Locating Carbon Bonds from INADEQUATE Spectra using Continuous Optimization Methods and Non-Uniform K-Space Sampling

Locating Carbon Bonds from INADEQUATE Spectra using Continuous Optimization Methods and Non-Uniform K-Space Sampling

By
Sean Colin Watson B.Eng

A Thesis
Submitted to the School of Graduate Studies
in Partial Fulfillment of the Requirements
for the Degree
Master of Applied Science

McMaster University

MASTER OF APPLIED SCIENCE(2011) McMaster University
COMPUTING AND SOFTWARE Hamilton, Ontario

TITLE: Locating Carbon Bonds from INADEQUATE Spectra using Continuous Optimization Methods and Non-Uniform K-Space Sampling

AUTHOR: Sean Colin Watson B.Eng(McMaster University)

SUPERVISOR: Dr. Christopher Anand

NUMBER OF PAGES: xi, 34

Abstract

The 2-D INADEQUATE experiment is a useful experiment for determining carbon structures of organic molecules known for having low signal-to-noise ratios. A non-linear optimization method for solving low-signal spectra resulting from this experiment is introduced to compensate. The method relies on the peak locations defined by the INADEQUATE experiment to create boxes around these areas and measure the signal in each. By measuring pairs of these boxes and applying penalty functions that represent a priori information, we are able to quickly and reliably solve spectra with an acquisition time under a quarter of that required by traditional methods. Examples are shown using the spectrum of sucrose. The concept of a non-uniform Fourier transform and its potential advantages are introduced. The possible application of this type of transform to the INADEQUATE experiment and the previously explained optimization program is detailed.

Acknowledgments

I would like to thank my supervisor, Dr. Christopher Anand, for his help and support and for never being too busy to answer any questions. I would also like to thank Dr. Alex Bain for his assistance with all things chemistry that appear in this thesis.

I would like to thank my fellow students, in particular Zhenghua Nie and Kevin Browne for their experience and advice as well as Qiong Wu for her unwaveringly positive attitude that helped me get through some of the most frustrating issues encountered in writing this thesis.

I would also like to thank my parents, family and friends for their support and encouragement.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Document Structure

We begin with an introduction to the study of Nuclear Magnetic Resonance by giving a basic summary of its effects, governing equations and some of its most common uses. We then introduce the INADEQUATE experiment, which is the focus of this thesis in order to describe the use of the experiment and give a basic idea of what can be done to improve its results. The functions used in the optimization, including the penalties and constraints, are described and an example of this method is given. The idea of non-uniform k-space sampling is introduced along with the non-uniform Fourier transform and proof-of-concept plots are given. Conclusions are then drawn and required future work is discussed.

Chapter 2

Background

2.1 Principles of NMR

Nuclear Magnetic Resonance (NMR) is a technique used to measure the aggregate magnetic field generated by the spins of one or more nuclei in a sample. This static magnetic field generates magnetization in the sample and aligns the nuclear spins in the direction of this magnetic field. This magnetization is then perturbed by a radio frequency (RF) pulse with a certain intensity and duration (pulse width). The pulse width and intensity of the pulse are calculated in order to perturb the spins by a specific angle as required for the current experiment. The recovery of the spins from this perturbation is what generates the signal measured by the spectrometer.

2.1.1 Bloch Equations

The most basic model of the spin dynamics encountered is that the moment, M, precesses in the magnetic field, B, according to

$$\frac{dM}{dt} = \gamma M \times B \tag{2.1}$$

where γ is the gyromagnetic ratio and is a physical constant. It can be noted that this same equation governs the motion of a gyroscope, giving an idea of the type of motion being modelled. This can model the dynamics of a single spin for very short time frames (a few microseconds), but in any longer time frames quantum interactions cause the net magnetization to return to thermodynamic equilibrium, $M_0 = (0, 0, \delta)^T$ where δ is the proton density. As a result we need to add a relaxation term to Eq. (2.1), by which we arrive at

the Bloch equation:

$$\frac{dM}{dt} = \gamma M \times B - R(M - M_0) \tag{2.2}$$

Where R is the relaxation matrix. R can be further broken down into two relaxation rates, $\frac{1}{T_2}$ for relaxation in the xy plane and $\frac{1}{T_1}$ for relaxation on the z-axis. $\frac{1}{T_1}$ and $\frac{1}{T_2}$, also known as spin-lattice relaxation and spin-spin relaxation, respectively, vary based on the nucleus. This equation is typically split into the x, y and z components as follows:

$$\frac{d}{dt}\begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} = \gamma \begin{bmatrix} 0 & B_0 & 0 \\ -B_0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} - \begin{bmatrix} \frac{1}{T_2} & 0 & 0 \\ 0 & \frac{1}{T_2} & 0 \\ 0 & 0 & \frac{1}{T_1} \end{bmatrix} \begin{bmatrix} M_x \\ M_y \\ M_z - M_0 \end{bmatrix}, \tag{2.3}$$

where B_0 is the static field along the z-axis. This equation is simplified by transformation to a rotating frame of reference such that the xy plane rotates around the z-axis at the Larmor frequency, $\omega_0 = \gamma B_0$. This equation then becomes:

$$\frac{d}{dt}\begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} = \begin{bmatrix} 0 & \omega_0 & 0 \\ -\omega_0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} - \begin{bmatrix} \frac{1}{T_2} & 0 & 0 \\ 0 & \frac{1}{T_2} & 0 \\ 0 & 0 & \frac{1}{T_1} \end{bmatrix} \begin{bmatrix} M_x \\ M_y \\ M_z - M_0 \end{bmatrix} \tag{2.4}$$

These equations are often re-written so that there is a single entry for the xy-plane, $M_{xy} = M_x + iM_y$ for simplicity:

$$\frac{d}{dt}\begin{bmatrix} M_{xy} \\ M_z \end{bmatrix} = \begin{bmatrix} i\omega_0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} M_{xy} \\ M_z \end{bmatrix} - \begin{bmatrix} \frac{1}{T_2} & 0 \\ 0 & \frac{1}{T_1} \end{bmatrix} \begin{bmatrix} M_{xy} \\ M_z - M_0 \end{bmatrix} \tag{2.5}$$

The solution of this differential equation models the appearance of the time domain output from any NMR experiments. This time domain data is made up of two exponentially decaying sine waves representing M_x and M_y as well as an exponential curve tending to the value of M_0 representing M_z as in Fig. 2.1, showing simulated data. The data in Fig. 2.1 only represents the signal from a single nucleus, but the data for molecules is not much more complex and is simply the sum of many of these single nucleus signals; an example is shown in Fig. 2.2. We notice that there is no line for the M_z component in the second example. This is because it gives us fairly little information compared to the

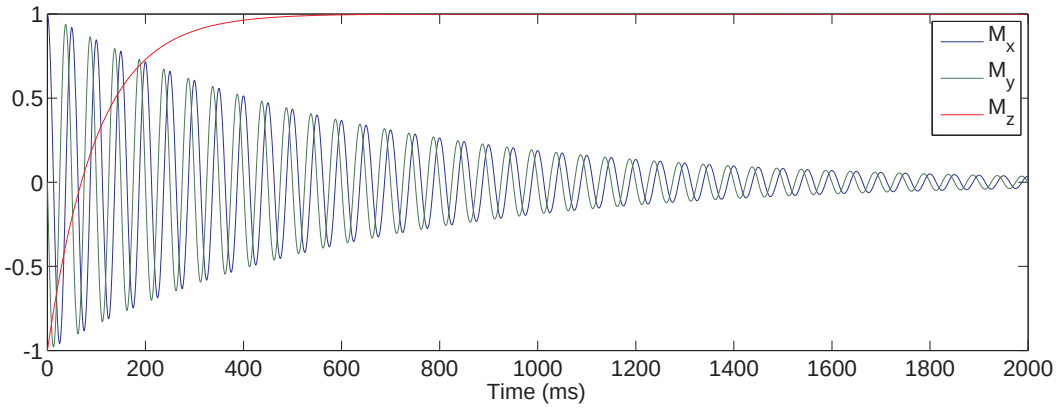

Figure 2.1: Sample simulated time domain data for a single nucleus.

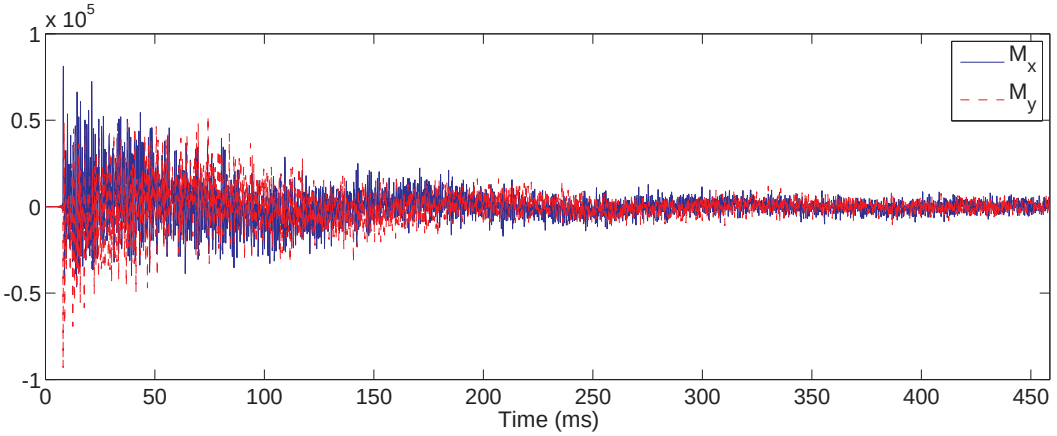

Figure 2.2: Time domain data after a single pulse, using a sample of sucrose.

M_x and M_y components as well as being more difficult to detect, meaning that in practice it is typically excluded.

The frequency domain representation of this data is where most of the data is recovered from NMR experiments. An example of this is shown in Fig. 2.3 and is the elementwise magnitude of the Fourier Transform of the M_{xy} data shown in Fig. 2.2. From this data we can extract the different chemical

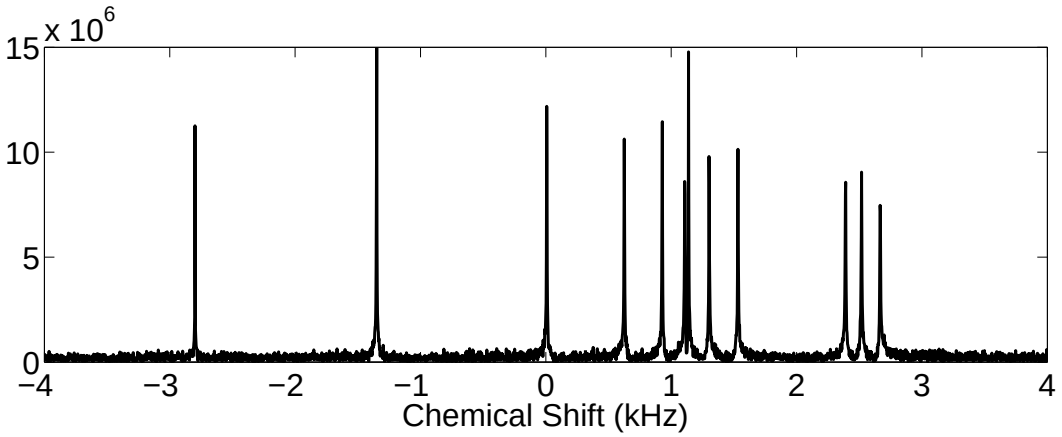

Figure 2.3: An NMR spectrum showing the twelve carbon peaks of sucrose from which we determine their chemical shifts in Hz.

shifts of every nucleus in the sample molecule.

2.1.2 Single-Dimensional Pulse Experiments

The simplest single-dimensional NMR spectrum, and the one we concern ourselves with in this thesis, displays a series of peaks representing the chemical shifts of the nuclei in the sample in Hertz. The nucleus appearing in the spectrum is specifically chosen for an experiment. Any nucleus with non-zero spin can be chosen, but most commonly either ^{1}H or ^{13}C are used [SH87a].

The experiment used to show the chemical shifts in a molecule is a single-pulse experiment, meaning we only use one pulse of RF energy before we collect the signal [SH87b]. More complex experiments are generated with multiple pulses, generally made up of a series of 90° and 180° pulses spaced apart with specific delay times. A one-dimensional experiment will have one variable delay, while multi-dimensional experiments have several. Two of the most common one-dimensional experiments are the inversion-recovery and spin-echo experiments for finding spin-lattice and spin-spin relaxation times, respectively [SH87c]. Important to note is that the experiments given so far only identify or determine information for a single type of nucleus. The interactions between nuclei require more information than can be provided with these experiments and so an extra dimension is added [ABE76, BF81]. There are many types of two-dimensional experiments, and the methods of each ex-

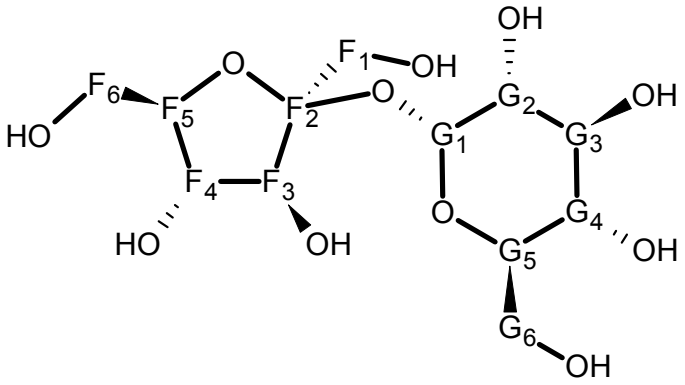

Figure 2.4: The structure of sucrose, showing the six-membered glucose ring(G) and five-membered fructose ring(F) with numbered carbons.

periment may differ, though the purpose is fairly similar: to determine the structure of a molecule. These experiments include correlation spectroscopy (COSY), nuclear Overhauser effect spectroscopy (NOESY) and the incredible natural abundance double quantum transfer experiment (INADEQUATE) which is the experiment we concern ourselves with in this thesis and describe in detail in the following section.

2.2 INADEQUATE Experiment

INADEQUATE [BF80, BFK80a, BFK80b, BFFL81, BM83] is a specific two-dimensional experiment for finding the complete carbon skeleton of an organic molecule [BFF81, FFR82, NL82, LKB89, ZSK+09]. This is done by finding all the bonded pairs of carbon atoms and tracing out this set of pairs. The reason this experiment was chosen for our optimization is because these pairs can be found and traced using very little information. We demonstrate this using the example of sucrose (the structure of which is shown in Fig. 2.4).

Fig. 2.5 shows the same one-dimensional experiment as Fig. 2.3, with labels to identify the twelve carbon atoms of sucrose by their chemical shifts. We note that this experiment is relatively quick and low-noise and that it results from tests to calibrate the spectrometer pulse width before carrying out more complex experiments. The chemical shifts shown in this spectrum are used to identify the atoms later in the two-dimensional INADEQUATE

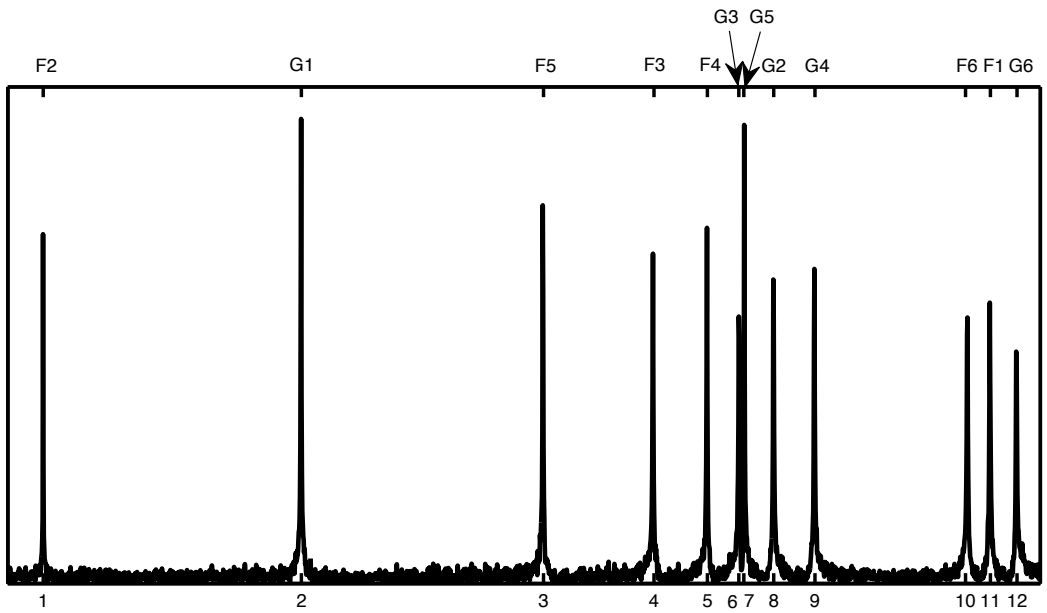

Figure 2.5: The numbering of the sucrose peaks with the structural numbering on the top and the numbering by chemical shift as used in our method on the bottom.

spectrum.

An example of the INADEQUATE spectrum of sucrose is shown in Fig. 2.6. What's important to notice in this spectrum is the way the various peak pairs line up vertically or horizontally. The locations at which these peaks can appear along the F_2(horizontal) and F_1(vertical) axes are predicted by the chemical shifts found in the one-dimensional experiment shown by Fig. 2.3. Along the F_2 axis, the peak pairs can only appear centred at locations that line up with one of the chemical shifts of a carbon; in the example of sucrose, there are exactly 12 possible locations on F_2 where we can see these peaks centred. The frequencies on this axis are referred to as the single-quantum frequencies, as they refer to the signal generated from each single carbon nucleus. Along F_1 the peaks represent the links between these carbons and occur at the sum of chemical shifts. These frequencies are known as double-quantum frequencies and are signals generated from bonded pairs. In general, the single-quantum signal will dwarf the double-quantum signal, however INADEQUATE uses a pulse program designed specifically to eliminate the single-quantum signals and

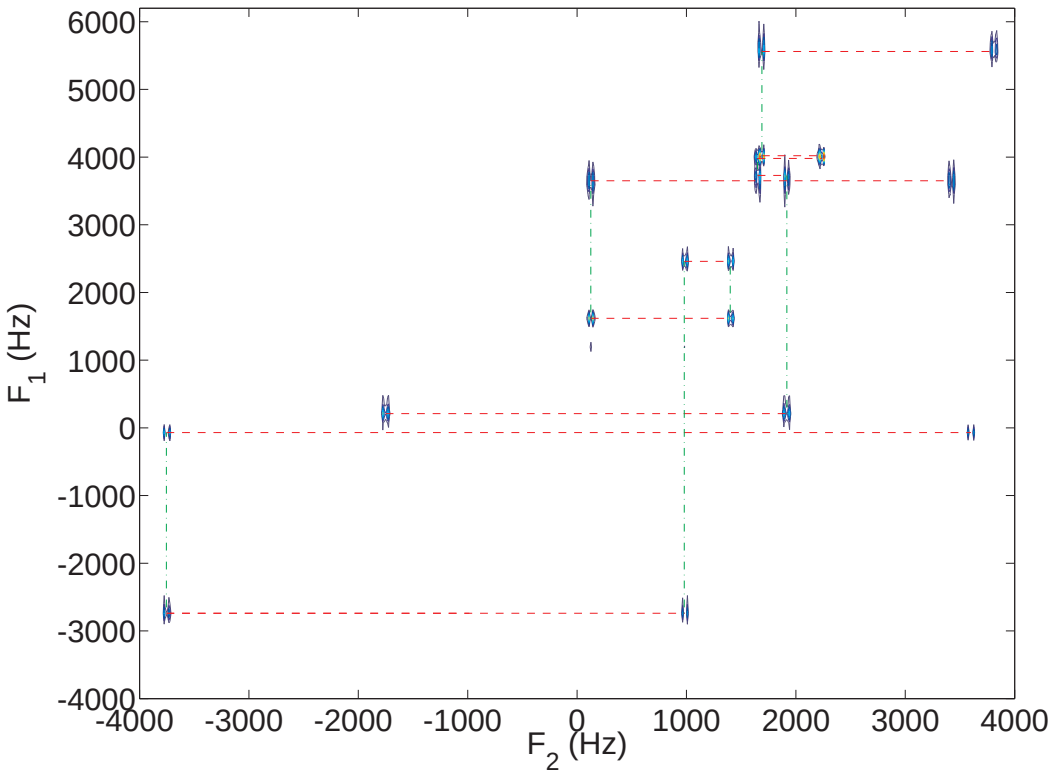

Figure 2.6: An INADEQUATE spectrum showing the carbon bonds of sucrose and how we trace them. The red lines indicate bonds between different nuclei while the green lines connect multiple peaks referring to the same nucleus.

enhance the double-quantum [BF80]. In this example we have a bond between a carbon with a chemical shift of $\omega_1 = -2800$Hz and one with a chemical shift of $\omega_2 = 630$Hz which is represented by two peaks at $(-2800, -2170)$ and $(630, -2170)$.

We can see this being applied in Fig. 2.6, where the bonds are traced. Bonds are shown connecting, with red lines, from left to right and top to bottom: F2 with F1; F2 with F3; G1 with G2; F5 with F6; F5 with F4; F3 with F4; G3 with G4; G3 with G2; G5 with G4 and G5 with G6. While the green lines indicate that the two peaks shown correspond to the same carbon, and so there are multiple bonds to that nucleus. This gives the two carbon rings as shown in Fig. 2.4, which gives the basic molecular structure of the of the sample. Due to the fairly low resolution of the spectrum shown here, it is

difficult to see some of the lines drawn in the upper right

The issue with this experiment is that it is inherently insensitive, requiring either extremely long acquisition times, or very high concentrations of sample [BB87, OWDM88, BLLO99, SFF$^+$82, KV86, Pod90]. The spectrum given in Fig. 2.6 is unusually clear and represents over six days of acquisition time. Usable, but less clear, spectra for a fairly simple molecule such as this one can take around two days of acquisition time. Even worse, these times are taking into account the ease of obtaining sucrose and the ease of creating concentrated solutions for testing. Most molecules of interest will be at extremely low concentrations, meaning that an extremely lengthy acquisition time is required to compensate for the lack of signal. This insensitivity is caused by the reliance of the experiment on ^{13}C, which accounts for $\sim 1.1\%$ of carbon molecules at natural abundance. This figure is then made even worse by the fact that the experiment looks for bonded pairs of ^{13}C, which occur at a rate of ~ 1 in 10000 pairs which is why the times given for collecting these spectra are so high. As a result, there is a lot of work that has gone into computer-assisted methods of reading these spectra in which less data is required than for a human operator to trace the bonds by eye. Many of the existing techniques developed for INADEQUATE spectra rely on complex line-shape fitting and large numbers of assumptions on the form of the experimental results [SE85, KULB92, LB93, DMC$^+$90, DMPG92, RAW81, FMD$^+$92, ODH$^+$95]. Our technique is based on the magnitude of the signal, making this technique insensitive to phase errors in the experiments and uses comparably few assumptions about the form of the experiment in order to simplify our model.

The following assumptions that we are able to make are strong, and based on the well-defined nature of the experiment's resulting spectra. Before the beginning of an INADEQUATE experiment, we know the single quantum frequencies (hence, all possible double-quantum frequencies); it is their assignment that is unclear. The correlations that we see in the INADEQUATE plots all appear as simple two-spin AB correlation spectra. Since both parts of a correlation come from the same double-quantum coherence, the two parts of the correlation should be mirror images [NSF96], although pathological offset effects [BHA$^+$10] may slightly distort the symmetry. These key facts allow us to build penalty functions to guide the optimization toward the global minimum. The process makes use of image regularization techniques to smooth the lineshapes as the optimization progresses. The combination of these allows us to find a complete set of carbon bonds from an INADEQUATE spectrum too noisy for traditional methods to be applied, significantly reducing the required experiment time. While the examples shown in this thesis use unprocessed

data, meaning the Fourier transforms are run on raw data in MATLAB, this method will work equally well with pre-processed data as the processing will not change any of the assumptions made.

2.3 Process Goal

It has been shown in Fig. 2.6 how an ideal spectrum may look and how this can be traced out by hand. However, as stated in Section 2.2 this spectrum takes approximately 6 days to acquire. A spectrum acquired in slightly over 48 hours is shown in Fig. 2.7 and while it is possible for chemists to visually trace the bonds in this spectrum, the difference in quality between these examples is vast. We intend to introduce a method where not only will the carbon bonds be traced by computer, but there will be a modified spectrum produced with clarity on the order of that in Fig. 2.6 using a quarter of the information present in Fig. 2.7.

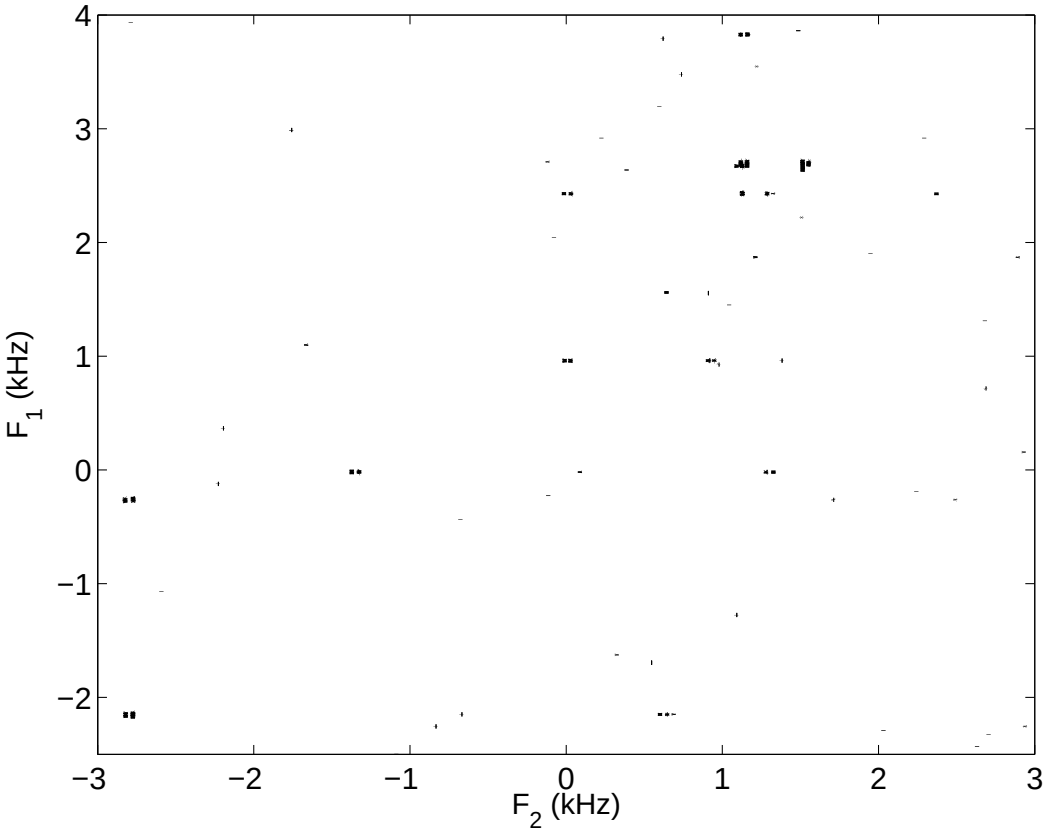

Figure 2.7: An INADEQUATE spectrum obtained over a 48 hour period, note the difficulty of choosing valid peaks when compared to Fig. 2.6.

Chapter 3

Optimization Model

Our approach is based, primarily, on the signal intensity present in a set of pre-defined locations based on the ^{13}C chemical shifts. In the INADEQUATE experiment, these locations are well-defined by the 1-D carbon spectrum. With a list of chemical shifts for all the carbons in the molecule we are testing, we are able to locate all possible peak locations and know which locations of peaks could indicate what bond. We draw boxes around all of these locations, large enough to contain the full peak and account for small shifts of location caused by the coupling constants and isotope effects (testing shows 32 points along the F_2 axis to be sufficient for an experiment using 4k points). In order to provide more intelligent estimates to compensate for high noise levels, we provide several penalties based on a priori knowledge of the experiment and basic chemistry. We penalize large differences between signal level in paired boxes as the INADEQUATE experiment will typically have equal signal in each of the two correlated boxes, and we force the number of bonds to any single carbon to be between one and four.

The method presented allows and requires some user input. We require an estimate of the level of connectivity in the molecule, as this controls the number of bonds we reject or accept as the algorithm runs. We allow an initial guess of the bonds that will be present. In small molecules this is not required, but in large molecules any prior knowledge will greatly improve the run-time of the method. The number of potential locations in a spectrum is $n(n-1)$, where n is the number of carbons in the molecule as each carbon could potentially be connected to any of the $n-1$ others.

3.1 Variables

We use two sets of variables that are intended to meet the goals of this optimization:

p: to determine the bonds present in a molecule, given a 2-D INADEQUATE spectrum; and

$\mathcal{S}$: to create a cleaner image of the given spectrum, clearly showing the locations of the peaks and bonds as determined.

The column vector, p, contains one entry for each potential bond (a total of n choose 2 entries, where n is the number of carbons present). Each entry contains a value ranging from 0 to 1 that represents the probability of its associated bond being present in the molecule. The entry of p relating to a specific bond, between carbons i and j, is given by the notation p_{ij}. This notation does not indicate the index of an entry, but rather indicates the contents by noting the two carbons being referenced. We use this non-standard convention in order to simplify the written problem. The entries in this vector are ordered with respect to the bonds they represent. For each carbon i, there is an ordered list of bonds from $i \leftrightarrow i+1$ to $i \leftrightarrow n$ (where $\leftrightarrow$ means "bonded to") where i ranges between 1 and $n-1$. The carbons themselves are numbered by their chemical shifts in increasing order, and an example of this is shown by Fig. 2.5.

The second set of variables, $\mathcal{S}$, is a sparse array with dimensions equal to those of the provided spectrum. $\mathcal{S}$ has the important structure of being zero-valued outside of the pre-determined boxes that may contain peaks (determined by the chemical shifts of the potentially bonded carbons), meaning that we eliminate noise from all regions that we know cannot contain a correlation. Each of these boxes is made up of a series of entries, referred to as s_k. This structure minimizes the number of variables we need to solve in order to create a new image of the spectrum without a loss of data. A small scale example of the structure as described is shown in Eq. (3.1). In order to simplify several later equations, we use the notation $\mathcal{S}_{ij}$ to refer to a sub-array of $\mathcal{S}$ corresponding to a box centred at the location $(\omega_i, \omega_i + \omega_j)$, where we could see one of the peaks indicating a bond of i to j. We show all the data for a potential bond between i and j with the two boxes $\mathcal{S}_{ij}$ and $\mathcal{S}_{ji}$.

$$\begin{bmatrix} 0 & s_1 & s_2 & s_3 & 0 & 0 & 0 & s_4 & s_5 & s_6 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ s_7 & s_8 & s_9 & 0 & 0 & s_{10} & s_{11} & s_{12} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \tag{3.1}$$

In the example in Eq. (3.1), the variables s_1, s_2, s_3 correspond to one box (if this represents a bond between carbons 1 and 2 we refer to it by $\mathcal{S}_{12}$). In a real spectrum we use wider boxes with many more variables per box.

3.2 Model

Our optimization problem minimizes the sum of the terms given in (3.2)-(3.7) and is constrained by (3.8) and (3.9). The solution of this minimization problem gives us both desired results: a low-noise image of the spectrum and a list of carbon bonds present in the molecule.

$$\min \|m - \mathcal{S}\|^2 \tag{3.2}$$
$$+ \lambda_1 ||\delta_x \mathcal{S}||^2 \tag{3.3}$$
$$+ \lambda_2 \sum_{ij} (1 - p_{ij})^2 ||\mathcal{S}_{ij} + \mathcal{S}_{ji}||^2 \tag{3.4}$$
$$+ \lambda_3 \sum_{ij} (||\mathcal{S}_{ij}||^2 - ||\mathcal{S}_{ji}||^2)^2 \tag{3.5}$$
$$+ \mu_1 \sum_{i} \left(2 - \sum_{j \neq i} p_{ij}\right)^4 \tag{3.6}$$
$$+ \mu_2 \sum_{ij} p_{ij} \tag{3.7}$$
$$\text{s.t. } p_{ij} \geq 0 \tag{3.8}$$
$$p_{ij} \leq 1 \tag{3.9}$$

Following is a quick summary of each term and an idea of its physical meaning:

Term (3.2) is a straightforward fit-to-data term attempting to minimize the difference between the two-dimensional Fourier transform of the measured spectrum (m) and the optimized image of the spectrum ($\mathcal{S}$). This gives us a final image whose peaks are located and shaped in the same way as those in the original spectrum as their size changes. This term also tends to minimize the difference in signal level between the two arrays, meaning that we redistribute the signal of the spectrum more so than we reduce or increase it as we modify the values in the boxes of $\mathcal{S}$.

For term (3.3), we define

$$||\delta_x \mathcal{S}||^2 = \sum_{\substack{i \\ \text{s.t. } s_i \text{ and } s_{i+1} \text{ are horizontally} \\ \text{adjacent in the same box}}} (s_i - s_{i+1})^2$$

as a short form for the L_2 difference norm in the x (F_2) dimension of the reconstructed spectrum, taking into account the sparsity of $\mathcal{S}$. This will tend to minimize the differences between adjacent points. Minimizing these differences promotes smoothness of the lines in this direction without trying to fit the peak to a particular line shape, which would require the incorporation of other variables. The line-shapes should be smooth due to the definition of the NMR peaks as given in Section 2.1.1. We do not regularize in the F_1 dimension as we do not generally expect smoothness in this direction.

Term (3.4) penalizes signals in the boxes of $\mathcal{S}$ if they are considered unlikely to correspond to a bond and penalizes the values in p if they correspond to low-signal areas of $\mathcal{S}$. These values will be increased in the opposite scenario. This promotes increasing signal in areas of likely peaks and a decrease of signal in areas of unlikely peaks. In the end, signal will be completely removed from locations we have decided will not contain peaks and will be concentrated in the areas of high certainty peaks.

Term (3.5) penalizes the signals in two paired boxes of $\mathcal{S}$ if the two boxes have a large difference in signal level. This means that a correlation will become less likely if its two related areas have very different signal levels.

Term (3.6) is a quartic penalty function designed to keep the total number of bonds for a single carbon between one and four, with two bonds being the most likely. The term achieves this by taking each carbon i and summing over the bond probability between i and every other carbon, j to get the total value of the bonds to i. This sum is then subtracted from two to shift the center of the quartic to $+2$ and the result of this is raised to the fourth power to generate our function. This function was chosen because it remains fairly flat in the centre and sharply rises afterward meaning that two bonds to a carbon is not penalized at all, one or three bonds are lightly penalized and outside that specified range, the penalty is strong.

Term (3.7) penalizes the sum of the probabilities in p, or the number of likely carbon bonds found. This function simply adds together all the bond probability values and minimizes this value. This is necessary because we do not have an a priori estimate for the peak heights, so we are only penalizing peaks which we do not believe correspond to bonds and large differences in paired peak heights. We would therefore expect a minimum in the objective function when many bonds are predicted, even ones with very low intensity peaks. This penalty prevents the optimizer from reaching such a non-descriptive minimum. In practice, investigators will have a good estimate of the total number of bonds in an unknown, or partially known, molecule and this information would be incorporated into this penalty.

3.3 Implementation

The problem is dominated by bi-quadratic terms; being quadratic in both $\mathcal{S}$ and p separately. This fact suggests that the problem can be split into two parts and solved using an alternating Gauss-Seidel approach to solve alternatively for $\mathcal{S}$ and for p.

In the first, unconstrained, problem (3.10) we solve for $\mathcal{S}$. The second part, (3.11) is a constrained problem that estimates the likelihoods of all possible bonds, thereby solving for p. The result of this split was a large increase in the speed of the solution. We implement this problem in a MATLAB program using standard routines included in the Mathworks optimization toolbox. We load the raw time series data into MATLAB using matNMR [vB07] and we run a standard MATLAB Fourier transform before normalizing the spectrum to make the largest value one in order to limit the size of the necessary constants. In initial tests using the sucrose spectra at a size of 128×4096, the solution using the single (combined) problem took between 10 and 12 hours, while the split problem took between 2 and 10 minutes to reach comparable solutions.

$$\begin{aligned} \min_{\mathcal{S}} \; & ||m - \mathcal{S}||^2 + \lambda_1 ||\delta_x \mathcal{S}||^2 + \lambda_2 \sum_{ij} (1 - p_{ij})^2 ||\mathcal{S}_{ij} + \mathcal{S}_{ji}||^2 \\ & + \lambda_3 \sum_{ij} (||\mathcal{S}_{ij}||^2 - ||\mathcal{S}_{ji}||^2)^2 \end{aligned} \tag{3.10}$$

The first (unconstrained) problem results in the redistribution of signal from the original spectrum which causes an increase of peak area in the possible bond locations and a decrease of peak area across the rest of the spectrum. This means that we will achieve a clearer image of the spectrum.

$$\begin{aligned} \min_{p_{ij}} \; & \sum_{ij} (1 - p_{ij})^2 ||\mathcal{S}_{ij} + \mathcal{S}_{ji}||^2 + \mu_1 \sum_i \left(2 - \sum_{j \neq i} p_{ij} \right)^4 \\ & + \mu_2 \sum_{ij} p_{ij} \\ \text{s.t. } & p_{ij} \geq 0 \\ & p_{ij} \leq 1 \end{aligned} \tag{3.11}$$

The second (constrained) problem estimates the bond probabilities.

This splitting is the most natural and it results in two subproblems with desirable performance properties. The original problem is a large, non-quadratic, constrained problem. When we split it into two however, the first

subproblem is a large, but mostly quadratic problem and the second, while non-quadratic and constrained, involves many fewer variables than the original problem. This relationship is maintained for all sizes of molecules.

The penalty parameters λ_i and μ_i are scaling factors for each of the terms we wish to minimize. Generally these constants do not need to be modified from default values where we assign: $\lambda_1 = 10, \lambda_2 = 15, \lambda_3 = 1, \mu_1 = 4$. However, the experiment is rather sensitive to the value of μ_2, which represents the degree of connectivity, and this value must be changed according to the amount of signal and the connectivity of a molecule.

3.4 Comparison to Existing Methods

We compare the method described above to the approach used by Dunkel et al. in [DMC+90] and [DMPG92]. There are two major points of comparison, firstly, our method examines only the magnitude of signal in specific locations, rather than the full lineshape. Secondly, we use only two assumptions about the form of the experiment: namely the number of bonds per carbon and that two peak pairs corresponding to a bond should have equal energy. On the other hard, the Dunkel method generates line shapes defined by the expectation function (Equation [1] in [DMC+90]):

$$S(\nu_1,\nu_2) = \frac{T_2^{\mathrm{DQ}}}{1+\left[2\pi T_2^{\mathrm{DQ}}(\nu_1-\nu_{\mathrm{DQ}})\right]^2}\left\{\frac{I_A I_{A_1} T_2^{\mathrm{SQ}}}{1+\left[2\pi T_2^{\mathrm{SQ}}(\nu_2-\nu_{A_1})\right]^2} - \frac{I_A I_{A_2} T_2^{\mathrm{SQ}}}{1+\left[2\pi T_2^{\mathrm{SQ}}(\nu_2-\nu_{A_2})\right]^2} + \frac{I_B I_{B_2} T_2^{\mathrm{SQ}}}{1+\left[2\pi T_2^{\mathrm{SQ}}(\nu_2-\nu_{B_2})\right]^2} - \frac{I_B I_{B_1} T_2^{\mathrm{SQ}}}{1+\left[2\pi T_2^{\mathrm{SQ}}(\nu_2-\nu_{B_1})\right]^2}\right\} \tag{3.12}$$

where ν_1 and ν_2 are positions on the F_1 and F_2 axes, respectively. The two potentially bonded nuclei are referred to as A and B, with their chemical shifts notated as ν_A and ν_B, the sum of these shifts is ν_{DQ} and refers to the position of the peak doublets on the F_1(vertical) axis. The doublets themselves are ν_{A_k} and ν_{B_k} where k is either 1 or 2, these are separated from the central shifts ν_A and ν_B by an amount relating to what's known as the coupling constant, J. I_{A_k} and I_{B_k} are the relative intensities of the peak doublets while I_A and I_B are scaling factors to convert the relative intensities to those observed. T_2^{DQ}

and T_2^{SQ} refer to the transverse relaxation times in the F_1 and F_2 directions, respectively.

This formula is based on the generalization of the Bloch equation presented in Sec. 2.1.1 expanded to model the case of two bonded nuclei, resulting in a 16×16 system, rather than a 3×3 system. The variable parameters in Eq. (3.12) are $\nu_A, \nu_B, J, I_A, I_B, T_2^{\mathrm{SQ}}, T_2^{\mathrm{DQ}}$ and these parameters are solved for for every potential bond. Meaning that, in the example of sucrose, Eq. (3.12) will be solved 66 times in order to reach a solution. Moreover, while the variables in our objective appear quadratically or to the fourth power (both convex polynomial functions), the variables in the equivalent Dunkel objective also appear in the denominator, resulting in difficult nonconvex subproblems, which cannot be simplified by a Gauss-Seidel approach.

Anecdotally, we know that the Dunkel approach, while based on sound theory, has not been adopted by experimentalists. We hypothesize that this is in part because the numerical difficulties which would not trouble a numerical analyst, but which would be outside the expertise of an experimentalist. In contrast, our approach is based on convex quadratic and quartic subproblems and can be expected to be numerically robust and easy to solve, which matches our experience.

Chapter 4

Experimental Results

The experiment we performed was with a concentrated sucrose (see Fig. 2.4 and Fig. 2.5) solution (approximately 80mg of sucrose in 0.5mL of water), run on a Bruker AV 500MHz spectrometer equipped with a 5mm room-temperature inverse-geometry probe. The pulse sequence we used was the standard INADEQUATE experiment including gradients to enhance pathway selection. Instead of running a single experiment overnight, we run several smaller experiments sequentially in order to help prevent any computer or spectrometer errors from affecting the entire dataset. Each of these experiments result in what we term a block; when averaged, these blocks give the results of the full experimental time. We gathered a total of eight blocks of data, each measuring 512 points in T_1 by 4096 in T_2 using 32 scans with a scan delay of one second and a 90° pulse with width $15\mu s$ taking slightly over six hours to run. The optimization used the optimization toolbox of MATLAB 7.9.0 installed on a 2.6GHz dual-core AMD Opteron processor. For small problems like this (under 20 or so carbons), the MATLAB solver is sufficiently fast, however for much larger problems as we expect to encounter, we will need to explore options for faster solvers.

In Fig. 4.1, we see the Fourier transform of the averaged value of all eight blocks. We can visually determine seven of the ten bonds reliably, but we have two bonds where only one of the two peaks are visible and one where neither is visible. In Fig. 4.2 we see the result of running our method on the average of two blocks. We see that even with only a quarter of the data, the final result of our method is much clearer than what can be seen in Fig. 4.1. Using two of the eight blocks was the minimum we could reliably use, but there existed a few combinations (all including one specific block) of two blocks that we could not solve. The results for all other pairs of blocks look highly similar to the results shown here.

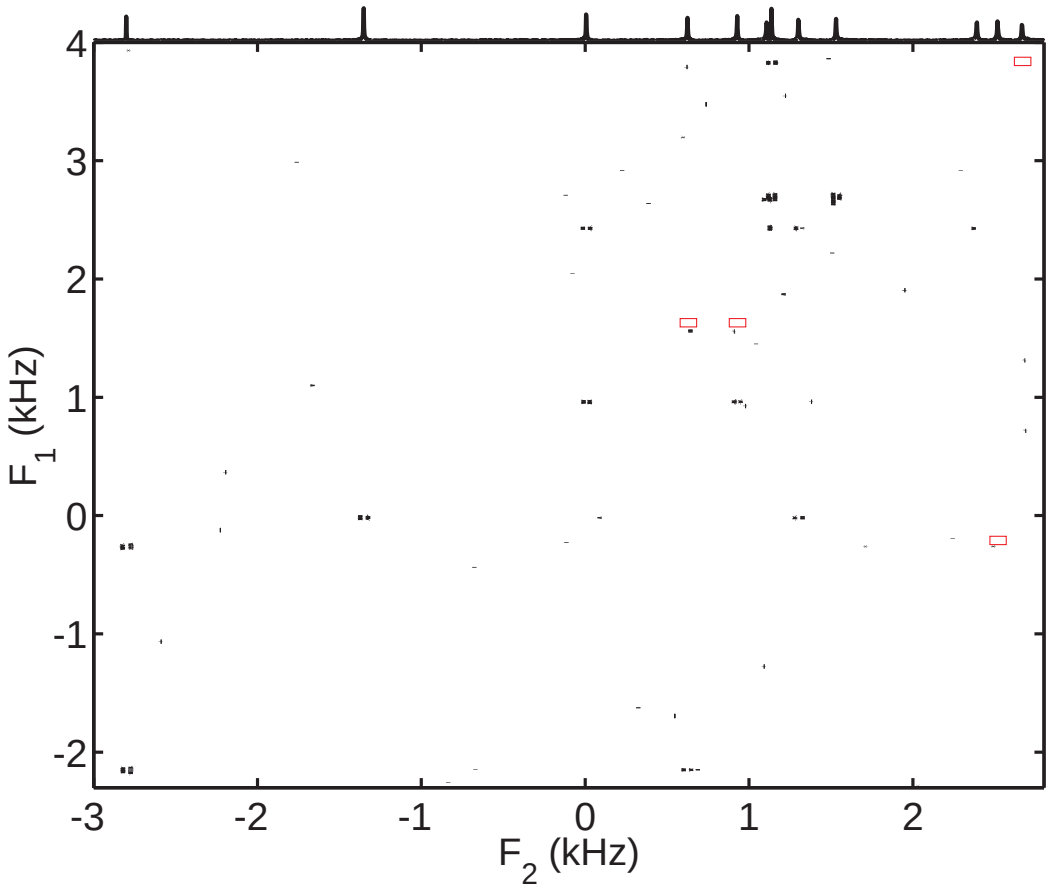

Figure 4.1: Averaged Fourier transform of all eight collected data sets. Here we can see most of the structure, but several bonds remain difficult to detect, most notably: $F2 \leftrightarrow F1, F3 \leftrightarrow F4$ and $G5 \leftrightarrow G6$. Red boxes indicate locations at which we should see peaks related to these bonds, but cannot.

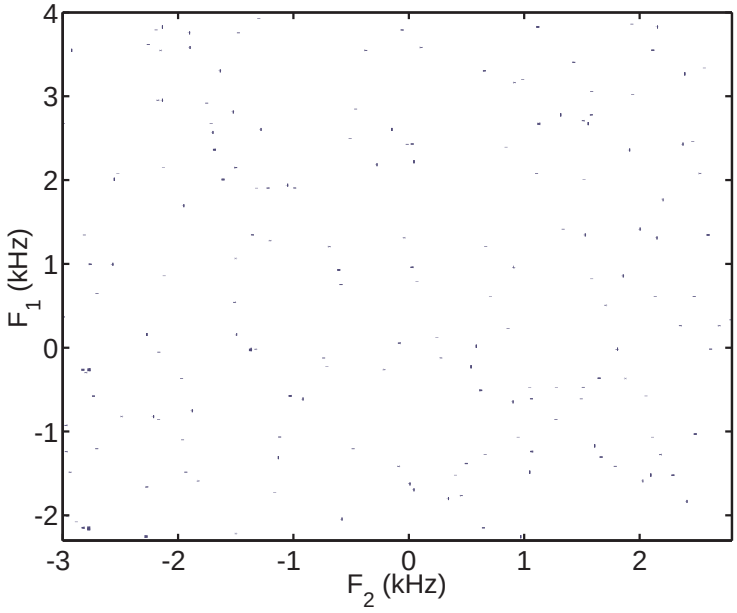

(a) Unmodified FT of two averaged data sets

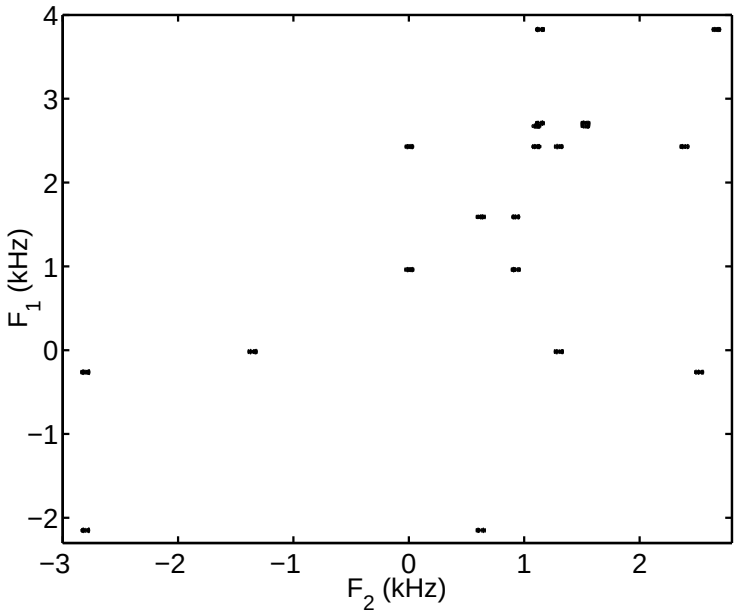

(b) Results of running our method on (a)

Figure 4.2: We can see here that running our method on two averaged data sets gives us better information than the average of all eight using traditional methods, constituting a decrease in experiment time of a factor of more than four.

Chapter 5

Non-Uniform K-Space Data

A possible method for lowering acquisition times for the spectrometers is simply to acquire less data. While most traditional methods use Fast Fourier Transforms to convert the time domain signal to the frequency domain, we suggest using non-uniform sampling along the F_1 axis so that less actual data is acquired.

While using these methods represents much slower processing than the traditional fast Fourier transform, we expect that the time saved in acquisition will more than make up for the increase in processing time.

The traditional Fourier transform is well known and given by,

$$X(m) = \sum_{n=1}^{N-1} x(n) e^{\frac{-2\pi i}{N} nm} \tag{5.1}$$

and its inverse given by,

$$x(n) = \frac{1}{M} \sum_{m=1}^{M-1} X(m) e^{\frac{2\pi i}{N} nm} \tag{5.2}$$

A non-uniform version of the frequency transform is given by,

$$\varphi(f_m) = \sum_{n=1}^{N-1} x(t_n) e^{-2\pi i t_n f_m} \tag{5.3}$$

where x is some function in the time domain and φ is its non-uniform frequency transform. The units of t are milliseconds and the units of f are kHz. The time transform is given by,

$$\psi(t_n) = \frac{1}{M} \sum_{m=1}^{M-1} y(f_m) e^{2\pi i t_n f_m} \tag{5.4}$$

where y is some function in the frequency domain and ψ is its non-uniform time transform. It is important to note that these two functions are not inverses of one another. In general, applying $\varphi(\psi(x))$ will not give us the original signal again. In the same manner as the uniform Fourier transform, the non-uniform version can be represented as a linear transform by multiplication with a matrix W, such that $\varphi = Wx$ with,

$$W = \begin{bmatrix} \omega(t_1, f_1) & \omega(t_1, f_2) & \omega(t_1, f_3) & \omega(t_1, f_4) & \dots & \omega(t_1, f_M) \\ \omega(t_2, f_1) & \omega(t_2, f_2) & \omega(t_2, f_3) & \omega(t_2, f_4) & \dots & \omega(t_2, f_M) \\ \omega(t_3, f_1) & \omega(t_3, f_2) & \omega(t_3, f_3) & \omega(t_3, f_4) & \dots & \omega(t_3, f_M) \\ \omega(t_4, f_1) & \omega(t_4, f_2) & \omega(t_4, f_3) & \omega(t_4, f_4) & \dots & \omega(t_4, f_M) \\ \vdots & \vdots & \vdots & \vdots & & \vdots \\ \omega(t_N, f_1) & \omega(t_N, f_2) & \omega(t_N, f_3) & \omega(t_N, f_4) & \dots & \omega(t_N, f_M) \end{bmatrix} \tag{5.5}$$

where $\omega(t_n, f_m) = e^{-2\pi i t_n f_m}$. The non-uniform version of this matrix does not have the exact same properties as the uniformly sampled version—notably that, in general, $W^{-1} \neq W^*$ and W is not necessarily invertible. The two-dimensional non-uniform transform can be defined as $\varphi = W_1 x W_2^*$ where W_1 is the column transform and W_2 is the row transform.

The differences in notation from the standard discrete Fourier transform are minor, we only note the dependence on the actual values of the time t_k and the frequency f_n rather than their indices. This is because during the transform, the points being constructed actually must be specified as we are not constructing the entire series each time. This gives us the advantage of being able to only construct the points we desire, and by concentrating the signal in only these areas where we know it can occur, we eliminate the rest of the signal to achieve better signal-to-noise ratios. As an example, in the one-dimensional spectrum of sucrose as in Fig. 2.3, we see exactly twelve peaks. This means that we can construct only twelve points from the spectrum, going from a uniform spectrum of 4096 points to a non-uniform frequency spectrum of twelve points.

5.1 Using the Non-Uniform Transform

We demonstrate the use of this method with a one-dimensional example. In order to find this transform as accurately as possible, we introduce an optimization function as follows,

$$\min W^* m - W^* W \varphi \tag{5.6}$$

$-\pi$	$-e$	-1	0	1	e	π
(peak)	(null)	(peak)	(null)	(peak)	(null)	(peak)

Table 5.1: The seven points used in our Fourier transforms, including 4 peaks and 3 null points.

where m is the original signal in the time domain, ψ is the vector of independent variables in the frequency domain and W is the matrix of non-uniform Fourier transform coefficients. The solution ψ of this problem provides the selected frequency points from the non-uniformly sampled time data given.

For this example we create a simple spectrum where peaks are placed at $\pm\pi$ and ± 1 using the function

$$f(x) = \sin(2\pi x) + \cos(2\pi^2 x) \tag{5.7}$$

We used 256 evenly distributed time points between 0 and 25ms to be able to create the fast Fourier transform of this function, shown in Fig. 5.1. 50 sets of noise were generated by MATLAB with entries ranging between -1 and 1, representing approximately 50% of the maximum and minimum points of the time-domain function $f(x)$. Running MATLAB's fast Fourier transform function `fft` on each of these sets of data and selecting seven points as in Tbl. 5.1, representing four peaks and three zero valued points, we created an error plot shown in Fig. 5.2. The mean value of the standard deviations among the 50 noisy sets over the seven selected points was 9.0462.

A comparison was then made between an `fft` and non-uniform Fourier transform of a 128 point time series, with time values between 0 and 12.5ms, of f as in Eq. (5.7). We see the `fft` solution in Fig. 5.3, which was done in the same manner as the 256 point example above; 50 time series with added noise were transformed and the seven points in the list shown by Tbl. 5.1 were selected from the transform, and we show the resulting mean and standard deviation for each point. We can see a problem that has occurred with this approach, notably that the peaks at -1 and 1 are now much smaller than in the above example as a result of having less data in the transform. The frequencies at which these peaks appear do not exist in this 128 point spectrum, as opposed to the original 256 point spectrum. We also show a non-uniform transformation in Fig. 5.4 in which only the 7 points shown were constructed from a time series of 128 random points from between 0 and 25ms. We see that we have error bars of approximately the same width as in Fig. 5.3, but

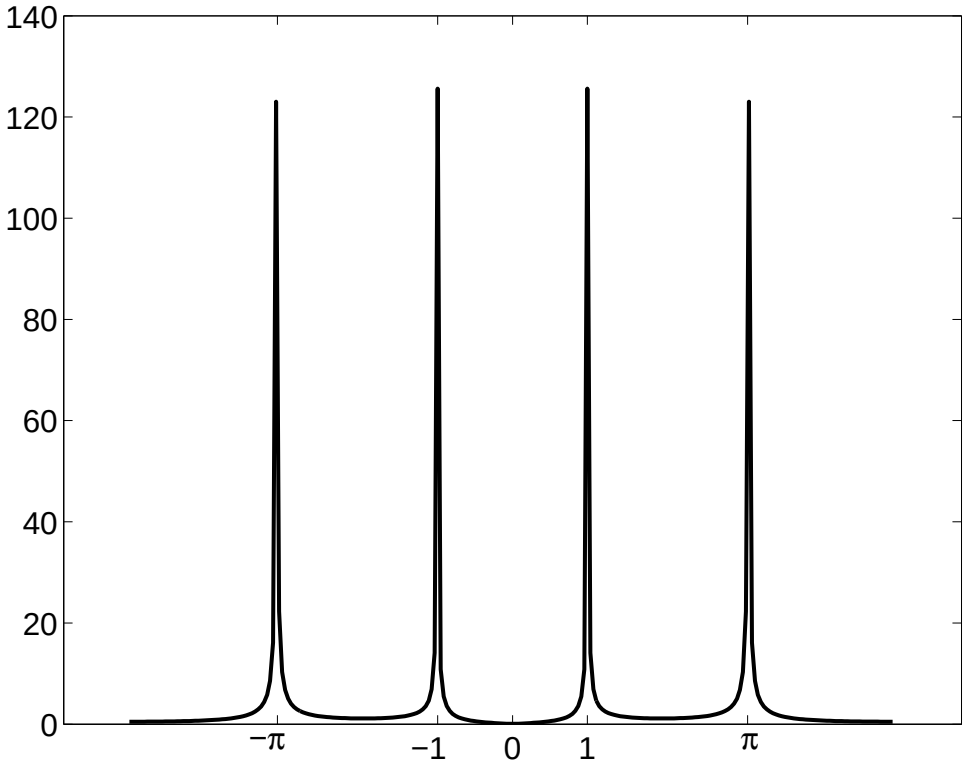

Figure 5.1: Full spectrum of the fast Fourier transform of function f.

the mean values of each of these points are much closer to those shown by Fig. 5.2.

This example shows that by using this non-uniform transform, we are able to gain flexibility in how we acquire data to process, both in the number of points we acquire and the time spacing between them. With the standard fast Fourier transform, we are unable to specify the frequency points we generate and therefore are likely to miss exact points where peaks occur as demonstrated in Fig. 5.3. This is related to the fact that we are tied to sampling times that will give the required spectral width. The first problem is solved by the choice of specific frequency points that we will solve through using the non-uniform transform. The second is solved because the sampling times required are based on the average distance between frequency points; with a standard fast Fourier

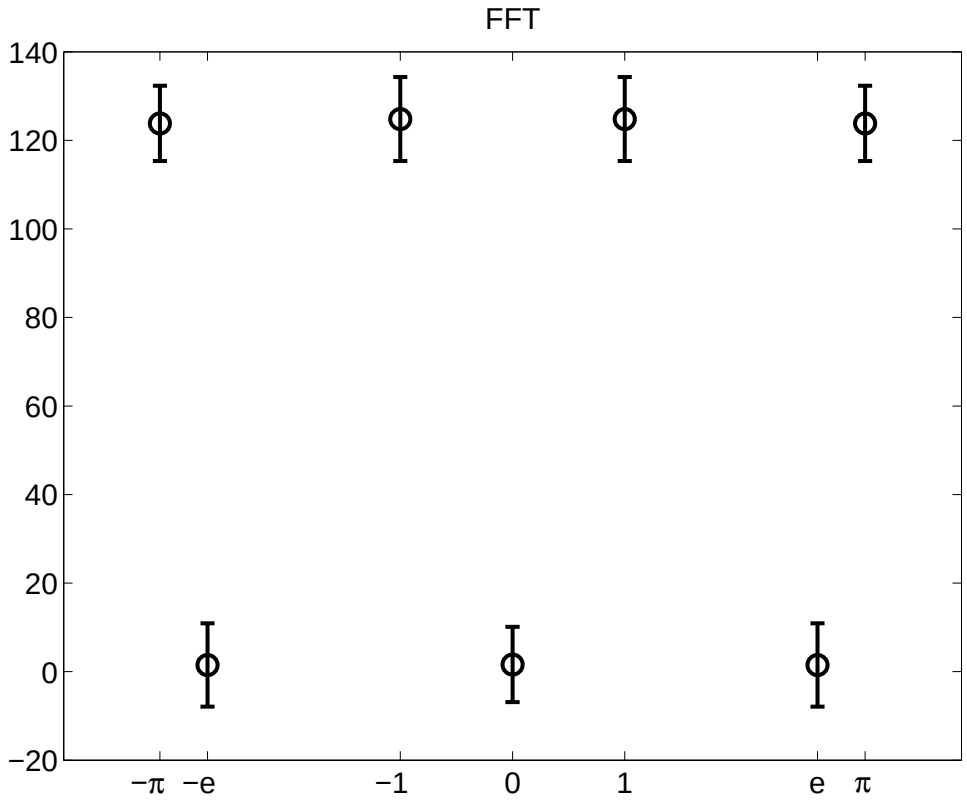

Figure 5.2: Error bar plot showing the mean value and standard deviation at each of seven selected points based on 50 different noisy functions after applying MATLAB's `fft` function. The mean value of each selected point is shown by a circle while the bars at each point represent a single standard deviation.

transform, the frequency steps are derived from the sampling times. Using a non-uniform transform, however, selects only a few frequencies with much larger average distances between them, making the required time sampling significantly less strict for acquiring useful results. We see this in Fig. 5.4 as the results are better than those in Fig. 5.3, but with sample times, on average, twice as far apart.

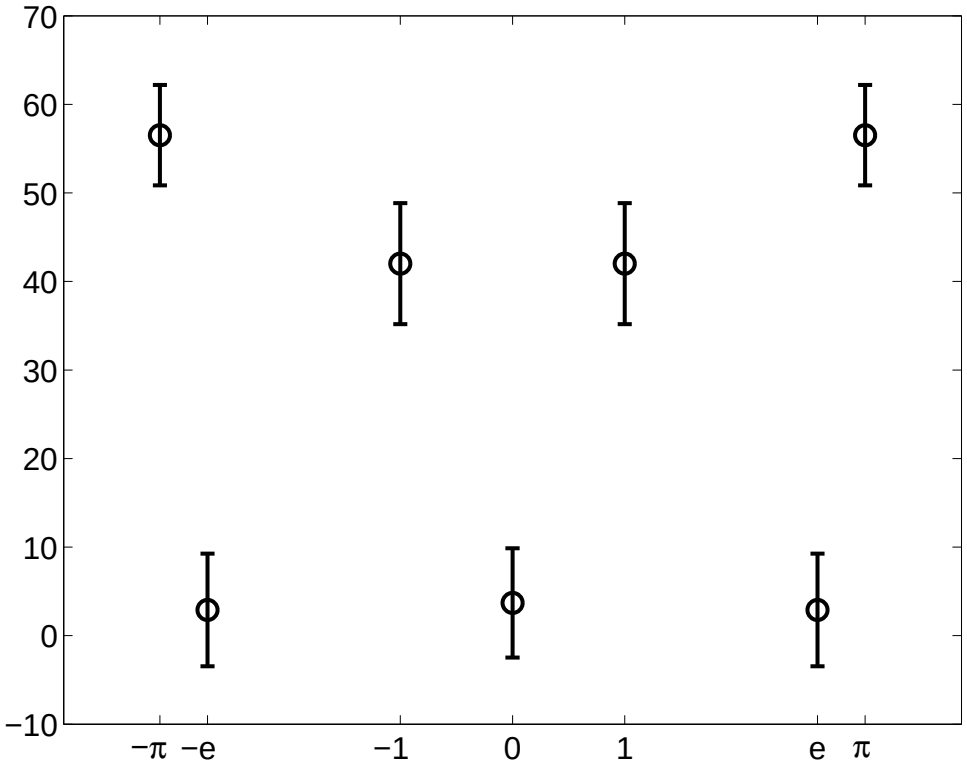

Figure 5.3: Error bar plot showing the mean value and standard deviation at each of seven selected points based on 50 different noisy functions after applying MATLAB's `fft` function to a 128-point function. The mean value of each selected point is shown by a circle while the bars at each point represent a single standard deviation.

5.2 Application to INADEQUATE

We can use the above method in order to minimize the sampling required for solving INADEQUATE spectra. This is accomplished by replacing the objective of Eq. (3.2) by Eq. (5.6) and by using time domain data rather than transformed data as our measurement. In the example of sucrose spectra at a size of 512×4096, as with the examples given in Sec. 4. However, we require only 24(corresponding to the double peaks at each of the 12 chemical shifts of sucrose) of the original 4096 columns and, in each of these columns,

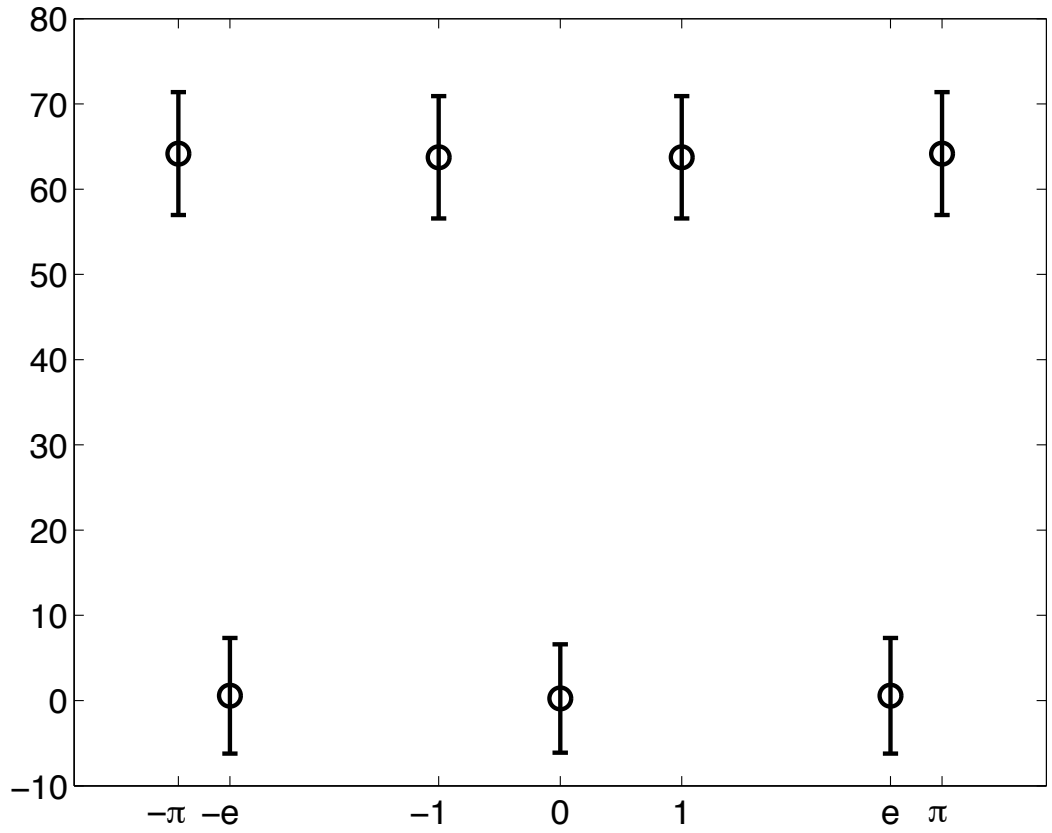

Figure 5.4: Error bar plot showing the mean value and standard deviation at each of seven selected points based on 50 different noisy functions after applying our non-uniform Fourier transform function to a 128-point function. The mean value of each selected point is shown by a circle while the bars at each point represent a single standard deviation.

only 11(corresponding to each possible pairing for each of the 12 carbons) of the original 512 rows. This means that by using the non-uniform Fourier transform, we will significantly reduce the number of variables that are used throughout the optimization process. In addition, we will reduce the amount of data that needs to be collected in order to generate the desired frequency points as was discussed previously in Sec. 5.1.

This means that by using this method we will be able to cut down on the data collected in the t_1 dimension, which adds much more to the length of

the experiment than does extra data collected in the t_2 dimension.

Chapter 6

Conclusion and Future Work

In conclusion, we have demonstrated in this thesis a method of solving for carbon bonds from INADEQUATE spectra by using image regularization and optimization techniques. Despite the limited amount of testing so far, the early results are promising. We have shown that we can reduce the experiment time required to determine the carbon backbone of sucrose from over 48 hours using traditional methods to around 12 using our method. We plan to expand our testing to a wide variety of molecules as soon as possible and we expect similar speed-ups in all cases. We have also shown a reliable method of finding a non-uniform Fourier transform of a single-dimensional NMR spectrum and communicated that using this method will allow greater flexibility in the amount of data required for these spectra. This technique is extendable to a two-dimensional spectrum by using two separate transforms, fast in the directly-acquired dimension and non-uniform in the indirect dimension.

It would be desirable to eventually translate all of the extant MATLAB code into C so that we will have more flexibility in the choice of solver and faster performance. Currently the code is all in MATLAB form simply to make use of the ease of obtaining data using the matNMR [vB07] tool, so a translation of the code into C will require the development or discovery of a tool that can import Bruker two-dimensional binary FID files into a C array to use with our method.

It would also be desirable to develop a full integration between the two currently separate parts presented in this thesis, namely the two-dimensional non-uniform Fourier transform method and the INADEQUATE bond solver. This will require extensive testing to ensure the non-uniform transform is progressing properly at each iteration of the full solver, and ensure that the reduction of points does not introduce new problems over the old method of using a sparse matrix at full size as our intermediate variables. We do, how-

ever, expect this integrated program to provide results with significantly better signal-to-noise ratios than the example presented in Sec. 4.

After these additional tasks are completed we expect that this will be an extremely versatile system for analysing INADEQUATE spectra.

Bibliography

[ABE76] W. P. Aue, E. Bartholdi, and R. R. Ernst. Two-dimensional spectroscopy. application to nuclear magnetic resonance. *Journal of Chemical Physics*, 64(5):2229–2246, 1976.

[BB87] J. Buddrus and H. Bauer. Direct identification of the carbon skeleton of organic compounds using double quantum coherence ^{13}C-NMR spectroscopy. the INADEQUATE pulse sequence. *Angewandte Chemie International Edition in English*, 26:625–642, 1987.

[BF80] A. Bax and R. Freeman. Investigation of 13C-13C couplings in natural-abundance samples: the strong coupling case. *Journal of Magnetic Resonance*, 41:507–511, 1980.

[BF81] A. Bax and R. Freeman. Investigation of complex networks of spin-spin coupling by two-dimensional NMR. *Journal of Magnetic Resonance*, 44(3):542–561, 1981.

[BFF81] A. Bax, R. Freeman, and T. A. Frenkiel. An NMR technique for tracing out the carbon skeleton of an organic molecule. *Journal of the American Chemistry Society*, 103:2102–2104, 1981.

[BFFL81] A. Bax, R. Freeman, T. A. Frenkiel, and M. H. Levitt. Assignment of carbon-13 NMR spectra via double quantum coherence. *Journal of Magnetic Resonance*, 43:478–483, 1981.

[BFK80a] A. Bax, R. Freeman, and S. P. Kempsell. Investigation of 13C-13C long-range couplings in natural-abundance samples. *Journal of Magnetic Resonance*, 41:349–353, 1980.

[BFK80b] A. Bax, R. Freeman, and S. P. Kempsell. Natural abundance 13C-13C coupling observed via double-quantum coherence. *Journal of the American Chemistry Society*, 102:4849–4851, 1980.

[BHA+10] Alex D. Bain, Donald W. Hughes, Christopher K. Anand, Zhenghua Nie, and Valerie J. Robertson. Problems, artifacts and solutions in the INADEQUATE NMR experiment. *Magnetic Resonance in Chemistry*, 48(8):630–641, 2010.

[BLLO99] Daniel W Bohman, Zheng Li, Du Li, and Noel L Owen. The analysis of mixtures using the 2-D INADEQUATE NMR technique. *Journal of Molecular Structure*, 480-481:125 – 132, 1999.

[BM83] Ad Bax and T. H. Mareci. Practical aspects of carbon-13 double quantum NMR. *Journal of Magnetic Resonance (1969)*, 53(2):360 – 363, 1983.

[DMC+90] R. Dunkel, C. L. Mayne, J. Curtis, R. J. Pugmire, and D. M. Grant. Computerized analysis of 2D INADEQUATE spectra. *Journal of Magnetic Resonance (1969)*, 90:290–302, 1990.

[DMPG92] R. Dunkel, C. L. Mayne, R. J. Pugmire, and D. M. Grant. Improvements in the computerized analysis of 2D INADEQUATE spectra. *Analytical Chemistry*, 64(24):3133–3149, 1992.

[FFR82] R. Freeman, T. A. Frenkiel, and M. B. Rubin. Structure of a photodimer determined by natural-abundance 13C-13C coupling. *Journal of the American Chemisty Society*, 104:5545–5547, 1982.

[FMD+92] M. P. Foster, C. L. Mayne, R. Dunkel, R. J. Pugmire, D. M. Grant, J.-M. Kornprobst, J.-F. Verbist, J.-F. Biard, and C. M. Ireland. Revised structure of bistramide a(bistratene a): application of a new program for the automated analysis of 2D INADEQUATE spectra. *Journal of the American Chemisty Society*, 114:1110 – 1111, 1992.

[KULB92] K. E. Kövér, D. Uhrin, T. Liptaj, and G. Batta. Easy implementation of selective INADEQUATE and H-C-C relay experiments using heteronuclear chemical shift filtering. *Magnetic Resonance in Chemistry*, 30:68 – 72, 1992.

[KV86] Paul J. Keller and Katherine E. Vogele. Sensitivity enhancement of INADEQUATE by proton monitoring. *Journal of Magnetic Resonance (1969)*, 68(2):389 – 392, 1986.

[LB93] J. Lambert and J. Buddrus. Sensitivity enhancement of two-dimensional 13C, 13C-INADEQUATE spectroscopy by considering symmetry and isotopic shifts. *Journal of Magnetic Resonance, Series A*, 101(3):307 – 312, 1993.

[LKB89] J. Lambert, H. J. Kuhn, and J. Buddrus. Complete identification of C-C connectivities by two-dimensional INADEQUATE NMR spectroscopy with composite pulses. *Angewandte Chemie International Edition English*, 28:738–740, 1989.

[NL82] A. Neszmelyi and G. Lukacs. Natural abundance one-bond C-13 - C-13 coupling constants in monosaccharide derivatives and in sucrose. *Journal of the American Chemisty Society*, 104:5342–5346, 1982.

[NSF96] T. Nakazawa, H. Sengstschmid, and R. Freeman. Enhancement of INADEQUATE spectra according to symmetry criteria. *Journal of Magnetic Resonance*, 120:269 – 273, 1996.

[ODH+95] A. M. Orendt, R. Dunkel, W. J. Horton, R. J. Pugmire, and D. M. Grant. Computerized analysis of 2D INADEQUATE spectra to assign chemical shifts in aromatic compounds. *Journal of Magnetic Resonance*, 33:803–811, 1995.

[OWDM88] B. H. Oh, W. M. Westler, P. Darba, and J. L. Markley. Protein carbon-13 spin systems by a single two-dimensional nuclear magnetic resonance experiment. *Science*, 240(4854):908–911, 1988.

[Pod90] I. S. Podkorytov. Seven-pulse sequence DEPT-INADEQUATE. *Journal of Magnetic Resonance (1969)*, 89(1):129 – 132, 1990.

[RAW81] R. Richarz, W. Amman, and T. Wirthlin. Computer-assisted determination of carbon connectivity patterns in organic molecules from natural-abundance $^{1}J_{CC}$ data. *Journal of Magnetic Resonance*, 45:270 – 283, 1981.

[SE85] S. W. Sparks and P. D. Ellis. DEPT polarization transfer for the INADEQUATE experiment. *Journal of Magnetic Resonance (1969)*, 62(1):1 – 11, 1985.

[SFF+82] Ole W. Sørensen, Ray Freeman, Thomas Frenkiel, Thomas H. Mareci, and Regina Schuck. Observation of 13C—13C couplings

with enhanced sensitivity. *Journal of Magnetic Resonance (1969)*, 46(1):180 – 184, 1982.

[SH87a] Jeremy K.M. Sanders and Brian K. Hunter. *Modern NMR Spectroscopy*, chapter 1. Oxford University Press, 1987.

[SH87b] Jeremy K.M. Sanders and Brian K. Hunter. *Modern NMR Spectroscopy*, chapter 2. Oxford University Press, 1987.

[SH87c] Jeremy K.M. Sanders and Brian K. Hunter. *Modern NMR Spectroscopy*, chapter 3. Oxford University Press, 1987.

[vB07] Jacco 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.

[ZSK+09] Z. Zhou, J. C. Stevens, J. Klosin, R. Kümmerle, X. Qui, D. Redwine, R. Cong, A. Taha, J. Mason, B. Winniford, P. Chauvel, and N. Montanez. NMR study of isolated 2,1-inverse insertion in isotactic polypropylene. *Macromolecules*, 42:2291–2294, 2009.