the measurements then it is likely to be reconstructed, and although it is possible, it is unlikely that genuine peaks will be suppressed that would not have already been suppressed by the previously proposed algorithms for

Compressed Sensing in NMR.

In this work two algorithms were used for finding the approximate solution to the ℓ_1 -norm minimisation problem: IHT, and IST. It was generally observed that IST was superior in terms of its stability and response to the incorporation into the reconstruction of statistical structural information. Although IHT in some circumstances yielded a lower reconstruction error, it often was unable to yield the correctly minimised solution for cases of high downsampling. For this reason it is advisable that in the context of multidimensional NMR, when seeking high downsampling and stable reconstruction, the use of IST over IHT is recommended. This also produces reconstructions that are less sensitive to the parameters of the structure being modelled (in this case transition probabilities and thresholds).

The results are nice in that they show, when the threshold is set correctly, the transition probabilities follow a “common sense” form. When the threshold is set to eclipse artefacts, but include enough large coefficients to make belief propagation a useful refinement, the transition probabilities imply diagonal coefficients that have been matched with cross peak coefficients should be large. This means that there is low probability a low coefficient in a cross peak (which may have been incorrectly estimated) will also be low in the corresponding diagonal peak. In addition, it means that a coefficient that is large in a cross peak implies the corresponding diagonal peak coefficient is also large.

Practically speaking, the results imply that it is possible to reconstruct spectra of comparable quality to the fully sampled case, in much shorter acquisition time. For instance, in the case of the COSY nonactin spectrum the reconstructed spectrum using only 15% of the measurements, was of better quality than that reconstructed using previous techniques, in addition to being comparable to the original spectrum. The number of measurements in a multidimensional NMR experiment directly determines the experiment duration and for experiments that can take up to a week to complete, and 85% reduction in time has significant practical consequences.

Multidimensional NMR spectra, through their predictable structural relationships, are good targets for model-based reconstruction techniques. The methods here used relatively simple statistical relationships between the coefficients and still managed to demonstrate enhanced reconstruction results. The utilisation of more complex belief networks, such as a loopy network between two diagonal peaks, and their cross peaks, is expected to further increase reconstruction quality, and influence the number of measurements required to obtain a usable spectrum. Also extension to three, and more, dimensional experiments through techniques very similar to that use for two-dimensional experiments here would yield even better results, due to high-dimensional data sets often being even sparser. As a further application, the tendency of the proposed algorithm to bring out peaks and suppress artefacts holds potential for reconstructing NMR spectra with low SNR. A notable difference in the spectra reconstructed by the proposed algorithm is that they contain less artefacts, so it is expected that would suppress noise for experiments were sample concentration, or other factors have reduced resulted in a high noise level. This opens further avenues for compressed sensing, and model-based techniques to be applied to NMR spectroscopy, as measurements could be made for samples in low concentration. This isn't easily achievable for some species due to reactions they undergo when in high concentration. The proposed algorithm is the beginning of an area in reconstruction that previously hasn't been tackled in the design of reconstruction algorithms but holds promise for the reduction of acquisition time, and more.

The algorithm presented here exploits the statistical structure inherent in COSY and many other types of NMR spectra to enhance the reconstruction of downsampled two-dimensional NMR experiments. The results indicate that incorporating this structure is advantageous in terms of improving the reconstruction error level, and reducing the strength of artefacts. The technique developed is readily extended

to higher dimensions and other experiments types such as TOCSY and NOESY. Thus, through the decrease in measurements required, the proposed algorithm has the potential to significantly decrease the measurement time for multidimensional NMR experiments whose spectra exhibit predictable structure, as well as pave the way for further application that are ultimately beneficial for NMR, and the structural determination process as a whole.



References

- [1] Lactose – biological magnetic resonance data bank. http://www.bmrb.wisc.edu/metabolomics/mol_summary/index.php?molName=lactose&id=bmse000065&whichTab=0. [Accessed: November 20, 2011].
- [2] Nonactin – biological magnetic resonance data bank. http://www.bmrb.wisc.edu/metabolomics/mol_summary/?molName=nonactin. [Accessed: November 20, 2011].
- [3] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde. Model-based compressive sensing. *Information Theory, IEEE Transactions on*, 56(4):1982–2001, 2010.
- [4] M. Billeter and D. K. Staykova. *Rapid Multidimensional NMR: Decomposition Methods and their Applications*. John Wiley & Sons, Ltd, 2007.
- [5] T. Blumensath and M.E. Davies. Iterative thresholding for sparse approximations. *Journal of Fourier Analysis and Applications*, 14(5):629–654, 2008.
- [6] R. N. Bracewell. Strip integration in radio astronomy. *Aust. J. Phys.*, 9(2):198–217, 1956.
- [7] E. J. Candès and M. B. Wakin. An introduction to compressive sampling. *Signal Processing Magazine, IEEE*, 25(2):21–30, 2008.
- [8] E. J. Candès, M. B. Wakin, and S. P. Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14(5):877–905, 2008.
- [9] J. Cavanagh. *Protein NMR spectroscopy: principles and practice*. Academic Press, 1996.
- [10] B. E. Coggins and P. Zhou. PR-CALC: A program for the reconstruction of NMR spectra from projections. *Journal of biomolecular NMR*, 34(3):179–195, 2006.
- [11] B. E. Coggins and P. Zhou. <http://zhoulab.biochem.duke.edu/software/pr-calc/index.htm>, December 2009. [Accessed December 2, 2011].
- [12] I. Daubechies, M. Defrise, and C. De Mol. An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on pure and applied mathematics*, 57(11):1413–1457, 2004.
- [13] M. A. Davenport, M. F. Duarte, YC Eldar, and G. Kutyniok. Introduction to compressed sensing. *Preprint*, 93, 2011. To appear in *Compressed Sensing: Theory and Applications*, Cambridge University Press, 2012.
- [14] F. Delaglio, S. Grzesiek, G.W. Vuister, G. Zhu, J. Pfeifer, and A. D. Bax. NMRPipe: a multidimensional spectral processing system based on UNIX pipes. *Journal of biomolecular NMR*, 6(3):277–293, 1995.
- [15] D. L. Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, 2006.

- [16] D. L. Donoho, A. Maleki, and A. Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.
- [17] I. Drori. Fast l_1 minimization by iterative thresholding for multidimensional NMR spectroscopy. *EURASIP Journal on Advances in Signal Processing*, 1, 2007.
- [18] M. F. Duarte, M. B. Wakin, and R. G. Baraniuk. Wavelet-domain compressive signal reconstruction using a hidden markov tree model. In *Acoustics, Speech and Signal Processing, 2008. ICASSP 2008. IEEE International Conference on*, pages 5137–5140. IEEE, 2008.
- [19] R. Freeman and E. Kupče. New methods for fast multidimensional NMR. *Journal of biomolecular NMR*, 27(2):101–114, 2003.
- [20] R. Freeman and E. Kupče. Distant echoes of the accordion: Reduced dimensionality, GFT-NMR, and projection-reconstruction of multidimensional spectra. *Concepts in Magnetic Resonance Part A*, 23(2):63–75, 2004.
- [21] L. He and L. Carin. Exploiting structure in wavelet-based bayesian compressive sensing. *Signal Processing, IEEE Transactions on*, 57(9):3488–3497, 2009.
- [22] L. He, H. Chen, and L. Carin. Tree-structured compressive sensing with variational bayesian analysis. *Signal Processing Letters, IEEE*, 17(3):233–236, 2010.
- [23] J. C. Hoch and A. S. Stern. *NMR data processing*. Wiley-Liss New York, 1996.
- [24] D. J. Holland, M. J. Bostock, L. F. Gladden, and D. Nietlispach. Fast multidimensional nmr spectroscopy using compressed sensing. *Angewandte Chemie International Edition*, 50(29):6548–6551, 2011.
- [25] F. V. Jensen. *Bayesian networks and decision graphs*. Statistics for engineering and information science. Springer, 2001.
- [26] S. M. Kay. *Fundamentals of Statistical Signal Processing: Detection theory*. Prentice Hall signal processing series. Prentice-Hall PTR, 1998.
- [27] K. Kazimierczuk and V. Y. Orekhov. Accelerated NMR spectroscopy by using compressed sensing. *Angewandte Chemie International Edition*, 50(24):5556–5559, 2011.
- [28] J. Keeler. *Understanding NMR spectroscopy*. John Wiley & Sons, 2011.
- [29] S. Kim and T. Szyperski. GFT-NMR, a new approach to rapidly obtain precise high-dimensional NMR spectral information. *Journal of the American Chemical Society*, 125(5):1385–1393, 2003.
- [30] F. R. Kschischang, B. J. Frey, and H. A. Loeliger. Factor graphs and the sum-product algorithm. *Information Theory, IEEE Transactions on*, 47(2):498–519, 2001.
- [31] E. Kupče and R. Freeman. Projection-reconstruction technique for speeding up multidimensional NMR spectroscopy. *Journal of the American Chemical Society*, 126(20):6429–6440, 2004.
- [32] E. Kupče and R. Freeman. The radon transform: a new scheme for fast multidimensional NMR. *Concepts in Magnetic Resonance Part A*, 22(1):4–11, 2004.
- [33] M. H. Levitt. *Spin dynamics*. John Wiley & Sons : Chichester, UK, 2001.

- [34] R. J. McEliece, D. J. C. MacKay, and J. F. Cheng. Turbo decoding as an instance of Pearl’s “belief propagation” algorithm. *Selected Areas in Communications, IEEE Journal on*, 16(2):140–152, 1998.
- [35] G. A. Morris and J. W. Emsley. *Multidimensional NMR: an Introduction*. John Wiley & Sons, Ltd, 2007.
- [36] D. Moskau. Application of real time digital filters in NMR spectroscopy. *Concepts in Magnetic Resonance*, 15(2):164–176, 2002.
- [37] K. Murphy, Y. Weiss, and M. I. Jordan. Loopy belief propagation for approximate inference: An empirical study. In *Proceedings of Uncertainty in AI*, volume 467475, 1999.
- [38] D. Needell and J. A. Tropp. Cosamp: Iterative signal recovery from incomplete and inaccurate samples. *Applied and Computational Harmonic Analysis*, 26(3):301–321, 2009.
- [39] V. Y. Orekhov, I. V. Ibraghimov, and M. Billeter. Munin: A new approach to multi-dimensional NMR spectra interpretation. *Journal of Biomolecular NMR*, 20(1):49–60, 2001.
- [40] J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 1988.
- [41] R. M. Pyndiah. Near-optimum decoding of product codes: Block turbo codes. *Communications, IEEE Transactions on*, 46(8):1003–1010, 1998.
- [42] X. Qu, D. Guo, X. Cao, S. Cai, and Z. Chen. Reconstruction of self-sparse 2D NMR spectra from undersampled data in the indirect dimension. *Sensors*, 11(9):8888–8909, 2011.
- [43] C. D. Ridge and V. A. Mandelshtam. On projection-reconstruction NMR. *J. Biomol. NMR*, 43:151–159, 2009.
- [44] P. Schniter. Turbo reconstruction of structured sparse signals. In *Information Sciences and Systems (CISS), 2010 44th Annual Conference on*, pages 1–6. IEEE, 2010.
- [45] S. Som, L. C. Potter, and P. Schniter. Compressive imaging using approximate message passing and a markov-tree prior. In *Signals, Systems and Computers (ASILOMAR), 2010 Conference Record of the Forty Fourth Asilomar Conference on*, pages 243–247. IEEE, 2010.
- [46] A. S. Stern, D. L. Donoho, and J. C. Hoch. NMR data processing using iterative thresholding and minimum l1-norm reconstruction. *Journal of Magnetic Resonance*, 188(2):295–300, 2007.
- [47] T. Szyperski and H. S. Atreya. Principles and applications of GFT projection NMR spectroscopy. *Magnetic Resonance in Chemistry*, 44(S1):S51–S60, 2006.
- [48] D. S. Taubman and M. W. Marcellin. *JPEG2000: image compression fundamentals, standards, and practice*. Number V. 1 in Kluwer international series in engineering and computer science. Kluwer Academic Publishers, 2002.
- [49] E. L. Ulrich, H. Akutsu, J. F. Doreleijers, Y. Harano, Y. E. Ioannidis, J. Lin, M. Livny, S. Mading, D. Maziuk, Z. Miller, E. Nakatani, C. F. Schulte, D. E. Tolmie, R. K. Wenger, H. Yao, and J. L. Markley. Biomagresbank. *Nucleic Acids Research*, 36(suppl 1):D402–D408, 2008.
- [50] J. D. van Beek. matNMR: A flexible toolbox for processing, analyzing and visualizing magnetic resonance data in matlab®. *Journal of Magnetic Resonance*, 187(1):19–26, 2007.

Summary

Exploitation of Structure for Enhanced Reconstruction of Multidimensional NMR Spectra

Multidimensional NMR의 주요 문제점의 하나는 데이터의 획득 시간이며 이미 Compressive Sampling을 이용하여 데이터 획득 시간을 줄인 NMR application이 존재한다. 본 연구의 동기는 데이터가 sparse하다는 가정에 데이터의 prior knowledge를 안다는 가정을 덧붙인다. 이는 NMR 실험에서 cross peaks와 diagonal peaks 간의 구조적인 관계를 연구하는데 도움이 된다.

Correlated Spectroscopy (COSY) 는 multidimensional NMR 실험에서 구조를 결정하기 위해 빈번하게 사용되는 방법이다. NMR 실험들은 diagonal peak cross peak 구조를 갖는다. 예를 들면, cross peak은 두 개의 diagonal peak 사이에만 존재한다. 그렇지 않다면 magnetization transfer 가 없게 된다. 따라서 diagonal peak는 cross peak의 존재 가능한 위치를 정의하게 된다. 만약 cross peak가 diagonal peaks 사이에 존재하지 않더라도, 그것이 존재할 위치는 한정되어 있으며 그 위치 역시 diagonal peaks에 의해 결정이 된다. 만약 cross peak가 존재한다면 그에 대응하는 두 개의 diagonal peaks의 위치를 결정할 수 있다. 또한 COSY spectra는 diagonal 에 대해 대칭이다. 이는 cross peaks에 대한 정보가 두 개의 diagonal peaks의 위치를 결정할 뿐만 아니라 opposing cross peak의 위치까지 결정한다는 것을 의미한다. 이러한 구조 연구에 Compressive Sampling을 이용하여 NMR spectra의 복원을 강화한 것이 본 연구의 의미라고 하겠다.

KAIST