

Simulation of coherence selection by pulsed field gradients in liquid-state NMR using an auxiliary matrix formalism

Luke J. Edwards^{a,*}

^a*Department of Chemistry, University of Oxford, Inorganic Chemistry Laboratory, South Parks Road, Oxford, OX1 3QG, UK.*

Abstract

An algorithm for simulating coherence selection due to a pulse sequence element consisting of two pulsed field gradients separated by a short collection of pulses and delays is introduced. This algorithm involves computation of the matrix exponential of an auxiliary matrix twice the size of the system Liouvillian, a dimensional increase smaller than is required with other known computational methods. Approximations valid for most simulations of liquid-state NMR spectra are involved in the derivation. Diffusion is omitted, but could be treated in an approximate way as a damping over the pulse sequence element. Several NMR pulse sequences using gradients for coherence selection have been implemented, making use of the functionality of *Spinach* (<http://spindynamics.org/Spinach.php>). Example simulations testing these implementations are presented, and the extent to which the formal dimensional reduction can lead to a speedup in simulation time discussed. It is found that the previously known methods can be made competitive with the auxiliary matrix method by making use of their embarrassingly parallel nature. Cases where the relative dimensional reduction of the auxiliary matrix method is very large, or where efficient parallelization of the simulation independent of the nature of the algorithm used exists, are found to lead to situations beneficial for the auxiliary matrix algorithm in this comparison.

Keywords: PFG; NMR; simulation

*Phone: +44 1865 275416

Email address: luke.edwards@chem.ox.ac.uk (Luke J. Edwards)

1. Introduction

The utilization of pulsed field gradients (PFGs) has now become routine in NMR experiments [1, 2]. Among the many ways in which they are employed [1], one of their most important uses is in selecting specific coherences [2, 3]. It is this use that we shall examine in this work.

While analytical rules can be derived for the effect of PFGs in special cases (e.g. in the absence of spin–spin couplings), computational techniques have an important rôle to play in determining the effects of PFGs in more general systems. The reasons for wanting to simulate such effects include, among other possibilities, creating new experiments, optimizing old experiments, and checking spectral assignments.

Several methods for simulating the effects of field gradients computationally are known in the literature; these can be separated neatly into two groups: those that use explicit slicing of the sample into discrete blocks [4, 5], and those that use analytical coherence selection to propagate each coherence individually followed by recombination with appropriate weightings [5, 6]. Explicit slicing has been implemented in several software packages [5, 7], as has analytical coherence selection [5, 6].

As very few assumptions have to be made, explicit slicing is the more general method. It is also easily parallelizable, even in the presence of diffusion [8]. Unfortunately, the number of slices can be large, and increases as the number of spins increases [4, 5]. This can result in very long simulations, despite the parallelization opportunities.

Analytical coherence selection can also be troublesome; while the simulation of each individual coherence could be run in a parallel manner, the number of coherences required quickly balloons [6]. One can guess which coherences will be important based upon known rules of thumb and thereby cut down the number of coherences that need to be simulated, but these rules could easily miss coherences that would contribute to the eventual spectrum.

In the next section we present an efficient algorithm that computes the effect of pairs of z -axis PFGs on NMR spectra in the absence of diffusion. Spin–spin couplings are included, and the use of a superoperator formalism means that, assuming the relaxation rates can be assumed independent of the gradient field, the inclusion of relaxation is trivial. This algorithm requires only a doubling of the matrix dimension, a much smaller increase as compared to the effective increases in matrix dimension provided by the methods mentioned above. Following this, the use of the algorithm in implementing explicit gradients in several NMR pulse

sequences is discussed, with example simulations shown.

2. Theory

A linear field gradient changes the strength of the magnetic field at a position z in the sample by an amount Gz , where G is the strength of the gradient per unit length. Within the usual approximations associated with liquid-state NMR (the high field approximation and assumption of complete averaging of non-isotropic interactions), which will be assumed throughout this report, this gives rise to an extra Zeeman term in the Liouvillian which, for a system containing N different types of heteronuclei, may be written [9]

$$-Gz \sum_{n=1}^N \gamma_n \sum_{m=1}^{M_n} \left(1 - \delta_m^{(n)}\right) \hat{I}_Z^{(m,n)} \quad (1)$$

where γ_n is the gyromagnetic ratio of the n -th type of nucleus, M_n is the number of nuclei of type n in the system, $\hat{I}_Z^{(m,n)}$ is the commutation superoperator representation [9] of $[\hat{I}_Z^{(m,n)}, \cdot]$ with $\hat{I}_Z^{(m,n)}$ the Z operator acting on spin m of type n , and $\delta_m^{(n)}$ is the isotropic chemical shift of the m -th nucleus of type n .

The Liouvillian should properly be modified to include the effects of diffusion as well. Unfortunately, such a modification would require the dimensions of the matrices used to be increased many-fold [10], likely beyond practicability for general-purpose simulations. Treating diffusion in an approximate way, e.g. as a simple damping for the duration of the pulse sequence element containing the gradients as described in Reference [6], would be possible, but for clarity of exposition we ignore the effects of diffusion in the rest of this report and assume that the length of the pulse sequence element containing gradients is short enough that such effects are negligible.

We now make the approximation that all terms in Equation (1) containing the chemical shift can be discarded. This is a reasonable approximation as long as we can assume that both

1. the gradients are weak relative to the static field.
2. the length of the pulse sequence element containing the gradients is short as compared to the reciprocal of the largest angular frequency arising from the terms omitted.

The first point describes the usual case in NMR. The second point is often true in NMR experiments using gradients for coherence selection as keeping the length

of such sequence elements short minimizes losses due to diffusion [3]. We restrict ourselves to cases in which such bounds hold; see the simulations in Section 3 for discussions of the validity of these assumptions.

The above truncation leads to the following superoperator,

$$\hat{G}(G, z) = -Gz \sum_{n=1}^N \gamma_n \sum_{m=1}^{M_n} \hat{I}_Z^{(m,n)}. \quad (2)$$

Importantly, this superoperator commutes with the high-field Liouvillian [9]. It should be noted that a secular relaxation superoperator [9] can be included as part of this high-field Liouvillian, assuming that it is to a good approximation independent of the gradient strength; this assumption is reasonable in NMR.

Figure 1 shows schematically the sort of pulse sequence element that we are attempting to simulate. Expressed mathematically this element may be written

$$\hat{S}(z) = \exp\{-i[\hat{L} + \hat{G}_2(z)]\tau_2\} \hat{P} \exp\{-i[\hat{L} + \hat{G}_1(z)]\tau_1\} \quad (3)$$

in the absence of diffusion. In this equation \hat{P} is a propagator representing pulses and delays between the two gradients. The two gradients are written as $\hat{G}_i(z)$ which is a shorthand for

$$\hat{G}_i(z) = \hat{G}(s_i G_i, z), \quad (4)$$

a gradient pulse with strength G_i and shape factor s_i . Each gradient pulse has envelope $A_i(t)$; this allows for pulses that are not rectangular. The shape factors appearing in Equation (4) are defined as the average of these pulse envelopes over the respective gradient pulse lengths τ_i [3, 11]:

$$s_i = \tau_i^{-1} \int_0^{\tau_i} A_i(t) dt. \quad (5)$$

Because of the commutativity of \hat{G}_i and \hat{L} the propagator in Equation (3) simplifies to

$$\exp\{-i\hat{L}\tau_1\} \hat{Q}(z) \exp\{-i\hat{L}\tau_2\}, \quad (6)$$

where

$$\hat{Q}(z) = \exp\{-i\hat{G}_2(z)\tau_2\} \hat{P} \exp\{-i\hat{G}_1(z)\tau_1\}. \quad (7)$$

We wish to compute [3]

$$l^{-1} \int_{-l/2}^{+l/2} \hat{S}(z) dz, \quad (8)$$




Figure 1: A schematic representation of the gradient-containing pulse sequence element explored in this article. The hatched rectangle represents pulses (applied to spins I and/or S) and delays in between the gradients. The gradients themselves are represented by sine-bell curves on the line labelled G, though the shape is arbitrary. Each gradient pulse is labelled with its shape factor [3, 11] s_i , its strength G_i , and the time τ_i over which it is applied.

the propagator averaged over the sensitive portion of the sample. This portion is of length l , and for definiteness we follow the NMR convention [3] that the field is zero in the centre of this portion. Examining Equation (6) we see that the only term depending on z is $\hat{Q}(z)$; the integration thus reduces to computing the integral

$$l^{-1} \int_{-l/2}^{+l/2} \hat{Q}(z) dz, \quad (9)$$

along with some straightforward matrix propagation under the Liouvillian. As mentioned in the Introduction, such an integration would usually be performed by slicing the simulation into a large number of explicit slices [4, 5]. In the following, we demonstrate how this integral can be calculated from the exponential of an auxiliary matrix, following work on the integration of matrix functions similar to that found in Equation (7) by Van Loan [12].

Defining the upper triangular block auxiliary matrix

$$\hat{A} = \begin{pmatrix} -\hat{G}_2(l)\tau_2 & +i\hat{P} \\ & +\hat{G}_1(l)\tau_1 \end{pmatrix}, \quad (10)$$

computation of the matrix exponential

$$\exp\{-i\hat{A}\} = \begin{pmatrix} e^{+i\hat{G}_2(l)\tau_2} & \hat{B} \\ & e^{-i\hat{G}_1(l)\tau_1} \end{pmatrix} \quad (11)$$

reveals a block, \hat{B} , given by

$$\hat{B} = e^{+i\hat{G}_2(l)\tau_2} \int_0^1 e^{-i\hat{G}_2(l)\tau_2\alpha} \hat{P} e^{-i\hat{G}_1(l)\tau_1\alpha} d\alpha. \quad (12)$$

A proof of this relation is given in Appendix A. Making the substitution

$$z = l \left(\alpha - \frac{1}{2} \right) \quad (13)$$

and remembering that the $\hat{\hat{G}}_i$ are linear gives

$$\hat{\hat{B}} = e^{+i\hat{\hat{G}}_2(l/2)\tau_2} \left[l^{-1} \int_{-l/2}^{+l/2} \hat{\hat{Q}}(z) dz \right] e^{-i\hat{\hat{G}}_1(l/2)\tau_1}, \quad (14)$$

where the portion in square brackets may be seen to be the integral we want to compute, Equation (9). This means that computation of $\hat{\hat{B}}$ through use of the auxiliary matrix $\hat{\hat{A}}$ results in a matrix that is only a few short propagations away from the integral that we want.

The results of this section suggest Algorithm 1 which propagates an initial density matrix $\hat{\sigma}_0$ to the end of the pulse sequence element under investigation. Importantly, we do not need to build any explicit propagators except for that representing $\hat{\hat{P}}$; we can use Krylov-like methods that instead only perform matrix–vector multiplications [13, 14].

The reliance on having to build the propagator $\hat{\hat{P}}$ explicitly could be a problem in large spin systems, or where $\hat{\hat{P}}$ contains long periods of propagation. In large spin systems the problem is that the dimension of the matrices that must be exponentiated increases with the size of the spin system, rapidly becoming prohibitive. Fortunately, recently developed state space restriction tools [13, 15] can mitigate this to a large extent, as demonstrated by the simulations presented in Figures 5 and 6. Evolution over long periods of time is problematic as the propagator representation is likely to have few non-zeroes; this means that sparse matrix methodology [14] is of little benefit, severely restricting the size of simulation that can be performed. We show in Section 3.3 how we can circumvent this problem in special cases by removing the period of propagation from the propagator $\hat{\hat{P}}$.

The result of this section is that we have computed the integral over the propagator, within the approximations set out above, with just a doubling of the dimensions of the matrices involved. Achieving the same level of precision using slicing methods could require tens of slices to be computed [5]; such a simulation would have an effective matrix dimension many times the dimension of the original superoperators.

Algorithm 1 Auxiliary matrix algorithm.

```

 $\hat{\sigma} \leftarrow \hat{\sigma}_0$ 
 $\hat{\sigma} \leftarrow \exp\{-i\hat{L}\tau_1\}\hat{\sigma}$ 
 $\hat{\sigma} \leftarrow \exp\{+i\hat{G}_1(l/2)\tau_1\}\hat{\sigma}$ 
 $\hat{A} \leftarrow \begin{pmatrix} -\hat{G}_2(l)\tau_2 & +i\hat{P} \\ & +\hat{G}_1(l)\tau_1 \end{pmatrix}$ 
 $\hat{\Sigma} \leftarrow \begin{pmatrix} 0 \cdot \hat{\sigma} \\ \hat{\sigma} \end{pmatrix}$ 
 $\hat{\Sigma} \leftarrow \exp\{-i\hat{A}\}\hat{\Sigma}$ 
 $\hat{\sigma} \leftarrow \text{top half of } \hat{\Sigma}$ 
 $\hat{\sigma} \leftarrow \exp\{-i\hat{G}_2(l/2)\tau_2\}\hat{\sigma}$ 
 $\hat{\sigma} \leftarrow \exp\{-i\hat{L}\tau_2\}\hat{\sigma}$ 
return  $\hat{\sigma}$ 

```

3. Results

The formalism presented in Section 2 has been used to implement gradients in the pulse sequences that are shown in Figure 2. Perfect RF pulses [9] are assumed throughout for convenience. These implementations, described below, make use of the functions contained within the *Spinach* toolbox [13] for MATLAB, available from <http://spindynamics.org/Spinach.php>. All simulations have been performed on a workstation with four Intel Nehalem-cores and 16 GB of RAM. The source code for the simulations shown in this section may be found in the Supplementary Information.

3.1. Homospoil pulse

The simplest application of PFGs is the homospoil pulse [3] which is used to dephase all but terms with coherence order zero [9] present in a given state. It is equally simple to implement using the auxiliary matrix formalism: the auxiliary matrix required is

$$\begin{pmatrix} 0_{\dim\{\hat{L}\}} & \mathbb{I}_{\dim\{\hat{L}\}} \\ & +i\hat{G}_1(l)\tau_1 \end{pmatrix} \quad (15)$$

where $0_{\dim\{\hat{L}\}}$ is a square matrix of zeroes with dimension equal to that of \hat{L} , and $\mathbb{I}_{\dim\{\hat{L}\}}$ is the identity matrix with the same dimension. An example of the effect




Figure 2: Pulse sequence diagrams for the sequences described in the text. a) homospoil pulse [3]. b) homonuclear multiple-quantum filter [3]. c) phase-sensitive gradient-selected DQF-COSY [16, Experiment 12.3]. d) gradient-selected absorption mode HSQC [17]. The gradient lines, marked G, are notated as in Figure 1. Dotted boxes enclose the regions of pulse sequences c) and d) over which we use Algorithm 1. Rectangles represent RF pulses: thin rectangles are $\pi/2$ pulses and thick rectangles π pulses; I and S denote the spins to which the pulses are applied. Pulses are applied about X unless otherwise labelled. The pulse marked ϕ in c) is phase cycled $\pm X$ with associated receiver phases $\pm X$, and is also cycled to generate separately the sine and cosine modulated spectra for frequency-discrimination [3]. The selection of either P- or N-coherence in d) is performed by changing the sign of the shape function, s_2 , of the dashed gradient pulse [3, 17].




Figure 3: Demonstration of the damping rates of different coherence orders caused by a homospoil pulse [3] simulated using Algorithm 1. The spin system consists of three randomly coupled ^1H with isotropic chemical shifts in the range $[0, 10]$, spin-spin couplings between 0 Hz and 20 Hz, and a magnetic field of 10 T. The initial state vector is chosen randomly, and each set of coherences is normalized to a different initial value (seen at time zero) in order to allow their differentiation on the plot. The gradient strength $G_1 = 20 \text{ G cm}^{-1}$, and the shape of the pulse is taken to be rectangular [11]: $s_1 = 1$. Source code for this simulation, which documents all of the parameters used, may be found in the Supplementary Information.

of this pulse sequence element on the different coherence orders of a three-spin system is shown in Figure 3.

3.2. Multiple quantum filter

Pulse sequence elements of the form shown in Figure 1 can be used to construct multiple quantum filters [3]. An explicit example is shown in Figure 2b; the auxiliary matrix corresponding to this pulse sequence element is

$$\begin{pmatrix} -\hat{G}_2(l)\tau_2 & +i\hat{I}_x\pi/2 \\ & +\hat{G}_1(l)\tau_1 \end{pmatrix}. \quad (16)$$

Shown in Figure 4 is a simulation using Algorithm 1 that demonstrates how we can use this filter to select specific coherences.

3.3. Double quantum filter in DQF-COSY

An example of a pulse sequence containing the element shown in Figure 1 is phase-sensitive gradient-selected double-quantum filtered- (DQF-) COSY [16, 18]. This pulse sequence is shown in Figure 2c, where the pulse sequence element of interest is demarked by the dotted-box. Examining that figure we see that the




Figure 4: Demonstration of how changing the length of the second gradient pulse relative to the length of the first in the pulse sequence element shown in Figure 2b changes which coherences are refocused, simulated using Algorithm 1. The spin system consists of three randomly coupled ^1H with parameters in the same ranges as those used to generate Figure 3. Prior to the application of the pulse sequence element the state vector consists entirely of $+2$ coherence, randomly selected and normalized to unity. The gradient strengths $G_1 = G_2 = 20 \text{ G cm}^{-1}$, the shapes of the pulses are taken to be rectangular [11]: $s_1 = s_2 = 1$, and the length of the first gradient pulse is $\tau_1 = 1 \text{ ms}$. Source code for this simulation, which documents all of the parameters used, may be found in the Supplementary Information.

pulse sequence contains an evolution period in between the two gradient pulses. The explicit propagator for this period is likely to have very few non-zeroes; this means that sparse matrix techniques [14] are not likely to be of benefit, restricting the size of system that we can simulate using Algorithm 1. We will show, however, that in the case where all the spins of interest have a spin quantum number of $1/2$ we can rearrange the propagator such that this period of evolution is simulated separately, easing the computational burden.

To begin we write out the propagator for the part of the sequence within the dotted box in Figure 2c at a given value of z :

$$e^{-i(\hat{G}_2(z) + \hat{L})\tau_2} e^{-i\hat{I}_X\pi} e^{-i\hat{L}\tau_2} e^{-i\hat{I}_X\pi/2} e^{-i(\hat{G}_1(z) + \hat{L})\tau_1} \quad (17)$$

where \hat{I}_X denotes the commutation superoperator representation [9] of $[\hat{I}_X, \cdot]$ that acts on the spins I . Making use of the commutativity of \hat{L} and $\hat{G}_i(z)$, as well as the relation involving matrix exponentials that for a square matrix B and invertible matrix A [19]

$$Ae^B A^{-1} = e^{ABA^{-1}}, \quad (18)$$

the above propagator is equivalent to

$$e^{-i\hat{L}\tau_2}e^{-i\hat{L}_X\pi}e^{-i\hat{G}'_2(z)\tau_2}e^{-i\hat{L}\tau_2}e^{-i\hat{L}_X\pi/2}e^{-i\hat{G}_1(z)\tau_1}e^{-i\hat{L}\tau_1}. \quad (19)$$

The modified gradient propagator $\hat{G}'_2(z)$ appearing in this equation is

$$\hat{G}'_2(z) = e^{-i\hat{L}_X\pi}\hat{G}_2(z)e^{+i\hat{L}_X\pi} \quad (20)$$

which, because of the definition of $\hat{G}_2(z)$ in terms of $\hat{I}_Z^{(k)}$, is for spin-1/2 particles given by [9]

$$\hat{G}'_2(z) = -\hat{G}_2(z). \quad (21)$$

Inserting this into the propagator, Equation (19), and remembering that $\hat{G}_2(z)$ commutes with \hat{L} results in

$$e^{-i\hat{L}\tau_2}e^{-i\hat{L}_X\pi}e^{-i\hat{L}\tau_2} \left[e^{+i\hat{G}_2(z)\tau_2}e^{-i\hat{L}_X\pi/2}e^{-i\hat{G}_1(z)\tau_1} \right] e^{-i\hat{L}\tau_1} \quad (22)$$

which is in the correct form for easy calculation using the auxiliary matrix method. Explicitly, the auxiliary matrix is

$$\begin{pmatrix} +\hat{G}_2(l)\tau_2 & +ie^{-i\hat{L}_X\pi/2} \\ & +\hat{G}_1(l)\tau_1 \end{pmatrix}. \quad (23)$$

A portion of the DQF-COSY spectrum of the 22 ^1H -spin system of strychnine, simulated by making use of Algorithm 1 with Equation (22), is shown in Figure 5. The validity of the assumptions made in moving from Equation (1) to Equation (2) may be confirmed by noting that the largest angular frequency arising from the dropped terms in this simulation comes from a spin with $\delta = 8.1$ ppm, leading to an angular frequency, given the simulation parameters in Figure 5, of 3.25 rad s^{-1} . The time period over which the effect of the gradients is to be simulated, 6 ms, is much smaller than the reciprocal of this frequency, 0.3 s. Simulation of the portion of the simulation containing gradients takes about three hours, and this portion must be run twice to allow for frequency discrimination in F_1 [3].

We can put this time in perspective by using the recommendations made in Reference [5] and numerical simulation to estimate the minimum time required to simulate this portion of the simulation using the slicing method. According to this reference the minimum number of slices needed to calculate the effect of one of the gradient pulses is 44. Test simulations, making the same approximations

in defining the gradient superoperator of each slice as were made in deriving the auxiliary matrix formalism in Section 2 to allow a fair comparison, show that the mean time taken for a given slice is about five minutes. This implies that when using a single-processor machine (or only a single-core of a multi-processor machine) the auxiliary matrix method will generate this spectrum more quickly. The easy parallelization of the slicing method, however, means that that method can easily be faster when performing this simulation on a multi-processor machine.

This prediction does however require some caveats. Firstly, the minimal number of slices may not give an accurate spectrum; more slices may be needed (though in principle the number of processors can be increased to offset the extra slices). The auxiliary matrix algorithm neatly sidesteps such convergence issues. Secondly, memory utilization is a concern; the number of slices that can be computed in parallel is limited by the amount of storage available as well as the number of processors. It should also be considered that such storage issues can lead to deviations from predicted parallelization speed ups [20]. Thirdly, parallelization of the matrix operations involved in the use of the auxiliary matrix algorithm is possible, though will not necessarily be as efficient as parallelization over the slices. Together, these caveats suggest that the auxiliary matrix algorithm may not always be slower on a multicore machine when performing this calculation.

3.4. *P-/N-detection in HSQC*

The arrangement of gradients and RF pulses that allowed the rearrangement of the propagator described in Section 3.3 is not unique to the DQF-COSY sequence examined in that section. Another gradient-containing pulse sequence that contains portions that can be similarly rearranged is the HSQC sequence shown in Figure 2d. In this pulse sequence the gradients are used to select either the P- or N-coherence during F_1 [3, 17]. Importantly, the presence of heteronuclei increases the minimal number of slices that would be needed to simulate this pulse sequence using the slicing method [5].

An example simulation which uses Algorithm 1 and Equation (22) (*mutatis mutandis*) to generate a P-/N-selected ^{13}C - ^1H HSQC spectrum is shown in Figure 6; this simulation took about 10 minutes. It should be noted that we make the assumption that the ^{13}C nuclei are dilute; this means that we perform a separate simulation for each ^{13}C . In this simulation $\Delta = (2 \cdot 140)^{-1}$ s, and the length of the first gradient is 1 ms. The ratio of the gyromagnetic ratios of the heteronuclei determines the length of the second gradient [3]; in this case this means that the second gradient is a quarter of the length of the first. This means that the period over which we need to use Algorithm 1 is ~ 3 ms. The largest frequency




Figure 5: Simulation of the phase-sensitive gradient selected DQF-COSY spectrum [16, Experiment 12.3] of strychnine performed using Algorithm 1 with the IK-2(1) basis [13] at a field of 11.74 T. The image has been cropped for comparison with the experimental spectrum presented in Reference [16]. Relaxation has been included by building a secular relaxation superoperator using the capabilities of *Spinach* [21] with an isotropic correlation time of 200 ps. Chemical shifts and couplings are taken from Reference [16], and the coordinates are those of the major conformer of strychnine from Reference [22]. Gradient parameters G_i , τ_i and s_i were chosen following Reference [16]: $G_1 = G_2/2 = 10 \text{ G cm}^{-1}$, $\tau_1 = \tau_2 = 2 \text{ ms}$, and $s_1 = s_2 = 2/\pi$ to represent sine bell shapes [11]. The length of the sensitive portion of the sample was set at 2 cm. Minor differences between the lineshapes in this simulation and the experimental spectrum of Reference [16] can be explained by differences in spectral processing (as the reference does not give explicit details), the basis set not necessarily including all the terms required to build an accurate relaxation superoperator, and the neglect of diffusion. Source code for this simulation, which documents all of the parameters used, may be found in the Supplementary Information.

that would be neglected may be determined from the spin system parameters to be $\sim 9.7 \text{ rad s}^{-1}$, the reciprocal of which is significantly larger than 3 ms, as we require.

In this simulation any comparison between the time taken for this simulation using Algorithm 1 and the slicing method is complicated by the possibility of efficient parallelization of the auxiliary matrix method: looping over the simulations corresponding to the individual dilute ^{13}C . It should also be noted that simulations using the slicing method that include heteronuclei can require more slices to be computed than in the homonuclear case [5]. The benefits of the slicing method as compared to the auxiliary matrix method are thus to an extent eroded in this simulation, benefitting the auxiliary matrix algorithm in terms of likely speedup.

4. Conclusions

A novel algorithm for simulating a pulse sequence element consisting of two gradients separated by pulses and/or delays has been presented. This algorithm makes use of an auxiliary matrix that has dimension twice the size of the dimension of the system Liouvillian, a dimensional increase much smaller than that required for prior known methods, though it is found that this does not always lead to a faster simulation. The use of this auxiliary matrix is made possible by the neglect of the effect of the gradient on chemical shift terms. Utilization of this algorithm to implement several common NMR pulse sequences has been demonstrated and the validity of ignoring the effect of the gradient on the chemical shifts discussed, with the source code available in the Supplementary Information.

Acknowledgements

I am grateful to Ilya Kuprov for helpful comments on the manuscript, to Andi Kolmer for help sourcing the strychnine spin system parameters, and to Deborah Potter for help with proof reading. This project is supported by the EPSRC (EP/F065205/1, EP/H003789/1, EP/J013080/1).

Appendix A. Proof of Equation (12)

In order to not clutter up the text with more equations than are necessary, the proof of the result shown in Equation (12) has been relegated to this appendix. While Equation (12) can be proved as a corollary of Theorem 1 in Reference [12], for clarity and completeness we here employ the method of proof used in that reference to prove the particular case under investigation.




Figure 6: Simulation of the HSQC spectrum of a 22 ^1H -spin, 12 ^{13}C spin system with gradient P-/N-selection [17] performed using Algorithm 1 with the modification suggested by Equation (22). The plot has been truncated to the densely populated region of the spectrum. The IK-2(1) basis set [13] is utilized, and the spin system parameters were taken from the file `sucrose.log` provided with *Spinach* [13]. It should be noted that these parameters come from a DFT calculation, and are likely not accurate representations of the true values for sucrose. Assuming that the ^{13}C nuclei are dilute allows the simulation to be treated as 12 individual simulations, each containing one ^{13}C nucleus, that are summed to give the full spectrum. Analytical decoupling of the ^{13}C nuclei from the ^1H nuclei is used during the final detection period. The delay Δ is taken as $\Delta = (2 \cdot 140)^{-1} \text{ s}$, and the gradient strength for both pulses is 20 G cm^{-1} with a 1.5 cm sensitive sample length; the pulses are both taken to have the sine bell shape [11], i.e. $s_1 = s_2 = 2/\pi$. Selection of the appropriate coherences [17] requires that the first gradient has length 1 ms , and the second 0.25 ms . Source code for this simulation, which documents all of the parameters used, may be found in the Supplementary Information.

The exponential of a matrix is defined by its Taylor series [23]. Multiplication of two upper triangular matrices results in another upper triangular matrix, as may be shown by examining the simple 3×3 case. As this series consists of just matrix multiplications and additions, the exponential of an upper triangular matrix will also be upper triangular. Recalling the definition of $\hat{\hat{A}}$, Equation (10), we may compute the exponential

$$\exp\{-i\hat{\hat{A}}\Delta\} = \begin{pmatrix} \hat{\hat{E}}_1(\Delta) & \hat{\hat{E}}_2(\Delta) \\ & \hat{\hat{E}}_3(\Delta) \end{pmatrix} \quad (\text{A.1})$$

where the $\hat{\hat{E}}_i(\Delta)$ are to be determined. It is important that the case where $\Delta = 1$ corresponds to Equation (11).

The equations defining the differentiation of a matrix with respect to a parameter mean that the derivative of the right hand side of Equation (A.1) may be written [20]

$$\frac{d}{d\Delta} \begin{pmatrix} \hat{\hat{E}}_1(\Delta) & \hat{\hat{E}}_2(\Delta) \\ & \hat{\hat{E}}_3(\Delta) \end{pmatrix} = \begin{pmatrix} d\hat{\hat{E}}_1(\Delta)/d\Delta & d\hat{\hat{E}}_2(\Delta)/d\Delta \\ & d\hat{\hat{E}}_3(\Delta)/d\Delta \end{pmatrix}, \quad (\text{A.2})$$

and the derivative of the left hand side of Equation (A.1) may be written [19]

$$\frac{d}{d\Delta} \exp\{-i\hat{\hat{A}}\Delta\} = -i\hat{\hat{A}} \exp\{-i\hat{\hat{A}}\Delta\}. \quad (\text{A.3})$$

Equating the right hand sides of these two equations gives

$$\begin{pmatrix} d\hat{\hat{E}}_1(\Delta)/d\Delta & d\hat{\hat{E}}_2(\Delta)/d\Delta \\ & d\hat{\hat{E}}_3(\Delta)/d\Delta \end{pmatrix} = -i\hat{\hat{A}} \exp\{-i\hat{\hat{A}}\Delta\}. \quad (\text{A.4})$$

The right hand side of this equation can be written in block matrix form by substituting in Equations (A.1) and (10) and performing the block matrix multiplication [20] to yield

$$\begin{pmatrix} d\hat{\hat{E}}_1(\Delta)/d\Delta & d\hat{\hat{E}}_2(\Delta)/d\Delta \\ & d\hat{\hat{E}}_3(\Delta)/d\Delta \end{pmatrix} = \begin{pmatrix} +i\hat{\hat{G}}_2(l)\tau_2\hat{\hat{E}}_1(\Delta) & +i\hat{\hat{G}}_2(l)\tau_2\hat{\hat{E}}_2(\Delta) + \hat{\hat{P}}\hat{\hat{E}}_3(\Delta) \\ & -i\hat{\hat{G}}_1(l)\tau_2\hat{\hat{E}}_3(\Delta) \end{pmatrix}. \quad (\text{A.5})$$

The blocks on each side of the equation are of equal size; we may thus equate respective blocks resulting in three equations:

$$\frac{d}{d\Delta} \hat{E}_1(\Delta) = +i\hat{G}_2(l)\tau_2\hat{E}_1(\Delta), \quad (\text{A.6})$$

$$\frac{d}{d\Delta} \hat{E}_2(\Delta) = +i\hat{G}_2(l)\tau_2\hat{E}_2(\Delta) + \hat{P}\hat{E}_3(\Delta), \quad (\text{A.7})$$

$$\frac{d}{d\Delta} \hat{E}_3(\Delta) = -i\hat{G}_1(l)\tau_2\hat{E}_3(\Delta). \quad (\text{A.8})$$

The definition of the matrix exponential as a Taylor series means that

$$\hat{E}_1(0) = \hat{E}_3(0) = \mathbb{I}_{\dim\{\hat{E}_1(\Delta)\}}, \quad (\text{A.9})$$

the identity matrix with the same dimension as $\hat{E}_1(\Delta)$, such that the solutions of the first and last equations are

$$\hat{E}_1(\Delta) = e^{+i\hat{G}_2(l)\tau_2\Delta}, \quad \hat{E}_3(\Delta) = e^{-i\hat{G}_1(l)\tau_1\Delta}. \quad (\text{A.10})$$

The second equation, Equation (A.7), is more interesting; the matrices of Equation (A.10) can easily be generated without recourse to an auxiliary matrix. Differentiating the matrix

$$\hat{F}(\Delta) = e^{-i\hat{G}_2(l)\tau_2\Delta}\hat{E}_2(\Delta) \quad (\text{A.11})$$

gives

$$\frac{d}{d\Delta} \hat{F}(\Delta) = e^{-i\hat{G}_2(l)\tau_2\Delta} \hat{P} e^{-i\hat{G}_1(l)\tau_1\Delta}, \quad (\text{A.12})$$

where we have substituted in $\hat{E}_3(\Delta)$ from Equation (A.10). As $\hat{E}_2(\Delta)$ is off the diagonal of the block matrix, the definition of the matrix exponential means that

$$\hat{E}_2(0) = 0_{\dim\{\hat{E}_2(\Delta)\}}, \quad (\text{A.13})$$

where $0_{\dim\{\hat{E}_2(\Delta)\}}$ is a square matrix of all zeroes with the same dimension as $\hat{E}_2(\Delta)$. Examining the definition of $\hat{F}(\Delta)$ then reveals that

$$\hat{F}(0) = 0_{\dim\{\hat{E}_2(\Delta)\}} \quad (\text{A.14})$$

also. Integration of Equation (A.12) thus gives the result

$$\hat{F}(\Delta) = \int_0^\Delta e^{-i\hat{G}_2(l)\tau_2\alpha} \hat{P} e^{-i\hat{G}_1(l)\tau_1\alpha} d\alpha \quad (\text{A.15})$$

which can be written in terms of $\hat{E}_2(\Delta)$ as

$$\hat{E}_2(\Delta) = e^{+i\hat{G}_2(l)\tau_2\Delta} \int_0^\Delta e^{-i\hat{G}_2(l)\tau_2\alpha} \hat{P} e^{-i\hat{G}_1(l)\tau_1\alpha} d\alpha. \quad (\text{A.16})$$

As mentioned above, when $\Delta = 1$ Equation (A.1) corresponds to Equation (11); we may therefore make the correspondence

$$\hat{B} = \hat{E}_2(1), \quad (\text{A.17})$$

and note that we have now proved Equation (12).

References

- [1] T. Parella, Pulsed field gradients: a new tool for routine NMR, *Magnetic Resonance in Chemistry* 36 (7) (1998) 467–495. doi:10.1002/(SICI)1097-458X(199807)36:7<467::AID-OMR325>3.0.CO;2-S.
- [2] T. D. W. Claridge, *High-resolution NMR Techniques in Organic Chemistry*, 1st Edition, Tetrahedron organic chemistry series, Elsevier, 1999.
- [3] J. Keeler, *Understanding NMR Spectroscopy*, 2nd Edition, Wiley, 2013.
- [4] G. H. Meresi, M. Cuperlovic, W. E. Palke, J. Gerig, Pulsed field gradients in simulations of one- and two-dimensional NMR spectra, *Journal of Magnetic Resonance* 137 (1) (1999) 186–195. doi:10.1006/jmre.1998.1665.
- [5] P. Nicholas, D. Fushman, V. Ruchinsky, D. Cowburn, The virtual NMR spectrometer: A computer program for efficient simulation of NMR experiments involving pulsed field gradients, *Journal of Magnetic Resonance* 145 (2) (2000) 262–275. doi:10.1006/jmre.2000.2108.
- [6] A. Jerschow, N. Müller, Efficient simulation of coherence transfer pathway selection by phase cycling and pulsed field gradients in NMR, *Journal of Magnetic Resonance* 134 (1) (1998) 17–29. doi:10.1006/jmre.1998.1491.

- [7] M. Veshtort, R. G. Griffin, SPINEVOLUTION: A powerful tool for the simulation of solid and liquid state NMR experiments, *Journal of Magnetic Resonance* 178 (2) (2006) 248–282. doi:10.1016/j.jmr.2005.07.018.
- [8] C. A. Waudby, J. Christodoulou, GPU accelerated Monte Carlo simulation of pulsed-field gradient NMR experiments, *Journal of Magnetic Resonance* 211 (1) (2011) 67–73. doi:10.1016/j.jmr.2011.04.004.
- [9] R. R. Ernst, G. Bodenhausen, A. Wokaun, *Principles of Nuclear Magnetic Resonance in One and Two Dimensions*, International Series of Monographs on Chemistry, Clarendon Press, 1990.
- [10] D. S. Grebenkov, NMR survey of reflected Brownian motion, *Rev. Mod. Phys.* 79 (2007) 1077–1137. doi:10.1103/RevModPhys.79.1077.
- [11] D. Sinnaeve, The Stejskal–Tanner equation generalized for any gradient shape—an overview of most pulse sequences measuring free diffusion, *Concepts in Magnetic Resonance Part A* 40A (2) (2012) 39–65. doi:10.1002/cmr.a.21223.
- [12] C. F. Van Loan, Computing integrals involving the matrix exponential, *Automatic Control, IEEE Transactions on* 23 (3) (1978) 395–404. doi:10.1109/TAC.1978.1101743.
- [13] 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, *Journal of Magnetic Resonance* 208 (2) (2011) 179–194. doi:10.1016/j.jmr.2010.11.008.
- [14] Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd Edition, Society for Industrial and Applied Mathematics, 2003.
- [15] A. Karabanov, I. Kuprov, G. T. P. Charnock, A. van der Drift, L. J. Edwards, W. Köckenberger, On the accuracy of the state space restriction approximation for spin dynamics simulations, *Journal of Chemical Physics* 135 (8) (2011) 084106. doi:10.1063/1.3624564.
- [16] S. Berger, S. Braun, *200 and more NMR experiments: a practical course*, Wiley-VCH, 2004.

- [17] A. L. Davis, J. Keeler, E. D. Laue, D. Moskau, Experiments for recording pure-absorption heteronuclear correlation spectra using pulsed field gradients, *Journal of Magnetic Resonance* (1969) 98 (1) (1992) 207–216. doi:10.1016/0022-2364(92)90126-R.
- [18] A. L. Davis, E. D. Laue, J. Keeler, D. Moskau, J. Lohman, Absorption-mode two-dimensional NMR spectra recorded using pulsed field gradients, *Journal of Magnetic Resonance* (1969) 94 (3) (1991) 637–644. doi:10.1016/0022-2364(91)90154-L.
- [19] W. Rossmann, *Lie Groups: An Introduction Through Linear Groups*, Oxford graduate texts in mathematics, Oxford University Press, 2002.
- [20] G. H. Golub, C. F. Van Loan, *Matrix Computations*, 3rd Edition, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, 1996.
- [21] I. Kuprov, Diagonalization-free implementation of spin relaxation theory for large spin systems, *Journal of Magnetic Resonance* 209 (1) (2011) 31–38. doi:10.1016/j.jmr.2010.12.004.
- [22] C. P. Butts, C. R. Jones, J. N. Harvey, High precision NOEs as a probe for low level conformers—a second conformation of strychnine, *Chem. Commun.* 47 (2011) 1193–1195. doi:10.1039/C0CC04114A.
- [23] C. Moler, C. F. Van Loan, Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later, *SIAM Review* 45 (1) (2003) 3–49. doi:10.1137/S00361445024180.