Available online at www.sciencedirect.com
ScienceDirect
Journal of Magnetic Resonance 185 (2007) 110–117
JMR Journal of Magnetic Resonance
www.elsevier.com/locate/jmr

Magnetic field homogeneity: A new approach to orthogonalizing and optimizing shim gradients

Carl A. Michal *

Department of Physics and Astronomy, The University of British Columbia, 6224 Agricultural Rd., Vancouver, BC, Canada V6T 1Z1

Received 30 September 2006; revised 7 December 2006
Available online 26 December 2006

Abstract

A new approach to optimizing shim coil currents for magnetic resonance magnets is presented. The new approach orthogonalizes the shim coil gradients to allow a simple one-dimensional optimization for each orthogonalized "composite shim." The technique demands no specialized equipment, requiring only the acquisition of simple one-dimensional NMR spectra. Examples from two high-resolution NMR spectrometers are presented, where the shim currents found by the new algorithm provide higher resolution than was obtained by the spectrometer vendor's installation engineers using field-mapping techniques. The examples shown demonstrate the advantages of the technique for high-resolution NMR, but we expect the approach will also find application in a broad variety of areas including imaging and *in vivo* spectroscopy.

Keywords: Shimming; Field mapping; Lineshape; Orthogonalization

1. Introduction

Magnetic field homogeneity is crucial in many magnetic resonance experiments. To achieve as homogeneous a field as possible, modern magnet systems are typically equipped with a number of magnetic field gradient coils, known as shim coils, each producing a different magnetic field gradient. The current passed through each coil can be adjusted individually, and the goal of "shimming" is to choose the currents that best cancel the inhomogeneity of the unshimmed magnetic field.

The set of shim coils may consist of more than 45 separate gradients, and while it is common to choose the spatial dependence of the gradients based on the spherical harmonics [1] in an attempt to make them orthogonal, the fact that samples are typically not spherically shaped and that susceptibility effects of the probe, sample container, and sample itself may perturb the field, it is not possible to make the gradients produced by the coils orthogonal under all conditions. This fact, combined with a potentially large number of shim gradients, has the potential of making the optimization of the shim currents both time-consuming and challenging. Often when magnet systems are installed, the main magnet field as well as the fields produced by the shims are mapped with a robotic field mapping system. This mapping process allows the calculation of the optimal current settings over the active volume of the system. This robotic approach is very effective, but requires expensive specialized equipment, and cannot be used on a routine basis because it requires removal of the probe and sample from the magnet and so does not take into account the details of the sample shape and the field perturbations due to susceptibility effects.

The advent of gradient shimming methods has in many cases dramatically simplified shimming [2,3], but gradient shimming is not available in many cases, particularly on older spectrometers, as it requires gradients that can be switched rapidly.

In this work we introduce a new approach to magnet shimming that can be fully automated and provides the optimal shim current settings for the sample at hand. It

* Corresponding author. Fax: +1 604 822 5324.
E-mail address: michal@physics.ubc.ca.

1090-7807/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmr.2006.12.006

requires no specialized equipment: neither pulsed field gradients nor a field-frequency lock are required. For a reasonable number of shim gradients the method is extremely fast. For a complete optimization with no assumptions made on the orthogonality of the various gradients, of order $2N^2$ free induction decays are necessary (where N is the number of gradients to be optimized), but where reasonable guesses can be made or prior knowledge is available as to which shims will interact with each other, this number can be reduced substantially. As part of the procedure, the gradients produced by the coils are orthogonalized, producing "composite shims" that do not interact with each other and can thus be optimized individually. The method is in principle non-iterative, although for best results a small number of iterations may be required.

The method has relatively few requirements; all that is needed is the ability to collect one-dimensional NMR spectra with a relatively high signal to noise ratio of a reasonably well resolved single peak.

While the examples that will be presented here are from two high-resolution solution NMR spectrometers at 400 and 600 MHz, where a large number of shim gradients (23 and 39, respectively) are optimized to provide high resolution spectra, we expect the method will also find application in imaging or *in vivo* spectroscopy, where a smaller number of shims is generally available and the speed of optimization is crucial. In our examples, the present algorithm finds shim current settings that provide higher resolution than was produced when the systems were recently moved and reinstalled, at which time the magnets and shims were mapped by the vendor's installation engineers.

We begin with an explanation of the mathematics underlying the method, followed by a description of the implementation and a number of practical issues and the strategies used to deal with them. Two examples are then presented and discussed, and we close with an offer of source code to make the method widely available.

2. Theory

To simplify the mathematics involved, we make the assumption that the NMR sample consists of a single visible species, resulting in an NMR spectrum of a single peak. In practice this is not necessary, though a spectrum containing a dominant, and ultimately (i.e. after shimming) well-resolved peak is preferred.

The z-component of magnetic field produced by the main magnet, as perturbed by the probe and sample, in the region of the sample can be expressed as:

$$B(\vec{x}) = B_0 + \sum_{i=0}^{\infty} a_i g_i'(\vec{x}), \tag{1}$$

where the a_i are coefficients of a set of orthonormal functions of position $g_i'(\vec{x})$. By orthonormal, we mean:

$$\int g_i'(\vec{x}) g_j'(\vec{x}) W(\vec{x})\, \mathrm{d}^3x = \delta_{ij}, \tag{2}$$

in which $W(\vec{x})$ is a weighting function accounting for the distribution of sample spins in the sample, the efficiency of the excitation coil, and the receive profile of the detection coil. The integral is over all space, though $W(\vec{x})$ will be 0 outside of the sample. One can think of the $g_i'(\vec{x})$'s as the (orthogonal) field gradients produced by a set of idealized shim coils. Ultimately, these functions will correspond to the gradients of the orthogonalized, composite shims. Real shim sets cannot possibly satisfy this orthonormalization condition for all samples, as it is sample dependent through the weighting factor $W(\vec{x})$. In practice, most high-resolution NMR shim sets have at least some gradients that do not come at all close to being orthogonal to the remaining shims in the set. If the shims could be made truly orthogonal, then shimming would be trivial, as each could be optimized independently of the rest.

The shim set provides a set of gradients, referred to as $g_i(\vec{x})$, that are almost certainly not mutually orthogonal. Assuming that this set contains no linear dependencies, the first N of the $g_i'(\vec{x})$ can be chosen so that they span the same space as the shim gradients. When optimized, the shims will be able to cancel the first N terms in the summation of Eq. (1). Our goal is to produce a set of orthogonal gradients, the $g_i'(\vec{x})$, as linear combinations of the $g_i(\vec{x})$, and to find the N coefficients a_i that provide the optimal correction to the magnetic field.

Once the functions $g_i'(\vec{x})$ are known, the coefficients can be found through an experimental determination of:

$$a_i = \int B(\vec{x}) g_i'(\vec{x}) W(\vec{x})\, \mathrm{d}^3x. \tag{3}$$

This value can be determined as follows. Consider the integral:

$$I = \int S(\omega)\omega^2\, \mathrm{d}\omega, \tag{4}$$

in which $S(\omega)$ is the intensity of the NMR spectrum at frequency ω. In the limit of an infinitely sharp NMR spectrum, $S(\omega)$ is given by:

$$S(\omega) = \int W(\vec{x})\delta(\omega - \gamma B(\vec{x}))\, \mathrm{d}^3x, \tag{5}$$

in which $\delta(x)$ is the Dirac delta function. Substituting Eq. (5) into Eq. (4) and exchanging the order of integration, we find

$$I^0 = \int W(\vec{x})\gamma^2 B^2(\vec{x})\, \mathrm{d}^3x, \tag{6}$$

where the superscript on I^0 indicates that this integral is obtained with no shim gradients applied. Now if shim gradient i is applied with a coefficient d_i' during acquisition, the field becomes $B(\vec{x}) + d_i' g_i'(\vec{x})$, and the integral of Eq. (4) becomes

$$I_i^+ = \int W(\vec{x})\left(\gamma^2 B^2(\vec{x}) + 2d_i' g_i'(\vec{x})\gamma B(\vec{x}) + d_i'^2 g_i'^2(\vec{x})\right) \mathrm{d}^3x, \tag{7}$$

where the notation I_i^+ now indicates gradient i applied with a positive offset. Repeating the experiment with gradient i applied with a coefficient of $-d_i'$, we find

$$I_i^- = \int W(\vec{x})\left(\gamma^2 B^2(\vec{x}) - 2d_i' g_i'(\vec{x})\gamma B(\vec{x}) + d_i'^2 g_i'^2(\vec{x})\right) \mathrm{d}^3x. \tag{8}$$

Taking the difference between these two results and dividing by $4d_i'$, we arrive at:

$$\frac{I_i^+ - I_i^-}{4d_i'} = \int W(\vec{x}) g_i'(\vec{x}) B(\vec{x})\, \mathrm{d}^3x = a_i. \tag{9}$$

This simple procedure generally fails to find the correct adjustment for shim i however, for the non-orthogonal gradients $g_i(\vec{x})$. Fortunately, we can orthogonalize the gradients through the well-known Gram-Schmidt orthogonalization procedure [4]. All that is necessary is the measurement of the quantities:

$$C_{ij} = \int W(\vec{x}) g_i(\vec{x}) g_j(\vec{x})\, \mathrm{d}^3x. \tag{10}$$

This can be done straightforwardly by acquiring spectra with pairs of gradients i and j with coefficients d_i and d_j applied simultaneously during acquisition. Under these conditions, the integral of Eq. (4) becomes:

$$I_{ij}^{++} = \int \Big(\gamma^2 B^2(\vec{x}) + d_i^2 g_i^2(\vec{x}) + d_j^2 g_j^2(\vec{x}) + 2\gamma B(\vec{x})(d_i g_i(\vec{x}) + d_j g_j(\vec{x})) + 2d_i d_j g_i(\vec{x}) g_j(\vec{x})\Big) \mathrm{d}^3x. \tag{11}$$

Spectra are also acquired with the other three combinations of positive and negative d_i and d_j, and combined to yield:

$$\frac{I_{ij}^{++} - I_{ij}^{+-} - I_{ij}^{-+} + I_{ij}^{--}}{8d_i d_j} = \int W(\vec{x}) g_i(\vec{x}) g_j(\vec{x}) d^3x = C_{ij} (i \neq j). \tag{12}$$

The coefficient C_{ij} for all pairs of shims that are expected to have significant interactions can then be measured, and the Gram-Schmidt orthonormalization procedure used to construct a set of composite shims that are orthogonal. This procedure is more practically carried out not by measuring all the C_{ij}, but rather by measuring the correlations between the gradient produced by a single shim coil, $g_i(\vec{x})$ and one of the previously orthogonalized composite shims $g_j'(\vec{x})$ (for $j < i$). We use the notation $C_{ij'}$, where the prime on the subscript j indicates that the second shim of the pair has already been orthogonalized. We now have all the ingredients necessary for an algorithm to orthogonalize and optimize the shims. As prescribed by the Gram-Schmidt orthogonalization algorithm, each composite shim $g_i'(\vec{x})$ is constructed according to

$$g_i'(\vec{x}) = \frac{g_i(\vec{x}) - \sum_{j<i} C_{ij'} g_j'(\vec{x})}{N_i}, \tag{13}$$

in which N_i is given by

$$N_i^2 = \int W(\vec{x}) \left(g_i(\vec{x}) - \sum_{j<i} C_{ij'} g_i'(\vec{x})\right)^2 \mathrm{d}^3x, \tag{14}$$

which, when expanded into useful form becomes:

$$N_i^2 = C_{ii} - \sum_{j<i} C_{ij'}^2. \tag{15}$$

C_{ii} is given by

$$C_{ii} = \frac{I_i^+ + I_i^- - 2I^0}{4d_i^2} = \int W(\vec{x}) g_i(\vec{x}) g_i(\vec{x})\, \mathrm{d}^3x. \tag{16}$$

In practice, determining N_i through Eq. (15) is problematic. If one of the gradients is nearly linearly dependent upon its predecessors, then measurement noise can cause Eq. (15) to suggest a negative value for N_i^2. A more robust alternative is to measure N_i directly. A temporary composite shim is constructed from:

$$g_{iT}'(\vec{x}) = g_i(\vec{x}) - \sum_{j<i} C_{ij'} g_j'(\vec{x}), \tag{17}$$

and then spectra are acquired with offsets $\pm d_{Ti}'$, where the subscript T indicates a quantity corresponding to the temporary composite shim. N_i is then given by

$$N_i^2 = \frac{I_{Ti}^+ + I_{Ti}^- - 2I^0}{4d_{Ti}^2}. \tag{18}$$

In the above calculations we have assumed a delta function lineshape. If the delta function is replaced with a physically reasonable lineshape, the values of all of the integrals are increased by a lineshape dependent constant amount, but since the quantities calculated from the integrals are all differences, this constant drops out of Eqs. (9), (12), (16), and (18).

In practice, symmetric positive and negative offsets to a shim current setting may not always produce symmetric changes in the corresponding gradient strength, due to non-linearities in the shim current supplies, heating of the shim coils, or magnetic properties of probe components. If this asymmetry is large, it may impact the accuracy to which the C_{ij}'s and the a_i's can be determined. But because the severity of this issue will diminish as the offsets d_i, are reduced, iterations of the procedure will converge to the correct settings.

While the orthogonalization of the shims is sample dependent because $W(\vec{x})$ depends on the sample size, shape, and susceptibility, we expect that for similarly shaped samples the coefficients C_{ij} will be similar. Once a shim gradient has been orthogonalized with respect to its predecessors, finding the optimum current setting for it is straightforward, and is given by Eq. (9). The procedure leading up to Eq. (9) can be understood intuitively as follows. The integral I varies parabolically with the gradient coefficient

(d_i'). This is true for each of the gradients $g_i'(\vec{x})$, regardless of the detailed spatial variation of the gradient. Three points are required to fully specify the parabola, and from them the minimum may be located. This is precisely what Eq. (9) provides. In the absence of noise, each orthogonalized gradient could be optimized and never adjusted again. In practice, small errors are difficult to detect in the presence of larger errors, and so a small number of iterations may be required. A number of other practical issues inform the implementation of a useful algorithm, and the strategy we have chosen to follow is detailed below.

3. Implementation

We begin with the first shim of the set: $g_0(\vec{x})$ (which is typically Z1, the linear gradient along the direction of the main field), and first measure $I_0^{\max}$ and $I_0^{\min}$, values of the integral of Eq. (4) when the shim current is set to the maximum and minimum values available. This is done simply to choose an appropriate value for d_0, which is chosen so that I_0^+ and I_0^- will vary from I^0 by about a factor of two, specifically we choose:

$$d_i = \sqrt{I_i^0/A_i}, \tag{19}$$

where A_i is the coefficient of the quadratic term in the fit of I to $I = A_i d_i^2 + B_i d_i + C_i$, and is given by

$$A_i = \frac{(I_i^0 - I_i^{\max})d_i^{\min} - (I_i^0 - I_i^{\min})d_i^{\max}}{d_i^{\max}(d_i^{\min})^2 - (d_i^{\max})^2 d^{\min}}, \tag{20}$$

in which $d_i^{\min}$ and $d_i^{\max}$ are the maximum and minimum settings of the shim gradient, and $I_i^{\max}$ and $I_i^{\min}$ are the integrals measured with those extreme settings.

With d_0 chosen, measurements are made for I_0^+ and I_0^-. Eq. (15) then gives the normalization for this shim. Eq. (9) immediately gives the correction current.

For each of the subsequent shims, i, in turn, we begin the same way by measuring I_i^0, $I_i^{\max}$ and $I_i^{\min}$ to select d_i, but then proceed with four measurements for I_{ij}^{++}, I_{ij}^{+-}, I_{ij}^{-+}, and I_{ij}^- for each $j < i$ where an interaction between shims i and j is expected. Eq. (12) gives the $C_{ij'}$'s.

The temporary composite shim of Eq. (17) is then constructed. At this point, the composite shim is set to its extreme limits to select a d_i' value in the same fashion d_i was selected above, and finally I_i^+ and I_i^- are measured to provide N_i and the correction a_i via Eqs. (18) and (9).

It may be noticed immediately that difficulties may arise with the procedure as outlined above. In particular, the suggested value of d_i given by Eq. (19) may be beyond the limits of the power supply for that gradient. In our implementation, as long as d_i can be set to at least 50% of the target value, the algorithm proceeds. The d_j's applied to other shims to be orthogonalized at this point are then also reduced by a similar factor so that the d_i and d_j values produce similar perturbations in I. If d_i cannot be set to at least 50% of the target, this shim is skipped for now and saved until all of the other shims have been optimized.

Another immediate question is how best to deal with the situation that arises when a correction determined by Eq. (9) would set one or more of the shims making up the composite shim beyond the limits of its power supply. Here we have chosen to set the composite shim consistently, by reducing a_i as necessary so that the shim set can produce the modified correction for all the components of the composite shim.

In principle, the order in which the various shims are attempted should not affect the final result. However, the choice of ordering may affect the number of iterations required for the algorithm to converge to the optimal shim currents. While we have not investigated the effects of the ordering in any detail, we have found that ordering the shims according to the magnitude of the inhomogeneity they can produce generally appears to make the algorithm run smoothly. Following this prescription, the low order shims are typically first, followed by increasingly higher-order shims. This may not, however be the optimal ordering. If the main source of higher-order field inhomogeneities is the error terms of the lower order shim coils, it might be advantageous to begin with the higher order shims so that the lower-order shims are orthogonalized last. Then the (large) corrections applied by the lower-order shims will have the higher-order errors associated with them corrected immediately, rather than by higher order corrections done later. This ordering might give "purer" gradients, in the sense of the spherical harmonics, however it is not at all obvious that it would improve the convergence. A compelling difficulty with this strategy is that the higher order shims often produce field perturbations that only negligibly affect the lineshape, at least until after the low order shims have been corrected. We believe the algorithm will converge most rapidly if the shims can be ordered so that the largest corrections, as measured by their impact on the second moment, are made first.

3.1. Calculating the integrals

A key practical issue is how best to decide how much of the spectrum to include in the integral of Eq. (4). We began with an extremely simple minded approach that finds the point in the spectrum with the highest intensity, and then searches to each side until the intensity has dropped by a fixed factor from the peak. A factor of 75 turned out to work very well and was used in all of the examples shown below.

3.2. Z0

The values of the integrals of Eq. (4) clearly depend on the position of the rf carrier relative to the NMR line in the spectrum. To carry out the algorithm as described, the Z0 gradient (constant field offset) must be included in the set of shims and included in the orthogonalization. Because each of the shims can be expected to produce some Z0, including Z0 would require that each of the other shims include it in

its orthogonalization. As an alternative that saves a significant amount of time, we chose to omit Z0 from the set of shims, and instead replaced all the integrals with the second moment of the spectrum, calculated according to

$$I = \frac{\int S(\omega)\omega^2 \,\mathrm{d}\omega}{\int S(\omega)\,\mathrm{d}\omega} - \left(\frac{\int S(\omega)\omega \,\mathrm{d}\omega}{\int S(\omega)\,\mathrm{d}\omega}\right)^2. \quad (21)$$

3.3. Iterations

Because of the finite signal to noise of any real experiment, simply orthogonalizing and optimizing each gradient once will not always be sufficient. After the orthogonalization has been performed once however, a great deal of information about the shim set is known, and may be used to dramatically speed up subsequent iterations. In the demonstrations below, we employ three different procedures. The first procedure is the orthogonalization and optimization procedure described above. This procedure requires $7N + 4M$ fid's, where N is the number of gradients, and M is the number of suspected interacting pairs included. With no assumptions made on which pairs of coils interact, M would be $N(N-1)/2$.

The second type of iteration is a "quick optimization" that simply measures I^+, I^-, and I^0 for each of the previously orthogonalized composite shims and calculates the position of the bottom of the parabola. This type of iteration requires only $3N$ fid's and is therefore much faster than the full procedure.

The third type of iteration is a "re-orthogonalization," in which the gradients are orthogonalized, but rather than starting with the single gradients supplied by the coils, the starting functions are the composite shims determined previously. In this type of iteration, appropriate d values can be determined without setting the shims to their extreme values, and the number of fid's required is reduced to $3N + 4M$. A further reduction might also be available here, as measurements of the overlap between shims have already been made. Pairs that were included unnecessarily could be automatically eliminated at this point, reducing the value of M.

4. Demonstration

The method was implemented for Varian Inova (Varian Inc., Palo Alta, CA) spectrometers running Vnmr. The implementation consists of two relatively short (<1800 lines total) programs written in C along with a simple macro that runs within Vnmr to allow control of the spectrometer by the C code. The algorithm was primarily developed on a system based on a 400 MHz (8.4 T) wide-bore (89 mm) magnet. After initial testing of the algorithm, the high-resolution solution probe that had been in place was replaced with a second high-resolution probe (Varian inverse detection triple resonance), and all of the shim currents set to 0. A doped sample containing 0.1 mg/ml $GdCl_3$, 0.1% DSS in 1% H_2O in D_2O (Isotec sample #0190185501, Isotec, Miamisburg, OH) was inserted into the probe. This initial sample has a relatively short T_2 of 0.16 s, which would result in an NMR spectrum in a perfectly homogeneous magnet of a Lorentzian line having full-width at half maximum (FWHM) of 2 Hz. We begin with this sample for two reasons: the first is that spectra may be acquired from it relatively quickly ($T_1 = T_2$ for this sample) so the initial orthogonalization can be achieved rapidly, and secondly because adjustments in the spectrum that might appear only very close to the baseline away from the center of the peak, which might be excluded from the integration, are convolved with the sample linewidth. This ensures that the effects of these changes are visible to the algorithm. Of course, a sharper signal could be used and the fid simply convolved with a broadening function, but then the advantage of the rapid relaxation is forfeited.

The initial spectrum with this sample (with all shims set to 0) is illustrated in Fig. 1a, and spans a width of more than 270 Hz. Much of this width is due to a large error in Z1, and the FWHM can be reduced to <20 Hz by adjusting this one shim.

The full orthogonalization procedure was then carried out with $N = 23$ gradients and $M = 74$ of the possible 253 pair-wise interactions included. This first orthogonalization was completed in about 50 min. This is about a factor of eight longer than would be required to simply acquire the necessary number of fully relaxed spectra from this sample. Extra time is included to allow the gradients to settle before acquiring new spectra. The result of this first iteration is shown in Fig. 1b.

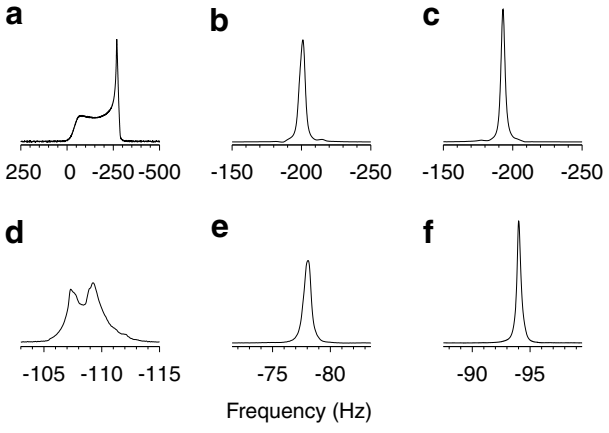

Fig. 1. Spectra acquired during shimming of a 400 MHz wide-bore (89 mm) magnet. (a) Initial spectrum with all shim currents set to 0. (b) Result of a single iteration of the orthogonalization/optimization procedure. (c) Same as (b), with one additional optimization pass. (d) Same shim current settings as in (c), but with the "lineshape" sample replacing the doped sample. (e) Same as (d) after one re-orthogonalization/optimization pass. (f) Final spectrum after one additional optimization. The horizontal scales in (b) and (c) are identical as are those in (d), (e) and (f). The vertical scales of (b) and (c) are identical, as are those of (e) and (f). Note that the horizontal scale in (d), (e) and (f) is magnified by approximately a factor eight compared to that in (b) and (c).

The specifications for resolution of the high-resolution NMR probes used here are given in terms of the full width of the peak at three locations: at half-height, at 0.55% of the peak (the height of the ^{13}C satellite peaks in the ^{1}H spectrum), and at 0.11% of the peak. For this probe, the manufacturer's resolution specification is given by 0.65/10/20 Hz. In all cases reported in this work, the line widths were determined using the Varian-supplied "res" macro in Vnmr. After the first iteration, the peak widths are 4.9/38/87 Hz, but it must be remembered that due to the broad line of the sample, the best shim settings possible would produce a spectrum having linewidths of 2/27/60 Hz (assuming a Lorentzian lineshape).

A second iteration was then performed, this time a "quick iteration" as described above in which the orthogonalization was omitted. This second iteration was therefore much faster, requiring only 3 fid's per shim, done here in about 10 min. The result of this iteration is shown in Fig. 1c, and has linewidths of 3.4/37/63 Hz.

At this point, the sample was changed to a "lineshape" sample, Isotec #968120-89, consisting of 1% chloroform in 99% deuteroacetone. This sample was found to have a T_1 of approximately 80 s, and so, assuming $T_1 \approx T_2$, the natural linewidth of the sample contributes negligibly to the observed lineshape. After simply switching the samples, the linewidth immediately improved to that shown in Fig. 1d, having widths of 3.2/9.6/11.7 Hz, showing that much of the width near the baseline of the earlier spectra was indeed due to the sample itself and not due to field inhomogeneity. A "re-orthogonalization" was then carried out with 395 fid's in 135 min. At this point, the linewidths were 0.67/5.0/9.2 Hz, and the spectrum appeared as shown in Fig. 1e. This spectrum very nearly meets the probe resolution specification, and is only a little too broad at half-maximum. One additional quick iteration requiring 94 fid's in 42 min resulted in the final spectrum shown in Figs. 1f and 2, having linewidths of 0.39/3.4/6.4 Hz. This is substantially better than the spectrum produced by the vendor's installation engineer using a robotic field mapping system (0.52/5.7/9.8 Hz), and much better than the installation specification.

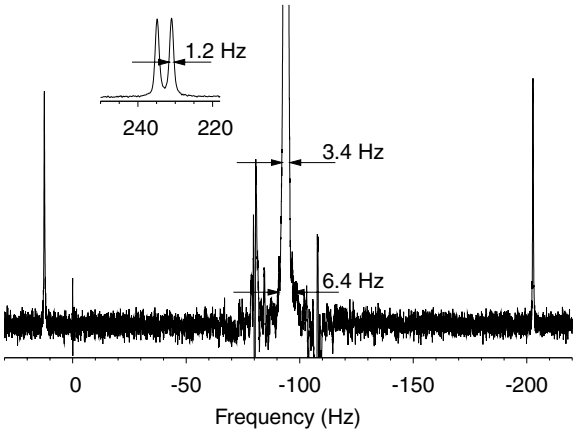

Fig. 2. Final spectrum from the 400 MHz spectrometer expanded to show the ^{13}C satellites and linewidths at 0.55% and 0.11% of the peak height. 0.55% of the peak height does not line up with the height of the ^{13}C satellites, presumably because the spectrum is not fully relaxed and the satellites can be expected to relax more quickly than the main peak. The noise in the baseline surrounding the peak is most likely due to vibrations in the room. The inset shows the doublet at 5.4 ppm in a 2 mM sucrose sample.

All spectra so far were collected without sample rotation. This probe has a second set of specifications for spinning samples: 0.45/5/10 Hz. The shim currents determined here beat the spinning specification without spinning the sample. Simply spinning the sample, without any further adjustments, resulted in further reductions in the linewidths to 0.34/2.3/4.9 Hz. The entire procedure required a total of 4 h of spectrometer time to go from completely unknown shim settings to a fully optimized set of 23 shim current values. Operator intervention was only required to change samples once.

Another commonly used criteria for judging field homogeneity is the splitting of the 5.4 ppm peak of sucrose in water. Following the shimming procedure described above, a sample containing 2 mM sucrose in 90% H_2O, 10% D_2O was inserted (Isotec sample #0190185505). At this point the probe was detuned several MHz away from resonance to minimize radiation damping effects, and the algorithm allowed to optimize the Z1–Z3, X1 and Y1 gradients using the water peak. The probe was then retuned and spectra acquired using a 2 s long presaturation pulse to suppress the H_2O peak. The 5.4 ppm sucrose doublet is shown in the inset in Fig. 2. Each peak has a full-width at half maximum of approximately 1.2 Hz.

A similar procedure was undertaken on a Varian Inova 600 spectrometer equipped with a cryoprobe and a shim set having 39 separate gradients. In this case, 207 of the 741 possible pair-wise correlations were selected to be included in the orthogonalization. After the first orthogonalization with the doped sample, the FWHM was 11 Hz (having started at ~150 Hz). After two more quick-optimization iterations linewidths of 4.2/37/80 were observed. At this point the sample was changed to the high-resolution "lineshape" sample, and a single reorthogonalization run was done, resulting in linewidths of 2.4/15/23 Hz. After four more optimization passes, the linewidth specification (1/10/20 Hz) was achieved, with widths of 0.88/9.1/17 Hz. The entire procedure thus far had taken approximately 12 h, most of which was completely unattended.

After one more re-orthogonalization and several further quick optimization iterations, the best linewidths obtained were 0.41/4.8/11.4 Hz (shown in Fig. 3), compared to the installation engineer's (with field mapping) results of 0.53/6.3/12.2 Hz. The whole procedure was carried out in about a day. For comparison, when this spectrometer system was reinstalled, it required about 12 h for the vendor's installation engineer to map the field and the 39 room temperature shims.

5. Discussion

The algorithm presented requires the selection of a number of parameters and sub-procedures. We have not

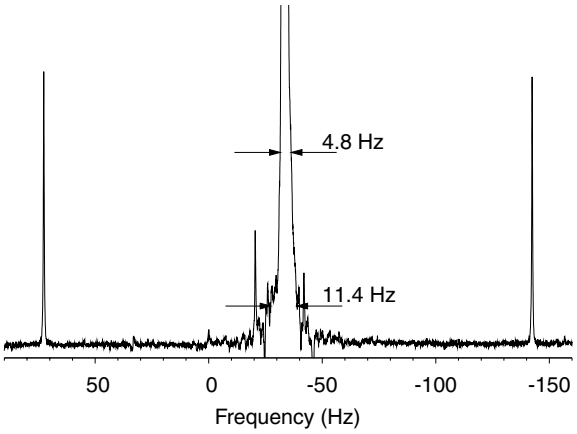

Fig. 3. Final spectrum from the 600 MHz spectrometer with cryoprobe, expanded to show the ^{13}C satellites and linewidths at 0.055% and 0.011% of the peak height. The noise in the baseline surrounding the peak is most likely due to vibrations in the room.

attempted an exhaustive study of the effect of different choices, and expect that further experimentation will improve the current implementation.

The choice of samples is one area that could certainly be optimized to speed up the process. The 2 Hz doped sample is a useful first sample for reasons explained above, but the lineshape sample with its very long T_1 is almost certainly not an optimal choice for the next stage. This lineshape sample was not allowed to relax fully (~20 s between acquisitions), however a second sample having a T_1 in the range of 5–10 s would allow more rapid collection of spectra in the later parts of the procedure. In the examples above, spectra were acquired with at least 5 s between scans to allow the gradient power supplies to fully settle. In practice, a significant fraction of this 5 s is taken up by communication between the various program modules and overhead in setting up and starting the spectrometer and so could not be eliminated easily. If the gradient supplies could be relied upon to settle quickly as can be expected in many modern systems, and the spectrometer start-times improved, the initial orthogonalization could be carried out extremely rapidly, in as little as 5 min for the 23 coil shim set, or about 10 min for the 39 coil shim set. Ultimately, the time taken will be limited by the final resolution desired, as the final sample linewidth should be no greater than the resolution desired.

Choices are also clearly required in the selection of which shim pairs should be included in the orthogonalization. In the case of the 400 MHz system, it appears that all of the important pairs were included as the algorithm quickly converged, but in the case of the 600, several iterations were required, and it appears that some non-negligible interactions were omitted. There is an important distinction between this outcome and what could happen if more important pairs had been omitted. If even one "important enough" pair was missing, iterating would not improve the homogeneity. The meaning of "important enough" is well defined: the coefficients $C_{ij'}/N_i$ in Eq. (13) indicate the contribution of one composite shim (j) in a subsequent physical shim (i). If a coefficient that would be greater than 1 (or a number of coefficients such that $\sum_j (C_{ij'}/N_i)^2 > 1$) is missing, then iterating will not allow the procedure to converge since the adjustments necessary to optimize the homogeneity will require making the second moment temporarily larger. As a specific example, if adjusting the current in the Z2 coil actually produces more (in the sense of change in the second moment of the line) Z1 than Z2, then simple iterations will never converge. As soon as enough Z1 is removed from Z2 so that the Z2 adjustment produces more Z2 than Z1, then iterations will be successful.

The pairs included in the above examples were selected simply by intuition, and so it is perhaps not surprising that some non-negligible interactions were missing. All of the experimentation on the 600 MHz system was performed over one weekend with no advance preparation. It is very likely that the convergence on this system could be improved. It should be possible to analyze correlations in the corrections made during the iterations on this system to add the shim pairs that are missing, and at the same time delete pairs that were included unnecessarily.

Which pairs of shims interact, and the interaction coefficients (the C_{ij}'s) will depend on probe geometry. We do not expect that the orthogonalization found with a 5 mm probe would provide the same composite shims as an orthogonalization of the same system with a 15 mm probe for example. In fact, even a significant change in the height of the sample can be expected to change the makeup of some of the composite shims dramatically.

The above examples began with essentially no knowledge of the interactions amongst the shims. For more routine use, it is expected that which shims interact, and which of those interactions can be expected to change with sample, could be well characterized, and only the portions of the algorithm necessary be run. This would allow a very much faster implementation for sample to sample changes.

Most of the samples used in the examples above are dilute enough in ^1H content that radiation damping causes no observable line broadening. For samples abundant in ^1H spins, where radiation damping can severely broaden the solvent spectra, simply detuning the probe well away from resonance is one option for using an abundant solvent peak, as was done above with the sucrose/H_2O sample. If the solvent T_2 does not allow a sufficiently sharp line, solvent suppression could be used, and the algorithm set to employ some more suitable peak. In this case, it might be necessary to include Z0 in the optimization so that adjustments to the other shims do not alter the resonance frequency of the solvent and interfere with the solvent suppression.

The simple requirements of this method should allow it to find broad application. With further optimizations such as in sample selection and pair list inclusion as discussed above, we expect the method to be significantly faster, and could become a preferred method of shimming new

probe installations, especially on older spectrometers. The method may be suited to areas such as imaging and *in vivo* spectroscopy, where there are often fewer relevant shim gradients, and the resolution required (for imaging anyway) may not be as high. For localized spectroscopy, the simple one pulse acquisition could be replaced by a volume selective sequence, effectively truncating $W(\vec{x})$ outside of the region of interest. In this context the method could be extremely fast, and would orthogonalize the available shims only in the region of interest, taking into account the particularities of the sample.

The source code used in the examples shown above has been made available for download at http://www.phas.ubc.ca/~michal/Shimming. The code has been published under the GNU General Public License and may be modified and redistributed. It is our hope that this code will provide a starting point for spectrometer vendors to incorporate the method into their respective spectrometer control software systems.

6. Conclusions

In conclusion, we have presented a method for optimizing magnetic field homogeneity. The method has few hardware requirements and produces orthogonal composite shims for the sample at hand. We expect it will find application in many areas of magnetic resonance. The method has been demonstrated to shim high-resolution solution NMR probes to homogeneities much better than achieved through more elaborate field mapping techniques. The method may be well suited to imaging and *in vivo* spectroscopy, where the number of relevant shim coils may be much smaller and the orthogonalization of the relevant shims may change dramatically from sample to sample.

Acknowledgments

This work was supported by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada. The Varian spectrometers were purchased with the assistance of the Canada Foundation for Innovation and the BC Knowledge Development Fund. The author thanks Mark Okon for helpful conversations and Lawrence McIntosh for access to the Varian 600.

References

[1] F. Romeo, D.I. Hoult, Magnet field profiling: analysis and correcting coil design, Magn. Reson. Med. 1 (1984) 44–65.
[2] M.G. Prammer, J.C. Haselgrove, M. Shinnar, J.S. Leigh, A new approach to automatic shimming, J. Magn. Reson. 77 (1988) 40–52.
[3] P.C.M. Van Zijl, S. Sukumar, M. Johnson, P. Webb, R.E. Hurd, Optimized shimming for high-resolution NMR using three-dimensional image-based field mapping, J. Magn. Reson. A 111 (1994) 203–207.
[4] G. Stran, Linear Algebra and its Applications, third ed., Harcourt Brace Jovanovich Inc., San Diego, 1988.