Journal of Magnetic Resonance 310 (2020) 106641

Contents lists available at ScienceDirect

Journal of Magnetic Resonance

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

De novo NMR pulse sequence design using Monte-Carlo optimization techniques

Joel Lapin, Alexander A. Nevzorov *

Department of Chemistry, North Carolina State University, 2620 Yarbrough Drive, Raleigh, NC 27695-8204, United States

ARTICLE INFO

Article history:
Received 24 September 2019
Revised 3 November 2019
Accepted 4 November 2019
Available online 6 November 2019

Keywords:
NMR pulse sequence design
Automation
Simulated annealing
GPU computing
Separated local field experiments
Dipolar couplings
PISEMA
SAMPI4

ABSTRACT

Separated Local Field (SLF) experiments have been routinely used for measuring ^{1}H-^{15}N heteronuclear dipolar couplings in oriented-sample solid-state NMR for structure determination of proteins. In the on-going pursuit of designing better-performing SLF pulse sequences (e.g. by increasing the number of subdwells, and varying the rf amplitudes and phases), analytical treatment of the relevant average Hamiltonian terms may become cumbersome and/or nearly impossible. Numerical simulations of NMR experiments using GPU processors can be employed to rapidly calculate spectra for moderately sized spin systems, which permit an efficient numeric optimization of pulse sequences by the Monte Carlo Simulated Annealing protocol. In this work, a computational strategy was developed to find the optimal phases and timings that substantially improve the ^{1}H-^{15}N dipolar linewidths over a broad range of dipolar couplings as compared to SAMPI4. More than 100 pulse sequences were developed *de novo* and tested on an N-acetyl Leucine crystal. Seventeen distinct pulse sequences were shown to produce sharper mean linewidths than SAMPI4. Overall, these pulse sequences have more variable parameters (involving non-quadrature phases) and do not involve symmetry between the odd and even dwells, which would likely preclude their rigorous analytical treatment. The top performing pulse sequence, termed ROULETTE-1, has 18% sharper mean linewidths than SAMPI4 when run on an N-acetyl Leucine crystal. This sequence was also shown to be robust over a broad range of ^{1}H carrier frequencies and various crystal orientations. The performance of such an optimized pulse sequence was also illustrated on ^{15}N Leucine-labeled Pf1 coat protein reconstituted in magnetically aligned bicelles. For the optimized pulse sequence the mean peak width was 14% sharper than SAMPI4, which in turn yielded a better signal to noise ratio, 20:1 vs. 17:1. This method is potentially extendable to *de novo* development of a variety of NMR experiments.

1. Introduction

Orientationally dependent heteronuclear dipolar couplings (DC) provide direct experimental input for structural elucidation in solid-state NMR of membrane proteins. Along the protein backbone the possible DCs between NMR active nuclei include ^{1}H-^{15}N, $^{13}C_{\alpha}$-$^{1}H_{\alpha}$, and ^{13}C-^{15}N DCs. The value of a DC is determined by the angle that the internuclear bond of interest makes with the external magnetic field, B_0, which provides angular restraints for structure calculations. Efficient pulse sequences correlating DCs between three or more spins are necessary for structure determination of membrane proteins in their native-like lipid environments. High-resolution Separated Local Field (SLF) pulse sequences such as PISEMA, SAMPI4, HIMSELF, and the more recent 2_n-SEMA [1–9], yield sufficiently sharp linewidths over a broad range of dipolar interactions, and have been shown to be robust over a wide range of proton carrier frequencies. For dilute low spins (^{15}N) surrounded by a "spin bath" of protons (^{1}H), the main goal in developing such sequences is to selectively evolve the heteronuclear DCs, while simultaneously decoupling the ^{1}H-^{1}H homonuclear DCs. In order to accomplish the evolution of heteronuclear DCs, as in PISEMA [1], the low-spin channel (S/^{15}N channel) has one phase alteration, i.e. a 180° x-pulse followed by a −180° x-pulse. On the high-spin channel (I/^{1}H channel), PISEMA uses off-resonance Lee-Goldburg irradiation [10,11] via tilting (spin-locking) the ^{1}H spins at the magic angle, $\theta = 54.7^{\circ}$. The subsequently developed Sandwich-Assisted Magnetic MicroscopY pulse sequence (SAMMY) [4] keeps the protons spin-locked either along the x-axis or along the z-axis with on-resonance pulses, which makes this sequence less sensitive to proton carrier offsets than PISEMA. This is accomplished by using a $(\pi/2)_y$ pulse, followed by a $-(\pi/2)_y$ pulse, between which the power is turned off on the ^{1}H channel. The SAMPI4 pulse sequence further improves

* Corresponding author.
E-mail address: alex_nevzorov@ncsu.edu (A.A. Nevzorov).

https://doi.org/10.1016/j.jmr.2019.106641
1090-7807/

on SAMMY by decreasing the number of phase transients on the low spin channel, and subtracting π/4 pulse duration from each subdwell to compensate for the finite π/2 pulses [6]. For both PISEMA and SAMPI4 the dwell time and the scaling factors can be derived theoretically. These pulse sequences have been shown to dramatically decrease dipolar linewidths (by approximately an order of magnitude) thus increasing sensitivity in 2D heteronuclear NMR experiments of uniaxially aligned samples.

All previous SLF pulse sequences only utilize quadrature phases. That is, the aforementioned pulse sequences, PISEMA and SAMPI4 utilize only phases of 0°, 90°, 180°, or 270°. Other heteronuclear DC experiments involve more pulses in their architecture, such as BLEW-12 [12] which consists of twelve back-to-back pulses, and the related HIMSELF sequence [3]. Still, those sequences are restricted to quadrature phases, which may limit their full potential. Average-Hamiltonian theory has been recently utilized [13] to optimize the pulse phases and rotation angles to find the radiofrequency (rf) pulse cycles which decouple the homonuclear dipolar couplings while evolving the chemical shift and heteronuclear dipolar interactions. However, analytical treatment of the Hamiltonian transformations involving multiple non-quadrature phases and arbitrary rotation angles can become prohibitively complex. In this case, average-Hamiltonian treatment to find the optimal pulse sequence parameters becomes a rather complicated task. To overcome the challenges of a rigorous analytical treatment, numerical spin simulations can easily and accurately predict spin evolution under any pulse sequence architecture and associated pulse parameters, thus permitting exploration of new NMR experiments involving non-quadrature phases and arbitrary rotation angles for spin magnetization. For instance, the decoupling pulse sequence DUMBO [14] represents a notable example of computer-aided pulse sequence design.

Herein we describe a protocol to find *de novo* pulse sequences using numerical optimizations performed on a relatively large 8-spin system, derived from the coordinates of a polyalanine amino acid sequence. The simulation method consists of the construction of the Hamiltonian, propagation of the density matrix, and scoring of the resulting Fourier transformed spectra. The optimizations are carried out by the Monte-Carlo Simulated Annealing (MCSA) protocol, which allows for sampling over a large sequence parameter space that includes pulse phases and timings. The automated algorithm is instructed to find a pulse sequence that maximizes peak sharpness, while correctly rendering a pre-defined set of three (3) dipolar splittings. This constitutes a *de novo* approach for pulse sequence design from first principles. The top scoring pulse sequences were run on an N-acetyl Leucine (NAL) crystal in a 300 MHz NMR spectrometer, and the spectral linewidths were compared to that for the SAMPI4 pulse sequence. The best performing pulse sequence was then re-optimized for a sample of ^{15}N Leucine labeled Pf1 coat protein reconstituted in magnetically aligned bicelles at 500 MHz ^{1}H NMR frequency. This method of numerical optimization is greatly versatile and potentially applicable to any NMR experiment that can be simulated with reasonable accuracy and speed.

2. Analytical framework

In general, the outcome of any pulse sequence can be simulated using the Liouville-von Neumann equation describing the evolution of the density matrix ρ for a spin system evolving under the Hamiltonian, H_n. One can write for each subdwell, n:

$$\frac{d\rho(t)}{dt} = -\frac{i}{\hbar}[H_n, \rho(t)] \tag{1}$$

The density matrix $\rho(t)$ can be successively propagated after each time subdwell Δt_n using the unitary transformations of the initial density matrix:

$$\rho(t_n) = e^{-iH_n\Delta t_n}\rho(t_{n-1})e^{iH_n\Delta t_n} \tag{2}$$

The real part of the time-domain signal can be calculated from the density matrix as:

$$S(t) = Trace(S_x\rho(t)) \tag{3}$$

where S_x is the detector matrix for transverse magnetization of the low (S)-spins. The Hamiltonian for a system of low (S/^{15}N) and high (I/^{1}H) spins along the backbone of a protein can be separated into single-spin terms involving the Zeeman interactions of the spins with the main B_0 and radiofrequency B_1 fields and the two-spin dipolar terms. The single-spin rf Hamiltonian terms for an arbitrary phase of each pulse can be written as:

$$H^S_{rf} = \omega^S_i[\cos(\phi^S_i)S_x + \sin(\phi^S_i)S_y] \tag{4a}$$

$$H^I_{rf} = \omega^I_i[\cos(\phi^I_i)I_x + \sin(\phi^I_i)I_y] \tag{4b}$$

$$H_{rf} = \sum_{spins}\left(H^S_{rf} + H^I_{rf}\right) \tag{5}$$

Let us consider for simplicity a system of one low spin and (N – 1) high spins. The two-body terms for the heteronuclear dipolar interactions are given by:

$$H^{(IS)} = \sum_{i=2}^{N} a_{1i}S_Z I^{(i)}_Z \tag{6}$$

and the homonuclear dipolar interactions among the high spins are:

$$H_{zz} = \sum_{i<j}^{N}\frac{a_{ij}}{2}\left[3I^{(i)}_z I^{(j)}_z - \left(I^{(i)}I^{(j)}\right)\right] \tag{7}$$

Here the dipolar coupling constants between spin pairs can be written as (in the units of rad × s^{-1}):

$$a_{ij} = \frac{\mu_0}{4\pi}\frac{\gamma_i\gamma_j\hbar(3\cos^2\theta - 1)}{r^3_{ij}} \tag{8}$$

where θ is the angle between the interspin vector and z-axis.

In most general terms, the goal of any rational SLF pulse sequence design is to find a periodic train of pulses with specific phases and timings that would ensure the evolution of the heteronuclear terms yielding accurate dipolar doublet splittings, a_{1i} (cf. Eq. (6)), while concomitantly canceling the homonuclear terms (Eq. (7)). Theoretical design of such pulse sequence usually starts with considering average Hamiltonians [15–17] up to the second-order approximation, followed by practical considerations including the probe efficiency, spectrometer dead time, magnetic field homogeneity, and other experimental limitations. Additional considerations may include proton frequency offsets, such as in the case of the PISEMA pulse sequence, which would result in spin-locking the protons away from magic angle, or phase transients, which generally preclude the use of too many subdwells due to accumulation of error over longer t_1 dwells.

Numerical spin simulations are an invaluable tool for the prediction and rationalization of experimental results, owing to the well-developed mathematical formalism which accurately simulates many-body spin systems. In addition to the available software packages for carrying out spin simulations [18–20], Eq. (2) can be relatively easily custom-coded in a high-level programming language in order to simulate a particular NMR experiment and optimize its parameters. However, the computational cost for the density matrix representation of spin systems exponentially increases with the system size, with the dimensions of the density matrix and the Hamiltonians scaling as $2^N \times 2^N$ with the number of spins N. This restricts the practical spin system size that can be used in the calculations due to both memory and time constraints.

Thus, in the present work, we have restricted ourselves to 8 spins total, with the proton coordinates constrained within a 3 Å sphere centered at a ^{15}N atom in a polyalanine molecule. A dual-GPU (NVidia 1080Ti) processor has been employed to perform the most computationally intensive matrix operations, which reduces the time required to simulate a single experiment by more than an order of magnitude as compared to a CPU calculation.

The ability to rapidly simulate individual NMR experiments permits an exhaustive search over a large parameter space. First, a desired outcome function must be defined, so that the resulting spectrum can be objectively scored (as defined in the subsequent section). The search then seeks to optimize the parameters for that outcome. Desirable characteristics for an experiment probing ^{1}H-^{15}N dipolar couplings include peak height and/or peak sharpness, i.e. the full width at half height (FWHH), and the appearance of the dipolar doublets at the locations either equal or proportional to the "true" couplings between the single ^{15}N spin and the directly bonded ^{1}H. The latter quantity is closely related to the so-called scaling factor (*scf*) arising in complex pulse sequences, and its robustness is critical for the accurate determination of the angular restraints in the SLF experiments. The major parameters defining an NMR pulse are the pulse phases, timings, and the rf power amplitudes. Optimization commences by selecting the initial sequence parameters for each subdwell, simulating the experiment, measuring the peak heights and locations for three different dipolar couplings, changing the parameters and repeating the process until the best pulse sequence is found. The optimization is carried out by the MCSA protocol [21]. MCSA is a broadly applicable optimization technique that can efficiently search over numerous parameters that may vary in the extent of their effect on the outcome. To avoid spurious solutions for a narrow range of DCs, additional steps need to be taken to ensure the robustness of the pulse sequence. This can include simulating dipolar spectra with different maximum DC values at each optimization step and verifying their relative positions along the frequency axis. Such an approach would eliminate pulse sequences that would be optimized to evolve dipolar couplings within a narrow frequency range. Another strategy to ensure robustness is to include random chemical shift perturbations (σ) for each of the protons (and also for the ^{15}N spins). This is added to the Hamiltonian as a sum over the individual chemical shift terms, viz.:

$$H_\sigma = \omega_0 \sum_i^N \sigma_i I_z \tag{9}$$

where ω_0 is the carrier frequency in MHz, and σ_i are the individual chemical shift perturbations. It is important to emphasize that including different chemical shifts and several dipolar couplings in the simulation is critical in order to ensure robustness of the final sequence.

Most of the previous pulse sequences were designed with symmetry considerations to resemble the canonical transformation for a matrix A, i.e. $A' = TAT^{-1}$. That is, each pulse that forms a subdwell would include a mirrored, counterpart pulse that would differ by a 180° phase. The same phase mirroring strategy is applied here for even vs. odd total t_1 dwells. For example, the odd dwell on the low S-spin channel in SAMPI4 starts with a 0° pulse, and ends with a 180° pulse, while the even dwell starts with a 180° pulse and ends with a 0° phase. Symmetry is reflected in the timings as well, with the first subdwell having the same timing as the last subdwell, the second subdwell the same as second to last, etc. In the current simulations we have imposed either full symmetry (FS), as described for SAMPI4 above, or partial symmetry (PS) in which the first and the last subdwell pulse phases would still differ by 180°, but the even and odd dwells were essentially treated independently from one another. In general, the simulation can let all parameters float freely within a certain range, taking no symmetry into consideration. The more symmetry that is invoked, however, the fewer free parameters there are in the pulse sequence, thus making it faster for the MCSA protocol to find a reasonably good scoring solution. The phase space can be searched with a mixture of symmetry, by forcing some channels/subdwells to be mirror images of each other, while letting others assume phase values independently.

3. Methods

3.1. Software and hardware

All calculations were carried out on an Alienware Aurora RB desktop computer with an Intel® Core™ I7-9700K processor operating at 3.60 GHz. The computer was custom equipped with two NVIDIA GeForce GTX 1080 Ti GPUs, on which the spectral simulations were carried out. It should be noted that without GPU computing, the simulations would have been prohibitively slow and severely limited in the sampling capability. For an 8-spin system in which the matrix dimensions were 256 × 256, sampling of a single SLF experiment was ~20 times faster on GPU than on the CPU (the relative speedup over CPU is expected to increase even more with the matrix size).

The script responsible for administering the MCSA loop and spectral simulation will be provided upon request to the authors. The code was specifically optimized for calculations on a PC equipped with dual GPU processors, which are concurrently used in the simulations when assembling the even and odd dwell Hamiltonians, as well as when propagating the density matrix. The simulation was GPU optimized with help from Python's Cupy library, which provides the functionality of Python's popular NumPy library, but with functions that are GPU compatible. All Python scripts were executed using the open-source data science platform Anaconda©.

3.2. Simulation details

As stated in the Introduction, the DC constants have been calculated using the coordinates of a polyalanine alpha-helix. Distant ^{1}H spins have been truncated beyond a 3 Å radius of the backbone amide Nitrogen, resulting in a system of 8 spins (with matrix dimension size 256 × 256). A series of arbitrary chemical shifts were included in the H_σ term of the Hamiltonian to account for the possible proton carrier offsets. The carrier frequency ω_0 was set to 500 MHz for protons and 50.7 MHz for the nitrogen spins.

The program was optimized at a ω_{rf} field amplitude of 58.14 kHz, in order to match the experimental value calibrated for the NAL single crystal sample. The rf amplitudes for each subdwell were modulated by the binary W_i variable, which in an experiment would correspond to irradiation being either power off, 0, or power on, 1. The full Hamiltonian for each subdwell is the sum of Eqs. (5)–(7) and (9).

$$H_{tot} = H_\sigma + H_{rf} + H^{(IS)} + H_{zz} \tag{10}$$

After being assembled via Eq. (10), the Hamiltonian, H_{tot}, is diagonalized to yield the eigenvalue matrix, D, and the eigenvector matrix, V. This decomposition is used to create the propagator matrix, P_n for each subdwell:

$$P_n = V_n\, e^{-iD_n \Delta t_n}\, V_n^\dagger \tag{11}$$

and then the overall propagator is assembled as a matrix product over all subdwells:

$$P = \prod_n P_n \tag{12}$$

Note that the even and odd dwells can have their own propagators. The manipulations with the Hamiltonian were GPU optimized due to the numerous linear algebra operations, such as diagonalization and matrix-matrix multiplication, necessary to implement Eqs. (11) and (12). The simulation then proceeds to the propagation of the FID (also GPU optimized), where the FID is calculated via Eqs. (1)–(3). Prior to the Fourier transformation, the FID vector, G, is multiplied elementwise by an exponential apodization function:

$$G'(t) = G(t)\exp(-\pi\Delta\nu t) \tag{13}$$

The linebroadening parameter $\Delta\nu$ was set to 50 Hz for all simulations. The Fourier transform of G' is matched to a frequency range f, based on the total dwell time, which is treated as the sum of all individual subdwell durations:

$$dwell = \sum_n t_n \tag{14}$$

In each simulation, 256 t_1 dwells have been propagated and the calculated FID was zero-filled to 1024 points to provide sufficient digitization of the Fourier-transformed spectra. The frequency range is given by the corresponding re-scaled Nyquist frequencies (in Hz) as:

$$f_{Nyquist} = \pm\frac{1}{2 * dwell * scf} \tag{15}$$

where *scf* is the apparent scaling factor for the pulse sequence that needs to be calibrated by comparing the apparent DCs to the known SLF sequences (see simulation results). The magnitudes of the apparent DCs are then evaluated as to whether they would yield the ratios of their relative positions corresponding to the predefined couplings (set to 20, 10, and 5 kHz or 4:2:1 in the present work). The relative intensity of each peak (which is inversely related to its linewidth) is also evaluated and included in the overall scoring function as described in the next section.

3.3. The MCSA algorithm

The MCSA protocol uses the dipolar spectrum simulation described in the previous section as the fundamental block for sampling the parameter space. More specifically, the peak amplitudes in the calculated spectrum and their locations are used to score a pulse sequence. The parameter search is then conducted using the Metropolis criterion in order to find a global minimum for the scoring function as defined below, always accepting a smaller proportion of higher energy moves as temperature decreases. The flowchart of Fig. 1 shows the logic flow for the algorithm.

In the present work, the program is initiated by setting manually the general pulse architecture and amplitude scheme (e.g. starting from the original SAMPI4 ordering of pulses), which in turn determines the variable space used in the simulation (i.e. the number of phases and timings). For PS, the total number of parameters varied during the MCSA run would then be one-half the number of subdwells times 5; 80% of the parameters belong to the phases and the rest to timings (for FS there are half as many varied phase parameters). For phases there are two channels, one for the high spins I and one for the low spins S. The two channels have both even and odd dwells. The symmetry is defined prior to starting the MCSA optimization process, which is represented in step 3 of Fig. 1a. While the phase parameters can in general have either FS (step 9a) or PS (step 9b), timings can only be PS. In principle, the rf frequency amplitudes for each channel and subdwell could also be considered parameters. Due to the inherent rise/fall times of the NMR amplifiers and their non-linearity, however, continuously varying the rf power between the subdwells would be impractical. Therefore, the values for the rf amplitudes were chosen to be either 0 or ω_{rf}. In addition, continuously varying the rf amplitudes would often result in a search that would never converge to a global minimum. Therefore, for the majority of the simulations we used the fixed power scheme such as used by SAMPI4, in which the power is turned off during the third and fourth subdwells on the ^{1}H-channel only.

Although the calculation of the individual terms in the Hamiltonian is a rather time consuming operation, it only needs to be completed once (step 2), allowing the simulation to change the parameters and continue using largely the same Hamiltonian matrices. The only Hamiltonian term in Eq. (10) that changes with the pulse sequence is H_{rf}; however, the matrices I_x, S_x only need to be calculated once, initially, and then H_{rf} is simply scaled per Eqs. (4a) and (4b). When one considers only a single dipolar doublet between the ^{15}N and its nearest ^{1}H, optimizing the pulse sequence carries a risk of finding a singularity for a specific dipolar doublet. This problem is assuaged here by running the simulation consecutively for three different DCs, thus ensuring that all resonances are properly evolved and decoupled. To calculate an intermediate score in the MCSA protocol, the simulation is, therefore, run three times for each random step. The three different precalculated $H^{(IS)}$ matrices are used in step 10, where the NMR experiment is simulated.

In this work the three maximum couplings were arbitrarily set to 5, 10, and 20 kHz. Since the generated pulse sequences may all have different scaling factors, the apparent DC frequencies could deviate from their true values. Thus, instead of their exact dipolar frequencies, only the ratios of the DC frequencies for the three peaks are factored into the score. The intensities from the three simulations (ph_i) are multiplied together and raised to the one-third power, which is considered as the raw score, termed *pmraw*. Using the product instead of the mean ensures that no single peak is optimized at the exclusion of the others. The final scoring function, termed *pm*, then becomes:

$$pm = -(ph_1 ph_2 ph_3)^{\frac{1}{3}}\left[1 - z\left(\left|2 - \frac{pf_1}{pf_2}\right| + \left|2 - \frac{pf_2}{pf_3}\right|\right)\right] \tag{16}$$

The ratios of the peak frequencies are accounted for by the pf_i terms. Parameter z (set to 0.5 in this work) is used to increase or decrease the penalty for violating the expected peak ratios. The score in Eq. (16) will be lowest (most negative) when the peak intensities are collectively high and the both ratios of the successive dipolar frequencies are as close to 2.0 as possible.

The starting parameters for an MCSA run must be initialized, either at random or from a fixed set of parameters that was previously saved to a file. If the initialization of the parameters is done randomly, the starting values are chosen within the same range where parameters would be sampled. For phases this range is [0,360°], which represents all possible angles that a pulse phase can assume. For timings this range was chosen to be 0–5, in the multiples of 90-degree pulse durations, i.e. $t90 = \frac{\pi}{2\omega_{rf}}$ (in microseconds). This range was chosen based on the previous SAMMY and SAMPI4 subdwells [6].

All MCSA protocols have two main loops: one for the "temperature" decrements and the other for random steps taken at each temperature. The temperature loop sets a unitless "inverse temperature" parameter, W^T, to be used in the Metropolis criterion when accepting or rejecting new steps. Since W^T is proportional to inverse temperature, it is increasing throughout the simulation, corresponding to cooling of the system. In this work, given the relative magnitudes of the scoring function (typically ranging from 0 to −30) it was sufficient to increment W^T from 1 to 8. The strategy for finding the best pulse sequences often stressed running numerous shorter, rather than fewer longer, simulations. Simulations were initially run with 100 temperature increments and

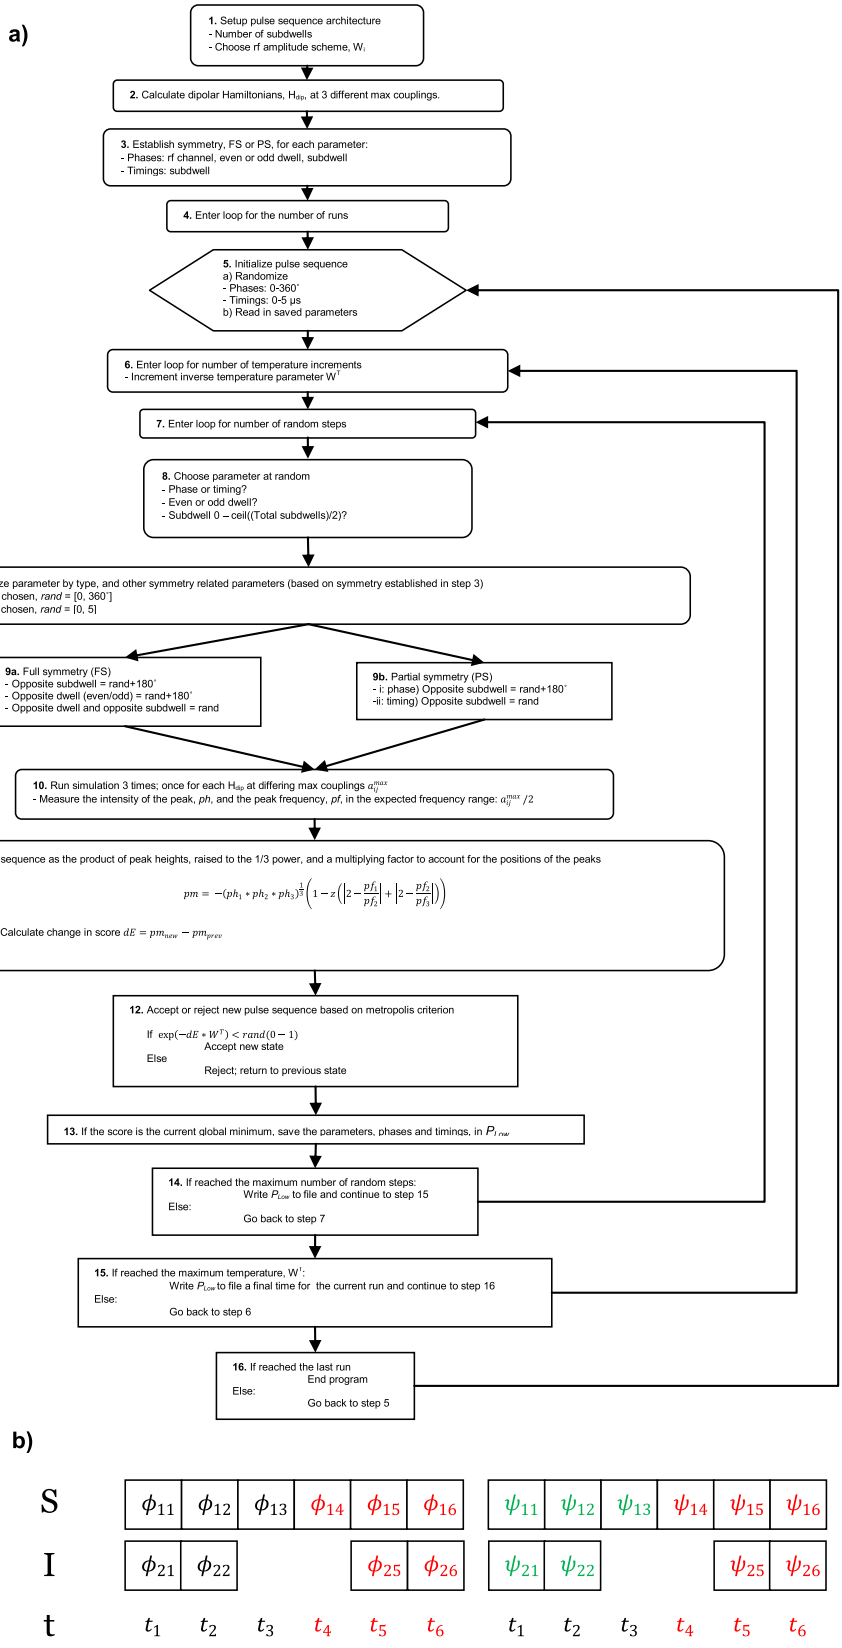

Fig. 1. (a) Flowchart for MCSA algorithm applied to pulse sequence optimization. (b) Block diagram for a pulse sequence architecture showing the individual subdwell parameters. Each row belongs to a specific spin channel, S (^{15}N) or I (^{1}H), and across the channel are the individual subdwells. Parameters ϕ,ψ are the phases for the odd and dwells, respectively. For FS symmetry, only parameters denoted by the black text are varied throughout the MCSA run; for PS symmetry black and green text are variable; the red text parameters are always fixed by minimal symmetry considerations. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

200 random steps per temperature, for a total of 20,000 random steps. Further refinement, when desired, was then run starting from each MCSA run's saved parameters for 10,000 additional steps.

Random steps are taken by choosing a particular parameter and then randomizing its value from a uniform distribution of its designated range. Two random integers are generated to index a specific parameter. The first random integer, *ind1*, is from 0 through 4. This integer represented one of the possible five (5) choices; 0: low-spin S-channel, odd dwell; 1: high-spin I-channel, odd dwell; 2: low-spin S-channel, even dwell; 3: high-spin I-channel, even dwell; 4: timings. The second random integer, *ind2*, chooses the subdwell to be sampled. Due to the symmetry restrictions across subdwells, the maximum value for *ind2* only needs to be half the number of subdwells (rounded up if there is an odd number of subdwells). At a minimum each parameter's symmetry was designated as PS. This means that setting the value for a subdwell in the first half of the dwell would also set the value for the mirrored subdwell in the second half of the dwell. Running all parameters with PS cuts down the true number of variables from 30 to 15, or 5 times the maximum value for *ind2*. For FS, the true number of variables becomes 9 since there is the same number of variables for timings (3), but half as many for phases (6). The variables for phases are further decreased because FS relates the subdwell variables between the even and odd dwells for each spin channel.

Once a parameter is specified by *ind1* and *ind2*, it is assigned a new value *rand* (step 9). As previously stated the range of *rand* is specific to its parameter type (phase or timing), and is drawn from a uniform distribution. The symmetry related parameters are automatically varied as well (step 9a,b). For phases, the mirrored subdwell is set to *rand* + 180°. If FS, then the opposite dwell for that channel (even or odd) is also set to *rand* + 180°, and lastly the parameter that has both opposite dwell and mirrored subdwell is set to *rand*. For timings the mirrored subdwell is set to *rand*.

The simulation is repeated three times for the new set of parameters, each simulation having a different heteronuclear dipolar Hamiltonian $H^{(IS)}$ that corresponds to a specific maximum coupling, and then is assigned a score. All the other steps that follow are standard MCSA procedure. Step 12 applies the Metropolis criterion for *dE*, which is the difference between the new score and the previous score. The scoring parameters are saved to file after every W^T increment completes. After all MCSA runs are completed the program terminates. The parameters can either be used for further refinement or converted into the final pulse sequence format for implementation on an NMR spectrometer.

3.4. NMR experiments

NMR experiments for the NAL crystal (18 mg) were carried out on a Bruker 300 MHz spectrometer operated by Topsin 3.2 software and using a static Doty Scientific™ 7 mm coil HX probe. The B_1 field was calibrated at 58.14 kHz and 2 scans were accumulated for each of 200 t_1 points, each having the architecture shown in Fig. 1b with the individual subdwell phases and timings calculated by the optimization, cf. Table 1 below. Experiments for the selectively Leu-labeled Pf1 coat protein in bicelles (ca. 4 mg) were carried out on a Bruker 500 MHz spectrometer operated by Bruker 2.0 software and using a static 5 mm coil HCN Bruker E-free™ probe. The B_1 field was calibrated at 49.5 kHz and 256 scans were co-added. Spectra were processed using NMRPipe [22]. Linewidth calculations were performed using the signal function in Python's Scipy library. The acquisition parameters were held constant between all experiments runs.

4. Results

4.1. MCSA simulations

MCSA trials were often run in succession, with the final scores typically ranging from −10 to −30. The resulting family of computationally optimized pulse sequences has been termed "ROULETTE" (Random Optimization Using the Liouville Equation Tailored To Experiment). Compared to SAMPI4, whose simulated *pmraw* score was −10.59, the top experimentally performing ROULETTE pulse sequence, termed ROULETTE-1, had a *pmraw* score of −14.43. A direct comparison between the simulations of this pulse sequence and SAMPI4 is provided in Fig. 2.

In silica, ROULETTE-1 has a *pmraw* score roughly 33% greater than that of SAMPI4. Since the scaling factors can vary widely between pulse sequences, retaining the correct relative ratios of the DC frequencies is most important.

The scaling factors for the ROULETTE sequences, *scf,* were calculated as the slope in a linear regression between the expected frequencies (from the maximum DCs used in the simulations) and the actual measured (or apparent) frequencies resulting from the simulated spectrum. With the dwell defined per Eq. (14), the numerical *scf* for SAMPI4 was 0.709 (which is to be compared with the theoretical result of $8 \times 1.0927/12 = 0.7285$ from Ref. [6]) and for ROULETTE-1 *scf* = 0.587. This ensured that the resonances for a given maximum coupling would all appear at approximately the

Table 1
Phases, timings, and mean linewidths for ROULETTE optimized sequences vs. SAMPI4.

Exp. #	ϕ_{11}	ψ_{11}	ϕ_{12}	ψ_{12}	ϕ_{13}	ψ_{13}	ϕ_{21}	ψ_{21}	ϕ_{22}	ψ_{22}	t_1(μs)	t_2(μs)	t_3(μs)	μ(Hz)
SAMPI4	0	180	0	180	0	180	0	180	90	270	15.05	4.30	6.45	213
ROULETTE 1	60	255	168	314	118	358	89	139	184	47	15.74	4.58	4.57	178
2	209	64	294	147	126	138	259	323	355	227	15.16	5.05	2.75	168
3	157	164	68	234	10	285	132	215	41	112	14.03	4.98	1.94	173
4	339	120	43	230	47	53	354	61	267	149	7.41	5.56	1.40	183
5	208	70	294	143	190	148	188	259	281	167	14.42	4.94	2.73	187
6	76	292	340	224	332	213	26	35	313	104	7.55	4.77	3.40	188
7	153	336	79	72	295	182	105	15	3	301	13.78	5.43	1.66	191
8	195	88	279	157	188	265	165	233	258	138	15.16	5.17	2.86	193
9	195	70	261	149	91	261	189	259	284	162	14.95	5.38	2.46	195
10	181	1	135	315	163	343	127	307	222	42	16.21	4.89	6.31	197
11	235	36	300	106	17	205	137	64	47	155	15.36	5.11	3.37	199
12	149	303	80	239	342	121	161	83	71	172	15.34	5.16	3.49	199
13	342	115	43	231	56	120	354	63	267	149	7.44	5.56	1.18	202
14	25	221	293	287	119	51	45	129	323	31	13.68	5.39	1.55	204
15	162	149	102	213	319	275	145	230	68	128	14.25	5.72	1.99	205
16	36	243	140	296	81	21	49	118	139	29	15.68	4.94	4.04	212
17	26	146	149	8	332	206	21	225	113	305	16.30	5.14	5.40	212

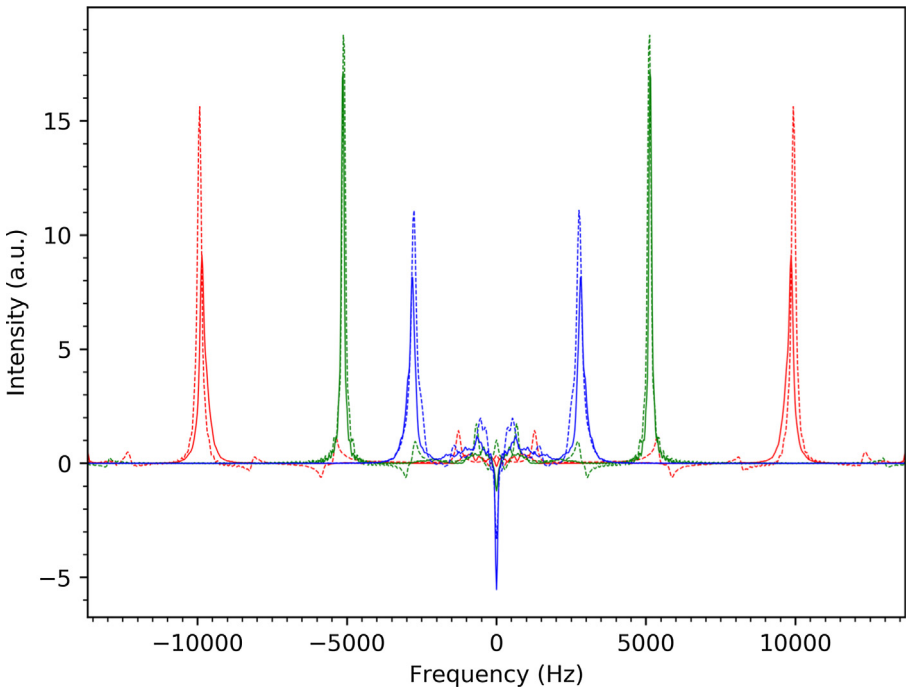

"True" frequency (Hz)	2500	5000	10000
SAMPI4	2830	5127	9853
ROULETTE-1	2755	5122	9926

Fig. 2. MCSA-optimized simulation for the ROULETTE-1 pulse sequence (dashed lines), vs. the simulation for the SAMPI4 pulse sequence (solid lines) shown across the ^{1}H-^{15}N DC dimension. The sets of three different NH DCs set in each simulation are color-coded, with red corresponding to 20 kHz, green for 10 kHz, and blue for 5 kHz full DC. The raw scores of SAMPI4 vs. ROULETTE-1 are −10.59 and −14.43, respectively. The empirically determined scaling factors for the SAMPI4 and ROULETTE-1 are 0.709 and 0.587, respectively. The accompanying table displays the calculated values for the apparent frequencies for each simulated experiment with the scaling factors taken into account. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

same locations across the different pulse sequences, cf. Fig. 2. The simulations also predict that, while the two pulse sequences match the expected frequencies for the intermediate couplings (e.g. 10 kHz), the ROULETTE-1 resonances are closer to the "true" couplings for both the higher (20 kHz) and lower (5 kHz) couplings than for SAMPI4.

4.2. Experimental results for NAL crystal at 300 MHz ^{1}H frequency.

More than one hundred (100) ROULETTE-family pulse sequences were tested experimentally using a single NAL crystal. The top seventeen pulse sequences that improved on the mean dipolar linewidths, μ's, are shown for comparison with SAMPI4, and each other, in Table 1.

In Table 1, only the phases and timings are shown for subdwells 1–3; the values for subdwells 4–6 are determined by the symmetry relationships as described in the Methods section and Fig. 1b. Furthermore, since the rf amplitude on the ^{1}H-channel is off during subdwells 3 and 4, the corresponding phases need not be specified.

A side by side comparison of the interferograms and spectra for SAMPI4 and the best performing ROULETTE pulse sequence, i.e. ROULETTE-1, is presented in Fig. 3. The interferogram in Fig. 3a shows much more persistent dipolar oscillations for all peaks in ROULETTE-1 as compared to SAMPI4. For instance, for peak 4 the signal has completely dissipated in SAMPI4 after ~75 t_1 points, while the signal still remains prominent even at ~125 t_1 points in ROULETTE-1. In Fig. 3b the SLF spectra for the two experiments are shown superimposed, along with respective FWHH linewidth for each peak. As predicted by the simulations presented in Fig. 2, ROULETTE-1 yields greater observed dipolar couplings than SAMPI4 at the higher frequencies and slightly lesser couplings at the lower frequencies. The simulations also suggest that ROULETTE-1 measured couplings are closer to the true couplings of the backbone NH bonds. The distributions of linewidths for the individual DC peaks present in NAL are shown for the seventeen best ROULETTE experiments in Fig. 4. The linewidth values for SAMPI4 are shown as horizontal lines for comparison with the ROULETTE experiments.

While all experiments included in the plots had better *average* linewidths than SAMPI4, they do not necessarily improve over SAMPI4 for *every* peak. The optimized pulse sequences perform particularly well at high couplings, cf. Fig. 4 (c,d), while SAMPI4 better resolves the lower couplings Fig. 4(a,b). The mean linewidth amongst the 4 peaks for ROULETTE-1 was 178 Hz, which is 18% sharper as compared to 212.75 Hz for SAMPI4, with a marked improvement over SAMPI4 for the largest DCs, cf. Fig. 4d.

4.3. Robustness of the scaling factor with respect to ^{1}H frequency offsets

To further evaluate robustness of the ROULETTE optimized pulse sequences, a series of experiments were performed at various ^{1}H carrier offset frequencies. A common pitfall for composite-pulse sequences is the potential variability in the scaling factor, which could alter the measured DC values, thus affecting subsequent structure calculations. Previous experiments have shown that the scaling factor can be sensitive to ^{1}H carrier frequency [4,6,23,24]. Fig. 5a shows the sensitivity of the measured DCs to the proton carrier frequency for ROULETTE-1.

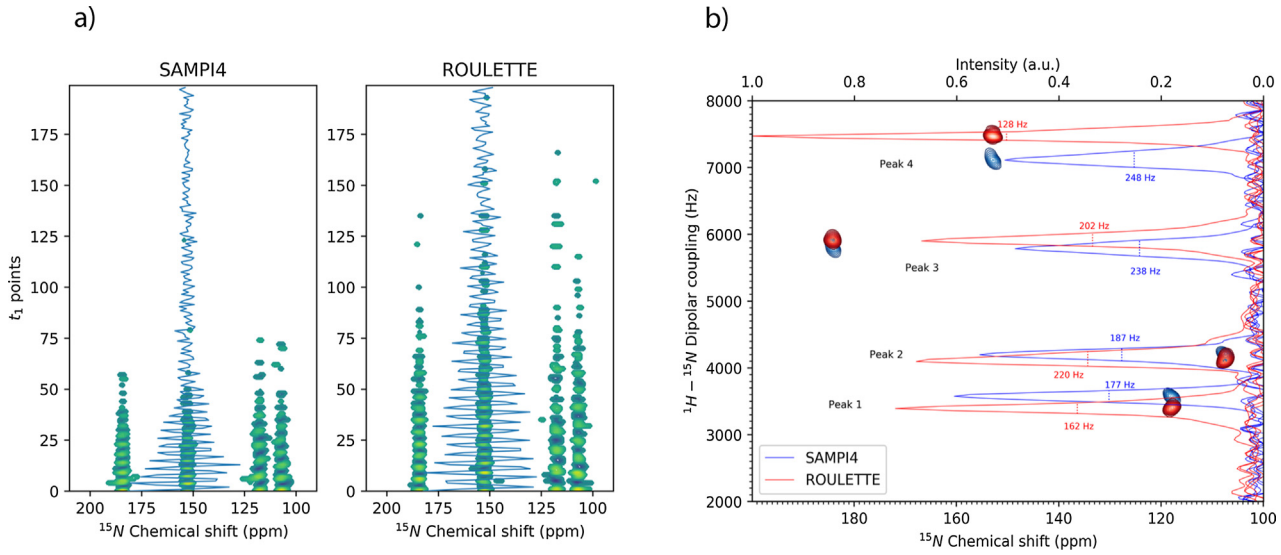

Fig. 3. (a) Side by side comparison of the interferograms for SAMPI4 and ROULETTE-1 experiments measured for the NAL crystal at 300 MHz ^{1}H frequency. The oscillations in the t_1 dimension through the 4th peak are shown in blue (b) An overlay of 2D SAMPI4 (blue) and ROULETTE-1 (red) spectra for all 4 NAL peaks. The additional horizontal axis on the top corresponds to the normalized ^{1}H-^{15}N DC spectra for each peak. The corresponding linewidths, calculated as FWHH, are listed beside each peak. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

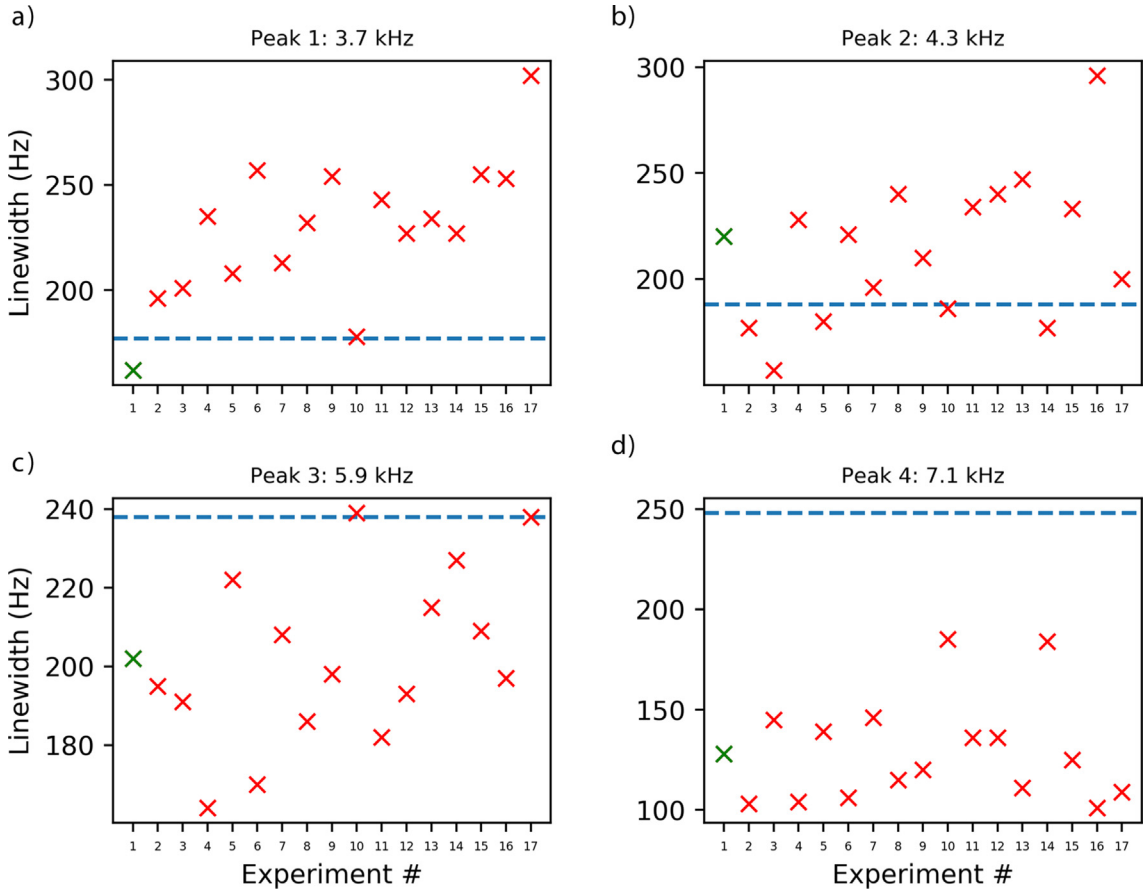

Fig. 4. Linewidths of the 17 best optimized ROULETTE pulse sequences for NAL crystal as compared to SAMPI4. Experiments are denoted by their relative number shown on the x-axis; ROULETTE-1 is highlighted by the green crosses. The individual plots (a–d) correspond to the peaks 1–4 in NAL. The linewidth values for SAMPI4 are represented by the horizontal dashed lines for comparison with the optimized experiments. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

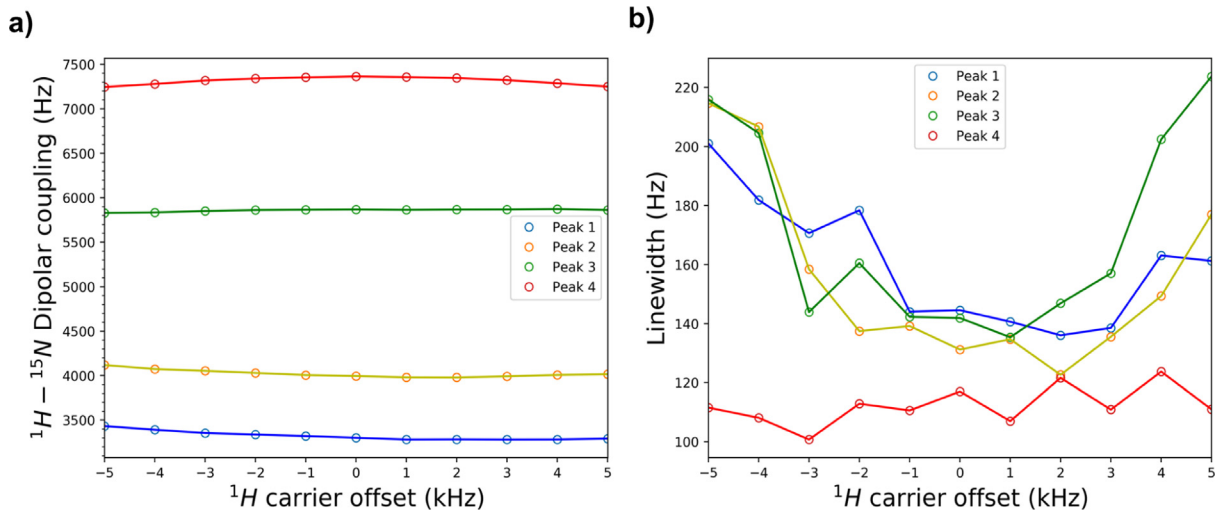

Fig. 5. Sensitivity of ROULETTE-1's DCs and linewidths to ^{1}H carrier frequency for 4 peaks of NAL crystal. The 0 kHz offset point corresponds to the optimum carrier frequency of 300.1267 MHz. The other 10 experiments were performed at ^{1}H carrier frequencies above and below the optimum center frequency, in increments of 1 kHz. For peaks 1–4 in (a), the ranges from the minimum to the maximum DC are 151, 140, 43, and 119 Hz, respectively (see legend for the color coding of peaks). (b) Sensitivity of the linewidths to the offset for the 4 peaks (same color coding for peaks as in a).

Relative to the DC range of roughly 4.5 kHz, the deviations in the measured coupling from the optimal carrier ^{1}H frequency of 300.1267 MHz are small, and comparable to the respective variances in SAMPI4 [6]. Fig. 5b shows the distribution of the measured linewidths at the same carrier frequencies. Peaks 1–3 tend to have the sharpest linewidths around the center carrier frequency, while the linewidths for peak 4 do not vary appreciably.

To verify that the optimized pulse sequence is still effective at other crystal orientations, the sample was rotated and the ROULETTE-1 sequence was compared to SAMPI4.

The second crystal orientation, shown in Fig. 6, is significantly more challenging than the one pictured in Fig. 3. Two of the four couplings have significantly lower DC frequencies than the original orientation, and are generally more difficult to evolve than moderate to high DCs. Nevertheless, it can be seen from Fig. 6 that ROULETTE-1 still yields much better signal to noise ratio than SAMPI4 at this crystal orientation, especially for the larger couplings.

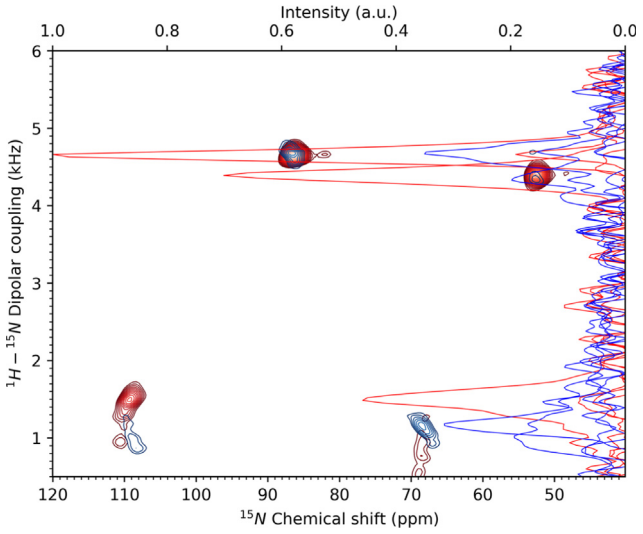

Fig. 6. ROULETTE-1 (red) vs. SAMPI4 (blue), for an alternate NAL crystal orientation. Compared to the previous orientation, shown in Fig. 3, the DCs for the 4 peaks have much lower frequencies. The lower x-axis corresponds to the chemical shift anisotropy, while the normalized intensities for dipolar slices, taken at the respective chemical shifts of the 4 peaks, are measured on the top x-axis. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4.4. Pf1 spectra at 500 MHz ^{1}H frequency

The SAMPI4 and refined ROULETTE-1 pulse sequence were then run on Leucine labeled Pf1 coat protein reconstituted in magnetically aligned bicelles at 500 MHz ^{1}H frequency. To conform with the calibrated pulses at 500 MHz, the sequence was refined starting from the previous phases and timings, but using $\omega_{rf} = 49.5$ kHz instead of 58.14 kHz. This resulted in a sequence with only small phase differences (within 20°) relative to the original ROULETTE-1 parameters that were optimized for 300 MHz ^{1}H frequency. The differences can be seen in Table 2 where the two pulse sequences are listed by their phases and timings.

The pulse sequence at 58.14 kHz has the same parameters as the ROULETTE-1 sequence from Table 1. The greatest deviation between the two pulse sequences is in the third odd subdwell for the ^{15}N channel (ϕ_{13}), where the subdwell phase at 49.5 kHz is ca. 18° more than the corresponding phase at 58.14 kHz. Naturally the timings are longer with a lower rf amplitude, but the ratios between the durations t_1, t_2, and t_3 remain roughly the same, albeit with a greater disparity between t_2 and t_3, as compared to 58.14 kHz. The comparison of SAMPI4 and the ROULETTE-1 sequence refined at $\omega_{rf} = 49.5$ kHz at 500 MHz ^{1}H frequency is pictured in Fig. 7.

The improvement in linewidths for ROULETTE-1 over SAMPI4 is slightly less for Pf1 than is observed for the NAL crystal (14% vs. 18%). However, this improvement in linewidths yields a corresponding increase in the signal to noise ratio from 17:1 in SAMPI4 to 20:1 for ROULETTE-1, as measured by NMRPipe. Moreover, the ROULETTE-1 lineshapes appear free of the additional artifacts such as "shoulders and "humps" that are otherwise present in the SAMPI4 dipolar slices for Pf1, cf. Fig. 7.

5. Discussion

The present work addresses the following three main aspects: application of MCSA to optimize NMR pulse sequences by

Table 2
Comparison of ROULETTE-1 pulse sequences, optimized at ω_{rf} = 49.5 kHz vs. 58.14 kHz.

Exp.	ϕ_{11}	ψ_{11}	ϕ_{12}	ψ_{12}	ϕ_{13}	ψ_{13}	ϕ_{21}	ψ_{21}	ϕ_{22}	ψ_{22}	$t_1(\mu s)$	$t_2(\mu s)$	$t_3(\mu s)$
ω_{rf} = 49.5 kHz	60	250	159	309	136	9	91	146	186	53	17.34	5.20	4.41
ω_{rf} = 58.14 kHz	60	255	168	314	118	358	89	139	184	47	15.74	4.58	4.57

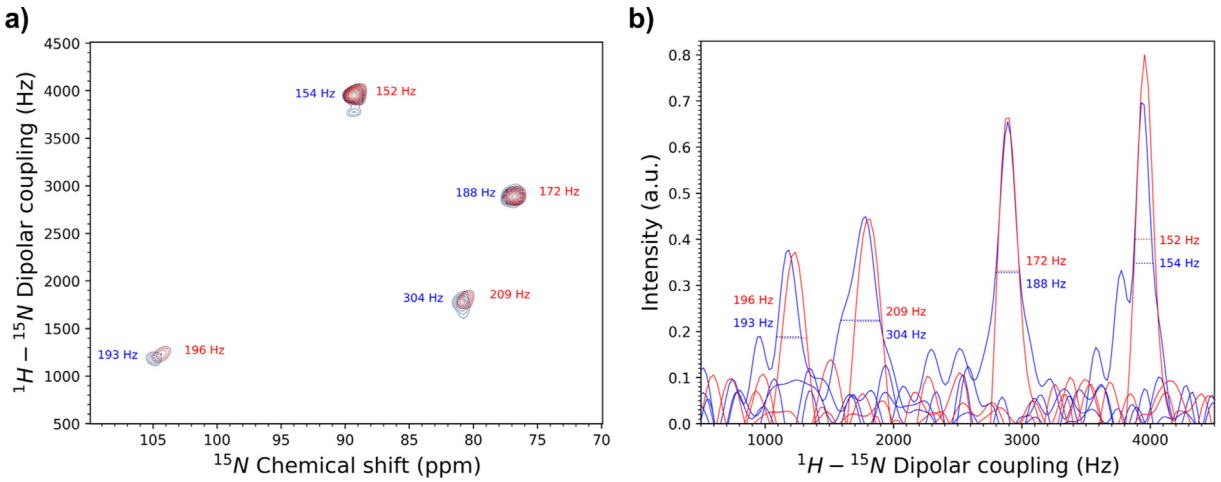

Fig. 7. Comparison of SAMPI4 (blue) and the re-optimized ROULETTE-1 (red) pulse sequence run on ^{15}N Leu-labeled Pf1 coat protein at 500 MHz ^{1}H frequency. (a) The 2D spectra with the ^{1}H-^{15}N DC linewidths labeled next to each peak. (b) Dipolar slices taken at the chemical shifts of each of the 4 peaks. Linewidth (FWHH) values are also displayed in Hz beside each peak. The mean peak width is 210 Hz for SAMPI4 and 182 Hz for the optimized pulse sequence, which is 14% sharper. This gain in resolution proportionally increases the signal-to-noise ratio from 17:1 for SAMPI4 to 20:1 for ROULETTE-1. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

computer simulations; experimental validation of the optimized pulse sequences using a static sample consisting of NAL crystal, and the performance of the pulse sequence on a much more dynamic biological system, i.e. Pf1 coat protein reconstituted in magnetically aligned bicelles. It was shown that an MCSA-guided parameter search can optimize a pulse sequence by minimizing an appropriately chosen scoring function (the product of peak intensities and the ratios of the DCs, cf. Eq. (16)). Rarely did a sufficiently long simulation fail to find a score that was better than that for SAMPI4 (often scores were greater than 1.5x that of SAMPI4). It was also rare for the simulation to converge to the same set of parameters in subsequent runs, which suggests that there are numerous local minima in the energy surface produced by the scoring function, Eq. (16). The best practice was to run the optimizations multiple times with different starting parameters so that the optimum parameter set would be largely determined by random chance, hence the pulse sequence name, "ROULETTE". For this purpose the program was written to run numerous MCSA simulations in succession, cf. step 4 in Fig. 1a. In that way multiple solutions would be obtained; the promising solutions could be further refined and the bad ones rejected. It is still important to note that a number of runs produced false positives, i.e. sharp peaks obtained in silica that failed to produce sharper resonances than SAMPI4 experimentally.

Running numerous shorter simulations helped identify the simulation parameters that will more likely yield sharper linewidths experimentally. The majority of the optimizations were run using the same architecture and rf amplitude scheme as SAMPI4. Increasing the number of subdwells not only increases the number of the parameters to be optimized in the simulations, but also provides more flexibility for the algorithm to find solutions with lower scores. Additional simulations contained as many as 15 subdwells, which allowed lower *pm* scores to be obtained (results not shown). Unfortunately, these results were not confirmed experimentally when run on NAL, often producing spectra with no discernable peaks. This is likely the result of the increased number of phase transients in between subdwells. It should be noted that increased number of subdwells can be of use when designing pulse sequences with continuous phase modulation [14]. This type of pulse sequence optimization is only effective when the difference in phases between the adjacent subdwells is small. The optimization method described in this paper, by contrast, allows adjacent subdwells to assume any phase, from 0 to 360°.

About 15% of the tested pulse sequences have succeeded in improving on SAMPI4's mean linewidth, and are listed in Table 1. Only one of the pulse sequences, ROULETTE-10, was constructed with FS symmetry, while the rest were only PS. A consistent pattern that can seen in the optimized pulse sequences arises from the timings, which tend to be roughly in a 3:1:1 ratio for subdwells 1–3, respectively (with a few exceptions). Amongst all sequences, each phase appears to be able to assume the entire range of values, i.e. 0–360°. Surprisingly, none of the optimizations have lead to the original SAMPI4 values, even when FS was imposed. Even more surprising is that there were no clusters of pulse sequences in which all phases were similar to each other (to within 10°). Overall, these results suggest that there may exist multiple ways of decoupling the ^{1}H-^{1}H dipolar interactions from ^{1}H-^{15}N interactions, especially when non-quadrature phases and arbitrary subdwell durations are employed.

The departures of the measured frequencies at low and, especially, high dipolar couplings between SAMPI4 and ROULETTE-1 (Fig. 3b) is largely in agreement with the simulations in Fig. 2. The simulations predict that ROULETTE-1 has greater fidelity in evolving the local NH couplings than SAMPI4, which is likely a consequence of the assumption of infinitely sharp refocusing 90-degree pulses utilized when deriving the scaling factor for SAMPI4 [6]. In addition, SAMPI4 is known to underestimate the values for larger (>7 kHz) dipolar couplings [6], especially at lower B_1 rf fields (<60 kHz). By contrast, the ROULETTE sequences have been optimized without such delta-pulse approximation. This feature of

the ROULETTE sequence could potentially result in more reliable measurements of the dipolar couplings for structure calculations with less uncertainty in determining the backbone folds of oriented membrane proteins [25,26].

Additional control experiments, cf. Section 4.3 in Results, further demonstrate the robustness of the ROULETTE-1 pulse sequence. A pulse sequence whose scaling factor is highly variable with the ^{1}H carrier frequency requires an extra optimization step in order to find the optimal frequency to achieve ^{1}H-^{1}H decoupling for all heteronuclear dipolar couplings, which may be difficult. It was shown that ROULETTE-1 has a largely consistent scaling factor for DCs in the relevant ^{1}H-^{15}N frequency range while maintaining sharp dipolar linewidths. The ROULETTE-1 pulse sequence was also tested at an alternate crystal orientation exhibiting much smaller DCs. Overall, the results show that ROULETTE-1 can evolve a wide range of DCs at various crystal orientations.

As an application to biological samples, the ROULETTE-1 pulse sequence was tested on ^{15}N Leucine labeled Pf1 coat protein reconstituted in magnetically aligned bicelles. The average improvement in linewidths over SAMPI4 was 14%, which is a slightly lesser improvement as compared to the tests on the NAL crystal. This may be due to the imperfections arising from the protein dynamics and greater B_1 field inhomogeneity over a larger sample volume (180 uL). Another possible explanation for the difference in results between the NAL crystal and Pf1 in bicelles is that ROULETTE-1 has not been experimentally optimized with the Pf1 sample at 500 MHz ^{1}H frequency. Whereas a single SLF experiment on the NAL crystal took only 30 min, a 1D experiment on the ^{15}N Leu-labeled Pf1 sample would take 44 h. This imposes practical limits on the feasibility of screening different optimizations/refinements using biological samples. However, the ROULETTE-1 spectrum of Pf1 coat protein still compares more favorably to SAMPI4 in the sense that in the former the dipolar peaks are more symmetric and are free of "shoulders". Although these features do not appear to appreciably influence the linewidths measured at half height, the extra intensity at the bases of the SAMPI4 peaks may potentially complicate overall spectral resolution and, therefore, the subsequent structural analysis.

6. Conclusion

We have presented a numerical strategy for optimizing NMR pulse sequences by minimizing a scoring function incorporating the desirable sequence outcomes (i.e. peak heights/widths and the relative DCs). The developed pulse sequence, termed ROULETTE-1 outperforms the existing SAMPI4 sequence, yielding sharper linewidths, better signal to noise ratio, while still producing robust DCs over a wide dipolar frequency range. It should be emphasized that ROULETTE-1 was not simply a refinement of SAMPI4, but rather a *de novo* pulse sequence optimization, i.e. the search was conducted largely free of assumptions about the phases of the subdwell pulses and their durations, except for the power application scheme of SAMPI4. The latter assumption has also restricted the overall number of subdwells, thus allowing for the more efficient and realistic search of the parameter space. The MCSA-optimized ROULETTE-1 pulse sequence was tested for two different samples, which yielded dipolar linewidths exceeding SAMPI4 over a broad ^{1}H carrier frequency range. The presented method is generally applicable to designing and optimizing NMR pulse sequences for a variety of experimental conditions and applications.

Declaration of Competing Interest

The authors declared that there is no conflict of interest.

Acknowledgements

The authors would like to thank Emmanuel Awosanya and Dr. Sergey Milikisiyants for providing the ^{15}N Leu labeled Pf1 sample and NAL crystal, respectively. This material is based upon work supported by the National Science Foundation under Grant No. MCB 1818240 and by U.S. Army Research Office under contract number W911NF1810363.

References

[1] C.H. Wu, A. Ramamoorthy, S.J. Opella, High-resolution heteronuclear dipolar solid-state NMR spectroscopy, J. Magn. Reson. A 109 (1994) 270–272.
[2] S.V. Dvinskikh, D. Sandström, Frequency offset refocused PISEMA-type sequences, J. Magn. Reson. 175 (2005) 163–169.
[3] S.V. Dvinskikh, K. Yamamoto, A. Ramamoorthy, Heteronuclear isotopic mixing separted local field NMR spectroscopy, J. Chem. Phys. 125 (2006) 034507.
[4] A.A. Nevzorov, S.J. Opella, A "Magic Sandwich" pulse sequence with reduced offset dependence for high-resolution separated local field spectroscopy, J. Magn. Reson. 164 (2003) 182–186.
[5] N. Sinha, C.V. Grant, S.H. Park, J.M. Brown, S.J. Opella, Triple resonance experiments for aligned sample solid-state NMR of C-13 and N-15 labeled proteins, J. Magn. Reson. 186 (2007) 51–64.
[6] A.A. Nevzorov, S.J. Opella, Selective averaging for high-resolution solid-state NMR spectroscopy of aligned samples, J. Magn. Reson. 185 (2007) 59–70.
[7] T. Gopinath, G. Veglia, Sensitivity enhancement in static solid-state NMR experiments via single- and multiple-quantum dipolar coherences, J. Am. Chem. Soc. 131 (2009) 5754–5756.
[8] S. Jayanthi, K.V. Ramanathan, 2(n)-SEMA-a robust solid state nuclear magnetic resonance experiment for measuring heteronuclear dipolar couplings in static oriented systems using effective transverse spin-lock, J. Chem. Phys. 132 (2010) 134501.
[9] S. Jayanthi, N. Sinha, K.V. Ramanathan, 2(4)-SEMA as a sensitive and offset compensated SLF sequence, J. Magn. Reson. 207 (2010) 206–212.
[10] M. Lee, W.I. Goldburg, Nuclear-magnetic-resonance line narrowing by a rotating Rf field, Phys. Rev. 140 (1965) 1261–1271.
[11] A. Bielecki, A.C. Kolbert, H.J.M. de Groot, R.G. Griffin, M.H. Levitt, Frequency-switched Lee-Goldburg sequences in solids, Adv. Magn. Reson. 14 (1990) 111.
[12] D.P. Burum, M. Linder, R.R. Ernst, Low-power multipulse line narrowing in solid-state NMR, J. Magn. Reson. 44 (1981) 173–188.
[13] J.Y. Cui, J. Li, X.M. Liu, X.H. Peng, R.Q. Fu, Engineering spin Hamiltonians using multiple pulse sequences in solid state NMR spectroscopy, J. Magn. Reson. 294 (2018) 83–92.
[14] D. Sakellariou, A. Lesage, P. Hodgkinson, L. Emsley, Homonuclear dipolar decoupling in solid-state NMR using continuous phase modulation, Chem. Phys. Letts. 319 (2000) 253–260.
[15] U. Haeberlen, J.S. Waugh, Coherent averaging effects in magnetic resonance, Phys. Rev. 175 (1968) 453–467.
[16] M. Leskes, P.K. Madhu, S. Vega, Floquet theory in solid-state nuclear magnetic resonance, Prog. Nucl. Magn. Reson. Spectrosc. 57 (2010) 345–380.
[17] A. Brinkmann, Introduction to average Hamiltonian theory. I. Basics, Concepts Magn. Resonance Part A 45A (2016).
[18] M. Bak, J.T. Rasmussen, N.C. Nielsen, SIMPSON: A general simulation program for solid-state NMR spectroscopy, J. Magn. Reson. 147 (2000) 296–330.
[19] M. Veshtort, R.G. Griffin, SPINEVOLUTION: a powerful tool for the simulation of solid and liquid state NMR experiments, J. Magn. Reson. 178 (2006) 248–282.
[20] H.J. Hogben, M. Krzystyniak, G.T.P. Charnock, P.J. Hore, I. Kuprov, Spinach - A software library for simulation of spin dynamics in large spin systems, J. Magn. Reson. 208 (2011) 179–194.
[21] N. Metropolis, S. Ulam, The Monte Carlo method, J. Am. Stat. Assoc. 44 (1949) 335–341.
[22] F. Delaglio, S. Grzesiek, G.W. Vuister, G. Zhu, J. Pfeifer, A. Bax, NMRPipe: a multidimensional spectral processing system based on UNIX pipes, J. Biomol. NMR 6 (1995) 277–293.
[23] K. Yamamoto, D.K. Lee, A. Ramamoorthy, Broadband-PISEMA solid-state NMR spectroscopy, Chem. Phys. Letts. 407 (2005) 289–293.
[24] Z. Gan, Spin dynamics of polarization inversion spin exchange at the magic angle in multiple spin systems, J. Magn. Reson. 143 (2000) 136–143.
[25] Y.Y. Yin, A.A. Nevzorov, Structure determination in "Shiftless" solid state NMR of oriented protein samples, J. Magn. Reson. 212 (2011) 64–73.
[26] J. Lapin, A.A. Nevzorov, Validation of protein backbone structures calculated from NMR angular restraints using Rosetta, J. Biomol. NMR 73 (2019) 229–244.