Available online at www.sciencedirect.com
ScienceDirect
Journal of Magnetic Resonance 192 (2008) 183–196
JMR Journal of Magnetic Resonance
www.elsevier.com/locate/jmr

Insights into homonuclear decoupling from efficient numerical simulation: Techniques and examples

Vadim E. Zorin [a], Matthias Ernst [b], Steven P. Brown [c], Paul Hodgkinson [a,*]

[a] Department of Chemistry, Durham University, South Road, Durham DH1 3LE, UK
[b] Laboratorium für Physikalische Chemie, Wolfgang-Pauli-Strasse, ETH Zürich, CH-8093 Zürich, Switzerland
[c] Department of Physics, University of Warwick, Coventry CV4 7AL, UK

Received 7 November 2007; revised 12 February 2008
Available online 16 February 2008

Abstract

A combination of techniques, including rational number synchronisation and pre-diagonalisation of the time-dependent periodic Hamiltonian, are described which allow the efficient simulation of NMR experiments involving both magic-angle spinning (MAS) and RF irradiation, particularly in the important special case of phase-modulated decoupling sequences. Chebyshev and conventional diagonalisation approaches to calculating propagators under MAS are also compared, with Chebyshev methods offering significant advantages in cases where the Hamiltonian is large and time-dependent but not block-diagonal (as is the case for problems involving combined MAS and RF). The ability to simulate extended coupled spin systems efficiently allows ^{1}H spectra under homonuclear decoupling to be calculated directly and compared to experimental results. Reasonable agreement is found for the conditions under which homonuclear decoupling is typically applied for rigid solids (although the increasing deviation of experimental results from the predictions of theory and simulation at higher RF powers is still unexplained). Numerical simulations are used to explore three features of these experiments: the interaction between the magic-angle spinning and RF decoupling, the effects of tilt pulses in acquisition windows and the effects of ''phase propagation delays'' on tilted axis precession. In each case, the results reveal features that are not readily anticipated by previous analytical studies and shed light on previous empirical observations.

Keywords: Solid-state NMR; H-1; Resolution; Numerical simulation; MAS; CRAMPS; Homonuclear decoupling

1. Introduction

Solid-state proton NMR spectra are typically poorly resolved due to the strength of the dipolar coupling network which, unlike in the solution-state, is not significantly averaged by molecular motion. Although fast magic-angle spinning (MAS) can substantially reduce ^{1}H linewidths, the dipolar linewidth (typically ~50 kHz for organic solids) is significant in comparison to spin rates achieved by currently available probes (up to 70 kHz).

^{1}H resolution can be greatly increased by using homonuclear decoupling sequences such as FSLG [1], PMLG [2] or DUMBO [3] to suppress the dipolar couplings, in combination with MAS to remove other anisotropic broadening terms, such as the chemical shift anisotropy (CSA). The efficient suppression of the dipolar couplings using homonuclear decoupling is critical for an increasing number of important solid-state NMR experiments such as those involving spin-diffusion between resolved ^{1}H sites [4], ^{1}H homonuclear double-quantum correlation [5,6], as well as heteronuclear correlation experiments [7,8].

Unfortunately understanding what resolution can be achieved in ^{1}H solid-state NMR is a formidable challenge, since it inevitably involves the multi-spin nature of the spin system. Although NMR lineshapes can be reasonably modelled by treating the effects of spin flips due to the couplings in stochastic terms [9], such models require empirical

* Corresponding author.
E-mail address: paul.hodgkinson@durham.ac.uk (P. Hodgkinson).

1090-7807/$ - see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmr.2008.02.012

factors ("spin diffusion" rates, etc.) which cannot be predicted *ab initio*. Moreover, analytical treatments of multi-spin systems rapidly become intractable, particularly if the time-dependence of the couplings due to sample spinning is included [10–12]. It has recently been shown, however, that numerical simulations can be made sufficiently efficient that the spectra from large spin systems (say 10 or more coupled spin -½) can be calculated in reasonable times on a desktop computer [13,14]. By combining numerical simulation with experimental data from model systems, a comprehensive picture of ^{1}H resolution under magic-angle spinning [15] could be derived.

In this paper, we consider the much harder problem of predicting ^{1}H lineshapes under simultaneous magic-angle spinning and multi-pulse decoupling. There are two separate aspects to this challenge. Firstly, the "experimental space" is considerably larger for the problem of combined sample rotation and RF decoupling. The experimental variables for magic-angle spinning are few (spinning rate and its stability, angle setting and its stability) and all are well controlled using current spinning technology. In contrast, decoupling using RF irradiation is subject to a variety of factors which are much harder to control and to characterise, such as variation of the B_1 field across the sample, transient responses associated with phase changes [16,17], etc. Each of these can potentially degrade the quality of the spectrum and the different factors may interact in non-intuitive ways.

The second aspect to the challenge is the difficulty of analysing experiments that involve two independent timescales, that is sample spinning and RF irradiation. Traditionally this problem has been avoided by ensuring that the RF decoupling has a much shorter cycle time, τ_c, than the period of sample rotation, τ_r. The MAS rate is just fast enough to suppress broadenings due to the CSA (which is typically modest for ^{1}H) and other "inhomogeneous" line-broadening factors. Hence the performance of the RF decoupling can be analysed in a "quasi-static" limit in which the time-dependence of the dipolar couplings from the spinning can be neglected. The other case in which significant analytical progress can be made occurs when the cycle times are simple multiples of each other. If, for instance, $3\tau_c = \tau_r$, then the complete spin Hamiltonian is periodic over the period τ_r, facilitating both analysis and numerical simulation. However such exact "synchronisations" of the rotor and RF periods are generally undesirable as they are often associated with "recoupling" conditions in which the time-dependence of the couplings due to spinning disrupts the averaging of the spin terms of the Hamiltonian from the RF decoupling [12], cf. also Fig. 7. But as we show below, there are sound experimental reasons to consider the general case of simultaneous RF decoupling and MAS away from the quasi-static limit while avoiding "resonances" from timescale synchronisation. Moreover, it is becoming increasingly important to find decoupling schemes that are effective at the faster MAS rates that are necessary at higher B_0 fields [18,19].

Progress can be made in analysing problems involving independent timescales using bimodal Floquet theory in which the Hamiltonian is expanded in distinct Fourier series for the sample spinning and RF decoupling time dependencies [20]; the phase-modulated Lee–Goldburg (PMLG) sequence has recently been analysed using this approach [21]. Bimodal treatments are highly complex, however, and it is difficult to relate general analytical formulae to the detailed behaviour of a specific experiment. Numerical simulation ought to provide an important bridge between exact analysis and experimental behaviour. Moreover it should be relatively straightforward to incorporate experimental deficiencies such as RF inhomogeneities into the simulations and gauge the effect of different experimental features on the effectiveness of decoupling. Indeed, simulations are potentially of great value in determining which potential factors are most important to describe experimental behaviour, and hence guide analytical work. However, the presence of multiple, incommensurate timescales is a severe challenge for simulation too. In contrast to "synchronised" cases where it is sufficient to calculate density matrix propagators for the evolution over a limited number of rotation or RF periods, it is necessary to integrate the evolution over the entire free induction decay to calculate evolution under an aperiodic Hamiltonian.

The first part of this article describes techniques which address the difficulty of simulating NMR problems with independent timescales. The combination of these individual "tricks" improves the calculation speed to the point where problems involving both MAS and RF can be calculated within times that are not much longer than those needed for problems involving MAS alone. The second part describes how these efficient simulations have been used to study multi-pulse homonuclear decoupling, focussing on three specific questions: the effect of RF inhomogeneity on homonuclear decoupling under static *vs.* spinning conditions, the effect of tilt pulses on the performance of windowed decoupling sequences and the effect of the "phase propagation delays" on the effective lock axis. In each case, the numerical simulations have provided fresh insights into the long-standing problem of achieving high resolution solid-state NMR spectra from abundant spins.

2. Techniques

The section describes the different techniques that have been used to improve the efficiency of calculation. Although some have been previously used individually, and are, for example, implemented in the SPINEVOLUTION program [14], the existing descriptions often lack the detail required for independent implementation. As the combination of all the techniques described are required to make simulations of multi-spin systems involving simultaneous RF and sample spinning a practical proposition, they are each described here in detail.

2.1. Rational number synchronisation

As discussed above, the treatment of problems involving independent periodicities is difficult for non-integral ratios x between the periods, τ_r for rotation and τ_c for the RF. In analytical work, bimodal Floquet approaches are required which do not restrict the values of x that can be considered. In numerical simulation, however, it is useful to consider values of x that correspond to rational numbers, i.e., the ratio between two integers, $x = N_r/N_c$. There is then a common period between the two time dependencies, $\tau_{sync} = N_r\tau_r = N_c\tau_c$, over which the evolution is uniquely defined. Provided that τ_{sync} is shorter than the time over which the evolution needs to be calculated then the efficiency of the calculation can be significantly improved by using well-established algorithms for computing the evolution under a periodic Hamiltonian [22,24,25].

Fig. 1 shows the permitted ratios $M : N$, $M < N$ for maximal Ns of 8 and 15 (where M and N correspond to N_r or N_c depending on whether $\tau_r > \tau_c$ or vice versa). As N_{max} increases, a greater number of values of x are permitted; the number of pairs is given by $N_{max}(N_{max} - 1)/2$ although a small proportion of these will correspond to the same ratio (whenever M and N have a common factor). The drawback of considering ratios that involve large integer multiples is that (in principle) these require computing the evolution over several cycle periods, proportionately increasing the time required to calculate the required propagation matrices. The lightness of the markers in Fig. 1 is thus related to the time required for calculation, with the simple ratios such as 1:2 and 1:3 being much quicker to evaluate than, say, 13:15.

The distribution of the rational numbers shown in Fig. 1 is far from uniform with some values of x lying much closer to a suitable ratio than others. Although the density of allowed ratios increases in general regions approximately quadratically with N_{max}, there are local regions around simple integer ratios where the gaps can be large, e.g., around $x = 0$ and $x = 1$. Here there is a "gap" of $1/N_{max}$, which only decreases relatively slowly with increasing N_{max}. Hence, although it is important for efficiency to select points at rational x, it may also be necessary to consider isolated points where this "synchronisation" cannot be achieved. A way of achieving reasonable computational efficiency even in this case is described at the end of this section.

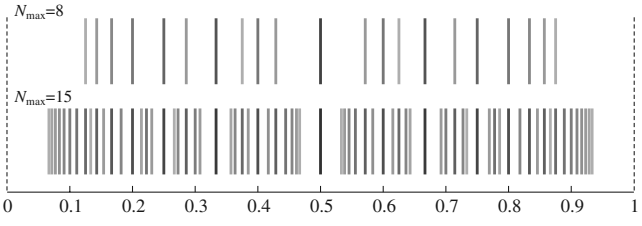

Fig. 1. Permitted rational numbers M/N for $M < N$ and $N \leqslant N_{max}$ for two values of N_{max}. The lightness of each marked ratio is proportional to the N for that ratio and so indicates the length over which the evolution needs to be evaluated, i.e., a lighter shade indicates a larger N and so a longer calculation time. M,N pairs which involve common factors are ignored. The distinct ratios for $N_{max} = 6$, for example, would be: 1:6, 1:5, 1:4, 1:3, 2:5, 1:2, 3:5, 2:3, 3:4, 4:5, 5:6.

The obvious concern with considering conditions that correspond to integer ratios of the cycle period is whether this re-introduces the potentially undesirable "resonances" between the timescales. It is clearly impossible experimentally to avoid all the integer ratios shown in Fig. 1—by definition, a given experimental ratio will be arbitrarily close to some rational number—but it could be argued that spatial and temporal experimental variations mean that an "unusual" matching condition would never be perfectly maintained for all the sample and for the complete experiment. As shown later, however, cf. Fig. 7, these "resonance" conditions are typically limited to a few simple ratios such as 1:N where N is an integer from one to four [20]. Hence using conditions which correspond to synchronised timescales is not expected to introduce systematic distortions.

2.2. Efficient simulation of phase-modulated RF sequences

The rotating frame Hamiltonian for RF of phase ϕ, and possibly including an offset from resonance Δ, can be derived from a corresponding Hamiltonian of another phase, say x, by a simple rotation about z:

$$H_{RF}(\phi) = \nu_{RF}[\cos\phi F_x + \sin\phi F_y] + \Delta F_z \quad (1)$$

$$= e^{-i\phi F_z} H_{RF}(0) e^{i\phi F_z} \quad (2)$$

where F_x, etc. are the sum operators for relevant nucleus. Since the spin system Hamiltonian at high field is invariant to z rotations, this means that the total spin Hamiltonian and exponentials of Hamiltonian (i.e., the density matrix propagators) can be derived by z rotations of corresponding matrices calculated for zero (x) phase [23]. The z rotation matrices, $\exp(i\phi F_z)$, are diagonal and so the transformation of Eq. (2) is computationally inexpensive. Moreover, the rotation matrices take up little storage space and can be conveniently cached for re-use.

This has important consequences. If, for instance, a pulse sequence consists of a series of RF pulses of the same duration that differ only in their phase, then the propagator for any step can be derived from a single propagator calculated with zero phase. Even if the system Hamiltonian is time-dependent but periodic, e.g., due to sample spinning, then it is still only necessary to compute the propagators of a suitable time step over the course of one rotor period and derive the required propagators by z rotation of the appropriate time-step propagators.

This is illustrated schematically in Fig. 2 for a 4:3 synchronisation between the RF and rotor cycle periods. The RF sequence involves steps with a common base duration, $\tau_c/4$, and the synchronisation between the RF and MAS periods implies that there is also a period, $\tau_{common} = \tau_c/12$, which is a common time base for both

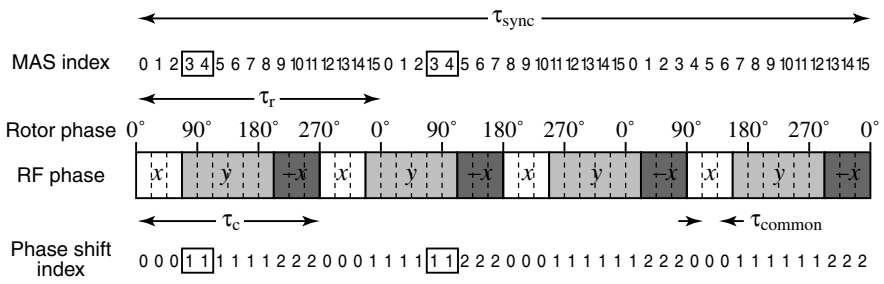

Fig. 2. Illustration of synchronisation between sample spinning and a periodic phase-modulated RF sequence. There is a common overall period of $\tau_{sync} = 3\tau_r = 4\tau_c$ and the phase changes occur at multiples of a base period $\tau_{common} = \tau_{sync}/48 = \tau_r/16 = \tau_c/12$. Rectangles mark an example of a repeated pair of MAS and phase shift indices (see text).

the phase modulation and rotation. Sixteen propagators can then be calculated over one rotation period in steps of τ_{common}, using zero phase for the RF. The propagator for a given τ_{common} step can then be quickly determined from the MAS propagator of the relevant index (0–15) shifted by the relevant RF phase (0–2). Propagators for arbitrary intervals (provided they are synchronised with τ_{common}) can be determined by accumulating these step propagators. As a result, even though calculation of the NMR signal requires the complete set of propagators over the full shared interval of three rotor periods, the time-consuming calculation of propagators is only required over a single rotor period. For simplicity we assume that the dwell time used for the observation of the NMR signal is commensurate with the base interval τ_{common} and the overall "synchronisation time", τ_{sync}. Since the spectral width can normally be chosen freely, this is not an unreasonable assumption and greatly simplifies the problem. A corresponding assumption is required by COMPUTE [24] and related algorithms [25,14] used for periodic problems.

It is occasionally possible that the accumulation of step propagators to determine the overall propagator for a time interval will involve multiplying the same pair of matrices. This potentially allows common matrix multiplications to be combined and further efficiency gains made. However, the original proposal [26] was made in the context of magic-angle spinning NMR in the absence of RF, in which case the propagators are uniquely defined by the MAS index. In the current situation where step propagators have both an MAS index and an RF index, these combinations can occur but are much rarer. In Fig. 2, for example, the pairing of propagator with MAS index 3 and phase index 1 followed by propagator with MAS index 4 and phase 1 occurs twice over the synchronisation period. However such combinations are rare and only occur here because the same phase (y) is used for more than one RF base period of $\tau_c/4$. Hence the time savings are expected to be minimal for non-trivial phase modulation periods and do not seem to justify the complexity of implementation.

The restriction of this technique to pure phase-modulated RF can be relaxed if the components of the sequence can be divided into a limited number of RF "states" and propagators over a rotor cycle are calculated for each of these states. For instance, windowed DUMBO [27] could be described in this framework by calculating two sets of propagators—one for the RF on and the other for zero RF—provided that there was a common time base between the phase modulation step and the duration of the windows. Clearly the efficiency advantage falls linearly as the number of "states" increases and applying this approach to, say, amplitude-modulated RF would clearly be inappropriate.

2.3. A quick-and-dirty method for continuous phase-modulated sequences

The discussion above suggests that calculations involving phase-modulated RF and magic-angle spinning can be evaluated with comparable degrees of efficiency since the propagation only needs to be evaluated over a single rotor cycle in both cases. There is, however, a penalty associated with cases where the common time base between the phase modulation and the rotation, τ_{common}, becomes very small. This will be associated with "poor" synchronisation conditions and/or phase modulations with short time steps. In these cases τ_{common} may be smaller than the time step used for propagator integration, Δt, which is typically of the order of 1 μs. In these cases, the total number of step propagators evaluated over the rotor period will exceed the number required to accurately capture the time-dependence of the Hamiltonian, leading to a loss of computational efficiency. As an example, consider a phase modulation such as DUMBO [3] which typically uses 128 phase steps in the complete RF cycle, τ_c. If the synchronisation condition is $7\tau_c = 2\tau_r$, then the largest possible common time base is $\tau_{common} = \tau_r/(7 \times 64)$. This corresponds to calculating propagators for intervals of 0.22 μs for an MAS rate of 10 kHz ($\tau_r = 10$ μs). If the integration of the propagator requires a time resolution of, say, 1 μs to accurately capture the time dependence of the Hamiltonian, then many more step propagators will be calculated over the single rotor period relative to a more conventional calculation.

However, much of the computational effort in calculating several propagators over periods very much smaller than the integration time step is redundant. If we are

assuming that the Hamiltonian is effectively constant over this interval, why re-diagonalise the Hamiltonian at each point? Hence we propose a modified algorithm which initially calculates and diagonalises the Hamiltonian over a set of points p, where τ_r/p is the time resolution. Whenever a propagator is required for an interval (t_1,t_2), the previously calculated eigenbasis and eigenvalues for the time point nearest the midpoint $(t_1+t_2)/2$ are recovered and used to calculate the propagator (cf. Eq. (3)), whose RF phase can then be adjusted as required. This avoids a time-consuming diagonalisation. The drawback to this approach is the introduction of "timing jitter"; the time point for the correct Hamiltonian eigenbasis and for the eigenbasis actually used differ by as much as $\tau_{jitter} = \tau_r/2p$ and this variation will differ from point to point of the calculated signal. This error can be compared to the error introduced by dividing the time evolution into finite integration time steps, although the nature of the error is different; in one case the time points at which the Hamiltonian is sampled are regularly spaced and the "error" arises from their finite spacing, while in the other case the time points "jitter" about their ideal values.

In practice, and as confirmed by detailed analysis [28], the calculated NMR spectra converge rapidly with decreasing integration time step (the error terms scale as the *square* of Δt). Fig. 3 illustrates the convergence of the calculation method for PMLG-17 homonuclear decoupling on a nine spin (three unit cell) geometry previously used to model ^{1}H lines shapes in regular 3D lattices such as adamantane [15]. The simulations were performed on Linux-based PC equipped with dual AMD Athlon 2800+ processors using an in-house simulation package specifically developed for handling complex multi-spin problems [29]. The calculation using the methodology discussed in the previous section, Fig. 3a, is extremely inefficient since the common timebase for the RF and magic-angle spinning, τ_{common}, is much smaller than the integration time step, Δt. In this case, the τ_r:τ_c ratio is 80:13 and the τ_{common} works out to be just 37 ns. Using the approach described above would lead to excessively long calculation times—much longer than the time required for a "conventional" calculation. Pre-diagonalising the Hamiltonian over one rotor cycle and allowing time points to move by up to 0.5 μs, Fig. 3b, reduces calculation times by a factor of about six. Fig. 3c shows how the spectrum approaches the "ideal" spectrum of Fig. 3a as this maximum time point jitter is reduced, at the cost of increasing the number of Hamiltonian diagonalisations and hence increased calculation time.

In many cases the dubious analytical characteristics of this method will not be significant, e.g., when comparing overall trends or crudely mapping out a problem. It would be less appropriate, however, to use this "quick-and-dirty" method when numerical accuracy is of paramount importance. For example, fitting data using gradient descent methods involves the numerical calculation of gradients, i.e., taking the difference of results calculated with slightly modified starting parameters. This difference calculation is very sensitive to numerical instabilities and so should not be combined with underlying algorithms with poorly defined convergence properties.

2.4. Calculation of propagators

Use of the methods described above helps to keep the number of propagator calculations required to a minimum. However, the resulting numerical simulations will only be practical if the time required for calculating propagators in multi-spin problems is reasonable.

The time-consuming step in most numerical simulations involves calculating the propagator, U, for a time interval, Δt, given the Hamiltonian at the midpoint of the interval, i.e., $U(\Delta t,0) = \exp(-iH(\Delta t/2)\Delta t)$. Although there are many ways to compute matrix exponentials [30], the con-

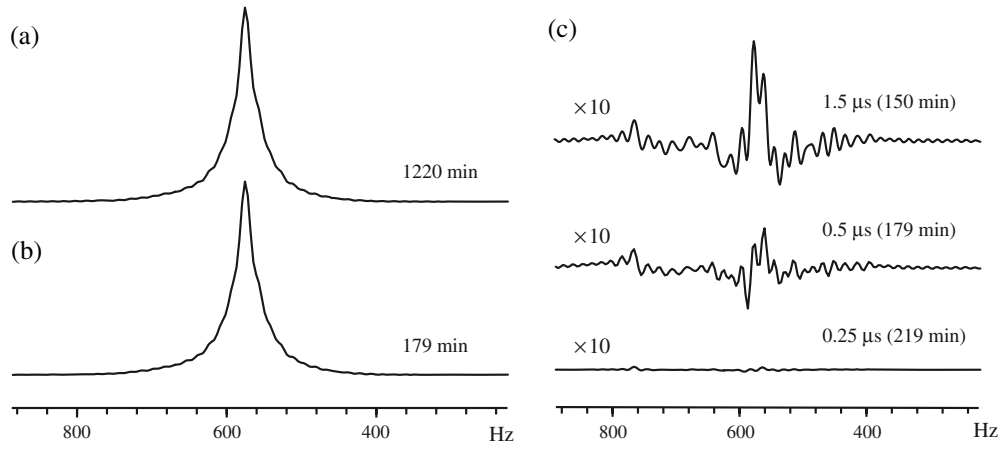

Fig. 3. (a) "Conventional" simulation of the ^{1}H spectrum of a nine spin adamantane-like spin system under 100 kHz PMLG-17 decoupling, (b) calculation in which time points are allowed to "jump" by up to 0.5 μs to the nearest time point for which the Hamiltonian has been pre-diagonalised. (c) Difference between the quick-and-dirty and conventional calculations as a function of the maximum jump time τ_{jitter} (figures in parentheses are total calculation times). 4096 data points were acquired, one per complete PMLG cycle. The spectrum was integrated over 50 crystallite orientations selected using a 3 Euler angle ZCW set.

ventional route in NMR simulations has been via matrix diagonalisation:

$$U(\Delta t, 0) = V^{\dagger} \exp(-\mathrm{i}\Lambda\Delta t) V \tag{3}$$

where V is the eigenvector matrix and Λ the diagonal matrix of eigenvalues of H. Since H is by definition Hermitian, the diagonalisation is numerically stable.

In some cases, however, it is clear that this route is not the most suitable. For example, when calculating Liouville space propagators, $\widehat{\widehat{U}} = \exp(\widehat{\widehat{L}}\Delta t)$, the Liouvillian, $\widehat{\widehat{L}}$, is often asymmetric (e.g., in the presence of exchange) and highly sparse. Diagonalisation of $\widehat{\widehat{L}}$ is both inefficient and numerically unstable. Calculation of the propagator using Chebyshev techniques is both more efficient and numerically more stable.

In "Chebyshev propagation" [31,28], the matrix exponential is expanded in terms of Chebyshev polynomials:

$$\mathrm{e}^{-\mathrm{i}H\Delta t} = \mathrm{e}^{\mathrm{i}H's} = \sum_{n}^{n_{\max}} a_n(s) T_n(H') \tag{4}$$

where s and H' are scaled versions of Δt and H, respectively (see below). The matrix terms of the expansion are defined by the recursive relationship

$$T_{n+1}(H') = 2H'T_n(H') - T_{n-1}(H') \tag{5}$$

with $T_1(H') = H'$ and $T_0(H') = \mathbf{1}$ (identity matrix). The coefficients of the expansion are

$$a_n(s) = (2 - \delta_{n0}) i^n J_n(s) \tag{6}$$

where $J_n(s)$ are the n-th order Bessel functions.

The utility of this expansion lies in the very rapid convergence of the coefficients $a_n(s)$—typically decreasing by an order of magnitude per iteration. This is only true however when $-1 < s < 1$, and so it is vital to scale the Hamiltonian such that the value of scalar s used for the polynomial expansion lies within this range. For the same fundamental reason, the expansion will only converge if $(||H\Delta t|| = ||H's||) < 1$ otherwise the $T_n(H')$ increase more rapidly than the $a_n(s)$ coefficients decrease. In other words, Chebyshev propagation is only stable for propagators calculated over relatively short time periods. This is not a significant restriction, however, since describing propagation under a time-dependent Hamiltonian by splitting the propagator into piecewise steps is only valid in this limit in any case.

Scaling of the Hamiltonian requires the range of eigenvalues of H to be known approximately. If $\omega_{\max}$ and $\omega_{\min}$ are the (approximate) maximum and minimum eigenvalues, then

$$H' = -\frac{H - \bar{\omega}\mathbf{1}}{\Delta\omega} \quad \text{and} \quad s = \Delta\omega\Delta t \tag{7}$$

where $\Delta\omega = (\omega_{\max} - \omega_{\min})/2$ and $\bar{\omega} = (\omega_{\max} + \omega_{\min})/2$. The mean of the eigenvalue range in principle contributes a phase factor of $\exp(-\mathrm{i}\bar{\omega}\Delta t)$ to the final propagator, but this phase can be neglected since it has no impact of the evolution of the density matrix. On a similar practical note, the target value of s can be chosen in advance (say 0.5) and the scale factor adjusted according (subject to the condition that it exceed $\Delta\omega$). This allows the a_n coefficients to be precalculated.

The eigenvalue range of H is efficiently estimated using the Gershgorin circle theorem. Specialised for Hermitian H, this states that the eigenvalues of H must lie within at least one of the ranges

$$H_{ii} \pm \sum_{j \neq i} |H_{ij}| \tag{8}$$

where i runs over the rows of H. The overall minimum and maximum values of Eq. (8) give $\omega_{\min}$ and $\omega_{\max}$.

The time-consuming step of Chebyshev propagation is the matrix–matrix multiplication of Eq. (5). Hence if H is a dense matrix, then the computational burden is still $O(N^3)$. Since several iterations are required for convergence, Chebyshev techniques tend then to be *less* efficient than classical diagonalisation methods, cf. Fig. 5. This is the case for free precession Hamiltonians involving homonuclear couplings. Provided that the diagonal block structure of the free precession Hamiltonian is exploited [22], then the diagonal blocks contain many non-zero elements due to the flip-flop terms of the homonuclear couplings. In contrast, the Hamiltonian for problems involving RF irradiation is no longer block diagonal, cf. Fig. 4. As a consequence, the eigenvector matrix, V, is full when the matrix is diagonalised and calculation of the step propagator via Eq. (3) must involve the multiplication of full matrices.

In contrast, Chebyshev propagation involves, in effect, the calculation of a power series in H. If H is sparse, then larger powers of H will become increasingly less sparse, but may not become completely full. If this sparsity is exploited when calculating the $T_n(H)$ matrices, then Chebyshev propagation can become competitive against diagonalisation. Fortunately simulations involving RF irradiation fall into this category; the block structure of H is "band diagonal" with the RF terms responsible for the ± 1 diagonal blocks. This is illustrated in Fig. 4. The $(H')^n$ (and $T_n(H')$) matrices steadily fill up with non-zero blocks as n increases. The matrices will be full once the number of Chebyshev iterations reaches the maximum coherence order supported by the irradiated spins. Before this point is reached, however, the presence of known empty blocks significantly reduces the time required for multiplication.

Previous applications of Chebyshev propagation to NMR problems in Liouville space have employed specialised sparse matrix methods for the storage and multiplication of Liouville space operators. In contrast to the highly sparse matrices of Liouville operators, the Hilbert space matrix representations used here are considerably denser and have a well-defined block structure (cf. Fig. 4). Moreover, the final propagators are relatively dense and so there is little advantage from the viewpoint of storage to the use of sparse matrices. Hence the matrices used here were stored as full matrices and the known block structure simply used to determine which blocks were non-zero and

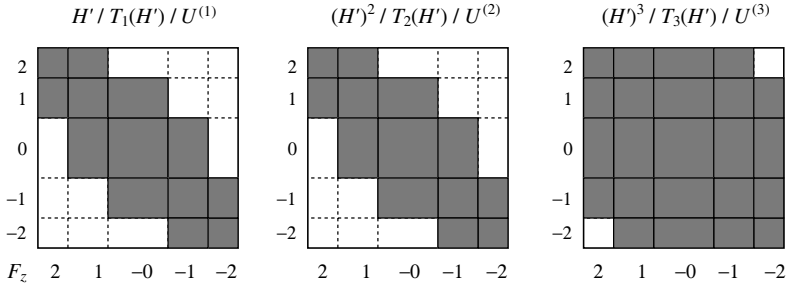

Fig. 4. Schematic illustration of the block structure of the Hamiltonian powers, $(H')^n$, Chebyshev matrices, $T_n(H')$, and n-th order approximation to the propagator, U^n, as a function of the number of iterations n for a 4 spin-½ system. The block structure is determined with respect to the total spin quantum number associated with the irradiated spin, with filled rectangles corresponding to non-empty blocks. The original (scaled) Hamiltonian matrix, H', has non-empty blocks along the diagonal (corresponding to the free-precession system Hamiltonian) plus blocks along the ± 1 diagonals corresponding to the RF Hamiltonian, i.e., the allowed NMR transitions, $\Delta F_z = \pm 1$. Higher diagonals are successively filled in higher powers of H'. The Chebyshev matrices and the approximation to the overall propagator, U, are constructed from linear combinations of $(H')^n$ and so share this block structure.

which empty. This allows conventional, but highly optimised routines, for dense matrix multiplication [32] to be used rather than the specialised routines required for the multiplication of sparse matrices.

Note that the algorithm described in the previous section cannot easily be combined with the Chebyshev approach to calculating propagators. When propagator calculation proceeds via diagonalisation, the eigenvalues and eigenvectors of the Hamiltonian at the different time points can be stored, allowing the rapid calculation of the propagator for an arbitrary interval based around the given time point. "Pre-calculation" in the Chebyshev approach would requiring storing the complete Chebyshev series, T_n, for each point and (expensive) calculation of the Chebyshev coefficients, $a_n(s)$, for arbitrary values of s.

Fig. 5 illustrates the relative merits of Chebyshev propagation *vs.* classical techniques using simulations of PMLG-11 decoupling. The network geometry was based on a periodic lattice with unit cells of two spins each repeated in one dimension—a geometry that has been previously used to model ^{1}H lineshapes in methylene groups [15]. Periodic boundary conditions imposed on the coupling network reduce "edge effects" associated with modelling an extended network in terms of a limited number of spins. In addition, a periodic lattice allows the problem to be expressed in a symmetry-adapted basis set, greatly improving the computational efficiency [13,33]. The figure shows the time taken as the number of unit cells is increased both when the symmetry-adapted basis is used ("periodic" calculation) or when using the conventional basis ("non-periodic" calculation). As previously demonstrated, exploiting periodic symmetry greatly improves calculation times and/or allows larger spin systems to be considered. The main interest, however, lies in comparing the performance of Chebyshev and classical propagation techniques. As anticipated above, Chebyshev techniques perform more poorly than diagonalisation for small spin systems, but are increasingly advantageous as the size of the spin system increases above about 4 coupled spin-½ nuclei. (Note that the timings for the Chebyshev, non-periodic calculations were very close to those obtained with an independent simulation programme [14] implying that the particular implementation of Chebyshev propagation used is entirely reasonable.) It is noticeable that the advantage of Chebyshev propagation is less marked for the specialised periodic calculations. The use of the symmetry-adapted basis set for a network of N unit cells divides the operator matrices into N smaller matrices. As already noted, the diagonalisation approach is most efficient with smaller matrices, hence the advantage of Chebyshev propagation is reduced if additional symmetry is exploited. However, it is not always possible or appropriate to exploit periodic symmetry and in these cases the Chebyshev approach provides substantial performance gains for larger spin systems.

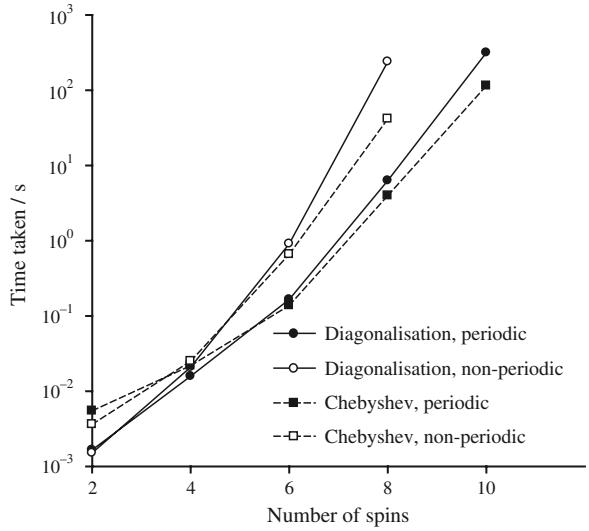

Fig. 5. Time taken per crystallite orientation for simulation of the ^{1}H spectrum under 100 kHz PMLG-11 homonuclear decoupling under MAS as a function of the number of spins (see text for calculation details). Calculations were performed in both conventional (open symbols) and symmetry-adapted basis sets (filled symbols) and using classical diagonalisation methods (solid lines) or Chebyshev methods (dashed lines) for calculating propagators. The MAS period ($\tau_r = 163.3$ μs) was matched to 10 times the PMLG cycle time.

3. Application to ^{1}H homonuclear decoupling

The design of efficient and robust pulse sequences for ^{1}H homonuclear decoupling is both an important but highly challenging topic of on-going research. Rather than attempt a comprehensive review of all the different factors that interact to determine spectral quality under homonuclear decoupling, the remainder of this paper will illustrate how numerical simulation can be used to shed light on the problem and address specific questions of interest.

We have previously shown that numerical simulation can effectively reproduce ^{1}H spectra under magic-angle spinning for a variety of network geometries without the need for adjustable or empirical parameters [15]. Fig. 6 performs a corresponding comparison between simulation and experiment for ^{1}H systems under homonuclear RF decoupling and magic-angle spinning. The plots compare full-width-at-half-height linewidths of numerically simulated spectra with results obtained from experiments in which a spin–echo was used to refocus decay due to inhomogeneous factors (chemical shift distributions, magnetic field inhomogeneity, etc.). The experiments were carried out on a Varian InfinityPlus spectrometer operating at 499.7 MHz for ^{1}H using microcrystalline powder samples packed into zirconia rotors with an outside diameter of 2.5 mm. Homonuclear decoupling was applied during the spin–echo delays and the intensity of the spin–echo signal measured as a function of the delay period. The decay was fitted to a simple decaying exponential to determine a nominal T_2' which was then expressed as a linewidth, i.e., $1/\pi T_2'$. T_2' is used to emphasise that the decay is not due to true T_2 relaxation [34] (and is often not simply exponential).

The numerical simulations used the 3×3 spin model previously described in association with Fig. 3, with the overall dipolar coupling strength, d_{rss}, matched to the corresponding system [33]. Since there are no inhomogeneous line-broadening factors in these idealised simulations, the spin–echo present in the experiments was omitted for simplicity of implementation; simulations on a test case confirmed that inclusion of an ideal π refocussing pulse had no effect on the decay of the simulated signal. Given that no adjustable parameters are present, the agreement between predicted and measured linewidths is reasonably good at low RF powers (as expressed by the nutation rate, ν_{RF}). The simulations also confirm that the linewidth under an "ideal" homonuclear decoupling sequence such as FSLG decreases as approximately ν_{RF}^{-2}. (The deviations from this dependence at lower RF rates can be associated with the "resonance condition" $\nu_{RF} \approx 4\nu_r$ cf. Fig. 7.)

However, the deviation between experimental and calculated values increases markedly as the RF power is increased. In particular, the experimental decoupling performance becomes worse at high RF nutation rates rather than improving as expected intuitively and confirmed by simulation. This breakdown occurs at much lower powers in the case of adamantane where the dipolar coupling is much weaker, while the cross-over occurs at nutation rates of about 100 kHz for the more typical organic solid. This difference reflects the intrinsically broader lines associated with stronger coupling networks; the increasing artefact level only becomes apparent when the dipolar linewidth has been sufficiently suppressed by high power RF decoupling. Although the origin of the deviations between experiment and simulation is not currently understood, the fact that very similar curves are observed in Fig. 6b for PMLG sequences with more or less coarse phase ramps suggests that they are neither directly related to increasingly short RF pulses (otherwise systematic trends would be observed going from PMLG-11 to PMLG-17 to PMLG-21 decoupling) nor due to subtle timing issues, i.e., whether pulse durations are simple multiples of internal clock times

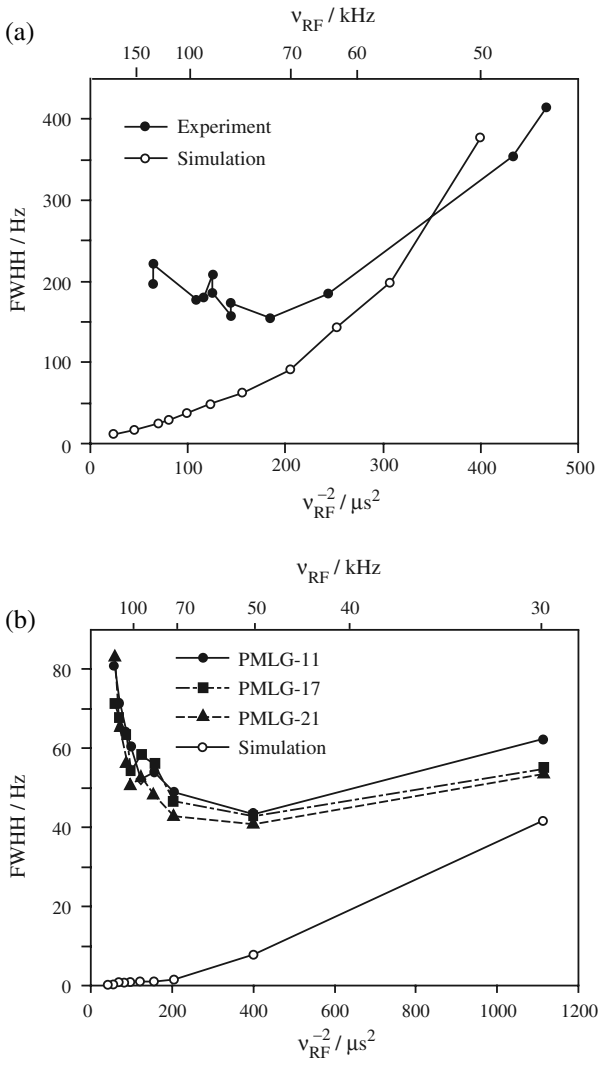

Fig. 6. Comparisons of ^{1}H linewidths under homonuclear decoupling and 12.5 kHz MAS determined experimentally (via spin–echo experiments) and from numerical simulation of PMLG-11 decoupling as a function of ν_{RF} for (a) the C–H resonance of alanine using PMLG-11 decoupling and (b) adamantane. The nine spin adamantane-like model was used in both cases but with an appropriate value of d_{rss}. Also shown in (b) are experimental results using different numbers of steps in the PMLG phase ramps.

(otherwise more erratic behaviour, depending on the matching of timing conditions, would be expected).

Although these multi-spin calculations provide theoretical reference points for comparison with experimental results in Fig. 6, they are not of themselves able to explain the observed experimental behaviour. In the examples below numerical simulation is used to address specific questions relevant to homonuclear decoupling. Under these conditions, the simulations are invaluable in rationalising experimental results.

3.1. Interaction between magic-angle spinning and homonuclear decoupling

The initial development of homonuclear decoupling sequences [35–37] focussed on suppressing dipolar linewidth in static samples. In Combined Rotation and Multiple-Pulse Spectroscopy (CRAMPS) experiments [38,39], magic-angle spinning is also applied to suppress other sources of line-broadening in solids, most notably the chemical shift anisotropy (CSA). But the magic-angle spinning rate was always modest to prevent interference between the periods of the rotation and RF irradiation, τ_c and τ_r, i.e., experiments were performed in the "quasi-static" limit in which $\tau_c \ll \tau_r$. Working in the quasi-static limit also greatly simplified the design of decoupling pulse sequences for CRAMPS since the magic-angle spinning could be effectively ignored. Analytical work also confirmed the destructive interactions between magic-angle spinning and homonuclear decoupling [20,40]. Although there have been isolated reports of decoupling techniques that take explicit account of magic-angle spinning [18,41], these approaches have not apparently been developed further.

Using the techniques described above, it is straightforward to examine the interaction between magic-angle spinning and homonuclear decoupling in simulation. Fig. 7 confirms that destructive interference occurs between FSLG homonuclear decoupling and MAS when τ_r/τ_c is a integer ratio less than five [20], and that efficient decoupling is only obtained if $\tau_r/\tau_c > 3$ (and the resonance condition at $\tau_r/\tau_c = 4$ is avoided). Note that multi-spin simulations are necessary to characterise this behaviour; calculations on simple two spin systems, for example, are unrevealing since the magic-angle spinning alone is sufficient to refocus the evolution under the dipolar coupling and so no sense of the performance of the RF sequence is obtained.

It is important to note, however, that these simulations only confirm that magic-angle spinning tends to degrade the performance of ideal homonuclear decoupling sequences that were developed in the quasi-static limit. This does not imply that MAS is always a negative influence. This can be seen clearly in Fig. 8b which compares FWHM linewidths for simulated ^{1}H spectra from the "alanine" model system under conditions of both FSLG and PMLG decoupling and static *vs.* MAS conditions. At low RF powers, the linewidths are indeed significantly worse in the presence of magic-angle spinning. As discussed above, however, this is associated with a resonance condition, and the performance of both FSLG and PMLG decoupling is noticeably *better* at higher RF powers than the corresponding results under static conditions. Under static conditions, where there is no interaction with sample spinning, the linewidth under FSLG scales perfectly with ν_{RF}^{-2}. This behaviour is not followed by PMLG, with the performance of PMLG-11 degrading significantly relative to an ideal phase ramp (i.e., FSLG) as the RF power increases. The most striking feature of Fig. 8 is that magic-angle spinning is effectively compensating for the approximate nature of the discrete phase ramp used in PMLG. This is consistent with experimental observations (obtained under MAS) where it was found that even very crude ramps, e.g., six phase steps in the complete PMLG cycle (PMLG-3) was sufficient to obtain good homonuclear decoupling [20]. Although detailed (and rather involved) multi-modal Floquet treatments may be required to elucidate the physical basis of this "cleaning up" of the imperfections associated

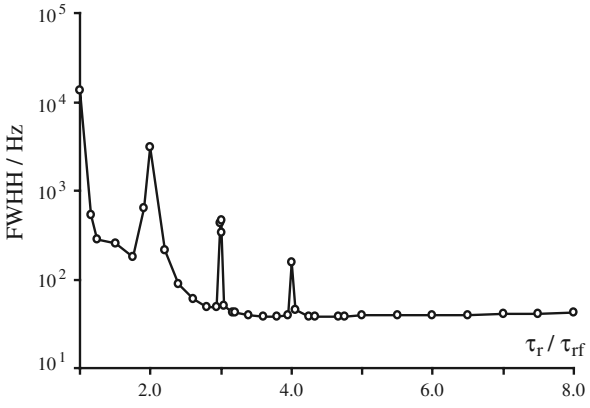

Fig. 7. Linewidth of the ^{1}H spectrum under 100 kHz FSLG decoupling as a function of MAS rate. Destructive interference occurs when τ_r/τ_c is a small integer ratio.

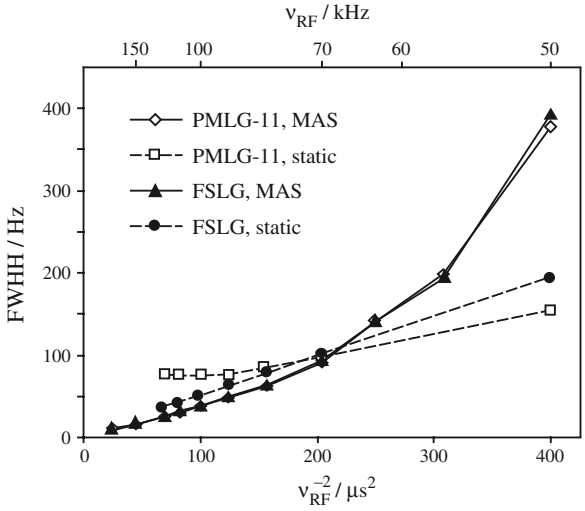

Fig. 8. Calculated linewidths for the alanine model system used in Fig. 6 under either PMLG-11 or FSLG homonuclear decoupling and under static or magic-angle spinning conditions (12.5 kHz) as a function of RF decoupling power, ν_{RF}.

with finite resolution phase, the simulations are invaluable in establishing the significance of interaction between sample spinning and RF irradiation.

Positive interactions between sample spinning and homonuclear decoupling are also demonstrated in Fig. 9. This considers the effects of RF inhomogeneity on the performance by FSLG decoupling by calculating the ^{1}H spectrum under FSLG decoupling as a function of the deviation of the actual RF nutation rate from a nominal value of 100 kHz. In the case of static samples, Fig. 9a, the spectrum is extremely sensitive to minor (5%) variations of the nutation rate. Hence inhomogeneity of the B_1 field would have a significant line-broadening effect. In contrast, the effect of deviations of the RF amplitude from its nominal value is much less marked under MAS conditions, Fig. 9b. Note that the narrowest lineshape is actually obtained when the actual RF nutation rate differs from its nominal value (at about 92 kHz).

These results have important practical consequences. Traditional CRAMPS experiments typically involved lengthy tune-up procedures to maximise resolution and minimise artefacts [42,43], including careful measurement of the RF nutation rate on static samples with narrow linewidths. Certainly in the case of FSLG decoupling, as shown in Fig. 9b, careful calibration of the nutation rate is relatively unimportant under MAS conditions and indeed the optimal conditions under MAS may be subtly different to those determined under static conditions. It is also worth noting the ease with which such valuable information about experimental parameters is derived from numerical simulation. Although FSLG (and the related sequence PMLG) has been analysed analytically using bimodal Floquet theory [20,21], the complexity of the analysis even for idealised conditions makes it difficult to determine how experiments will perform under less than ideal conditions.

3.2. Tilt pulses in windowed homonuclear decoupling sequences

A significant complicating factor in the use of homonuclear decoupling to improve ^{1}H resolution in the solid state is that the effective precession axis during the homonuclear decoupling period is tilted away from the z-axis, e.g., at the magic angle for Lee–Goldburg and frequency-switched Lee–Goldburg decoupling. The tilted plane of precession means that large quadrature images would be observed unless the precession plane is tilted into the xy plane prior to detection of the x and y components of the magnetisation. There are a limited number of decoupling sequences in which the precession axis is oriented along z [44–46], however, such sequences have not, to date, out-performed sequences which involve tilted axis precession, and they have not been widely adopted.

The use of tilt pulses either side of homonuclear decoupling periods in *indirect* dimensions is relatively straightforward, although the position of the effective precession axis needs to be clearly established if quadrature and zero frequency artefacts are to be minimised [27]. Tilt pulses are also required, in principle, if homonuclear decoupling is being applied during the acquisition dimension, i.e., the ^{1}H magnetisation is being sampled during gaps in a windowed sequence. It is found, however, that the optimised tilt pulses for detection in t_2 are generally very different from those optimal for t_1 [47,6]. In particular, the parameters describing the optimal tilt pulses in t_1 generally coincide with the values expected for the given decoupling sequence, while optimal results are obtained for detection in t_2 with much shorter pulses or by omitting tilt pulses completely.

Fig. 10 shows the results of simulations which examine the role of tilt pulses in the directly detected dimensions. As shown in Fig. 10b, an artefact free spectrum is obtained when using ideal (delta function) pulses of the "expected" duration, that is tilting through the magic angle. This simple picture changes, however, if finite-duration pulses (of the same RF amplitude as the homonuclear decoupling) are used, Fig. 10c. Although the quadrature artefact remains suppressed, the quality of the decoupling is degraded; the linewidth increases and so the signal height and sensitivity is also diminished. Decreasing the duration of the tilt pulses improves resolution at the cost of increasing the magnitude of the quadrature artefact, with the

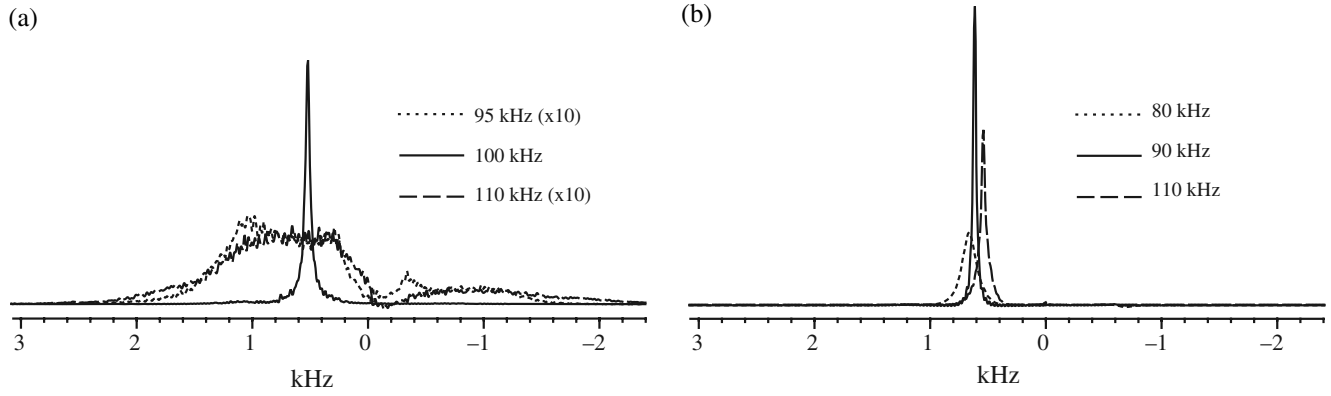

Fig. 9. Effect of deviation of the RF field strength from its ideal value on the ^{1}H spectrum of the adamantane-like spin system under FSLG decoupling: (a) static conditions, (b) under magic-angle spinning at 12 kHz. Timing parameters and RF offsets were set for $\nu_{RF} = 100$ kHz. The effect of deviations of the actual RF field rate from this nominal value is much more dramatic in the case of static samples.

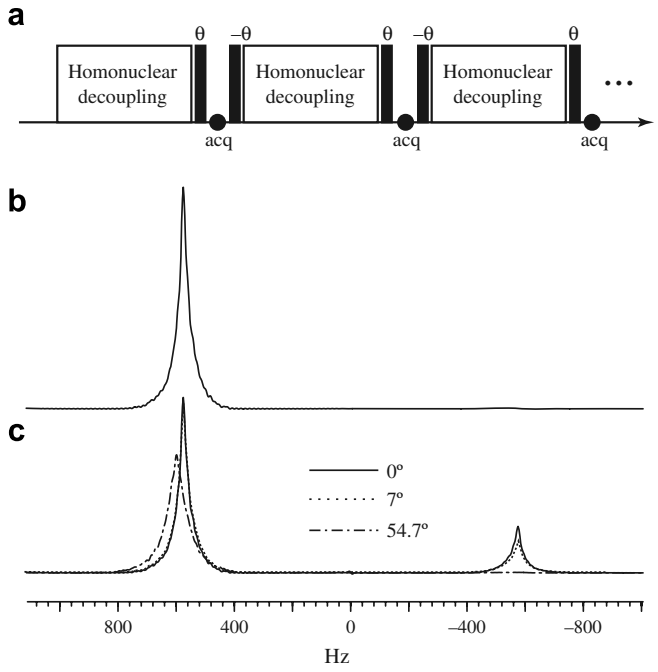

Fig. 10. Effect of tilt pulses on the ^{1}H spectrum under windowed 100 kHz PMLG-17 homonuclear decoupling under static conditions. (a) Schematic representation of the windowed decoupling experiment, (b) simulated spectrum using ideal (delta) magic-angle tilt pulses, (c) simulated spectra using finite-duration tilt pulses (of the same amplitude as the homonuclear decoupling) for different tilt angles.

results from a 7° tilt pulse (as used experimentally in Ref. [6] for example), being barely distinguishable from the results obtained without tilt pulses.

Clearly it is not possible to consider the effect of tilt pulses in directly detected dimensions independently of the decoupling sequence, and again, numerical simulation provides a straightforward and direct way to investigate the interaction. It is also worth emphasising that the presence of decoupling windows themselves can interact in interesting, and potentially useful, ways with the decoupling sequence. For example, the combination of acquisition windows and a coarse Lee–Goldburg phase ramp in the windowed-PMLG-5 decoupling sequence resulted in a net precession axis around z, avoiding the problems associated with tilted axis precession [47].

3.3. Phase propagation delays

As a consequence of non-linear elements (typically filters) in the RF generation pathway, there is often a small, but observable, lag between amplitude and phase changes. Hence if two pulses of the same duration but different phases are programmed back-to-back, the phase change will typically lag behind the amplitude changes associated with the start and end of the pulse pair. Effectively the first pulse (or rather the phase of the first pulse) will appear to be longer than the second pulse. The magnitude of these effects depends strongly on the particular spectrometer console, but delays of the order of about 200 ns have been observed in practice [17]. Hence in a train of constant-amplitude phase-modulated pulses, the first pulse will appear to be longer than expected and the last shorter than expected by the same amount due to this lag. The overall effect of the delay can thus be expressed as the ideal phase modulation sandwiched between a pair of short tilt pulses.

The impact of these delays on the performance of homonuclear decoupling depends strongly on the nature of the decoupling sequence. If the homonuclear decoupling is being applied as a windowless train of pulses, e.g., during an indirect dimension, then the effective addition of a pair of small tilt pulses either side of the homonuclear decoupling has minimal effect, and may be absorbed into the optimisation of any "magic angle" tilt pulses bracketing the homonuclear decoupling [27]. In contrast, the cumulative impact of phase propagation delays can be much larger for windowed sequences since each window effectively introduces an additional pair of tilt pulses.

These effects are illustrated in Fig. 11 which shows the precession plane traced out by the ^{1}H magnetisation vector for both PMLG and DUMBO-1 [3] decoupling for different values of the phase propagation delay. Since the overall precession is of most interest, it is sufficient to perform the simulations using a single nuclear spin with a non-zero frequency relative to the transmitter and trace out the components of the magnetisation in x, y and z as a function of the acquisition time. Since the effects of the propagation delays are associated with the windows of the decoupling sequences, the dwell times have been chosen to be approximately equal for PMLG and DUMBO. Despite this, the effects on the precession plane are very different for the two homonuclear decoupling sequences. In the case of PMLG, the propagation delays lead to a modest tilting of the precession plane which could be readily compensated by reducing the size of any observation tilt pulses (see above). In practice, any tilt pulses are generally optimised empirically and this modification of the precession plane is unlikely to be noticed. In contrast, the propagation delays have a much more

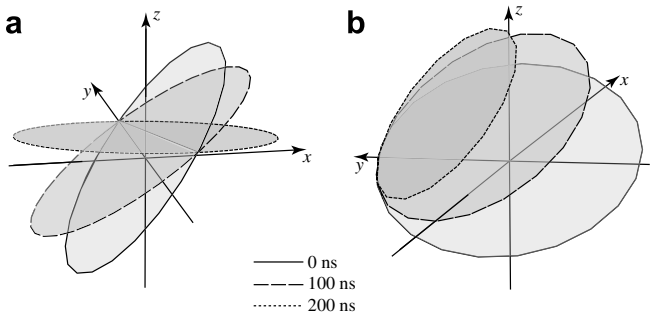

Fig. 11. Effect of phase propagation delays on the tilted axis precession using (a) PMLG-17 and (b) DUMBO-1 homonuclear decoupling at RF nutation rates of 100 kHz nutation rate for different values of the propagation delay: 0, 100 and 200 ns. The diagrams plot the simulated position of the rotating frame ^{1}H magnetisation vector for a single spin with a frequency of 1 kHz relative to the transmitter. The dwell time between observations was 81.7 μs for PMLG (5 complete cycles) and 90 μs for DUMBO (3 complete cycles).

marked effect on the precession under DUMBO decoupling. In this case, the size of the precession circle (and hence the overall signal) is reduced and a sizeable zero frequency artefact (associated with the overall shift away from the origin of the xy plane) would be expected.

The origin of this difference is relatively easily explained. For most phase-modulated homonuclear decoupling sequences, including both PMLG and DUMBO decoupling, the phase of the first and last pulses of the decoupling sequence differ by 180° (this ensures that the decoupling sequence is "cyclic", i.e., there is no net rotation of the axis system over the complete RF cycle). Since the phase propagation delay effectively lengths the first and shortens the last pulse of a train, its effects can be described in terms of a pair of additional short tilt pulses of the same phase as the initial pulse. The effect of these tilt pulses depends strongly on their phase relative to the spin–lock axis. The phase modulations used for the two decoupling sequences in Fig. 11 have been chosen such that the spin–lock axis lies in the xz plane. However, the phases of the initial pulses, with respect to x, are quite different. In the case of PMLG, Fig. 11a, the initial pulse has a phase close to x (exactly in the limit of large N) whereas in the case of DUMBO-1, the phase of the initial pulse is y. The net precessional motion is very different in the two cases. The conclusion is that aligning the initial RF phase with the effective spin–lock axis leads to decoupling sequences that are much more robust with respect to phase propagation delays.

It is important to stress, however, that the magnitude of the propagation delays are strongly spectrometer-dependent. Indeed on some current console designs (including the InfinityPlus used in this work), the delays were negligibly small and major source of pulse imperfections is likely to lie in the finite rise and fall times associated with amplitude and phase changes in any tuned circuit [16].

4. Conclusions

A number of algorithms that are relevant to the numerical simulation of problems in spin dynamics involving both magic-angle spinning and RF decoupling have been described. These have been implemented and evaluated using in-house software optimised for large spin systems. (This software and its sources is freely available [29], and the specific input files used are also available on request.) Chebyshev propagation is shown to be a useful alternative to conventional approaches to propagator calculation where the Hamiltonian is time-dependent and relatively sparse, as is the case for multi-spin systems under RF decoupling. A simple, if rather crude, algorithm for efficient simulations of problems involving phase-modulated RF decoupling has also been described which allows a single set of propagators over one rotor cycle to be used, irrespective of the degree of "synchronisation" of the RF and spinning time periods. Note that some of these algorithms are also implemented in SPINEVOLUTION [14] which is also designed to handle large and challenging problems in solid-state NMR, and which performs with similar efficiency for comparable problems. In contrast, the performance of programs such as SIMPSON [48] which do not exploit the sparse nature of Hamiltonian for many-spin problems (except for the special case of diagonal Hamiltonians) is significantly poorer [14].

The resulting efficient simulations have been used to investigate ^{1}H resolution under homonuclear decoupling. This is a formidable challenge due to the number of factors involved and the potential interactions between these factors. In general terms, resolution loss can either result from degraded decoupling efficiency, resulting in broader lines, or from variations in the effective precession axis or nutation rate (such variations can either be spatial or temporal). Variations of the precession axis can usually be determined in terms of a individual nuclear spin, but questions of resolution invariably involve considering multi-spin systems. Unfortunately, the dynamics of multi-spin systems are challenging to determine analytically and are rarely tractable outside special (and idealised) cases. Although numerical simulation may appear to offer little in the way of physical insight, it is proving to be a invaluable tool in unravelling the different factors that determine ^{1}H resolution in solids. Efficient algorithms, including those described here, allow arbitrarily complex and non-ideal experiments to be assessed in a straightforward and routine manner. This allows specific questions to be addressed and new insights into observed experimental behaviour that are not readily obtained from analytical work. In particular, we have shown that magic-angle spinning and homonuclear decoupling often interact constructively. While MAS may lead to unwelcome interactions when classic homonuclear decoupling sequences are applied under ideal conditions, it can significantly reduce artefacts associated with experimental imperfections, such as RF inhomogeneity, or limitations of the pulse sequence itself (e.g., PMLG decoupling applied with a limited number of phase steps). This confirms the interest of developing decoupling sequences by direct spectral optimisation under MAS conditions [19].

The complexity of experiments involving homonuclear decoupling is significantly increased if the ^{1}H signal is being acquired via windows inserted in the pulse sequence. We have shown that effects, such as phase propagation delays and finite-duration tilt pulses, which are of negligible significance in indirect dimensions, become very important in directly acquired dimensions. The simulations confirm empirical observations that much shorter tilt pulses are required in directly (as compared to indirectly) acquired dimensions. They also show that the impact of small time lags between phase and amplitude changes depends strongly on the decoupling sequence, providing additional guidelines for the design of effective and robust phase-modulated decoupling sequences.

Although the work described advances our understanding of these experiments, it is important to emphasise that there are still important discrepancies between observed

and predicted behaviour, cf. Fig. 6. This suggests that there is still significant room for improvement (either in hardware or experimental methodology), and the limits of ^{1}H resolution in solids have yet to be reached.

Acknowledgments

This work was initially supported under Grant GR/S56993/01 from the Engineering & Physical Sciences Research Council (EPSRC). VEZ is currently supported under EPSRC Grant EP/D057159. SPB also thanks the EPSRC for the award of an Advanced Research Fellowship.

References

[1] A. Bielecki, A.C. Kolbert, M.H. Levitt, Frequency switched pulse sequences: homonuclear decoupling and dilute spin NMR in solids, Chem. Phys. Lett. 155 (1989) 341–346.
[2] E. Vinogradov, P.K. Madhu, S. Vega, High-resolution proton solid-state NMR spectroscopy by phase-modulated Lee–Goldburg experiment, Chem. Phys. Lett. 314 (1999) 443–450.
[3] D. Sakellariou, A. Lesage, P. Hodgkinson, L. Emsley, Homonuclear dipolar decoupling in solid-state NMR using continuous phase modulation, Chem. Phys. Lett. 319 (2000) 253–260.
[4] B. Elena, L. Emsley, Powder crystallography by proton solid-state NMR, J. Am. Chem. Soc. 127 (2005) 9140–9146.
[5] P.K. Madhu, E. Vinogradov, S. Vega, Multiple-pulse and magic-angle spinning aided double-quantum proton solid-state NMR spectroscopy, Chem. Phys. Lett. 394 (2004) 423–428.
[6] S.P. Brown, A. Lesage, B. Elena, L. Emsley, Probing proton–proton proximities in the solid state: high-resolution two-dimensional ^{1}H–^{1}H double-quantum CRAMPS NMR spectroscopy, J. Am. Chem. Soc. 126 (2004) 13230–13231.
[7] B.-J. van Rossum, H. Förster, H.J.M. de Groot, High-field and high-speed CP-MAS ^{13}C NMR heteronuclear dipolar-correlation spectroscopy of solids with frequency-switched Lee–Goldburg homonuclear decoupling, J. Magn. Reson. 124 (1997) 516–519.
[8] A. Lesage, D. Sakellariou, S. Steuernagel, L. Emsley, Carbon–proton chemical shift correlation in solid-state NMR by through-bond multiple-quantum spectroscopy, J. Am. Chem. Soc. 120 (1998) 13194–13201.
[9] M. Mehring, Principles of High Resolution NMR in Solids, second ed., Springer, 1983.
[10] C. Filip, X. Filip, D.E. Demco, S. Hafner, Spin dynamics under magic angle spinning by Floquet formalism, Mol. Phys. 92 (1997) 757–771.
[11] C. Filip, S. Hafner, I. Schnell, D.E. Demco, H.W. Spiess, Solid-state nuclear magnetic resonance spectra of dipolar-coupled multi-spin systems under fast magic angle spinning, J. Chem. Phys. 110 (1999) 423–440.
[12] S. Ray, E. Vinogradov, G.-J. Boender, S. Vega, Proton MAS NMR spectra at high magnetic fields and high spinning frequencies: spectral simulation using Floquet theory, J. Magn. Reson. 135 (1998) 418–426.
[13] P. Hodgkinson, D. Sakellariou, L. Emsley, Simulation of extended periodic systems of nuclear spins, Chem. Phys. Lett. 326 (2000) 515–522.
[14] M. Veshtort, R.G. Griffin, SPINEVOLUTION: a powerful tool for the simulation of solid and liquid state NMR experiments, J. Magn. Reson. 178 (2005) 248–282.
[15] V.E. Zorin, S.P. Brown, P. Hodgkinson, Origins of linewidth in ^{1}H magic-angle spinning NMR, J. Chem. Phys. 125 (2006) 144508:1–144508:13.
[16] M. Carravetta, M. Edén, O.G. Johannessen, H. Luthman, P.J.E. Verdegem, J. Lugtenburg, A. Sebald, M.H. Levitt, Estimation of carbon–carbon bonds lengths and medium-range internuclear distances by solid-state nuclear magnetic resonance, J. Am. Chem. Soc. 123 (2001) 10628–10638.
[17] A.J. Vega, Controlling the effects of pulse transients and RF inhomogeneity in phase-modulated multiple-pulse sequences for homonuclear decoupling in solid-state proton NMR, J. Magn. Reson. 170 (2004) 22–41.
[18] S. Hafner, H.W. Spiess, Multiple-pulse line narrowing under fast magic-angle spinning, J. Magn. Reson. Ser. A 121 (1996) 160–166.
[19] B. Elena, G. De Paëpe, L. Emsley, Direct spectral optimisation of proton–proton homonuclear dipolar decoupling in solid-state NMR, Chem. Phys. Lett. 398 (2004) 532–538.
[20] E. Vinogradov, P.K. Madhu, S. Vega, A bimodal Floquet analysis of phase modulated Lee–Goldburg high resolution proton magic angle spinning NMR experiments, Chem. Phys. Lett. 329 (2000) 207–214.
[21] M. Leskes, P.K. Madhu, S. Vega, Proton line narrowing in solid-state nuclear magnetic resonance: new insights from windowed phase-modulated Lee–Goldburg sequence, J. Chem. Phys. 125 (2006) 124506.
[22] P. Hodgkinson, L. Emsley, Numerical simulation of solid-state NMR experiments, Prog. Nucl. Magn. Reson. Spectrosc. 36 (2000) 201–239.
[23] R.R. Ernst, G. Bodenhausen, A. Wokaun, Principles of Nuclear Magnetic Resonance in One and Two Dimensions, Clarendon Press, Oxford, 1987 (section 6.3).
[24] M.H. Levitt, M. Edén, Numerical simulation of periodic nuclear magnetic resonance problems: fast calculation of carousel averages, Mol. Phys. 95 (1998) 879–890.
[25] M. Hohwy, H. Bildsøe, H.J. Jakobsen, N.C. Nielsen, Efficient spectral simulations in NMR of rotating solids. The γ-COMPUTE algorithm, J. Magn. Reson. 136 (1999) 6–14.
[26] W.B. Blanton, J.W. Logan, A. Pines, Rational reduction of periodic propagators for off-period observations, J. Magn. Reson. 166 (2004) 174–181.
[27] A. Lesage, D. Sakellariou, S. Hediger, B. Elena, P. Charmont, S. Steuernagel, L. Emsley, Experimental aspects of proton NMR spectroscopy in solids using phase-modulated homonuclear dipolar decoupling, J. Magn. Reson. 163 (2003) 105–113.
[28] M. Veshtort, Numerical Simulations in Nuclear Magnetic Resonance: Theory and Applications, Ph.D. thesis, Massachusetts Institute of Technology, 2003.
[29] P. Hodgkinson, pNMRsim: a general simulation program for large problems in solid-state NMR. Available from: <http://www.durham.ac.uk/paul.hodgkinson/pNMRsim/>.
[30] C. Moler, C. Van Loan, Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later, SIAM Rev. 45 (2003) 3–49.
[31] R.S. Dumont, S. Jain, A. Bain, Simulation of many-spin dynamics via sparse matrix methodology, J. Chem. Phys. 106 (1997) 5928–5936.
[32] R.C. Whaley, A. Petitet, J.J. Dongarra, Automated empirical optimization of software and the ATLAS project, Parallel Comput. 27 (2001) 3–35.
[33] V.E. Zorin, S.P. Brown, P. Hodgkinson, Quantification of homonuclear dipolar coupling from magic-angle spinning ^{1}H NMR, Mol. Phys. 104 (2006) 293–304.
[34] G. De Paëpe, N. Giraud, A. Lesage, P. Hodgkinson, L. Emsley, Transverse dephasing optimized solid-state NMR spectroscopy, J. Am. Chem. Soc. 125 (2003) 13938–13939.
[35] J.S. Waugh, L.M. Huber, U. Haeberlen, Approach to high-resolution NMR in solids, Phys. Rev. Lett. 20 (1968) 180–182.
[36] P. Mansfield, Symmetrized pulse sequences in high resolution NMR in solids, J. Phys. C 4 (1971) 1444–1452.
[37] D.P. Burum, W.K. Rhim, Analysis of multiple pulse NMR in solids. III, J. Chem. Phys. 71 (1979) 944–956.
[38] B.C. Gerstein, R.G. Pembleton, R.C. Wilson, L.M. Ryan, High resolution NMR in randomly oriented solids with homonuclear

dipolar broadening: combined multiple pulse NMR and magic angle spinning, J. Chem. Phys. 66 (1977) 361–362.
[39] G.E. Maciel, C.E. Bronnimann, B.L. Hawkins, High-resolution ^{1}H nuclear magnetic resonance in solids via CRAMPS, Adv. Magn. Reson. 14 (1990) 125–150.
[40] C. Filip, S. Hafner, Analysis of multiple-pulse techniques under fast MAS conditions, J. Magn. Reson. 147 (2000) 250–260.
[41] P.K. Madhu, X. Zhao, M.H. Levitt, High-resolution ^{1}H NMR in the solid state using symmetry-based pulse sequences, Chem. Phys. Lett. 346 (2001) 142–148.
[42] D.P. Burum, M. Linder, R.R. Ernst, A new "tune-up" NMR pulse cycle for minimizing and characterizing phase transients, J. Magn. Reson. 43 (1981) 463–471.
[43] P. Jackson, R.K. Harris, A practical guide to combined rotation and multiple-pulse NMR spectroscopy of solids, Magn. Reson. Chem. 26 (1988) 1003–1011.
[44] M. Hohwy, N.C. Nielsen, Elimination of high order terms in multiple pulse nuclear magnetic resonance spectroscopy: application to homonuclear decoupling in solids, J. Chem. Phys. 106 (1997).
[45] H. Cho, Off-resonance multiple-pulse dynamics in solid-state NMR spectroscopy: a revised coherent averaging theory analysis, J. Magn. Reson. 141 (1999) 164–179.
[46] M. Leskes, P.K. Madhu, S. Vega, A broad-banded z-rotation windowed phase-modulated Lee–Goldburg pulse sequence for ^{1}H spectroscopy in solid-state NMR, Chem. Phys. Lett. 447 (2007) 370–374.
[47] L. Bosman, P. Madhu, S. Vega, E. Vinogradov, Improvement of homonuclear dipolar decoupling sequences in solid-state nuclear magnetic resonance utilising radiofrequency imperfections, J. Magn. Reson. 169 (2004) 39–48.
[48] 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.