Available online at www.sciencedirect.com
ScienceDirect
Journal of Magnetic Resonance 182 (2006) 38–48
JMR Journal of Magnetic Resonance
www.elsevier.com/locate/jmr

Gradient shimming with spectrum optimisation

Markus Weiger *, Thomas Speck, Michael Fey

Bruker Biospin AG, Industriestrasse 26, CH-8117 Faellanden, Switzerland

Received 18 April 2006; revised 6 June 2006
Available online 27 June 2006

Abstract

Shimming, i.e. homogenising the unavoidable distortion of the static magnetic field B_0 in NMR spectroscopy, is still an annoying, time consuming task. Although compared with conventional manual or computerised search methods gradient shimming initiated a new era in terms of operation and efficiency, there remain aspects that inhibit fully automated shimming with a result of guaranteed quality. The major reason for this limitation is that the judgement of the quality of the B_0 homogeneity takes place in the spatial domain, although the actual objects of interest are the lines in the spectral domain. In this work, this restriction is removed by the introduction of a new framework for gradient shimming that enables to directly aim at the spectrum quality. Based on the mapped B_0 field shimming is simulated and spectra are calculated for the virtual residual inhomogeneity. Using a suitable criterion to judge the spectrum quality an optimisation is performed, thus providing the predicted optimum spectrum and the corresponding residual B_0 field. This target field is then aimed at during the real, iterative shimming procedure. For the widely applicable case of optimising the shape of a single line a powerful quality criterion was developed using an envelope of the calculated lineshape spectrum. The whole procedure is demonstrated for adjusting the on-axis shim functions based on one-dimensional field map data and for both on- and off-axis shimming using three-dimensional data. The results are verified with ^{1}H NMR spectra acquired on standard NMR test samples.
© 2006 Elsevier Inc. All rights reserved.

Keywords: Gradient shimming; Field mapping; Spectrum optimisation; Lineshape; Envelope

1. Introduction

For optimum sensitivity and resolution in nuclear magnetic resonance (NMR) spectroscopy the static magnetic field B_0 throughout the sample volume seen by the radio-frequency (RF) coil should be as homogenous as possible. However, residual inhomogeneity in the B_0 field due to the magnet, probe, and sample very often add up to spatial field variations that would deteriorate the quality of an NMR spectrum to an unacceptable degree. Therefore, NMR spectrometers are commonly equipped with a set of shim coils that serve to reduce the remaining B_0 inhomogeneities. Each shim coil is designed to generate an individual static field distribution per unit current and is connected to a current source that can be driven independently. In a mathematical sense, the shim coils provide a set of basis functions, the shim functions, for the generation of a desired field distribution in the visible sample volume.

Shimming is now the procedure of finding the setting of the shim currents that compensates the B_0 field inhomogeneities the best. Mathematically, this is an optimisation problem defined by the B_0 field variations, the set of shim functions, and by the quality criterion judging the field homogeneity. While the first two aspects are predetermined for a certain setup, the criterion can be chosen such as to represent the desired spectrum quality. In most cases the solution of this optimisation task is not trivial, since even by careful design it can usually not be avoided that the shim functions interact with each other to a certain degree in the sensitive volume [1]. Hence, it is not sufficient to optimise each shim current independently and only once, and an appropriate optimisation algorithm has to be chosen.

* Corresponding author. Fax: +41 44 825 96 38.
E-mail address: markus.weiger@bruker-biospin.ch (M. Weiger).

1090-7807/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmr.2006.06.006

This leads to the class of iterative *search algorithms*, where after modification of the shim setting the value of the criterion is determined, and this procedure is repeated until an optimum has been found (Fig. 1). The total duration of a completed search is given by the number of iterations performed N_{iter} times the duration per iteration T_{iter}, the response time. Such a search can be performed either manually by the spectrometer user or by algorithms implemented on a computer, e.g. the simplex algorithm [1–3]. Of course, the latter approach is preferable in terms of convenience and cost; however, quantitative criteria are required, ideally in the form of a single numeric value.

Quality criteria are usually derived from NMR spectra obtained for the nucleus of interest or the lock nucleus. Ideally, various properties should be determined from the spectral lines to judge the B_0 homogeneity. However, very often only the maximum height is used, thus neglecting the shape of the lines. The choice of the quality criterion not only determines the achievable spectrum quality but also the efficiency of the optimisation procedure: T_{iter} is often dominated by the measurement time for the data used to derive the criterion, whereas N_{iter} depends on the optimisation algorithm that can be selected due to the given criterion.

Overall, shimming with search algorithms suffers from temporal inefficiency and from the danger of finding a local rather than a global optimum in the multi-dimensional parameter space spanned by the shim functions. Hence it is often tedious and a matter of hours.

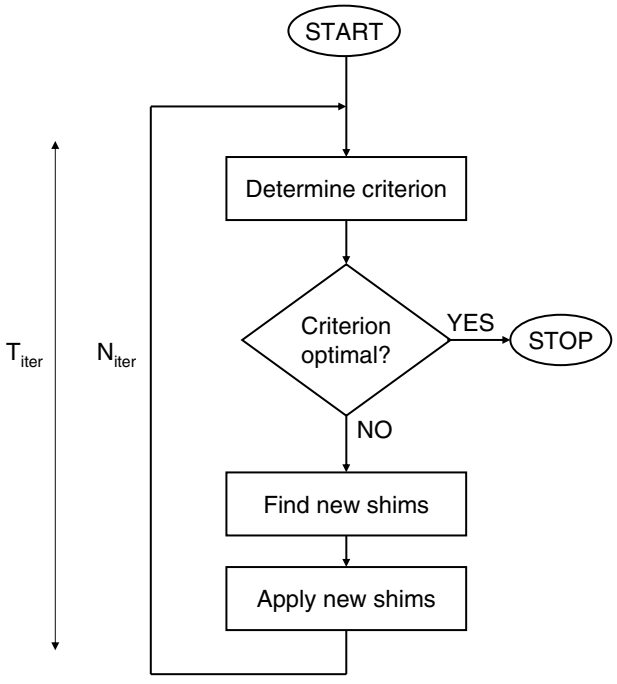

Fig. 1. Search algorithm for iterative shimming. For the initial state the criterion is determined and only in case of an already optimal result the procedure is left immediately. Otherwise the search loop is entered, new shim values are chosen and applied, and the criterion is determined again. This procedure of duration T_{iter} is repeated N_{iter} times until the criterion cannot be further improved.

With the introduction of gradient shimming to NMR spectroscopy the efficiency of shimming was increased dramatically [4–6]. By means of imaging techniques applying magnetic field gradients, the B_0 inhomogeneities are resolved spatially, hence their judgement takes place in the spatial rather than in the spectral domain. The crucial advantage is now that by knowledge of the spatial dependence of the available shim functions, the required shim setting can be derived in a single step based on the choice of a suitable criterion. The latter can e.g. be a sum-of-squares of the field deviations which enables direct calculation of the optimum solution via inversion of a system of linear equations. Other criteria as e.g. proposed in Refs. [1,7,8] require iterative optimisation analogue to that shown in Fig. 1, however, without application of any measurements but only by means of calculations.

Although in principle with gradient shimming the predicted homogeneity improvement should be realised in a single step, there are reasons to repeat the procedure several times. On the one hand, noise and systematic errors in the field map but also in the shim function data lead to imperfect results that can be improved from iteration to iteration. Moreover, after each iteration the quality of the field map data improves with the better homogeneity. Hence, iterative gradient shimming also follows the scheme of Fig. 1, however, for very different reasons than the search algorithms. Therefore, the number of necessary iterations N_{iter} is quite small, usually less than 10.

A remaining disadvantage of gradient shimming is due to the fact that the criteria proposed so far are determined in the spatial domain. Hence the criteria provide only an indirect connection with the spectrum and cannot guarantee optimum spectrum quality in terms of sensitivity and resolution. In this paper we propose to come back to using the spectrum of interest as a basis for the criterion, however, in the different context of gradient shimming. In this way, the efficiency of gradient shimming can be combined with the ideal quality criterion.

2. Methods

The basic idea of gradient shimming with spectrum optimisation is to simulate the shimming procedure based on initially acquired B_0 field map data. The shim quality is judged not with a measured but a calculated spectrum. The best simulated solution found is then actually realised. The method of gradient shimming with spectrum optimisation is explained while making use of a lineshape spectrum, i.e. a spectrum consisting of a single line. It is described how a spectrum is calculated based on B_0 field map and other experimental data, a quality criterion is derived from the spectrum, the quality criterion is used for optimisation of the spectrum, and how the spectrum is actually realised with a shimming procedure.

2.1. Spectrum calculation

In order to keep the spectrum calculation and evaluation simple, a trivial substance is chosen with a gyromagnetic ratio γ, a chemical shift of 0 ppm, no coupling, and a transverse relaxation time T_2. Furthermore, the most simple experiment is selected, consisting of an excitation pulse of flip angle α followed by the acquisition of a free induction decay. B_0 denotes the reference value for the z component of the static magnetic field, and setting the receiver frequency to $-\gamma B_0$ provides spectral intensities as a function of the rotating frame frequency Ω. For spatially constant B_0 the resulting spectrum $I(\Omega)$ of such an experiment and the corresponding calculation is, apart from an arbitrary scaling constant, a single Lorentz line of width $\lambda = 1/(\pi T_2)$ centred around zero frequency

$$I(\Omega) = L(\Omega, \lambda). \tag{1}$$

Note that with the chosen simple virtual setup it is sufficient to consider the phased real part of the spectrum, i.e. the absorption component of the Lorentzian.

A vector $\mathbf{r}$ of dimension n with $n = 1$, 2, or 3 now describes a spatial location in the sample. For a certain location the deviation $\Delta B_0(\mathbf{r})$ from the reference field B_0 shifts the resonance frequency to $\Omega_0(\mathbf{r}) = -\gamma \Delta B_0(\mathbf{r})$. Moreover, the magnitude of the sensitivity of the RF coil $B_1(\mathbf{r})$ may have spatial variations. Here, B_1 is normalised in the sense that for $B_1 = 1$ the nominal flip angle α is actually achieved. Assuming a sample with a spatially constant spin density the contribution to the spectrum supplied by a volume element $\mathrm{d}^n r$ located at $\mathbf{r}$ is then

$$I(\Omega, \mathbf{r}) = L(\Omega - \Omega_0(\mathbf{r}))w(\mathbf{r})\mathrm{d}^n r \text{ with } w(\mathbf{r}) = \sin(B_1(\mathbf{r})\alpha)B_1(\mathbf{r}), \tag{2}$$

where the shifted Lorentz line is weighted due to the local B_1 during excitation as well as during acquisition. Note that for static spins the phase of the RF sensitivity does not play a role, hence all quantities in Eq. (2) are real numbers. The parameter λ is omitted for convenience.

The full spectrum is calculated as the integral over the sample volume

$$I(\Omega) = \int_{\mathbf{r}} L(\Omega - \Omega_0(\mathbf{r}))w(\mathbf{r})\mathrm{d}^n r. \tag{3}$$

Expanding the expression with a delta function

$$I(\Omega) = \int_{\mathbf{r}} \int_{\Omega\prime} L(\Omega - \Omega')\delta(\Omega' - \Omega_0(\mathbf{r}))\mathrm{d}\Omega' w(\mathbf{r})\mathrm{d}^n r \tag{4}$$

enables rewriting it as a product of two functions of frequency

$$\begin{aligned} I(\Omega) &= \int_{\Omega'} L(\Omega - \Omega') \int_{\mathbf{r}} \delta(\Omega' - \Omega_0(\mathbf{r}))w(\mathbf{r})\,\mathrm{d}^n r\,\mathrm{d}\Omega' \\ &= \int_{\Omega\prime} L(\Omega - \Omega')H(\Omega')\,\mathrm{d}\Omega\prime, \end{aligned} \tag{5}$$

where

$$H(\Omega) = \int_{\mathbf{r}} \delta(\Omega - \Omega_0(\mathbf{r}))w(\mathbf{r})\,\mathrm{d}^n r \tag{6}$$

is the weighted histogram of the spatial variation of resonance frequencies. Hence, the spectrum in an inhomogenous B_0 field is the convolution of the natural, undisturbed line of Eq. (1) with the histogram of the inhomogeneities given by Eq. (6):

$$I(\Omega) = L(\Omega, \lambda) \otimes H(\Omega). \tag{7}$$

Practical spectrum calculation is based on non-continuous data creating various implications. Generally it can be said, that in order to limit computation time the number of points in a data set should be kept as small as possible while not degrading the usefulness of the final result. B_0 field maps and the corresponding frequency maps $\Omega_0(\mathbf{r})$ may have large steps in their values between adjacent data points causing discontinuities in the histogram. This can be avoided in a preparative procedure by interpolation of intermediate data points in order to bring down the maximum step size below a certain limit $\Delta\Omega_{\mathrm{M}}$. Interpolation of spatially equidistant data points based on the largest occurring frequency step is straightforward but can become time-consuming in particular for large, three-dimensional (3D) data sets. Better efficiency is obtained by local adaptation of the number of interpolated data points to just meet $\Delta\Omega_{\mathrm{M}}$. During this procedure care must be taken to maintain the weight of each volume element. Further parameters controlling the compromise between computational time and data quality are the bin size of the histogram $\Delta\Omega_{\mathrm{H}}$, and the frequency interval of the spectrum $\Delta\Omega_{\mathrm{S}}$. As the natural line width λ defines the size of the relevant frequency interval it can be used as a basis for selecting suitable values for the three described discretisation parameters.

2.2. Quality criterion

A powerful quality criterion must account for several aspects of a spectral line, such as e.g. height, width at half maximum height, width of the hump, shape, smoothness, and symmetry. The multiple aspects can be combined in some way to form a single criterion or enter a multi-objective optimisation algorithm. Moreover, on a level of already high quality there remains usually some freedom in choosing the "best" lineshape, which then becomes a matter of the needs of a specific NMR experiment or of the personal taste of the spectrometer user. Hence, the criterion should enable the user to control this remaining degree of freedom.

As a comprehensive solution to the above requirements, we propose to use an envelope of the spectrum as the basis for calculation of the quality criterion. First, the envelope is capable of intrinsically combining the individual criteria listed above as it reacts to different kinds of deviations from the ideal line shape in a similar way by simply becoming broader. In particular, this is also true for merely local

distortions and those at some distance from the main line. Second, a very useful single criterion can be derived from the envelope, which is the full width of the envelope at half maximum height of the spectrum (FWEHMS). Third, the choice of the shape of the envelope gives the user an intuitive means of influencing the shape of the spectral lines.

The chosen envelope shape should of course represent a sensible shape for a spectral line. On the other hand, the requested shape must also be a realistic goal for the given situation. Naturally, the basic shapes suitable for envelopes are Lorentz or Gauss curves. Different types of transitions between the two are provided by the Voigt or the Tsallis [9] curves, whereas the latter is easier to handle for practical calculations. The shape of the Tsallis curve is defined by a parameter $q > 1$, with the special cases of a Gauss function for $q \to 1$ and a Lorentz function for $q = 2$.

The calculation of the envelope is as follows: all envelope shapes are described by the three parameters position, width, and integral or amplitude, which need to be determined for a given spectral line under the condition of tightly enclosing the latter. Hence, the proposed calculation of the envelope $E(\Omega)$ is an optimisation itself using a penalty function

$$P = \sum_{\Omega} P(\Omega) \tag{8}$$

with

$$P(\Omega) = D^2(\Omega) \quad \text{for} \quad D(\Omega) \geqslant 0 \tag{8a}$$

$$P(\Omega) = (f|D(\Omega)|)^c \quad \text{for} \quad D(\Omega) < 0 \tag{8b}$$

and the difference between envelope and spectrum

$$D(\Omega) = E(\Omega) - I(\Omega). \tag{8c}$$

With this approach spectrum points below the envelope are penalised conventionally with the squared difference. For cases with the spectrum running above the envelope the parameters c and f serve to determine the coarse match and to fine tune the strictness of the envelope, respectively. Setting $c = 4$ was found working well in all cases. A suitable range for fine tuning the envelope strictness is $f = 100, \ldots, 500$, where a large value of f prevents spectral intensity in the side lobes of the spectrum at the price of less flexibility for shaping the main line. In this work, $f = 200$ was generally used. In order to find the envelope fitting best to the spectrum the three free envelope parameters are varied for minimisation of the penalty function, e.g. using the simplex algorithm [2,3].

As an example Fig. 2 shows an artificial spectral line enclosed by an envelope of Lorentzian shape. When using the criterion FWEHMS for an optimisation the buckle in the line would be penalised by the increased width *and* height of the envelope.

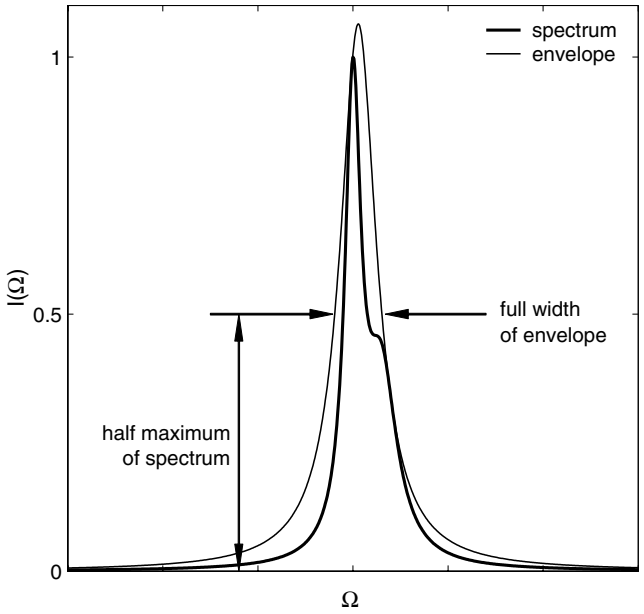

Fig. 2. Envelope of a spectral line. The artificial line is fully enclosed by the envelope of Lorentzian shape calculated by means of the penalty function in Eq. (8). As a criterion for judging the quality of the line the full width of the envelope at half maximum height of the spectrum (FWEHMS) is proposed as indicated by the arrows. Thus, in the shown example the buckle in the line is penalised by the increased width *and* height of the envelope.

2.3. Optimisation

The spectrum optimisation performed in complete analogy to Fig. 1 is based on a single measured field map $\Delta B_{0,\mathrm{M}}(\mathbf{r})$. A set of shim values a_i is chosen and the residual field map expected after application of the shim functions $S_i(\mathbf{r})$ is calculated according to

$$\Delta B_{0,\mathrm{R}}(\mathbf{r}) = \Delta B_{0,\mathrm{M}}(\mathbf{r}) + \sum_i a_i S_i(\mathbf{r}). \tag{9}$$

For this field map the spectrum and the quality criterion are determined as described above. A search algorithm is now required for optimisation of a single criterion by variation of a set of parameters, the shim values. In order to save computation time high efficiency is desired, i.e. the preferably global optimum should be found with as little trials as possible.

Of course, one of the common multi-parameter optimisation algorithms can be used. However, the application of such a method may require a large number of trials, in particular if a large set of shim functions is available. Therefore in this work a different solution is chosen. First, it is assumed that some kind of least-squares fitting of $\Delta B_0(\mathbf{r})$ with the shim functions is generally a sensible approach to minimise inhomogeneity and therefore obtain a narrow line. Second, knowledge about the specific geometrical situation of high-resolution liquid state NMR is utilised to introduce a weighting to the fitting procedure. With the usual, cylindrical sample geometry the axial ends of the sensitive range play a particular role. On the one hand, in these regions the B_0 field is very often disturbed by the sample ends and by reduced magnet homogeneity. On the other hand, the B_1 sensitivity drops down with a steepness characteristic for the individual RF coil. Hence, for

improved homogeneity in the sample centre it is often a good compromise to accept somewhat larger inhomogeneities in the border regions. This is what is usually done more or less explicitly with whatsoever shimming method. Based on these assumptions a spatial weighting function $W(\mathbf{r}, k)$ is defined where the free parameter k serves to scale the weighting of the different regions. An example for B_1-based weighting is $W(\mathbf{r}) = [B_1(\mathbf{r})]^k$. The weighting function for a certain k can then be applied during least-squares minimisation of $\Delta B_{0,\mathrm{R}}(\mathbf{r})$, thus providing the set of shim values associated with the used k-value. In this way the multi-parameter search is reduced to a single-parameter search, as the optimisation is now performed for the parameter k. Thus, only a subset of the possible solutions is tested, assuming that it approaches the optimum sufficiently close.

The output of the optimisation is the expected spectrum, the shim values a_i to be applied, and the corresponding residual field map $\Delta B_{0,\mathrm{R}}(\mathbf{r})$ according to Eq. (9).

2.4. Shimming procedure

The role of the shimming procedure is to actually realise the theoretically predicted spectrum quality. This task is equivalent to creating a field inhomogeneity identical with the corresponding residual field, making the latter the target field $\Delta B_{0,\mathrm{T}}(\mathbf{r}) = \Delta B_{0,\mathrm{R}}(\mathbf{r})$ of the shimming procedure. In principle, all information is available to obtain the target field directly by application of the associated shim values. However, as with conventional gradient shimming various sources of error usually inhibit the realisation of the target at the required accuracy in a single step and an iterative approach is applied. After setting the shim values a new field map is acquired, and the next shim values are calculated by least-squares minimisation of the difference of this map with the target map

$$\Delta B_{0,\mathrm{Diff}}(\mathbf{r}) = \Delta B_{0,\mathrm{M}}(\mathbf{r}) - \Delta B_{0,\mathrm{T}}(\mathbf{r}). \tag{10}$$

Also during this fit procedure it can be useful to apply some local weighting to account for a lower importance of certain regions. The procedure is repeated until the (weighted) standard deviation σ of the difference has converged or fallen below a given absolute threshold.

To account for poor initial B_0 homogeneity it is advantageous to repeat the whole procedure, including the determination of the target field during the spectrum optimisation, thus benefiting from the generally improved data quality during shimming. Consequently, the full iterative shimming procedure as shown in Fig. 3 consist of two nested loops, where the inner loop serves for error correction while the outer one takes advantage of improved B_0 homogeneity.

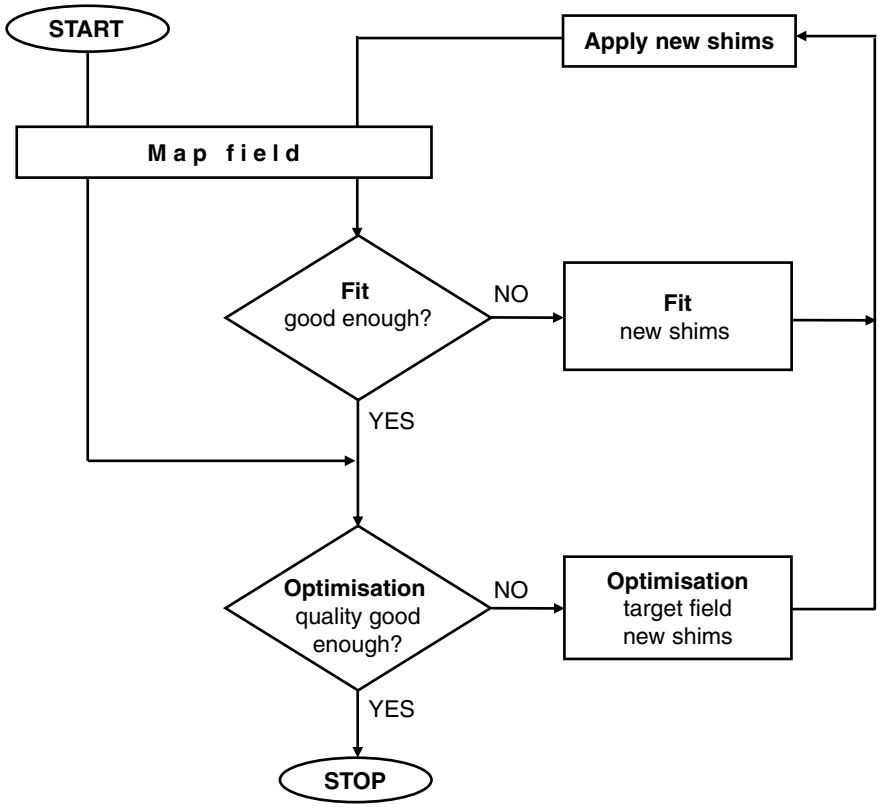

Fig. 3. Iterative gradient shimming procedure consisting of two nested loops. While passing the outer *optimisation loop* the spectrum optimisation takes place providing the target field, and with multiple runs through it advantage is taken of an improved B_0 homogeneity. The inner *fit loop* serves for error correction during realisation of the target map. Both loops are performed iteratively until the desired termination criteria are fulfilled. Each iteration requires acquisition of a B_0 field map, execution of the respective minimisation algorithm, and setting of the newly determined shim values.

2.5. Experimental details

Experiments were performed at 300 and 800 MHz on Bruker AVANCE spectrometers operated under TopSpin® 1.3. The standard bore magnets were equipped with shim systems BOSS1 (20 shim functions) and BOSS2 (34 shim functions), respectively, both providing axially symmetric (on-axis) polynomial shim functions up to the sixth order. At 300 MHz a 5 mm probe was used with an outer coil for ^{1}H and the ^{2}H lock (BBO). At 800 MHz a 5 mm probe was used with ^{1}H on the inner coil (TXI). The required ^{1}H RF sensitivity profiles were calculated from magnitude maps obtained with standard gradient echo imaging sequences with a small flip angle and a short repetition time executed on a water sample. Both probes were equipped with actively shielded Z gradient coils.

One-dimensional (1D) shimming was performed at 300 MHz on a standard lineshape test sample with 1% $CHCl_3$ in acetone-d_6, which was spinning at a rate of 20 Hz. 1D ^{2}H field maps [6] were obtained via acquisition on the lock channel using a standard gradient-spoiled gradient echo phase mapping sequence with an echo time difference of 163 ms, an excitation flip angle of 43°, a pulse repetition time of 348 ms, 256 complex data points acquired during 169 ms, and a gradient strength of 0.78 G/cm corresponding to a field-of-view of 3.0 cm. For a sufficiently high signal-to-noise ratio 128 averages were necessary, leading to a total acquisition duration of 89 s per field map. (Note that this duration can be much shorter e.g. at higher frequency, with cryogenically cooled probes, or for ^{1}H field maps.) For comparison with the simulated lines ^{1}H NMR spectra were acquired with a standard lineshape test procedure with 16 averages.

3D shimming was performed at 800 MHz on a non-spinning standard water-suppression test sample with 2 mM sucrose and 0.5 mM DSS in a solvent mixture of 90% H_2O and 10% D_2O. 3D ^{1}H field maps [5] were acquired with frequency encoding along z using the probe gradient coil and phase encoding along x and y, utilising the X and Y gradient coils of the shim system. Fast switching of the latter was accomplished with the standard Bruker RCB (real-time control board) functionality. 3D data sets were obtained with $128 \times 22 \times 22$ voxels corresponding to a volume of $3.06 \times 0.66 \times 0.66\ \mathrm{cm}^3$. With an acquisition time of 3.5 ms the Z gradient strength was 2.82 G/cm, the echo time difference was 25 ms, the flip angle was 4.5°, and a repetition time of 104 ms lead to a total duration of 100 s per 3D field map. After shimming a ^{1}H NMR spectrum was acquired with H_2O pre-saturation and eight averages.

For fully automatic shimming the algorithms described above were implemented in C for execution on the spectrometer PCs equipped with a 3 GHz Intel® Pentium® processor and 1 GB memory.

3. Results

3.1. 1D shimming

The whole optimisation and shimming procedure as described in Section 2 was used to adjust all six available on-axis shim functions. Due to their axial symmetry this can be performed with 1D data resolved along z, provided that a data point represents the local cross-section of the sample in a reasonable fashion. Hence, good B_0 homogeneity is required across each slice which is supported by spinning the sample. Accordingly, spectral lines simulated based on 1D data are suitable to predict spectra measured on a spinning sample, where a similar averaging effect occurs as during field mapping.

The input data of the spectrum optimisation is shown in Fig. 4 with the B_1 profile and the initially mapped unshimmed B_0 field plotted as the local ^{1}H resonance frequency $\Omega_{0,\mathrm{M}}(z)$. Note, that the displayed usable range along the z dimension of 2.2 cm is used in all following similar figures. For the four isotope lines of $CHCl_3$ [10] only a combined approximate width of the natural line can be given. Hence, in the present case at 300 MHz $\lambda = 0.1$ Hz was selected for spectrum calculation according to Eq. (7). The discretisation parameters were set to $\Delta\Omega_{\mathrm{M}} = \lambda$ and $\Delta\Omega_{\mathrm{H}} = \Delta\Omega_{\mathrm{S}} = \lambda/3$.

In order to demonstrate the influence of the choice of the envelope shape two different shapes were selected, a

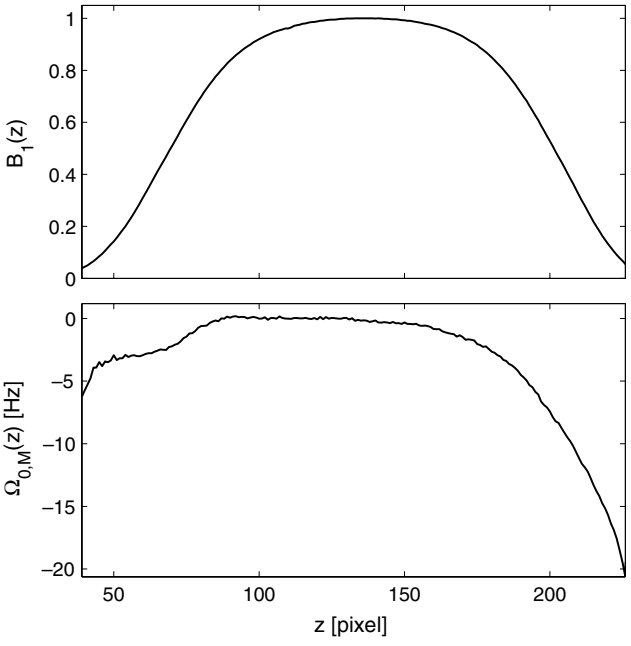

Fig. 4. Input data for the spectrum optimisation. The 1D profile $B_1(z)$ of the ^{1}H magnitude sensitivity of the RF coil is normalised in the sense that for $B_1 = 1$ the nominal flip angle α is performed. The ^{1}H frequency map $\Omega_{0,\mathrm{M}}(z)$ was calculated from the field $\Delta B_{0,\mathrm{M}}(z)$ of the initial inhomogeneity that was mapped with ^{2}H. The z axis is given in units of reconstructed pixels and the displayed usable region corresponds to a range of 2.2 cm.

pure Lorentz curve and a Tsallis function with $q = 1.75$. The spectrum optimisation was performed with the single-parameter approach using the weighting function $W(z) = [B_1(z)]^k$ and by stepping k from 0.2 up to 10.8 in steps of 0.2 ending up with totally 54 trials per envelope. Executed on a recent standard PC the mean duration of the calculation per trial was $T_{\text{iter}} = 0.14$ s adding up to 5–10 s per optimisation run. The resulting values of the criterion FWEHMS as shown in Fig. 5 exhibit a similar behaviour for both envelopes. A global minimum can be found for lower values of k, and an increase occurs for larger k while several local extrema are observable. However, the global minima for the two envelope shapes clearly differ. Detailed data of the calculations for the Tsallis shape are shown in Fig. 6 for the selected trials as indicated in Fig. 5. For trial 1 the lowest weighting was applied and accordingly the deviations in the field map increase only slightly towards the axial borders. The resulting shape of the spectrum is good and the envelope fits well, but the line is relatively broad. Trial 5 is the optimum found for the Tsallis shape. The stronger weighting allows for larger field deviations of the same sign at the two ends leading to a slight asymmetric increase of the intensity in the side lobes. However, in return the line becomes notably narrower and higher. Trial 14 was found being the optimum line with the envelope of Lorentzian shape. Again stronger weighting enables an improved width and height. However, the Tsallis shape cannot be fitted very tightly and the increased asymmetric hump intensity is penalised by the criterion.

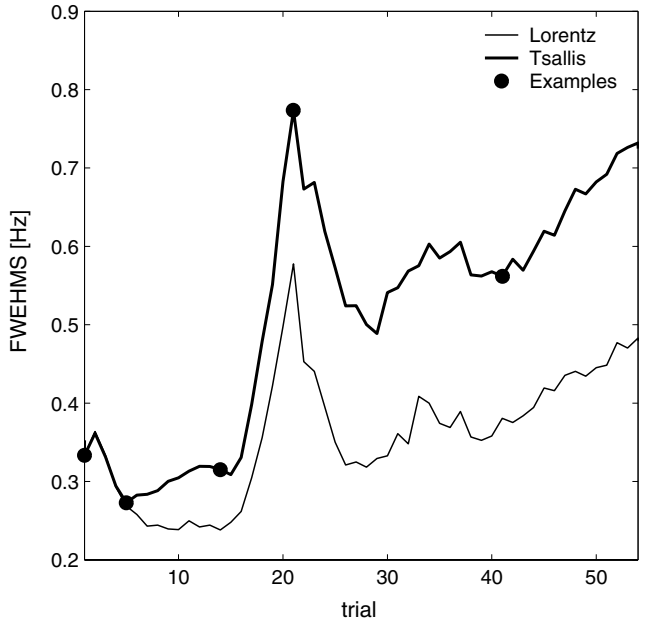

Fig. 5. Values of the quality criterion FWEHMS during optimisation. Each trial represents the calculation of a spectrum for a residual field obtained with a different spatial weighting. For each spectrum two envelopes were calculated, a Lorentz curve and a Tsallis function with $q = 1.75$, and the quality criterion FWEHMS was determined. The trial with the lowest criterion value is selected as the result of the optimisation procedure. Detailed data for the indicated examples are shown in Fig. 6.

Trial 21 is the worst case found during this optimisation. As indicated by the arrows, this example demonstrates how the envelope criterion reacts to the small additional peak away from the main line that is caused by the small plateau in the field map at about −2 Hz. Finally, trial 41 shows the data for a local minimum obtained with very strong weighting. Here, the large field deviations have opposite sign leading to a rather symmetric spectral line. However, despite a narrow line width the large hump is punished by the broad and badly fitting envelope.

Shimming was performed with the procedure described in Fig. 3 for both the target map obtained with the Lorentzian envelope and that of the Tsallis shape. As the initial situation was not badly inhomogeneous one run through the outer optimisation loop was sufficient. For both cases the fit loop converged after two iterations. Fig. 7 shows the measured field maps and their difference with the target map. Already after the first iteration the standard deviation of the difference map σ has improved by more than an order of magnitude. After the second iteration σ is on the order of the noise amplitude. Note that the latter is in the range of $\lambda = 0.1$ Hz which is just about acceptable to provide the necessary accuracy. In the centre the zoomed difference map after the second iteration exhibits only negligible deviations from zero. However, at the borders slightly larger variations are observable. This is due to the fact that the fit to the target map (and the calculation of σ) was performed with the local weighting $[B_1(z)]^2$.

The measured ^{1}H lineshape spectra for the two shim situations are shown in the bottom row of Fig. 8. Both lines are narrow and smooth, and, beside the hump, they are rather symmetric. The line obtained with the Lorentzian envelope is slightly narrower and higher but has a broader hump than the Tsallis result, as can be seen in the zoomed plot. This observation is supported by the full widths determined at 50, 0.55, and 0.11% of maximum height that are 0.18, 3.4, and 11.1 Hz for Lorentz, and 0.19, 2.9, and 6.7 Hz for Tsallis, respectively. The calculated lines in the top row of Fig. 8 clearly show the same tendencies as the measured spectra, despite the difference in the shape of the natural line.

3.2. 3D shimming

Using field maps resolved in all three spatial dimensions not only all available on-axis but also non-axially symmetric (off-axis) shim functions up to the fifth total order were employed for shimming with spectrum optimisation. Application of the Lorentzian envelope with the identical parameters as used in the 1D case lead to an optimum at $k = 5.6$ with FWEHMS = 0.48 Hz. On average each lineshape calculation took 1.4 s. Shimming to the corresponding target field converged after five iterations with $\sigma = 0.2$ Hz.

The water-suppressed ^{1}H NMR spectrum displayed in Fig. 9 was acquired after the described 3D shimming

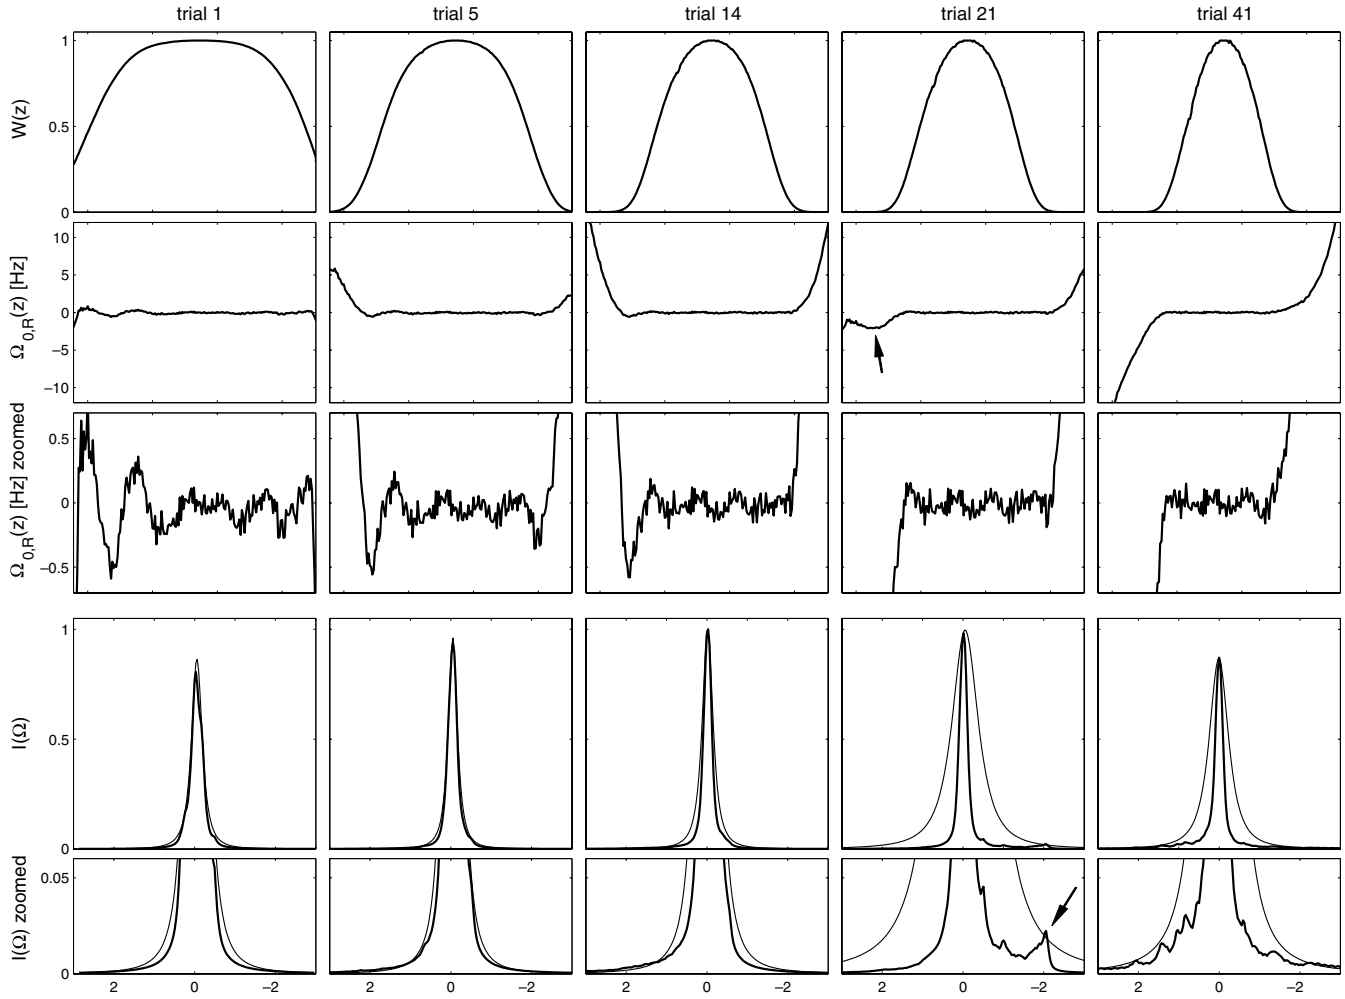

Fig. 6. Data obtained during spectrum optimisation. Examples are shown for the Tsallis-shaped envelope as indicated in Fig. 5. The weighting functions $W(z)$ applied during least-squares fitting and the obtained residual field maps with full and zoomed range are shown in the three upper rows. The range of the z axis is identical to that given in Fig. 4. The two bottom rows show the corresponding calculated spectra $I(\Omega)$ (bold line) along with the envelopes (thin line) where the zoomed display serves to observe the hump with more detail.

procedure and shows excellent water suppression and resolution. The narrow and simple DSS line was selected for comparison with the simulated result, and Fig. 10 shows a very good agreement of the two spectra.

4. Discussion

A new framework for gradient shimming was presented that enables to find a shim setting that is optimised directly for spectrum quality. Spectrum optimisation by means of simulations based on a field map leads to a shimming procedure that aims at the realisation of a target map rather than an impossible spatially constant field. For the widely applicable case of the optimisation of a single line the envelope quality criterion was developed. Features and usefulness of the approach were demonstrated for 1D and 3D shimming on standard NMR test samples.

As a major advantage of spectrum optimisation with subsequent target shimming the quality of the spectrum can be guaranteed. This is a prerequisite for full automation of shimming. In contrast to that, conventional gradient shimming aims at homogenising the field in real space and a satisfactory spectrum quality cannot always be guaranteed. Hence, multiple attempts with modified parameters are required, each followed by the acquisition of an NMR spectrum and judgement by the spectrometer user. With the presented approach this task is not only taken away from the user but also transferred from measurement to calculation, thus reducing the response time. Although for the used calculations computation time is still an issue the future speed improvement expected from computer performance is definitely larger than that from increased NMR sensitivity.

Several parameters are available to influence the optimisation result. Calculation of the spectrum incorporates the width of the natural line λ, and it is worth adjusting this parameter approximately to the expected line widths. This becomes clear when considering the natural line being a smoothing kernel for the histogram of the field inhomogeneity. Hence, distortions unimportant on the scale given by

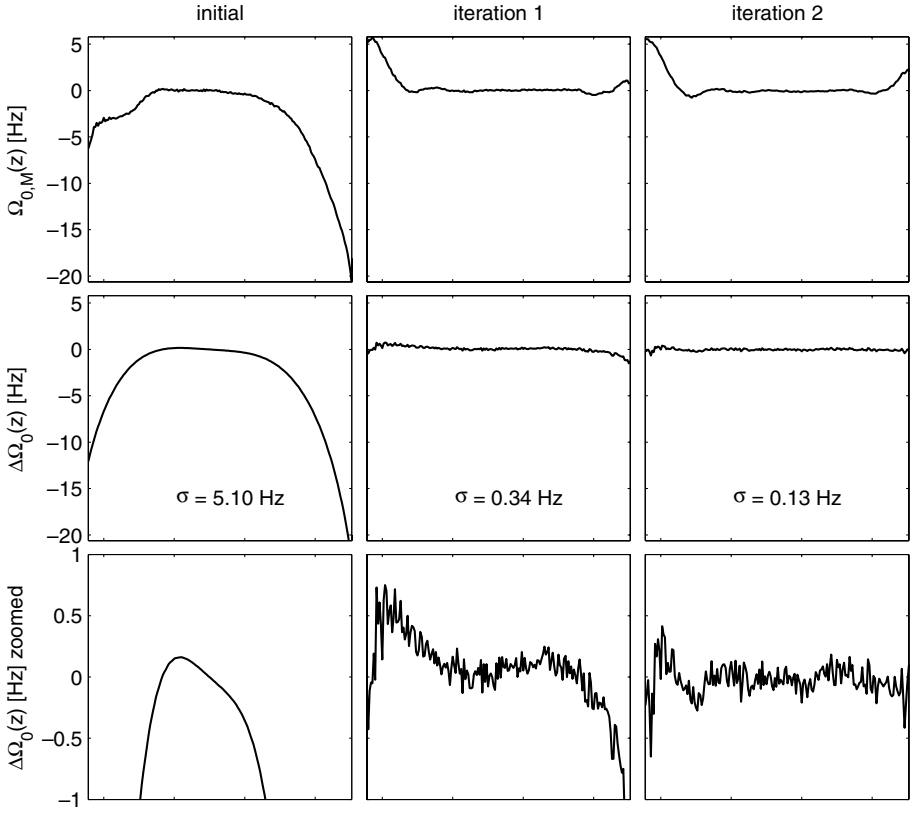

Fig. 7. B_0 field maps measured during the shimming iterations. The field map after an iteration and its difference with the target map are shown. The standard deviation of the difference σ as well as the zoomed display indicate that after the second iteration the residual deviations in the centre are on the order of the noise. Only at the borders the deviations are somewhat larger due to the weighting applied during the fit to the target map. The range of the z axis is identical to that given in Fig. 4.

λ will not be visible in the spectrum and during optimisation the limited resources of the shim system are not wasted in order to get rid of them.

A further parameter is the shape of the envelope that enables to influence the shape of the spectral lines to a certain extent depending on the user preferences. However, the optimisation has to select from the set of actually possible lineshapes, which is fully determined by the given setup. Envelope shapes set in the range between a pure Lorentz and a Tsallis curve with $q = 1.75$ as presented here proved to successfully accomplish this task during testing the optimisation with a couple of different probe types.

The single-parameter optimisation itself is determined by the chosen weighting function and the associated parameter k. With a B_1-based weighting function the useful range of k depends not only on the B_0 inhomogeneity but also on the variation of sensitivity of the RF coil and has to be found by experience. In this work, the optimum value of k was then determined by simply sampling a range of k values with a sufficiently small step size. However, the relatively well-structured behaviour of the criterion suggests that a more efficient optimisation is possible, e.g. by separation into a coarse and a fine search.

The pragmatic approach of reducing the optimisation from a multi- to a single-parameter search dramatically improves its efficiency. This is particularly true for an increasing number of shim functions, as the number of trials remains constant with the proposed technique but increases for a multi-parameter optimisation. Of course, due to the pre-selection of a single path in the multi-dimensional parameter space the global optimum may be missed. However, the assumptions for this approach are strongly supported by experience and the potentially slightly worse results come at a considerable gain in optimisation speed. Nevertheless, an intelligent and fast multi-parameter optimisation would be the ideal solution. In terms of reproducibility it is important to note that the found optimum does not depend on the original shim setting as long as the B_0 homogeneity allows the same spatial region to be mapped.

The spectrum simulation employs the spatially resolved magnitude B_1 sensitivity of the RF coil of the nucleus observed in the planned NMR experiment, which was ^{1}H in the presented example. Consequently, with the same probe optimisation for nuclei with different B_1 sensitivities results in different shim settings. In particular nuclei observed with different RF coils, potentially exhibiting differing length and/or shape, can favour unlike optima. Providing the necessary information about B_1 is not always trivial. For nuclei where sufficient NMR sensitivity can be provided easily, such as e.g. ^{1}H, ^{2}H, or ^{19}F, this can be accomplished by means of a simple NMR imaging sequence. The same B_1 information can usually also be

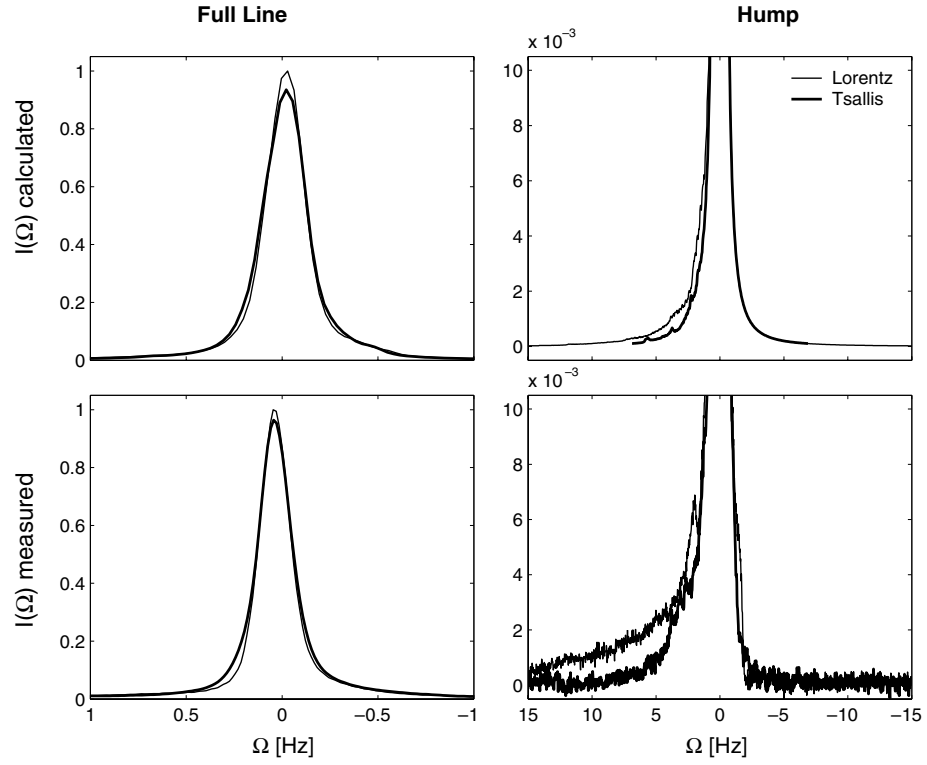

Fig. 8. ^{1}H lineshape spectra simulated and measured at 300 MHz after 1D shimming optimised with different envelopes. The top row shows the predicted calculated spectra during optimisation based on 1D data and the bottom row the corresponding spectra measured on $CHCl_3$ with a spinning sample. The full lines are plotted in the left column and the zoomed hump is shown in the right column. Note the good agreement of calculation and measurement despite the difference in the natural line.

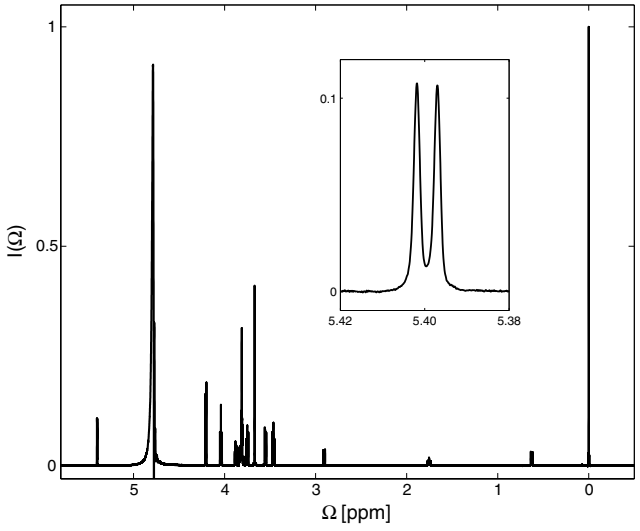

Fig. 9. Water-suppressed ^{1}H spectrum measured at 800 MHz without sample spinning after 3D shimming with spectrum optimisation. Four seconds of H_2O pre-saturation were followed by a 90° excitation and an acquisition of 3.4 s duration, and eight averages were performed. The widths of the water hump at 50 and 10% height of the DSS maximum are 15 and 45 Hz, and the splitting of the sucrose doublet shown in the zoomed region is 12%.

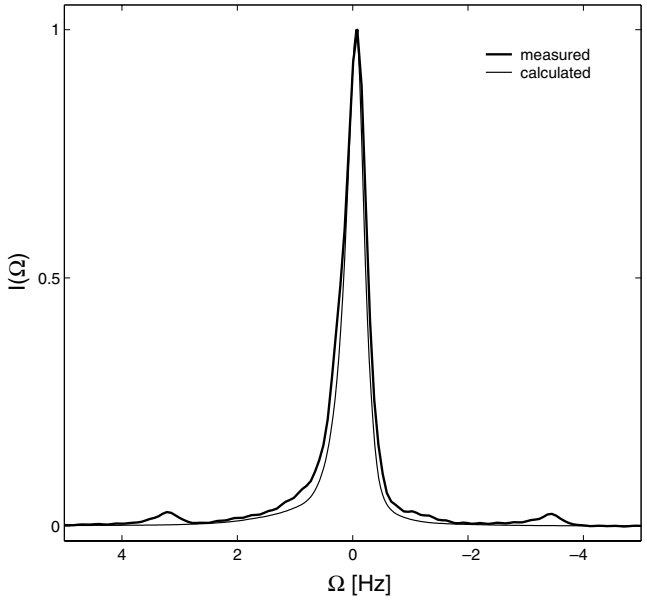

Fig. 10. Zoomed section of the measured spectrum in Fig. 9 showing the DSS line with a half width at full height of 0.47 Hz. The overlay with the line calculated from 3D field map data shows a very good agreement. Note that the satellites at ±3.3 Hz are part of the DSS spectrum and were not taken into account during the simulation.

used for less sensitive nuclei observed on the same coil. If none of these methods is applicable, more advanced NMR procedures, specific samples, or other means of measuring the sensitivity are required. Note that the B_1 sensitivity of the nucleus used for acquisition of the field maps, e.g. ^{2}H in the 1D case presented here, does not play

a role for the optimisation. However, the useful spatial range should sufficiently cover the sensitive region of the nucleus of interest.

The spectrum optimisation necessarily leads to shimming to a target field. In the proposed approach two nested iteration loops decouple the tasks of error correction and taking advantage of the improved B_0 homogeneity. In order to provide good convergence, parameters must generally not be modified while running through the inner loop. However, during the outer loop not only a new target field is determined but also parameters such as the set of used shim functions or the considered spatial volume can be adjusted based on the current B_0 homogeneity. Multiple runs through the outer loop were not required in the presented examples as the initial B_0 field was relatively homogeneous. However, in a routine employed for fully automatic shimming this approach enables shimming starting from a very bad initial B_0 homogeneity and ending with a high-end result.

The general framework described here is not limited to the optimisation of a single lineshape, although this case is widely applicable. The common situation of solvent suppression requires sufficient homogeneity at the sample borders while maintaining a good lineshape. This task is particularly demanding for specific sample geometries with discontinuities. With the available free parameters the presented method enables to satisfy such a setup by choosing a shape with Gaussian contribution (e.g. $q = 1.75$) and tight envelope (e.g. $f = 500$). Moreover, also complicated substances and sequences can be simulated and other criteria can be derived to obtain shim settings dedicated to a specific NMR experiment. E.g. for a substance leading to multiplets in the spectrum a criterion could be developed to judge the splitting of the neighbouring lines; or for a sequence containing a series of RF pulses the influence of B_1 and B_0 on the shape of the final lines could be taken into account.

Finally, the principles of the presented method are not limited to high-resolution NMR but can also be applied for localised NMR spectroscopy performed in vitro as well as in vivo.

Acknowledgments

The authors thank many of their colleagues in Switzerland and abroad for helpful input, fruitful discussion, and support of all kind.

References

[1] R.R. Ernst, Measurement and control of magnetic field homogeneity, Rev. Sci. Instr. 39 (1968) 998–1012.
[2] J.A. Nelder, R. Mead, A simplex method for function minimization, Comput. J. 7 (1965) 308–313.
[3] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical Recipes in C, Second ed., Cambridge University Press, Cambridge, 1992, Chapter 10.4.
[4] M.G. Prammer, J.C. Haselgrove, M. Shinnar, J.S. Leigh, A new approach to automatic shimming, J. Magn. Reson. 77 (1988) 40–52.
[5] 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.
[6] S. Sukumar, M.O. Johnson, R.E. Hurd, P.C.M. Van Zijl, Automated shimming for deuterated solvents using field profiling, J. Magn. Reson. 125 (1997) 159–162.
[7] R. Dunkel, Method for correcting spectral and imaging data and for using such corrected data in magnet shimming, US Patent 5,218,299 (1993).
[8] T.G. Reese, T.L. Davis, R.M. Weisskoff, Automated shimming at 1.5 T using echo-planar image frequency maps, J. Magn. Reson. Imag. 5 (1995) 739–745.
[9] D.F. Howarth, J.A. Weil, Z. Zimpel, Generalization of the lineshape useful in magnetic resonance spectroscopy, J. Magn. Reson. 161 (2003) 215–221.
[10] F.A.L. Anet, M. Kopelevich, Ultrahigh resolution in proton NMR spectra at 500 MHz: two-bond intrinsic chlorine and silicon isotope effects, J. Am. Soc. 109 (1987) 5871–5872.