Available online at www.sciencedirect.com
ScienceDirect
Journal of Magnetic Resonance 189 (2007) 241–250
JMR Journal of Magnetic Resonance
www.elsevier.com/locate/jmr

Polynomially scaling spin dynamics simulation algorithm based on adaptive state-space restriction

Ilya Kuprov *, Nicola Wagner-Rundell, P.J. Hore

Department of Chemistry, University of Oxford, Physical and Theoretical Chemistry Laboratory, South Parks Road, Oxford OX1 3QZ, UK

Received 15 June 2007; revised 8 September 2007
Available online 21 September 2007

Abstract

We report progress with an old problem in magnetic resonance—that of the exponential scaling of simulation complexity with the number of spins. It is demonstrated below that a polynomially scaling algorithm can be obtained (and accurate simulations performed for over 200 coupled spins) if the dimension of the Liouville state space is reduced by excluding unimportant and unpopulated spin states. We found the class of such states to be surprisingly wide. It actually appears that a *majority of states in large spin systems are not essential in magnetic resonance simulations and can safely be dropped from the state space.* In restricted state spaces the spin dynamics simulations scale polynomially. In cases of favourable interaction topologies (sparse graphs, e.g. in protein NMR) the asymptotic scaling is linear, opening the way to direct fitting of molecular structures to experimental spectra.

Keywords: NMR; EPR; Spin; Simulation; Polynomial scaling

1. Introduction

Spin dynamics simulations are a powerful and well developed tool used extensively in all areas of magnetic resonance spectroscopy [1–4]. The most popular approaches include the diagonalization [5,4], direct propagation [6,2], many-body theory [7,8], perturbation theory [9,10] and Campbell–Baker–Hausdorff series techniques [11,12], frequently adapted to include spatial [13], temporal [14] and permutation [15–17] symmetry, as well as a host of analytical [18,19] and numerical [1,20,21] techniques for pulse sequence design and spin system analysis. Such simulations have been successfully used for decades and excellent reviews exist, outlining the theory, implementation and programming in fine detail [2–4].

This tranquil picture, however, has one important problem: it only works for fewer than $n \approx 10$ spins. In numerical simulations, the dimension of operator matrices is $\geqslant 2^n$ and a propagation step requires $\geqslant 4^n$ multiplications, meaning that the computation becomes prohibitively large for more than ten, or with some tricks fifteen, spins. As n goes up, the situation steadily gets worse, and for systems with over 20 spins the Liouvillian matrix cannot currently even be stored, let alone diagonalized: we hit the notorious 'exponential scaling wall' [22], and in quantum spin dynamics there currently is no way around it.

We demonstrate in this communication that a polynomially scaling spin dynamics simulation algorithm can be obtained if certain assumptions are made about the structure of the state space in which the spins evolve. Specifically, it appears that a large number of states in big spin systems (very high multi-spin orders and orders linking the spins that are remote on the interaction graph) are not essential in magnetic resonance simulations and can be dropped from the state space. In restricted state spaces the matrix dimensions scale polynomially and accurate simulations may be performed for over 200 coupled spins.

* Corresponding author. Fax: +44 1865 275410.
E-mail address: ilya.kuprov@chem.ox.ac.uk (I. Kuprov).

1090-7807/$ - see front matter © 2007 Elsevier Inc. All rights reserved.
doi:10.1016/j.jmr.2007.09.014

2. The case for state-space restriction

Within the density operator formalism [23], the operators spanning the state space may be written in the form $P_1 \otimes P_2 \otimes \cdots \otimes P_n$, where each P_m is either a Pauli matrix of spin m or an identity operator. If we were to restrict the state space (thus far for no particular reason) and exclude entanglements of more than $k < n$ spins, the operators would take the form: $\cdots E \otimes P_1 \otimes \cdots \otimes P_2 \otimes E \otimes \cdots \otimes P_k \cdots$, so that k Pauli matrices are scattered among the identity operators. The dimension N of the new state space is

$$N = 4^k \frac{n!}{k!(n-k)!} = 4^k \frac{1}{k!} n(n-1)\cdots(n-k+1) = O((4n)^k) \tag{1}$$

(with trivial modifications if spins greater than $I = 1/2$ are present). The cost of the system propagation step in this restricted space is $O((4n)^{2k})$ multiplications, so the scaling is polynomial in the total number of spins n. This is a tantalizing possibility: if state-space restriction is a good approximation, it promises fundamental improvements in the simulation time. It is also clear that the procedure can be made adaptive, with the level of description depending on the local coupling density.

A number of arguments may be put forth to justify the exclusion of high-order entanglements (illustrated schematically in Fig. 1). We may remember that (with rare, but notable exceptions [24–27]) high spin orders will be relaxing or dephasing faster than low orders and therefore might not accumulate in large quantities. Another appealing (if retrospective) argument comes from the approximate methods introduced during the early days of magnetic resonance: product operator treatments of even very complicated NMR pulse sequences rarely included spin orders greater than four [28], suggesting that higher orders, describing the fine multiplicity details, may be relatively unimportant. Furthermore, Lee et al. have successfully obtained analytical expressions for anomalous dipole–dipole peaks in COSY and CRAZED type spectra by neglecting certain high-order coherences [29]. It is also known that in a linear chain of binary couplings, a two-spin order with a spin that is k couplings away will only accumulate in k^{th} order upwards in time-dependent perturbation theory (TDPT) [23]. That is, within the convergence radius of TDPT, the accumulation will be slower the further the two spins are away from each other on the interaction graph. For very long-range connectivities it might therefore be reasonable to suppress the corresponding entanglements altogether. Yet another argument comes from the success of the Bloch–Redfield–Wangsness relaxation theory [9,10]—if a second-order TDPT expression (which by definition only includes direct coherence transfer to at most the next nearest neighbor) is appropriately closed in time (the $\rho(0) \rightarrow \rho(t)$ substitution [9]), the resulting equations correctly describe the phenomena well beyond what might be expected from second-order perturbation theory [10].

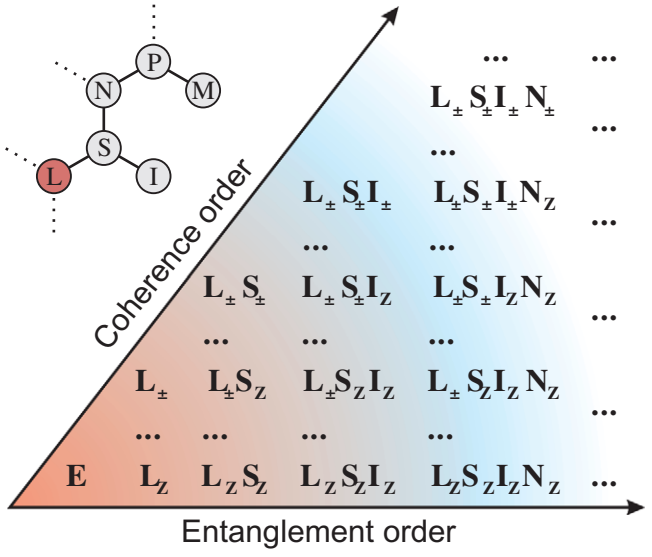

Fig. 1. Schematic illustration of the restricted state-space approximation. For a given spin, the most important states in the state space are likely to involve its nearest neighbors on the interaction graph. Very high coherences and entanglements are likely to be unimportant (see text for detailed analysis of these assumptions).

It thus appears that restricting the state space to include only low-order coherences and entanglements between nearby spins is (*a*) not entirely unreasonable, at least in the liquid state, and (*b*) promises fundamental improvements in the spin dynamics simulation times. We shall therefore proceed to a deeper algebraic analysis of the consequences of state-space restrictions.

3. Algebraic consequences of state-space restriction

For a given static Liouvillian superoperator $\hat{\hat{L}}$ the set G of all possible forward and backward propagators is a uniparametric Lie group [30]:

$$\begin{aligned} &\mathrm{e}^{-i\hat{\hat{L}}t_1}, \mathrm{e}^{-i\hat{\hat{L}}t_2} \in G \Rightarrow \mathrm{e}^{-i\hat{\hat{L}}t_1}\mathrm{e}^{-i\hat{\hat{L}}t_2} = \mathrm{e}^{-i\hat{\hat{L}}(t_1+t_2)} \in G, \\ &\forall \mathrm{e}^{-i\hat{\hat{L}}t_1} \in G \; \exists \mathrm{e}^{i\hat{\hat{L}}t_1} \in G \quad \text{s.t. } \mathrm{e}^{-i\hat{\hat{L}}t_1}\mathrm{e}^{i\hat{\hat{L}}t_1} = \hat{\hat{E}}, \end{aligned} \tag{2}$$

that is, a product of two propagators is another propagator and for every propagator there is a unique inverse. The set of elements $G(\hat{S}) = \{\mathrm{e}^{-i\hat{\hat{L}}t}\hat{S}, t \in \mathbb{R}\}$ is therefore a propagator group orbit of the initial state operator $\hat{S}$. The basis of the space containing $G(\hat{S})$ is the minimal set of operators required to describe the spin-system evolution exactly. Reduced state-space approximations should therefore be aimed at reducing (exactly[1] or approximately) the dimension of the space enclosing $G(\hat{S})$.

Simply truncating the basis set to low spin orders is not an option for the following reasons. Let the complete state space be Q and the desired reduced state space be K. An

[1] An example would be the common practice of treating uncoupled subsystems separately. For independent subsystems the inter-subsystem entanglement terms are usually outside the propagator group orbit.

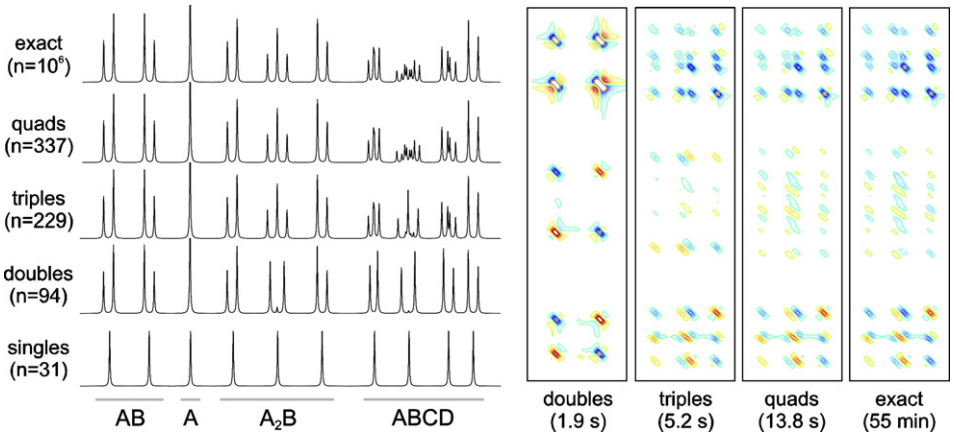

Fig. 2. State-restricted NMR spectrum simulation results for a model 10-spin system (composed of 4 independent subsystems, from left to right: AB, A, A_2B, ABCD). (Left panel) 1D pulse-acquire spectrum simulation results (*singles*, *doubles*, *triples* and *quads*refer to the maximum allowed spin-state order) and the corresponding state space dimensions. (Right panel) 2D COSY spectrum simulation convergence with the increase in approximation level and CPU time used by the simulation.

unsuitably chosen K is very likely not to be closed with respect to infinitesimal propagation—it would contain operators that leak out of it when the Hamiltonian superoperator is applied:

$$\exists \hat{\rho} \in K, \quad -i[\hat{H}, \hat{\rho}] \in Q \setminus K \tag{3}$$

Such operators $\hat{\rho}$ may be called "loose ends". Accepting the fact that the magnetization is lost when it is sent out of K would lead to something resembling relaxation and would amount to assuming that the states in $Q\backslash K$ (states in Q that are not in K) relax infinitely fast. This, however, implies that the propagator will be non-unitary, and therefore the effective Liouvillian non-Hermitian. This will lead to complex-valued energies and twisted phases and is best avoided. An alternative strategy would be to keep the magnetization from diffusing away from the restricted state space by renormalizing the truncated propagator, so that it remains a unitary operator. However, the recovered magnetization would still be returned into the system in a non-physical fashion: it will be smeared across the spin system; this is also best avoided. Therefore, for any reduced state space K it is a requirement that it is naturally closed with respect to the original or modified temporal propagation operator—for any state in K the result of its propagation must also be in K:

$$\forall \hat{\rho} \in K, \quad -i[\hat{H}, \hat{\rho}] \in K \tag{4}$$

One way of achieving this would be to avoid direct manipulation or truncation of the propagators (because of the above concerns), and manipulate the state space at the Liouvillian level. The propagator non-unitarity problem then does not appear so long as the modified Liouvillian stays Hermitian.

Because we know that small amounts of high-order multi-spin entanglement *are* generated in the exact simulations, it is clear that the operators we are about to neglect are present in the system orbit. The task is therefore narrowed to deflecting the orbit and preventing it from passing through high entanglements. It is also clear that this objective cannot be achieved without modifying the Hamiltonian. The conclusion therefore is that the state restriction should be accompanied by such modification of the system Hamiltonian as would satisfy condition (4) yet have a minimal impact on the simulation quality.

4. Restricted state-space simulation results

Before embarking on a technical description of the construction of restricted state spaces and Liouvillians satisfying conditions (2) and (4), we wish to present the results of some simulations that utilize the restricted state space approach. The mathematical and programming details are presented in full in the subsequent sections and the annotated program source code is included in the Supplementary information.

Historically the most challenging parameter in all types of approximations in magnetic resonance was the signal phase [31,32]. We are pleased to report that the state-restricted simulations get the phases precisely right at all approximation levels (Fig. 2, left panel). The second feature in the difficulty ranking—intensity patterns for strongly coupled nuclei—is also reproduced: once a coupling is included, the relative peak intensities come out correctly. The coupling patterns neatly approach the exact result as the approximation level is increased: first the doublets appear in their exact form, then triplets and finally the more complicated splittings. Similar things happen to COSY cross-peaks (Fig. 2, right panel). Finally, the signal integrals do not depend on the approximation level and always come out correctly.

It follows from both the definition of the restricted state-space and from the simulation behavior in Fig. 2 that the exact result is the natural limit of the approximation made. However, the reduction in computational complexity is

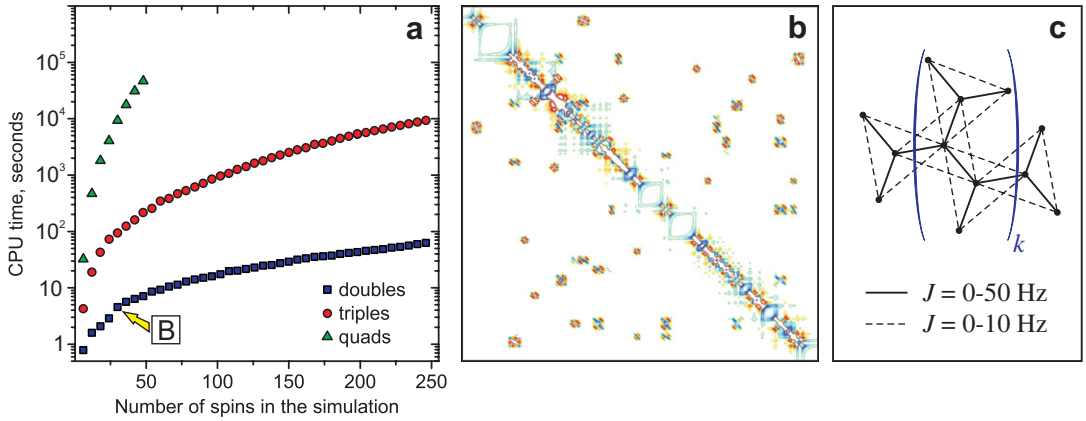

Fig. 3. (a) CPU time (AMD Opteron 265 / DDR400) required for the simulation of a 2D COSY spectrum of spin systems containing between 6 and 246 spins. (b) An example output of the simulation (indicated with yellow arrow in the left hand panel). (c) The coupling structure of the spin system used for these simulations. The couplings are randomly selected within the indicated intervals, the chemical shifts are randomly selected within the spectral window.

Table 1
State space and matrix density statistics for the full and state-restricted simulations of the spin systems in Fig. 3c

Number of spins	State-space dimension				Quads Liouvillian density (%)	Quads propagator density (%)	Quads, speed-up factor
	Full	Doubles	Triples	Quads			
6	4096	109	532	1747	0.63	1.2	~2
12	1.7×10^{7}	262	2134	15,913	0.068	0.10	$\sim10^{3}$
18	6.9×10^{10}	415	3943	37,981	0.028	0.045	$\sim10^{6}$
24	2.8×10^{14}	568	5752	60,985	0.017	0.024	$\sim10^{9}$

dramatic: with 4-spin clusters,[2] which reproduce the spectrum in Fig. 2 exactly, the simulation is done with just 337 states as opposed to around 10^6 in the exact formulation (Table 1). As a further demonstration, Fig. 3a reports the timings for the simulation of a series of much more complicated (and no longer disjoint) spin systems shown in Fig. 3c. With a restricted state space, an explicit time-domain density matrix propagation simulation of a 1024×1024 COSY spectrum of a strongly and densely coupled 54-spin system takes 10 s with doubles, 100 s with triples and about an hour with quads. The exact simulation of a 2D spectrum of a coupled 50+ spin system is beyond current computers.

The rule of thumb for the couplings is that including the spin states up to *n*-tuples will correctly reproduce multiplet structure arising from up to *n* coupled spins. Higher multiplicities will still be represented but with some distortions (Fig. 2, left panel). Given that the state-space restriction procedure (as well as the program listed in the Supplementary information) supports variable approximation levels, tight knots of *J*-couplings (e.g. a valine side chain) may be accommodated by a local increase in the number of spin states. In principle, just accepting these distortions is an option: in the experiments where those splittings are decoupled or not resolved, the relevant cross-peaks will still appear and will have the correct intensity.

Somewhat counterintuitively, restricted state space simulations can produce reasonably accurate results for quite densely coupled spin systems. Fig. 4 shows a triples-restricted simulation of the effect of a weak applied magnetic field on the product yield of backward electron transfer in a spin-correlated radical pair formed by photolysis of pyrene and *N*,*N*-dimethylaniline [33–35]. The circles show the experimental data (presented as the first derivative of the product yield with respect to the magnetic field intensity) which are essentially identical to an exact simulation (1.6×10^6 states) using literature values of the 18 largest hyperfine couplings [35]. The triples-restricted simulation (4008 states), using the same hyperfine couplings agrees with the data remarkably well in the practically interesting [34,35] "low field" region where the electron Zeeman interaction and the hyperfine interactions are comparable in strength. Under these conditions non-secular parts of the hyperfine interactions cannot be neglected, and this is therefore a clear example of a large spin system in which the traditional approximate approaches (e.g. the weak coupling assumption) cannot be employed, and in which conventional methods of analysis are seriously impeded by the enormous size of the spin system.

Another remarkable finding is that, even though we have restricted the state space to just the states which are supposed to be important and participate actively in creat-

[2] Here and below, we refer to 1-, 2-, 3-, 4- and *n*-spin clusters as, respectively, *singles*, *doubles*, *triples*, *quads* and *n-tuples*. It is important to note that the clustering procedure is only used to decide which spin states to include—the spin system is always simulated as a whole (see the algorithm description below).

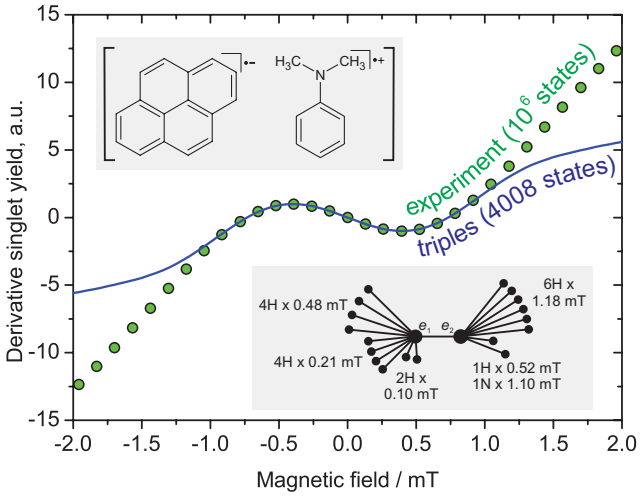

Fig. 4. Experimental data [35] (circles) showing the effect of a weak applied magnetic field on the backward electron transfer in the photo-generated pyrene / *N*,*N*-dimethylaniline radical pair and the result of the triples-restricted simulation (annotated program code included in the Supplementry information, for theoretical details see Ref. [48]). (Upper inset) Structure of the two radicals involved in the reaction. (Lower inset) Spin system topology and experimental values [33] of the isotropic hyperfine coupling constants used in the simulation. See text for further details.

ing the coupling structure, both the Liouvillian and the associated propagator are still very sparse (Table 1). On the one hand this makes the simulation very CPU- and memory-efficient, on the other this might be a hint that even the intelligent state space restriction described below is not the last step in improving the computational efficiency and reducing the dimension of the spectral simulation problem. Finally, an interesting result is that the restricted Liouvillian, even for quads, is actually *orders of magnitude smaller* than the full Hamiltonian (instead of having the square of its dimension). Given the accuracy of state-restricted simulations reported in Figs. 2–4 and the state space reduction reported in Table 1, these facts are astonishing. It appears that the majority of states in any large spin system are not essential in NMR simulations and can safely be dropped from the state space.

5. Restricted state-space construction

We will now outline the principal stages in the construction of the restricted state space and the effective Liouvillian satisfying condition (4) as well as the group property (2). The procedure is schematically illustrated in Fig. 5 and the annotated *Matlab* source code is included in the Supplementary information.

5.1. Interaction graph construction and connected subgraph expansion

For a given spin, the important states and coherences are likely to involve its close neighbors on the interaction graph. One can identify a subgraph (hereafter *cluster*) of spins, which are the nearest neighbors, next nearest neighbors, etc. Constructing the overall state space as a direct sum of state spaces of these clusters (Fig. 5) satisfies our requirements: the important states are included, and entanglements higher than the size of the biggest cluster are excluded. Furthermore, cluster size can be varied through the system to adapt to local changes in the coupling density.

Graph-theoretical algorithms, as applied to chemical problems, have recently been reviewed by Klein et al. [36]. Linear-time algorithms exist for the enumeration of sub-trees [37] and sub-graphs [38]. We adapted the depth-first path-tracing procedure [39]. In the simulation, the spin interaction matrices are supplied by the user as a part of the system specification. The adjacency matrix for the spin interaction graph is then a sparse logical array of the following form:

$$\texttt{conmatrix} = \begin{cases} 1 & \text{if } i \text{ and } j \text{ are coupled} \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

where `conmatrix` is the name of the corresponding variable in the enclosed *Matlab* code. Connected sub-graph expansion then results in a sparse logical array of cluster membership specifications:

$$\texttt{cluster\_list} = \begin{bmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{k1} & \cdots & a_{kn} \end{bmatrix};$$
$$a_{kn} = \begin{cases} 1 & \text{if spin } n \text{ is in cluster } k \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

While other ways, of course, exist to store the cluster membership list (for example, simply specifying the spin numbers), the form given in Eq. (6) is preferable because use can be made in subsequent stages of *Matlab*'s formidable logical and matrix indexing capabilities. The static dimensionality of the cluster list as given in Eq. (6) also permits straightforward implementation of adaptive clustering.

In the case when dipolar interactions are averaged to first order, the interaction graph is usually very sparse. When dipolar interactions are present in the static Hamiltonian, the non-zero band of the graph is broader, but the connectivity is still significantly sparser than "everything to everything" assumed in the scaling estimate in Eq. (1).

5.2. Local and global state list generation

Because each cluster is small enough for the full quantum mechanical treatment to be performed, the local operator basis is generated in the usual way, using direct products of propagator group generators and unit operators:

$$\hat{S}_m^{(k)} = \otimes_{q=1}^{n} \hat{L}_q; \quad \hat{L}_q = \hat{E}, \hat{\sigma}_z, \hat{\sigma}_+, \hat{\sigma}_- \ldots \tag{7}$$

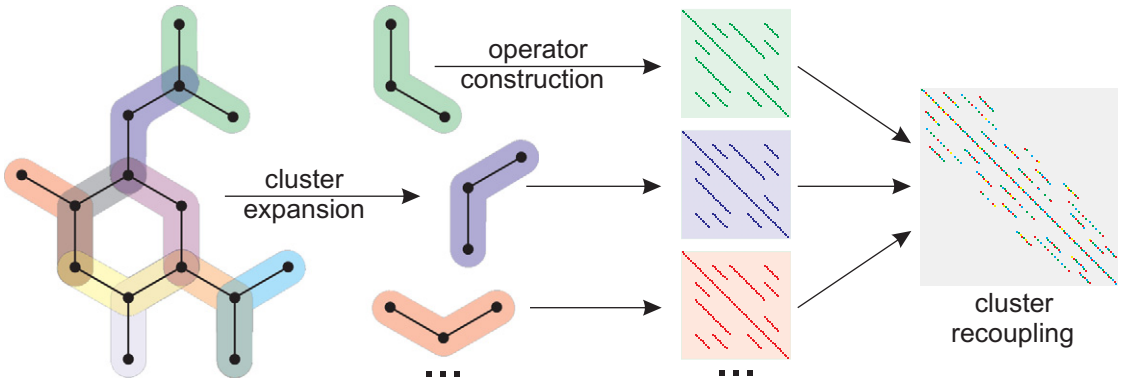

Fig. 5. Schematic illustration of the restricted state-space treatment. The interaction graph is expanded in a complete set of connected subgraphs (*spin clusters*) of the specified size, each cluster is treated quantum mechanically and the resulting Liouvillian superoperators are recoupled in such a way as to preserve the group theoretical properties of the resulting propagators. For detailed description see text.

Assuming the following indexing convention for the states in the direct product (7)

$$\texttt{cluster(k).state\_list(m)} = [s_{m1} \quad \cdots \quad s_{mn}];$$

$$s_{mq} = \begin{cases} 0 & \text{if} \quad \hat{L}_q = \hat{E} \\ 1 & \text{if} \quad \hat{L}_q = \hat{\sigma}_z \\ 2 & \text{if} \quad \hat{L}_q = \hat{\sigma}_+ \\ 3 & \text{if} \quad \hat{L}_q = \hat{\sigma}_- \\ \cdots & \end{cases} \tag{8}$$

(higher numbers will be required for spin quantum numbers larger than 1/2), the generation of a complete cluster state list is a simple combinatorial procedure involving the permutations of integers and zeros. Again, the sparse matrix indexing is preferred to direct enumeration for reasons of computational convenience and speed.

The global state list is obtained by vertical concatenation of the cluster state lists and elimination of repetitions. The resulting global state list does not contain any spin orders higher than the biggest cluster size and is consistent with the interaction topology of the system: there would be no cross-cluster entanglements unless several clusters contain the same subset of spins.

5.3. State-cluster cross-membership matrix generation

The recoupling stage described in the next section critically depends on efficient cluster and state indexing. For every cluster the local state list must be available, and for every state the list of all hosting clusters. The most convenient form of this index is the cross-membership matrix stored as a binary sparse array:

$$\texttt{xmm(i,j)} = \begin{cases} 1 & \text{if } \texttt{state (i)} \in \texttt{cluster (j)} \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

Computing this array can potentially be an expensive procedure, because the length of both the cluster list and the state list may be significant. We suggest a simple hashing procedure here, based on the fact that the sum of elements along the n index in the state specification (Eq. (8)) would not change under element-by-element multiplication by the hosting cluster specification line (Eq. (6)), but will be reduced in the case of a non-hosting cluster. By construction, this procedure utilizes fast binary and integer arithmetic and avoids an element-by-element state list trawl.

5.4. Building local and global Liouvillians

The clustering procedure described above is only used to generate the new state space—the spin system itself must, of course, be simulated as a whole. This section describes the 'recoupling' stage—generation of a single global Liouvillian from Liouvillians of individual clusters. This is the most important step in the algorithm, and the one that was not available in prior work on hierarchical coupling analysis [40].

Inside every cluster Cl the Hamiltonian $\hat{H}_{Cl}$ and Liouvillian $\hat{L}_{Cl}$ matrices for the overall evolution and the evolution under any RF pulses are computed exactly:

$$\begin{aligned} \hat{H}_{Cl} &= \sum_{n\in Cl} \hat{\vec{S}}_n \mathbf{g}_n \vec{B} + \sum_{n,m\in Cl} \hat{\vec{S}}_n \mathbf{A}_{nm} \hat{\vec{S}}_m + \cdots \\ \hat{\hat{L}}_{Cl} &= \hat{H}_{Cl} \otimes \hat{E} - \hat{E} \otimes \hat{H}_{Cl}^{\mathrm{T}} \end{aligned} \tag{10}$$

Importantly, the Liouvillian element linking two states $\hat{S}_1$ and $\hat{S}_2$ is the same in every cluster that hosts both $\hat{S}_1$ and $\hat{S}_2$:

$$\hat{S}_1, \hat{S}_2 \in Cl_1, Cl_2 \Rightarrow \langle \hat{S}_1 | \hat{\hat{L}}_{Cl_1} | \hat{S}_2 \rangle = \langle \hat{S}_1 | \hat{\hat{L}}_{Cl_2} | \hat{S}_2 \rangle \tag{11}$$

This is the consequence of the fact that if the two states are in a cluster, then their constituent spins are also in that cluster, meaning that the entire section of the Hamiltonian representing all interactions within that group of spins will also be there (Eq. (10)). The Hamiltonian and Liouvillian matrix elements describing those interactions will therefore be exact and therefore the same in every cluster.

With Eq. (11) in place, the global state-restricted Liouvillian is generated from cluster Liouvillians by the procedure that is best described as matrix blending. If a pair of

states (i,j) has no cluster in common, zero is written; if they do share a cluster, then the corresponding element is fetched from the Liouvillian matrix of that cluster:

$$L_{ij} = \bigcup_{Cl} \begin{cases} 0 & \text{if } \hat{S}_i \notin Cl \vee \hat{S}_j \notin Cl \\ L_{mk}^{(Cl)} & \text{if } \hat{S}_i \in Cl \wedge \hat{S}_j \in Cl \end{cases} \tag{12}$$

This is a critical step, which accomplishes seamless re-coupling of clusters for subsequent simulation of the spin system as a whole in the restricted state space. The union over clusters in Equation is applied in the following sense: the elements of the Liouvillian operator of every cluster in turn are re-indexed (states $\hat{S}_i$ and $\hat{S}_j$ on the global list are states $\hat{S}_m$ and $\hat{S}_k$ inside the cluster Cl) and assigned to the elements of the global Liouvillian. The non-zero elements are the same in every hosting cluster (Eq. (11)) and therefore do not change even if reassigned multiple times.

The computational complexity of every step in the procedure described above scales linearly with the number of clusters in the system. The resulting restricted state space and the set of global Liouvillians (for overall propagation and, perhaps, for pulses) are mutually correct in the sense of Eq. (4). Furthermore, because the manipulations had been performed at the Lie algebra level, the group property of the propagators (Eq. (2)) is also preserved, and the propagator group orbit of the initial state is correctly confined, without leakage, to a reduced state space, which is based on the actual interaction topology in the system.

6. Propagator construction and system propagation

The generation of theoretical magnetic resonance spectra requires either Hamiltonian diagonalization followed by the application of Fermi's golden rule, or numerical propagation of the density matrix followed by a trace operation and Fourier transformation [23]. The diagonalization procedure works well for small systems, but becomes impractical when the matrix dimensions exceed 10^4 because the eigenvector matrices are usually dense. For large spin systems, which are the subject of this paper, we are therefore confined to numerical propagation. Furthermore, the restricted state-space approximation is essentially a Liouville-space technique: once the space is set up, the operators no longer have the $\hat{H} \otimes \hat{E} - \hat{E} \otimes \hat{H}^{\mathrm{T}}$ form and cannot be projected back into Hilbert space. Nor is Hilbert space operation desirable, because it would neglect many relaxation pathways. We therefore need to set up numerical propagation in Liouville space, which is actually quite manageable now that its dimension has been reduced.

For 200+ spin systems even the restricted state-space dimension is large enough to make the direct calculation of the propagator $\exp(i\hat{\hat{L}}\Delta t)$ problematic: the standard scaling and squaring (S&S) technique (in the shape of *Matlab*'s `expm` function) [41] overflows 4 GB of RAM around $n = 50$. The alternative option of using the Krylov space technique [42] to compute $\exp(i\hat{\hat{L}}\Delta t)\hat{\rho}(t)$ directly from $\hat{L}$ and $\hat{\rho}(t)$ requires less memory but is much slower, because it has to be applied at every step. The situation might have been quite difficult, if it were not for another convenient property of restricted state spaces.

A remarkable feature of a state-restricted Liouvillian is its narrow eigenvalue spectrum: $|\lambda| < k \max_n\{\omega_n\}$, where k is the maximum cluster size and ω_n are Zeeman frequencies. This would compare with $|\lambda| \leqslant \sum_n |\omega_n|$ in the exact case. This means that the maximum eigenvalue of $i\hat{\hat{L}}\Delta t$ for a sub-Larmor time step Δt would be $|\lambda_{\max}\Delta t| \leqslant \pi k$, where k is a small integer. That in turn means that the power series for $\exp(i\hat{\hat{L}}\Delta t)$ will converge very rapidly. Because of this narrow eigenvalue range and the fact that the restricted Liouvillian is still very sparse (Table 1), in practical calculations the straightforward Taylor expansion is faster than Krylov and S&S[3] techniques by several orders of magnitude, permitting the calculation of exponentials for restricted state-space dimensions in excess of $N = 10^5$. The Krylov space technique may still be preferable in the cases where the number of time steps is smaller than the state-space dimension (i.e. in most 1D simulations). When the opposite is the case (most 2D and 3D experiment simulations), it is reasonable to pre-compute the propagators.

The computational scaling of the propagator calculation and system propagation will depend on the density of the matrices involved and vary between quadratic (with dense matrices) and linear (with diagonal and near-diagonal sparse matrices). Table 1 shows that, for a system in Fig. 3c, both the Liouvillian and the propagator are very sparse indeed, with the share of non-zeros well under 1% and decreasing with the spin system size. Direct inspection also reveals that the non-zeros are clustered along the diagonal. With linear belts of J-couplings (e.g. for proteins in the liquid phase) we should therefore expect the asymptotic scaling to be linear.

7. Parametric gradient and Hessian

An important application of a technique for simulating systems containing hundreds of spins would be direct fitting of macromolecular structures to experimental data (bypassing or improving upon the currently ubiquitous molecular dynamics stage [43,44]). Efficient minimizers require a gradient and (ideally) a Hessian of the simulation with respect to the simulation parameters [45]. While both can be obtained by finite difference calculations, such estimates are usually computationally inefficient. We will now demonstrate that the analytical gradient and Hessian of the simulation may be computed directly within the restricted state-space formalism.

Taking derivatives of the Liouville–von Neumann equation with respect to the Hamiltonian parameters, we obtain:

[3] Technically speaking, Taylor series is a special case of the scaling and squaring technique (S&S uses Padé approximants [41]). The memory problem might therefore be specific to *Matlab*'s `expm`, and an appropriately modified S&S technique may be faster than Taylor series.

$$\begin{cases} \frac{\partial\hat{\rho}}{\partial t} = -i\hat{\hat{L}}\hat{\rho} \\ \hat{\rho}(0) = \hat{\rho}_0 \end{cases} \xrightarrow{\partial/\partial J} \begin{cases} \frac{\partial\hat{\rho}'_J}{\partial t} = -i(\hat{\hat{L}}\hat{\rho}'_J + \hat{\hat{L}}'_J\hat{\rho}) \\ \hat{\rho}'_J(0) = 0 \end{cases}$$
$$\xrightarrow{\partial/\partial K} \begin{cases} \frac{\partial\hat{\rho}''_{JK}}{\partial t} = -i(\hat{\hat{L}}\hat{\rho}''_{JK} + \hat{\hat{L}}'_J\hat{\rho}'_K + \hat{\hat{L}}'_K\hat{\rho}'_J) \\ \hat{\rho}''_{JK}(0) = 0 \end{cases} \tag{13}$$

where ρ'_J denotes a directional derivative of the density matrix along the interaction parameter vector J. Note also that $\hat{\hat{L}}''_{JK} = 0$. Eq. (13) can be solved analytically:

$$\hat{\rho}'_J(t) = \left[\int_0^t \mathrm{e}^{i\hat{\hat{L}}(t'-t)} i\hat{\hat{L}}'_J \mathrm{e}^{-i\hat{\hat{L}}(t'-t)}\,\mathrm{d}t\right]\hat{\rho}(t)$$
$$\hat{\rho}''_{JK}(t) = -\left[\int_0^t \mathrm{d}t' \int_0^{t'} \mathrm{e}^{i\hat{\hat{L}}(t''-t)}(\hat{\hat{L}}'_J\hat{\hat{L}}'_K + \hat{\hat{L}}'_K\hat{\hat{L}}'_J)\mathrm{e}^{-i\hat{\hat{L}}(t''-t)}\,\mathrm{d}t''\right]\hat{\rho}(t) \tag{14}$$

Although Eqs. (14) do formally provide closed expressions for density matrix derivatives, the parametric differentiation operators in square brackets are very expensive to compute. Under special circumstances, namely if the derivative direction vector J is chosen in such a way as to ensure that $[\hat{\hat{L}}'_J, \hat{\hat{L}}] = 0$, Eq. (14) do simplify into

$$\hat{\rho}'_J(t) = i\hat{\hat{L}}'_J t\hat{\rho}(t)$$
$$\hat{\rho}''_{JK}(t) = -(\hat{\hat{L}}'_J\hat{\hat{L}}'_K + \hat{\hat{L}}'_K\hat{\hat{L}}'_J)\frac{t^2}{2}\hat{\rho}(t) \tag{15}$$

In that case, given a sufficiently large set of linearly independent parameter vectors along which the derivatives in Eq. (15) are taken, the gradient and the Hessian of the simulation may be recovered by a linear transformation. It is far from clear, however, whether the required parameter space directions exist in sufficient numbers.

A less computationally expensive avenue to the same goal would make use of the way that the density matrix is propagated in a state-restricted simulation:

$$\begin{aligned}
\hat{\rho}(t+\Delta t) &= \exp\{-i\hat{\hat{L}}\Delta t\}\hat{\rho}(t); \quad \hat{\rho}(0) = \hat{\rho}_0 \\
\hat{\rho}'_J(t+\Delta t) &= \frac{\partial}{\partial J}(\exp\{-i\hat{\hat{L}}\Delta t\})\hat{\rho}(t) \\
&\quad + \exp\{-i\hat{\hat{L}}\Delta t\}\hat{\rho}'_J(t); \quad \hat{\rho}'_J(0) = 0 \\
\hat{\rho}''_{JK}(t+\Delta t) &= \frac{\partial^2}{\partial J\partial K}(\exp\{-i\hat{\hat{L}}\Delta t\})\hat{\rho}(t) \\
&\quad + \frac{\partial}{\partial J}(\exp\{-i\hat{\hat{L}}\Delta t\})\hat{\rho}'_K(t) \\
&\quad + \frac{\partial}{\partial K}(\exp\{-i\hat{\hat{L}}\Delta t\})\hat{\rho}'_J(t) \\
&\quad + \exp\{-i\hat{\hat{L}}\Delta t\}\hat{\rho}''_{JK}(t); \quad \hat{\rho}''_{JK}(0) = 0
\end{aligned} \tag{16}$$

In the propagation equation for the derivative density matrix the second term is known completely (the propagator was computed in the main problem and $\hat{\rho}'_J(t)$ comes from the previous step of the derivative problem) and in the first term $\hat{\rho}(t)$ is known from the main problem. The only unknown term is therefore the propagator derivative, and it is easily computed using the property we noted in the previous section about the eigenvalue spectrum of $i\hat{\hat{L}}\Delta t$ and the fast convergence of the exponential series for the propagator. By the same argument, the Taylor expansions for the propagator derivatives:

$$\begin{aligned}
\frac{\partial}{\partial J}(\exp\{-i\hat{\hat{L}}\Delta t\}) &= \sum_{n=0}^{\infty}\frac{(-i\Delta t)^n}{n!}\sum_{k=1}^{n}\hat{\hat{L}}^{k-1}\hat{\hat{L}}'_J\hat{\hat{L}}^{n-k} \\
\frac{\partial^2}{\partial J\partial K}(\exp\{-i\hat{\hat{L}}\Delta t\}) &= \sum_{n=0}^{\infty}\frac{(-i\Delta t)^n}{n!} \\
&\quad \times \sum_{k=1}^{n}\left(\sum_{m=1}^{k-1}\hat{\hat{L}}^{m-1}\hat{\hat{L}}'_K\hat{\hat{L}}^{k-1-m}\hat{\hat{L}}'_J\hat{\hat{L}}^{n-k}\right. \\
&\quad \left. + \sum_{m=1}^{n-k}\hat{\hat{L}}^{k-1}\hat{\hat{L}}'_J\hat{\hat{L}}^{m-1}\hat{\hat{L}}'_K\hat{\hat{L}}^{n-k-m}\right)
\end{aligned} \tag{17}$$

are going to converge very rapidly. Efficient computation of first derivatives in Eq. (17) is made possible by the fact that all the relevant powers of $i\hat{\hat{L}}\Delta t$ were obtained when the main propagator was computed and the fact that $\hat{\hat{L}}'_J$ would in general have only a handful of non-zeroes. The second derivative is, of course, significantly more expensive, but in practical situations the exact second derivatives are rarely necessary as a Hessian of sufficient quality can be generated and updated by the many varieties of the BFGS formula [46].

The relative ease with which the gradient and Hessian of the simulation can be obtained owes much to the fast convergence of the Taylor series expansion for the propagator and its derivatives. This would not be the case outside the restricted state space, since the high coherences present in the full state space would lead to the emergence of large eigenvalues in $i\hat{\hat{L}}\Delta t$ and therefore convergence problems.

With Eqs. (16) and (17) implemented efficiently, the parametric derivatives of the density matrix may be co-propagated alongside the main simulation (an example

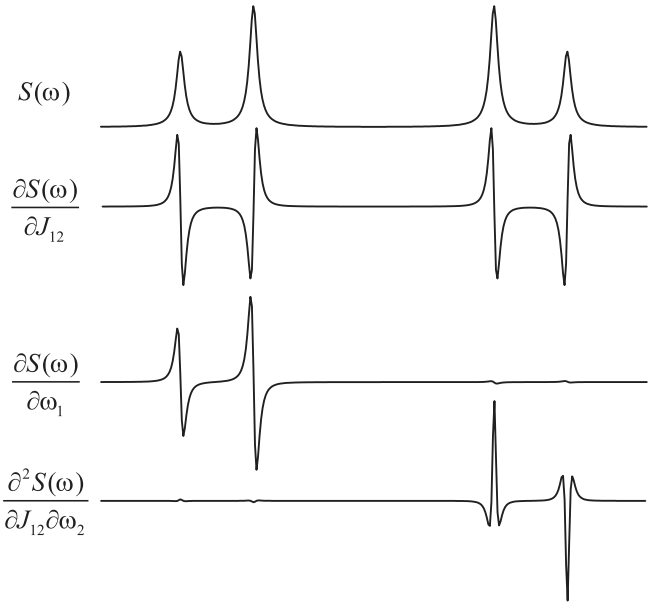

Fig. 6. An example of parametric spectrum derivatives computed for the 1D spectrum of a strongly coupled two-spin system using Eq. (16).

is given in Fig. 6). The computational expense of calculating simulation derivatives with respect to k parameters is smaller than k times the main simulation, because the intermediate results (e.g. Liouvillian powers and the main propagator) are re-used between the main and the derivative simulation. With compound pulse sequences, so long as the pulses are hard, no special care needs to be taken, because hard pulse propagator derivatives with respect to system parameters are zero. The propagators of soft pulses need to be differentiated in the way described above.

8. Relaxation and chemical kinetics

The cluster expansion described above was based on the topology of coherent couplings. A significant chunk of the relaxation superoperator (e.g. Zeeman tensor anisotropy, intra-cluster dipole and quadrupole interactions and their cross-correlations) will be accounted for by the relaxation superoperator obtained in the same way as was done for the global Liouvillian: by blending the cluster relaxation superoperators, which can be obtained exactly. However, there will be contributions beyond these, due to inter-cluster relaxation processes and chemical kinetics.

The inter-cluster relaxation processes of practical importance are dipolar cross-relaxation and DD-CSA (dipole anisotropy to chemical shift anisotropy) cross-correlation. While these processes would, in general, be non-secular between all longitudinal spin orders in the system, in practice they can be adequately described by including only the $\hat{I}_Z \leftrightarrow \hat{S}_Z$ and $\hat{I}_Z \leftrightarrow 2\hat{I}_Z\hat{S}_Z$ processes [47], that is, low spin orders are again sufficient. These relaxation processes would be accounted for if extra clusters are added for the cross-relaxing spins. The same procedure will also account for coherence transfer as a result of chemical transformations.

9. Conclusions and outlook

We propose a polynomially scaling algorithm for spin dynamics simulation and have illustrated its performance with NMR simulations of systems containing hundreds of coupled spins. The approximation involves restricting the state space to low spin-orders linking nearby spins and takes care to observe a number of restrictions imposed by the algebraic nature of the propagator group. A somewhat surprising find is that most spin states in large spin systems are not essential in magnetic resonance simulations and can safely be dropped from the state space. In cases of favourable interaction topologies (sparse graphs, e.g. in protein NMR) the asymptotic scaling of the simulation complexity is linear, opening the way to direct fitting of molecular structure to the experimental spectra.

We have identified the following issues as open for exploration. Firstly, the computational efficiency of the clustering algorithm depends on the coupling graph being sparse. This raises a question of what gains, if any, can be obtained in the cases where the coupling graph does not have this property, e.g. in solid state NMR and in EPR spectroscopy. Another important question is why, even though the state-space dimension has been reduced dramatically, the propagators are still sparse and whether there might be a possibility of further reducing the dimension of the system orbit (which is actually one-dimensional itself), perhaps through basis set transformation. The third open question is that of cluster isomorphism: there is just one type of 2-cluster, two types of 3-clusters etc, which can be pre-computed (or even pre-compiled) as building blocks, removing the need for intra-cluster treatment every time the simulation is run. Finally, the natural application of the algorithm described in this paper would be the (hitherto impossible) direct fitting of structural models to experimental spectra, possibly improving upon the molecular dynamics stage that is currently ubiquitous in NMR structure determination.

Acknowledgments

We are grateful to Dr Christopher Rodgers for enlightening discussions and the Oxford Supercomputing Centre for the generous allocation of CPU time. I.K. is a Fellow by Examination at Magdalen College, Oxford.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.jmr.2007.09.014.

References

[1] S.A. Smith, T.O. Levante, B.H. Meier, R.R. Ernst, Computer simulations in magnetic resonance. An object-oriented programming approach, J. Magn. Reson. 106 (1994) 75–105.
[2] P. Hodgkinson, L. Emsley, Numerical simulation of solid-state NMR experiments, Progr. NMR Spec. 36 (2000) 201–239.
[3] M. Mehring, V.A. Weberruss, Object-oriented Magnetic Resonance: Classes and Objects, Calculations and Computations, Academic Press, San Diego, CA, 2001.
[4] M. Eden, Computer simulations in solid-state NMR. I. Spin dynamics theory, Conc. Magn. Reson. A 17A (2003) 117–154.
[5] A. Kreiter, J. Huettermann, Simultaneous EPR and ENDOR powder-spectra synthesis by direct Hamiltonian diagonalization, J. Magn. Reson. 93 (1991) 12–26.
[6] T. Charpentier, C. Fermon, J. Virlet, Efficient time propagation technique for MAS NMR simulation: application to quadrupolar nuclei, J. Magn. Reson. 132 (1998) 181–190.
[7] A.A. Nevzorov, J.H. Freed, Dipolar relaxation in a many-body system of spins of 1/2, J. Chem. Phys. 112 (2000) 1425–1443.
[8] A.A. Nevzorov, J.H. Freed, Direct-product formalism for calculating magnetic resonance signals in many-body systems of interacting spins, J. Chem. Phys. 115 (2001) 2401–2415.
[9] A.G. Redfield, The Theory of Relaxation Processes, in: J.S. Waugh (Ed.), Advances in Magnetic Resonance, Academic Press, 1965, pp. 1–30.
[10] M. Goldman, Formal theory of spin–lattice relaxation, J. Magn. Reson. 149 (2001) 160–187.
[11] M. Hohwy, N.C. Nielsen, Systematic design and evaluation of multiple-pulse experiments in nuclear magnetic resonance spectroscopy using a

semi-continuous Baker–Campbell–Hausdorff expansion, J. Chem. Phys. 109 (1998) 3780–3791.
[12] T.S. Untidt, N.C. Nielsen, Closed solution to the Baker–Campbell–Hausdorff problem: Exact effective Hamiltonian theory for analysis of nuclear-magnetic-resonance experiments, Phys. Rev. E 65 (2002) 021108/1–021108/17.
[13] P. Hodgkinson, D. Sakellariou, L. Emsley, Simulation of extended periodic systems of nuclear spins, Chem. Phys. Lett. 326 (2000) 515–522.
[14] M. Hohwy, H. Bildsoe, H.J. Jakobsen, N.C. Nielsen, Efficient spectral simulations in NMR of rotating solids. The g-COMPUTE algorithm, J. Magn. Reson. 136 (1999) 6–14.
[15] J.I. Musher, Magnetic equivalence of nuclear spins, J. Chem. Phys. 47 (1967) 5460–5461.
[16] J.I. Musher, Equivalence of nuclear spins, J. Chem. Phys. 46 (1967) 1537–1538.
[17] S.M. Nokhrin, J.A. Weil, D.F. Howarth, Magnetic resonance in systems with equivalent spin-1/2 nuclides. Part 1, J. Magn. Reson. 174 (2005) 209–218.
[18] A. Jerschow, MathNMR: spin and spatial tensor manipulations in Mathematica, J. Magn. Reson. 176 (2005) 7–14.
[19] I. Kuprov, N. Wagner-Rundell, P.J. Hore, Bloch–Redfield–Wangsness theory engine implementation using symbolic processing software, J. Magn. Reson. 184 (2007) 196–206.
[20] 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, J. Magn. Reson. 145 (2000) 262–275.
[21] T.E. Skinner, T.O. Reiss, B. Luy, N. Khaneja, S.J. Glaser, Application of optimal control theory to the design of broadband excitation pulses for high-resolution NMR, J. Magn. Reson. 163 (2003) 8–15.
[22] W. Kohn, Nobel lecture: electronic structure of matter-wave functions and density functionals, Rev. Mod. Phys. 71 (1999) 1253–1266.
[23] R.R. Ernst, G. Bodenhausen, A. Wokaun, Principles of Nuclear Magnetic Resonance in One and Two Dimensions, Clarendon Press, Oxford, 1987.
[24] M. Carravetta, O.G. Johannessen, M.H. Levitt, Beyond the T1 limit: singlet nuclear spin states in low magnetic fields, Phys. Rev. Lett. 92 (2004) 153003/1–153003/4.
[25] M. Carravetta, M.H. Levitt, Long-lived nuclear spin states in high-field solution NMR, J. Am. Chem. Soc. 126 (2004) 6228–6229.
[26] M. Carravetta, M.H. Levitt, Theory of long-lived nuclear spin states in solution nuclear magnetic resonance. I. Singlet states in low magnetic field, J. Chem. Phys. 122 (2005) 214505/1–214505/14.
[27] G. Pileio, M. Concistre, M. Carravetta, M.H. Levitt, Long-lived nuclear spin states in the solution NMR of four-spin systems, J. Magn. Reson. 182 (2006) 353–357.
[28] O.W. Soerensen, G.W. Eich, M.H. Levitt, G. Bodenhausen, R.R. Ernst, Product operator formalism for the description of NMR pulse experiments, Progr. NMR Spec. 16 (1983) 163–192.
[29] S. Lee, W. Richter, S. Vathyam, W.S. Warren, Quantum treatment of the effects of dipole–dipole interactions in liquid nuclear magnetic resonance, J. Chem. Phys. 105 (1996) 874–900.
[30] J.H. Kristensen, H. Bildsoe, H.J. Jakobsen, N.C. Nielsen, Application of Lie algebra to NMR spectroscopy, Progr. NMR Spec. 34 (1999) 1–69.
[31] R.J. Ober, E.S. Ward, Correcting for phase distortion of NMR spectra analyzed using singular-value decomposition of Hankel matrixes, J. Magn. Reson. 114 (1995) 120–123.
[32] P. Koehl, Linear prediction spectral analysis of NMR data, Progr. NMR Spec. 34 (1999) 257–299.
[33] J.R. Woodward, C.R. Timmel, K.A. McLauchlan, P.J. Hore, Radio frequency magnetic field effects on electron-hole recombination, Phys. Rev. Lett. 87 (2001) 077602/1–077602/4.
[34] C.T. Rodgers, K.B. Henbest, P. Kukura, C.R. Timmel, P.J. Hore, Low-field optically detected EPR spectroscopy of transient photoinduced radical pairs, J. Phys. Chem. A 109 (2005) 5035–5041.
[35] C.T. Rodgers, S.A. Norman, K.B. Henbest, C.R. Timmel, P.J. Hore, Determination of radical re-encounter probability distributions from magnetic field effects on reaction yields, J. Am. Chem. Soc. 129 (2007) 6746–6755.
[36] D.J. Klein, D. Babić, N. Trinajstić, Enumeration in Chemistry, Chemical Modelling: Applications and Theory, Royal Society of Chemistry, Cambridge, 2002.
[37] W.G. Yan, Y.N. Yeh, Enumeration of subtrees of trees, Theor. Comp. Sci. 369 (2006) 256–268.
[38] P. Bonzini, L. Pozzi, Polynomial-Time Subgraph Enumeration for Automated Instruction Set Extension, Technical Report 1 2006/07, University of Lugano, Switzerland (2006).
[39] G. Ruecker, C. Ruecker, Automatic enumeration of all connected subgraphs, Comm. Math. Comput. Chem. 41 (2000) 145–149.
[40] A.V. Buevich, N.M. Sergeyev, V.A. Chertkov, Poster W144: Proceedings of the 12th EENC (Oulu, Finland) (1994).
[41] N.J. Higham, The scaling and squaring method for the matrix exponential revisited, SIAM J. Matr. Anal. Appl. 26 (2005) 1179–1193.
[42] M. Hochbruck, C. Lubich, On Krylov subspace approximations to the matrix exponential operator, SIAM J. Numer. Anal. 34 (1997) 1911–1925.
[43] B. Lopez-Mendez, P. Guentert, Automated protein structure determination from NMR spectra, J. Am. Chem. Soc. 128 (2006) 13112–13122.
[44] B. Richter, J. Gsponer, P. Varnai, X. Salvatella, M. Vendruscolo, The MUMO (minimal under-restraining minimal over-restraining) method for the determination of native state ensembles of proteins, J. Biomol. NMR 37 (2007) 117–135.
[45] D.M. Gay, Subroutines for unclonstrained minimization using a model/trust-region approach, ACM Trans. Math. Soft. 9 (1983) 503–524.
[46] M. Al-Baali, Improved Hessian approximations for the limited memory BFGS method, Numer. Alg. 22 (1999) 99–112.
[47] M. Goldman, Interference effects in the relaxation of a pair of unlike spin-1/2 nuclei, J. Magn. Reson. 60 (1984) 437–452.
[48] U. Till, C.R. Timmel, B. Brocklehurst, P.J. Hore, The influence of very small magnetic fields on radical recombination reactions in the limit of slow recombination, Chem. Phys. Lett. 298 (1998) 7–14.