

Dynamical effects in *ab initio* NMR calculations: Classical force fields fitted to quantum forces

Mark Robinson^{1,a)} and Peter D. Haynes²

¹*Theory of Condensed Matter, Cavendish Laboratory, J.J. Thomson Avenue, Cambridge CB3 0HE, United Kingdom*

²*Departments of Physics and Materials, Imperial College London, Exhibition Road, London SW7 2AZ, United Kingdom*

(Received 22 April 2010; accepted 12 July 2010; published online 31 August 2010)

NMR chemical shifts for an L-alanine molecular crystal are calculated using *ab initio* plane wave density functional theory. Dynamical effects including anharmonicity may be included by averaging chemical shifts over an ensemble of structural configurations generated using molecular dynamics (MD). The time scales required mean that *ab initio* MD is prohibitively expensive. Yet the sensitivity of chemical shifts to structural details requires that the methodologies for performing MD and calculating NMR shifts be consistent. This work resolves these previously competing requirements by fitting classical force fields to reproduce *ab initio* forces. This methodology is first validated by reproducing the averaged chemical shifts found using *ab initio* molecular dynamics. Study of a supercell of L-alanine demonstrates that finite size effects can be significant when accounting for dynamics. © 2010 American Institute of Physics. [doi:10.1063/1.3474573]

I. INTRODUCTION

Nuclear magnetic resonance (NMR) is a powerful tool for studying structure and dynamics on an atomic scale. It is now widely applied in the physical, chemical, and biological sciences. First principles quantum mechanical calculations of NMR parameters provide a link between the experimental spectra and the underlying atomic structure. Traditional quantum chemical approaches have shown great success assigning solution-state NMR spectra.¹ However their application to the solid-state has been restricted to modeling clusters which has been fraught with difficulties. It is only in the past decade that methodology exploiting the periodicity of crystals to calculate chemical shifts on truly extended systems has been developed.^{2–4} The method used in this work, known as the gauge including projector augmented wave method, is based on the density functional theory (DFT) plane wave pseudopotential framework.⁵ It deals with both the problem of the electron response to an externally applied magnetic field² and the need to use pseudopotentials to reduce the size of the plane wave basis set.^{6,7} Over the past few years these techniques have been successfully applied to a variety of problems.^{8–10}

However in some systems discrepancies between experimental and calculated shifts are still observed.¹¹ Dumez and Pickard¹² have shown that the dynamics of a system can be a contributing factor in organic solids, while De Gortari *et al.*¹³ have recently seen a similar effect for a peptide in the solid state. Calculations of the chemical shifts are particularly sensitive to dynamics as the shifts are nonlinear and anharmonic effects are significant. Since a NMR experiment studies a dynamic system over a period of milliseconds, it is therefore essential to account for the motion in a calculation. This has

been widely studied in the gas phase and in solution;^{14–16} however methods in the solid-state are not well developed. In the past dynamics has been accounted for by averaging over an ensemble of configurations. These can be generated using methods such as vibrational averaging or molecular dynamics (MD). While MD has the advantage of including anharmonic effects, the nuclei are governed by classical equations of motion and so any quantum dynamical effects are neglected. A comparison of the contributions from the two effects is given by Dumez and Pickard. Though protons still behave quantum mechanically at room temperature, the harmonic motion of the carbon atoms is well approximated classically, and the influence of anharmonicity was found to be far greater. In this study we therefore focus on the MD method.

In the work of Dumez and Pickard the MD trajectory from which configurations were sampled was generated using DFT. However DFT MD is computationally very expensive. For systems that exhibit dynamics on long time scales this can become prohibitive as configurations must be uncorrelated. We overcome this by using classical MD to generate configurations, while still using DFT to calculate the chemical shifts. This can provide a significant speed up since a DFT calculation is now only performed on each configuration, rather than at every time step in the MD simulation. However, this introduces a new problem. The atomic coordinates obtained with an off-the-shelf force field such as the GAFF (Ref. 17) or AMBER (Ref. 18) force fields can differ from DFT by over 0.1 Å. Thus the region of configuration space sampled is different from that sampled by DFT MD. This can cause substantial errors in the chemical shieldings, far larger than any improvement gained from time averaging. The force field must therefore be parametrized so that it replicates the structural properties found using DFT for the particular system under study.

^{a)}Electronic mail: mr419@cam.ac.uk.

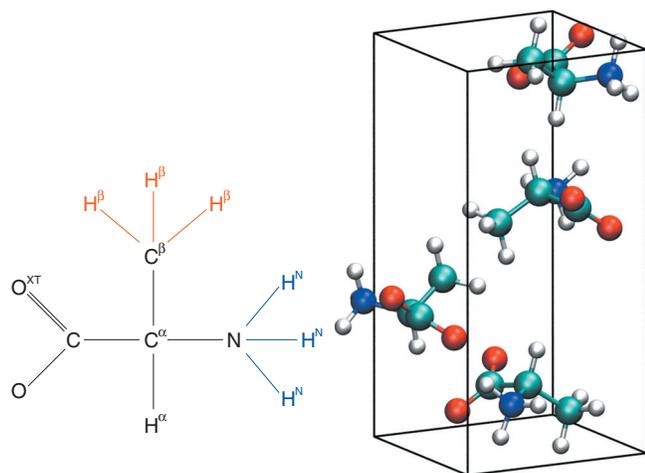


FIG. 1. The chemical structure of an L-alanine amino acid and a depiction of the L-alanine unit cell are shown. Chemical shieldings on the H^β and H^N atoms are averaged to compensate for the failure of MD to accurately describe their rotations.

In this work we follow two strategies to avoid correlation between configurations generated using MD. First we use Langevin dynamics implemented within the DFT code CASTEP (Ref. 19) to efficiently sample configuration space and converge the chemical shifts for the molecular crystal L-alanine. However the computational expense of DFT means that this does not scale well to larger systems. We therefore develop the methodology required to fit a force field to DFT calculated forces. The parametrized force field is used with the classical MD suite AMBER (Ref. 20) to reproduce the DFT MD result. A larger supercell of L-alanine is then studied to determine if finite size effects can also affect the chemical shifts when including dynamics.

II. ENSEMBLE AVERAGING

For comparison we study the same system as Dumez and Pickard, a molecular crystal of L-alanine. The chemical structure of the L-alanine amino acid and a representation of its crystal structure are shown in Fig. 1. The primitive unit cell contains four crystallographically equivalent alanine molecules.²¹ Under magic angle spinning conditions these four molecules are equivalent and their chemical shieldings can be averaged. In addition hydrogens on the methyl and amine groups are averaged to be consistent with experiment, where a single peak is observed. This is because the hydrogens jump between equivalent sites. Since this motion occurs on a time scale of a few hundred picoseconds,²² it is not accessed by our MD simulations. We therefore perform the averaging by hand.

DFT calculations were performed using the plane wave pseudopotential code CASTEP.¹⁹ A cutoff-energy of 550 eV was used to construct the plane-wave basis, with ultrasoft pseudopotentials²³ and the Perdew–Burke–Ernzerhof (PBE) exchange-correlation functional.²⁴ k -point sampling was performed using a Monkhorst–Pack grid²⁵ of $2 \times 1 \times 2$ for MD and $4 \times 2 \times 4$ for the NMR calculations. MD simulations were performed at 293 K with a 1 fs time step. The PBE functional is known to describe dispersion forces inaccu-

rately and overestimate lattice parameters. All calculations were therefore performed using the lattice parameters found in a neutron scattering experiment performed at room temperature.²¹ The lattice parameters found in that experiment were $6.025 \times 12.324 \times 5.783 \text{ \AA}^3$.

A. Langevin dynamics

To sample configuration space we use Langevin dynamics^{26,27} implemented within the CASTEP code.²⁸ Langevin dynamics regulates the temperature by modifying Newton's second law with drag and stochastic terms representing interactions with a fictional solvent,

$$m\mathbf{a} = \mathbf{F} - \gamma m\mathbf{v} + \mathbf{R}. \quad (1)$$

γ is a collision frequency representing drag caused by the solvent and \mathbf{R} is a random force associated with the solvents thermal motion. $\langle \mathbf{R}^2 \rangle$ is related to γ by the fluctuation-dissipation theorem.²⁹ The fictitious solvent acts like a thermal bath such that the canonical ensemble is approximated. Once a trajectory has been generated configurations are sampled from it at a regular time interval. To ensure efficient sampling this interval should be greater than the correlation time of the chemical shifts. Due to the expense of DFT MD a short correlation time is therefore beneficial.

The collision frequency γ provides control over the correlation time. Larger collision frequencies reduce the correlation time at the expense of less stable temperature regulation. The CASTEP default is for $\gamma = 10 \text{ ps}^{-1}$, but this results in a long correlation time. With $\gamma = 100 \text{ ps}^{-1}$ we find that the velocity is uncorrelated after around 48 fs. At this γ the temperature is still well controlled with a standard deviation of $30.4 \text{ }^\circ\text{C}$ over the trajectory. Using a block averaging algorithm²⁷ we find that the chemical shieldings have uncorrelated after around 8 fs. A shorter correlation time for the chemical shieldings than the velocity is expected due to their sensitivity to structure. In the following work we sample configurations every 16 fs.

Once an ensemble of N uncorrelated configurations has been obtained, DFT calculations are carried out to calculate the chemical shieldings. The shieldings for each configuration are then averaged to find the mean and standard deviation of the mean. Isotropic chemical shifts are related to the shielding by a reference shielding σ_{ref} ,

$$\delta = \sigma_{\text{ref}} - \sigma. \quad (2)$$

Using this equation the experimental chemical shifts are fitted to the computed chemical shieldings to obtain the reference shielding. This is a standard procedure to avoid any systematic error introduced from performing a NMR calculation on the reference molecule.¹¹ The computed chemical shifts can then be found and compared to experiment. Like Dumez and Pickard, we found systematic variations between methods. Thus, in order to capture the experimentally relevant relative variations, a different reference shielding has been found for each method.

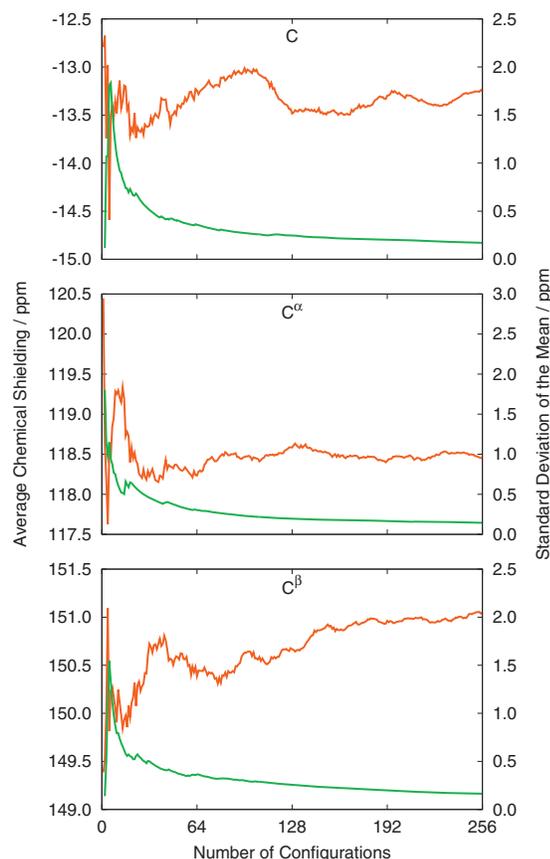


FIG. 2. Convergence of the chemical shieldings for carbon atoms in an L-alanine molecular crystal. The (red) mean chemical shielding and (green) standard deviation of the mean are shown. Configurations were generated using PBE DFT.

B. Results and discussion

Figure 2 shows the convergence of the averaged chemical shieldings for the carbon atoms as the number of configurations is increased. Until around 192 configurations there is still significant variation in the shieldings. Chemical shieldings for hydrogen atoms converge far quicker than those for the carbons. This is due to the more complicated dynamics exhibited by the carbon atoms. After 256 configurations the errors have become reasonably small. Carbons are converged to better than 0.2 ppm, and hydrogens to 0.02 ppm. After around 16 configurations the standard deviations decrease smoothly with $1/\sqrt{N}$, indicating well behaved convergence.

The reference shieldings were found to be $\sigma_{\text{ref}}=167.92$ for carbons and $\sigma_{\text{ref}}=30.11$ for hydrogens. Using these values the chemical shifts were found and are compared to experiment in Table I. The result for a fully geometry optimized structure is also shown. It is clear that including the dynamics has made a significant difference over the static structure. The C shift improves by over 1 ppm, and the rms chemical shift for all the carbons improves by 0.74 ppm. These changes are far greater than the errors due to the averaging. The hydrogen shifts are not affected by dynamics to the same extent due to the simplicity of their dynamics.

These results verify the importance of including dynamical effects in NMR calculations on organic solids. By using Langevin dynamics and careful adjustment of the collision

TABLE I. Experimental chemical shifts for carbon (Ref. 30) and hydrogen (Refs. 31 and 32) atoms are compared to those calculated using (static) a geometry optimized structure, and by (averaged) averaging over 256 configurations generated with DFT MD.

	Expt.	Static	Averaged	
C	176.8	+5.38	+4.37	± 0.17
C $^{\alpha}$	50.9	-1.71	-1.43	± 0.14
C $^{\beta}$	19.8	-3.68	-2.93	± 0.16
H $^{\text{N}}$	8.4	+0.55	+0.48	± 0.03
H $^{\alpha}$	3.6	-0.14	-0.13	± 0.02
H $^{\beta}$	1.2	-0.40	-0.35	± 0.02
rmsd C		3.89	3.15	
rmsd H		0.40	0.35	

frequency γ , we have improved the sampling efficiency of the MD trajectory. This avoided the problem of correlation between configurations and meant that the corresponding correction to the mean found in previous work was not required. However even using the extensive computing resources available to the authors, these calculations are at the limit of what is currently practical. Even doubling the L-alanine system size to check for any finite size errors would be prohibitive using DFT MD. Larger biological systems requiring nanosecond long trajectories are completely out of the question. To overcome this restriction we turn to classical MD to generate configurations.

III. FORCE FIELD FITTING

Classical MD was performed using the AMBER suite of programs.²⁰ Input files were generated using the ANTECHAMBER and LEAP programs. MD simulations were performed using SANDER at 293 K with a 1 fs time step and $\gamma=1 \text{ ps}^{-1}$. Configurations were sampled at 5 ps time intervals. Our force field parametrizing code was implemented within the SANDER code.

A. The force field

The force field used in this work is based on the AMBER force field.¹⁸ This is a simple two-body additive model aimed at accurately modeling conformational energies and intermolecular interactions in organic and biological systems. Our force field contains bond, angle, dihedral, electrostatic, and a Pauli-exclusionlike R^{-12} term,

$$\begin{aligned}
 E = & \sum_{\text{bonds}} K_r (r - r_{\text{eq}})^2 + \sum_{\text{angles}} K_{\theta} (\theta - \theta_{\text{eq}})^2 \\
 & + \sum_{\text{dihedrals}} \frac{V_n}{2} [1 + \cos(n\phi - \gamma)] + \sum_{i < j} \left[\frac{A_{ij}}{R_{ij}^{12}} + \frac{q_i q_j}{\epsilon R_{ij}} \right].
 \end{aligned} \quad (3)$$

Bonds and angles are treated as springs with force constants K_r and K_{θ} , and equilibrium lengths r_{eq} and θ_{eq} . Dihedrals are represented by a symmetry parameter n , a corresponding force constant V_n , and a phase γ . The R^{-12} term is described by the A_{ij} parameters. The nonbonded interactions are only evaluated for atoms separated by at least three bonds. For

atoms separated by exactly three bonds the terms are included but scaled. The PBE functional we are using is known not to capture dispersion effects. We have therefore not included the R^{-6} dispersion term found in the AMBER force field.

The number of parameters describing a system is determined by the number of atoms, bond types, angle types, and dihedrals. The R^{-12} term is difficult to fit as it is so short ranged. However this should also mean that the exact values of the parameters will have a minimal effect at room temperature. For simplicity we therefore retain the AMBER values for the A_{ij} parameters. This leaves a total of 94 parameters requiring optimization for the L-alanine system.

B. Newton's method of optimization

Our fitting procedure builds upon that found within the literature for parametrizing force fields for silica.³³ We minimize the function

$$Q(\alpha) = |\mathbf{F}^{\text{QM}} - \mathbf{F}^{\text{FF}}(\alpha)|^2, \quad (4)$$

where α is a vector containing all the force field parameters $\{\alpha_j\}$. Its length is the number of parameters in the force field, P . The three components of the DFT and classical forces on each atom are stored in the vectors \mathbf{F}^{QM} and \mathbf{F}^{FF} , respectively. They are therefore each of size $3NM$, where N is the number of atoms in the system and M is the number of configurations. Note the distinction between vectors of length P that are denoted with an underline and vectors of length $3NM$ that are in bold. The function Q is just the sum of the squared differences between DFT and classical forces. The classical forces and therefore Q are a function of the parameter set α that we seek to optimize.

Q is easily differentiated to find its first and second derivatives with respect to the parameters,

$$\frac{\partial Q}{\partial \alpha_n} = -2(\mathbf{F}^{\text{QM}} - \mathbf{F}^{\text{FF}}) \cdot \frac{\partial \mathbf{F}^{\text{FF}}}{\partial \alpha_n}, \quad (5)$$

$$\frac{\partial^2 Q}{\partial \alpha_n \partial \alpha_m} = -2(\mathbf{F}^{\text{QM}} - \mathbf{F}^{\text{FF}}) \cdot \frac{\partial^2 \mathbf{F}^{\text{FF}}}{\partial \alpha_n \partial \alpha_m} + 2 \frac{\partial \mathbf{F}^{\text{FF}}}{\partial \alpha_n} \cdot \frac{\partial \mathbf{F}^{\text{FF}}}{\partial \alpha_m}. \quad (6)$$

The simplicity of Eq. (3) means that these derivatives are trivial to find analytically. Hence the gradient $\underline{g} = \nabla Q$ and Hessian $\underline{\underline{H}} = \nabla^2 Q$ are easily constructed. This makes Newton's method an obvious technique to minimize Q .

Thus for a given parameter set, the gradient and Hessian are constructed, and the required change in parameters $\Delta \alpha$ found by solving

$$\underline{\underline{H}} \cdot \Delta \alpha = -\underline{g}. \quad (7)$$

In practice Q is not quadratic in the parameters due to the charge term and the collinearity of the force constants and equilibrium bond lengths and angles. Close to the minimum however Q is well approximated by a quadratic. In these regions we repeatedly find the gradient and Hessian, calculate the change in parameters, and take a step. As we approach the minimum and the quadratic terms dominate, we

enter superlinear convergence and rapidly move to the minimum. In regions of parameter space where the quartic terms are dominant, Q can have local minima and maxima that can confuse Newton's method. Here we use our knowledge of Q and its derivatives to fit a quartic to the search direction from which the stationary points can be found analytically. It is then trivial to evaluate Q at each stationary point and move to the lowest value of Q .

A measure of how good a fit has been obtained can be found by normalizing Q ,

$$\sqrt{\frac{Q}{N}} = \frac{|\mathbf{F}^{\text{QM}} - \mathbf{F}^{\text{FF}}|}{|\mathbf{F}^{\text{QM}}|}. \quad (8)$$

This can be used to evaluate improvements in the quality of our fitted force field.

C. Parameter constraints

Using the methodology laid out so far we have observed that for some initial parameter sets, nonphysical bond and angle force constants can result. In these cases Newton's method converges to a point where a subset of the force constants are zero. Consider the example of a one-dimensional system comprised of a single bond. The derivative of the spring term with respect to the equilibrium length,

$$\frac{\partial F^{\text{FF}}}{\partial r_{\text{eq}}} = 2K_r \quad (9)$$

is zero when $K_r=0$. We are not finding the solution where $\mathbf{F}^{\text{QM}} = \mathbf{F}^{\text{FF}}$, but a point where $\partial F^{\text{FF}} / \partial \alpha_n = 0$. Further information on this point can be found by constructing the Hessian for the spring using Eq. (6). When $K_r=0$ we find that the determinant of the Hessian is

$$\det \underline{\underline{H}} = -4(F^{\text{QM}})^2. \quad (10)$$

Since the DFT forces are calculated on configurations generated by an inconsistent force field, their average is not necessarily zero. The determinant is therefore less than zero and the stationary point a saddle point. This is expected since though $\partial F^{\text{FF}} / \partial r_{\text{eq}}$ is zero when $K_r=0$, $\partial F^{\text{FF}} / \partial K_r$ need not be.

To avoid these saddle points we write constrained versions of Eqs. (5) and (6),

$$\frac{\partial \tilde{Q}}{\partial \alpha_n} = -2(\mathbf{F}^{\text{QM}} - \mathbf{F}^{\text{FF}}) \cdot \left. \frac{\partial \mathbf{F}^{\text{FF}}}{\partial \alpha_n} \right|_{\alpha_{\text{ref}}}, \quad (11)$$

$$\frac{\partial^2 \tilde{Q}}{\partial \alpha_n \partial \alpha_m} = + 2 \left. \frac{\partial \mathbf{F}^{\text{FF}}}{\partial \alpha_n} \right|_{\alpha_{\text{ref}}} \cdot \frac{\partial \mathbf{F}^{\text{FF}}}{\partial \alpha_m}, \quad (12)$$

where \tilde{Q} is the constrained Q . Here we have replaced the true force derivatives with reference values fixed to their evaluation at the initial parameter set. The only point at which Eq. (11) is zero is now when $\mathbf{F}^{\text{QM}} = \mathbf{F}^{\text{FF}}$. Although the symmetry of the Hessian is now broken, it is weak enough that Eq. (7) still performs well.

We might expect this constraint to impede convergence. Consider the case where the force is linear in the parameter α ,

$$\mathbf{F}^{\text{FF}} \propto \alpha. \quad (13)$$

Integrating with respect to α gives

$$\int \mathbf{F}^{\text{FF}} d\alpha \propto \frac{1}{2} \alpha \mathbf{F}^{\text{FF}}. \quad (14)$$

Thus integrating the constrained derivative given in Eq. (11) and ignoring any constant term lead to an expression for \tilde{Q} ,

$$\tilde{Q} = - (2\mathbf{F}^{\text{QM}} - \mathbf{F}^{\text{FF}}) \cdot \left[\frac{\partial \mathbf{F}^{\text{FF}}}{\partial \alpha} \Big|_{\alpha_{\text{ref}}} \alpha \right]. \quad (15)$$

But expanding the brackets in Eq. (4) and removing the constant $[\mathbf{F}^{\text{QM}}]^2$ term,

$$Q = - (2\mathbf{F}^{\text{QM}} - \mathbf{F}^{\text{FF}}) \cdot [\mathbf{F}^{\text{FF}}]. \quad (16)$$

We see that when the force is linear in α , Q and \tilde{Q} are the same and convergence should not be affected. Of course, as we have seen above, the force is not linear in the parameters. However it appears to be good enough that to first order, Q and \tilde{Q} are the same for reasonably fast convergence. The approximation improves as we approach the minimum.

In our experience it is only the bond and angle terms that suffer from this problem, although in theory the same problem can afflict the charges. Optimizing charges using the constrained derivatives is also subject to a larger performance penalty. To get the best of both worlds we follow a slightly less elegant approach. We first optimize only the bond and angle parameters using the constrained derivatives. Then optimize the charges using the unconstrained derivatives. Finally, once the parameters are fairly close to the global minima, all parameters may be optimized using the unconstrained derivatives without fear of finding a nonphysical stationary point. This procedure has proved to be fast and stable. The 94 parameters describing the 52 atom L-alanine unit cell with 256 configurations converges to exactly the same minimum of Q in around 20 steps from a variety of different initial parameter sets.

IV. PARAMETRIZATION OF L-ALANINE

Using the methodology developed above we have parametrized our force field for L-alanine using the GAFF (Ref. 17) parameters as our starting point. We found that 256 configurations each containing 52 force vectors were enough to converge the 94 fitted force field parameters. This results in a converged value for $\sqrt{Q/N}$ of 0.16, close to that found for similarly fitted force fields in the literature.³³

A. Contribution by configuration

The parameters included in the optimization will have an effect on the quality of fit. This can be seen by plotting histograms where configurations have been binned by their contribution to Q . Four of these plots are compared in Fig. 3. The lower a configuration's contribution to Q , the better it is represented by the force field. These plots therefore provide a visual representation of how well the configuration space spanned by a system is represented in the force field. The lower the value of Q a peak is centered on, the better the

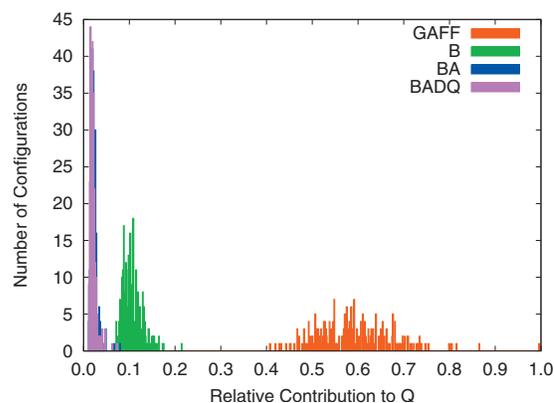


FIG. 3. Contributions to Q by configuration after fitting the force field via different parameters. Results are shown (GAFF) before fitting, and after fitting (b) bonds; bonds and angles (BA); bonds, angles, dihedrals, and charges (BADQ).

space is represented. The breadth of the peak indicates the range in quality of fit across the ensemble of configurations. A sharp peak implies that all configuration space is equally well represented. Our goal therefore is a sharp peak at low Q , where all configuration space is well represented.

As expected the quality of fit improves as parameters are added to the fit; the peak is sharper and at a lower value of Q . Bonds and angle parameters make the largest improvements to the fit. Dihedrals make very little difference, perhaps sharpening the peak slightly. This is unsurprising due to the small size of the alanine amino acid. Charges can have an effect on both bond lengths and angles. Although adding them to a fit that already includes bonds and angles makes a relatively small improvement, we have observed that fitting charges on their own does have a significant effect.

A similar analysis can be made by plotting the contribution to Q by atom. We find that it is the terminating atoms that are worst described by the force field. The two oxygen atoms are particularly poorly described by GAFF. However after fitting both bonds and angles they are as well represented as other atoms.

B. Convergence to DFT

An optimized force field should be self-consistent. That is, fitting to the DFT forces for an ensemble of configurations generated using an optimized force field should result in no change to the force field parameters. To obtain self-consistency we repeatedly apply the fitting procedure above. At each iteration a new generation of the force field is produced that matches DFT better than the previous. In Fig. 4 we visualize the convergence of the force field by plotting bond lengths, chemical shifts, and $\sqrt{Q/N}$ for each generation.

The value of $\sqrt{Q/N}$ after fitting is fairly constant, indicating a fundamental limit on the quality of fit that can be achieved using our force field. This can only be improved by adding more terms to the energy expression in Eq. (3) and fitting more parameters. The $\sqrt{Q/N}$ before fitting shows a fairly well behaved convergence toward the value after fitting. Hence at each generation the ensemble of configurations generated by the force field is closer to those that would be generated using DFT MD.

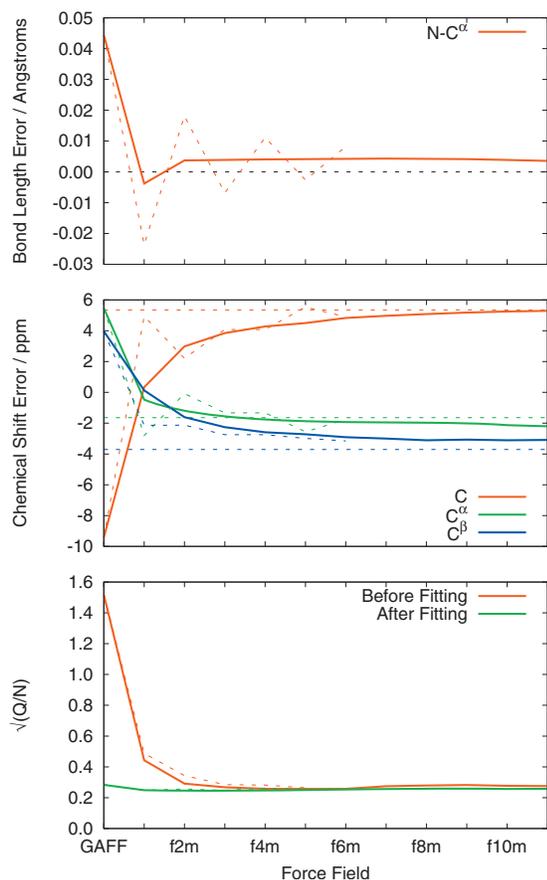


FIG. 4. The convergence of a force field toward self-consistency is shown with mixing (solid lines) and without mixing (dashed lines). At each iteration a set of configurations is generated and fitted to, creating a new generation of the force field. The bond length error is the difference in bond lengths of relaxed structures obtained with PBE DFT and the force field. Chemical shieldings were calculated on structures optimized using the respective force field and are shown relative to experiment. The dashed horizontal lines indicate the result using the PBE optimized structure. For generation n the mixed force fields have been labeled f_{nm} .

The bond length errors are given as the difference in bond lengths found by optimizing L-alanine using PBE and the force fields. Oscillations are apparent in many of the bond lengths including the N–C α bond shown in Fig. 4. This occurs because we fit to a set of configurations taken from a small localized region of configuration space, rather than an ensemble representing the entire configuration space. The effect of this is that the chemical shifts do not converge smoothly. This impedes convergence and makes it difficult to tell when the force field has converged.

We overcome this by linearly mixing each new generation with the previous. This is achieved by a mixing parameter m such that each parameter α_i becomes

$$\alpha_i = m\alpha_{i,n} + (1 - m)\alpha_{i,n-1}, \quad (17)$$

where $\alpha_{i,n}$ is the value of the parameter at the n th generation. We have verified that the bond lengths are linear in the parameters, and so this scheme works well. Empirically the oscillations in the bond length obey the recurrence relation

$$(y_n - y_\infty) = -\lambda(y_{n-1} - y_\infty), \quad (18)$$

from which a best guess at m can be found by fitting to the data in Fig. 4. Here y_n is the bond length error at generation n , y_∞ is the converged bond length error, and λ is a decay constant. The mixing parameter obtained varies depending on exactly which bonds are included in the fit. However most result in a mixing parameter between 0.6 and 0.7, and so for simplicity we choose a mixing parameter of $\frac{2}{3}$.

The convergence of the force field using this mixing parameter is also shown in Fig. 4. Oscillations in the bond lengths have clearly been damped. After a few generations even the C–N bond that exhibited the largest oscillations has converged. This has the effect that the chemical shifts converge far more smoothly than when no mixing is performed.

V. RESULTS AND DISCUSSION

Armed with a force field optimized for L-alanine under PBE DFT, we can begin the calculation of the chemical shieldings. The procedure is the same as in Sec. II, except now that the force field is used to generate the trajectory from which configurations are sampled. Chemical shieldings are calculated and averaged, and the reference shielding and chemical shifts are found by fitting to Eq. (2). The convergence of chemical shieldings calculated under each approach behave similarly, converging with $1/\sqrt{N}$. 256 configurations were found to provide nearly identical errors for both the DFT and classical methods.

A. Comparison with experiment

The averaged chemical shifts calculated are compared to experiment in Table II. As expected GAFF does relatively poorly. The differences in the forces produced by GAFF compared to PBE leads to sampling in a different region of configuration space, resulting in very inaccurate chemical shifts. GAFF performs even worse than the rms difference of 6.60 ppm suggests since the signs of the averaged shifts are opposite to those found using PBE. Remember that it is PBE we want to replicate using the force field, not the experimental result.

The poor performance of GAFF reinforces the need to parametrize a force field if structurally sensitive properties such as chemical shifts are to be found. In the work of De Gortari *et al.*¹³ chemical shifts of a peptide averaged over configurations generated using the OPLS force field³⁴ were found to perform better than DFT MD. This was attributed to the longer time scales spanned by classical MD. However, it could equally be due to fortuitous properties of the OPLS force field. By fitting to forces calculated within the same level of approximation as the chemical shifts, our approach reduces any error to that of a single methodology. This represents a significant advance in the controlled accuracy of dynamical effects.

By fitting the force field to PBE forces we are able to vastly improve the averaged chemical shifts. As well recovering the optimal geometry, the region of configuration space sampled is replicated, as shown schematically in Fig. 5. Using f11m the C shift is very good, within one standard de-

TABLE II. Comparison of carbon and hydrogen chemical shifts calculated by averaging over 256 configurations sampled from different MD trajectories. Results are shown for (PBE) DFT MD using the PBE exchange-correlation functional, (GAFF) using classical MD with the GAFF, and (f11m) using classical MD with a force field fitted to PBE DFT. Using the optimized f11m force field averaged shifts were found for both a $(1 \times 1 \times 1)$ unit cell of L-alanine and a $(2 \times 1 \times 2)$ supercell. Chemical shifts are given relative to the experimental result. The aim of the fitting is for the f11m $1 \times 1 \times 1$ result to replicate the PBE result.

	Expt.	DFT MD				Classical MD				
		Static	PBE		GAFF	f11m $1 \times 1 \times 1$	f11m $2 \times 1 \times 2$			
C	176.8	+5.38	+4.37	± 0.17	-9.25	± 0.14	+4.50	± 0.18	+4.94	± 0.13
C^α	50.9	-1.71	-1.43	± 0.14	+5.70	± 0.16	-1.90	± 0.15	-1.98	± 0.10
C^β	19.8	-3.68	-2.93	± 0.16	+3.55	± 0.16	-2.60	± 0.16	-2.97	± 0.11
H^N	8.4	+0.55	+0.48	± 0.03	+0.33	± 0.02	+0.72	± 0.02	+0.71	± 0.01
H^α	3.6	-0.14	-0.13	± 0.02	-0.14	± 0.02	-0.30	± 0.02	-0.30	± 0.02
H^β	1.2	-0.40	-0.35	± 0.02	-0.20	± 0.02	-0.41	± 0.02	-0.42	± 0.01
rmsd C		3.89	3.15		6.60		3.19		3.52	
rmsd H		0.40	0.35		0.24		0.51		0.51	

viation of the PBE result. Though the improvement is still significant, C^α and C^β shifts fare slightly worse. We notice that though these discrepancies are correlated with the errors seen in Fig. 4, they are smaller. The regions of configuration space sampled are similar enough that some of the errors in the minimum are averaged out. This is encouraging as it implies that a relatively simple force field such as that in Eq. (3) is capable of capturing the dynamics exhibited by L-alanine. Adding extra terms to increase the accuracy at the minimum would not necessarily provide a significant improvement in the averaged shifts.

Distributions of chemical shifts found using trajectories from GAFF and f11m are compared to DFT MD in Fig. 6. This provides a clear visual illustration of the improvement

achieved by parametrizing the force field. Where as the peak for the GAFF chemical shifts lies far from the PBE peak, the f11m peak is nearly identical in both position and breadth. This reflects the f11m force fields ability to replicate the region of configuration space sampled using DFT MD.

B. Finite size effects

So far we have only studied a single unit cell of L-alanine. However it is possible that there are dynamical effects operating on larger length scales that would fit within the unit cell. These are prohibited by the periodic boundary conditions imposed in the calculations above. To investigate we have used the f11m force field to generate configurations for a larger supercell of L-alanine. The unit cell of L-alanine is composed of four alanine molecules in a column (see Fig. 1). It is reasonable then to expect that any finite size effects will be most significant along the x and z axes, perpendicular to the column. We have studied a $2 \times 1 \times 2$ supercell, the results of which are shown in Table II.

Configurations for the supercell were generated using the same f11m force field used for the unit cell calculations. Since each configuration now contains four times as many alanine molecules the chemical shieldings converged faster. The $2 \times 1 \times 2$ results are averaged over 128 configurations resulting in slightly smaller errors than for the 256 configurations of the $1 \times 1 \times 1$ unit cell. We see that the carbon shifts change by as much as 0.44 ppm. These significant differences imply new dynamical effects not captured previously. Any new flexibility afforded by the larger supercell only affects the carbon shifts, the hydrogen shifts remain the same.

Using the f11m force field on the supercell is an extrapolation since it was fitted to dynamics on the unit cell. However the accuracy of this approximation can be tested. Refitting the f11m force field to the 128 configurations used in the NMR calculation above results only in small changes to the parameters, nothing larger than the changes seen between the f10m and f11m force fields. Thus the f11m force field is as good at replicating forces (and therefore dynamics) in the supercell as it is in the unit cell.

From the calculations presented here it is not possible to

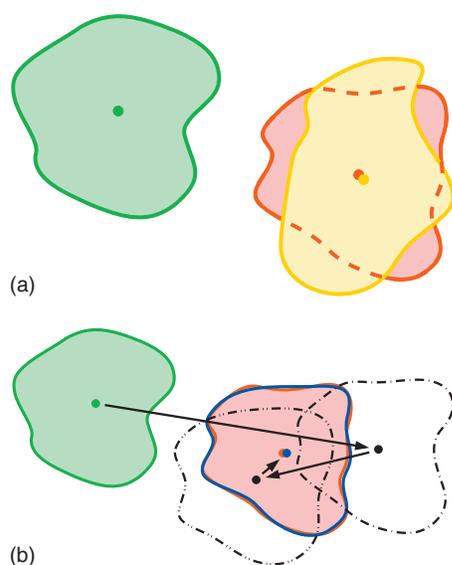


FIG. 5. Comparison of the configuration space sampled using different force fields. The region of configuration space sampled during a MD simulation is represented by the closed curve, while the optimized geometry is indicated by a dot. (a) First, a comparison is made between (red) DFT, (green) GAFF, and (yellow) a force field parameterized to reproduce the optimized geometry. Though the optimized geometry fits well to DFT, the region sampled is different. (b) Second, the convergence of the force field to (blue) f11m is shown. Since we fit to the forces, the shape of the region is replicated as well as the optimized geometry.

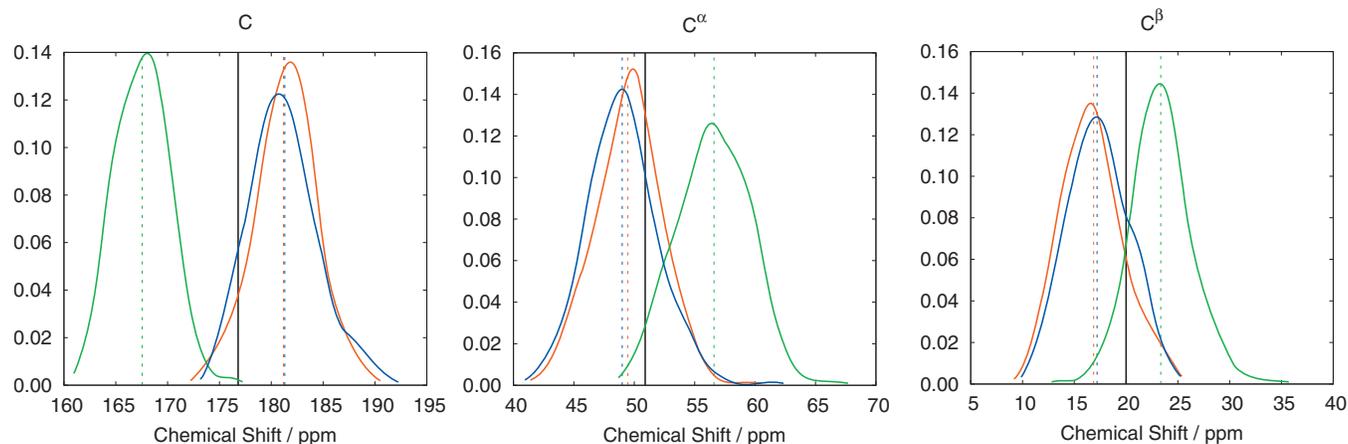


FIG. 6. Comparison of chemical shift distributions for configurations generated using (red) DFT MD, (green) GAFF, and (blue) f11m force fields. Dashed lines indicate the average chemical shift. The black line indicates the experimental shift. Distributions were created by summing Gaussians with standard deviation $\sigma=1$ ppm centered on each of the 256 data points and normalizing.

tell if the $2 \times 1 \times 2$ supercell is large enough to include all significant dynamical effects. Assuming that f11m is good for any sized supercell, it is of course very cheap to generate configurations. However the NMR calculations still scale as $\mathcal{O}(N^3)$ and become prohibitively expensive. This may become possible in the future with the implementation of NMR in linear scaling codes. For now though we cannot give a precise result for L-alanine. However, these results do make it clear that care must be taken over finite size effects when considering dynamics.

VI. CONCLUSION

We have demonstrated the use of MD to include dynamical effects in *ab initio* NMR calculations. First Langevin dynamics was used to efficiently sample the configuration space of the L-alanine molecular crystal. By careful use of the collision frequency we were able to remove the correlation between configurations encountered in previous work. This allowed us to fully converge the time averaged chemical shifts using DFT based MD. Including dynamics resulted in changes to the chemical shifts of up to 1 ppm, verifying previous work showing dynamics to be an important factor in organic solids.

The study of larger systems however is still impractical due to the computational expense of the PWP method. We have overcome this by using a classical force field based approach to generate configurations. Due to the sensitivity of the chemical shielding tensor to structure, it is important that the force field reproduces the DFT structures. This was achieved by fitting the force field to DFT calculated forces. The methodology required has been outlined including details of the constraints required to avoid saddle points in the parametrization. Using the optimized force field to generate configurations and DFT to calculate the chemical shieldings, we have been able to replicate the averaged chemical shifts found with DFT MD. A larger supercell of L-alanine could then be studied, which we found significantly changed the chemical shifts. Finite size effects are therefore an important consideration when including dynamical effects in NMR calculations.

Once finite size effects are included the only remaining significant source of error is the exchange-correlation functional. These results therefore provide a true picture of the contribution the exchange-correlation error makes to *ab initio* NMR calculations. It is found to be less than previously thought when performing calculations on a single geometry optimized configuration. As improved exchange-correlation functionals become available, the force field can be refitted. In this way the accuracy of the calculation is improved in a controlled and systematic fashion.

Over the past couple decades the size of systems studied using DFT has rapidly increased. However as the length scale increases, so also does the time scale of the dynamics that these systems exhibit. For as the size increases so does the number of degrees of freedom. This is particularly relevant for biological systems where weaker hydrogen and van der Waals bonds are prevalent, allowing greater dynamical freedom. As NMR calculations are performed on ever larger systems, it is therefore going to be increasingly necessary to include dynamical effects.

ACKNOWLEDGMENTS

The authors would like to thank Paul Tangney, Chris Pickard, and Jonathan Yates for helpful discussions. M.R. would like to thank the EPSRC for a Ph.D. studentship and P.D.H. the Royal Society for a University Research Fellowship. Computational facilities were provided by the University of Cambridge High Performance Computing Service.

- 1 T. Helgaker, M. Jaszunski, and K. Ruud, *Chem. Rev.* **99**, 293 (1999).
- 2 F. Mauri, B. G. Pfrommer, and S. G. Louie, *Phys. Rev. Lett.* **77**, 5300 (1996).
- 3 D. Sebastiani and M. Parrinello, *J. Phys. Chem. A* **105**, 1951 (2001).
- 4 T. Thonhauser, D. Ceresoli, A. A. Mostofi, N. Marzari, R. Resta, and D. Vanderbilt, *J. Chem. Phys.* **131**, 101101 (2009).
- 5 M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, *Rev. Mod. Phys.* **64**, 1045 (1992).
- 6 C. J. Pickard and F. Mauri, *Phys. Rev. B* **63**, 245101 (2001).
- 7 J. R. Yates, C. J. Pickard, and F. Mauri, *Phys. Rev. B* **76**, 024401 (2007).
- 8 R. K. Harris, S. A. Joyce, C. J. Pickard, S. Cadars, and L. Emsley, *Phys. Chem. Chem. Phys.* **8**, 137 (2006).
- 9 J. R. Yates, T. N. Pham, C. J. Pickard, F. Mauri, A. M. Amado, A. M. Gil, and S. P. Brown, *J. Am. Chem. Soc.* **127**, 10216 (2005).

- ¹⁰J. R. Yates, S. E. Dobbins, C. J. Pickard, F. Mauri, P. Y. Ghi, and R. K. Harris, *Phys. Chem. Chem. Phys.* **7**, 1402 (2005).
- ¹¹R. K. Harris, P. Hodgkinson, C. J. Pickard, J. R. Yates, and V. Zorin, *Magn. Reson. Chem.* **45**, S174 (2007).
- ¹²J.-N. Dumez and C. J. Pickard, *J. Chem. Phys.* **130**, 104701 (2009).
- ¹³I. De Gortari, G. Portella, X. Salvatella, V. S. Bajaj, P. C. A. van der Wel, J. R. Yates, M. D. Segall, C. J. Pickard, M. C. Payne, and M. Vendruscolo, *J. Am. Chem. Soc.* **132**, 5993 (2010).
- ¹⁴S. Tang and D. Case, *J. Biomol. NMR* **38**, 255 (2007).
- ¹⁵A. Bagno, F. D'Amico, and G. Saielli, *ChemPhysChem* **8**, 873 (2007).
- ¹⁶U. F. Rohrig and D. Sebastiani, *J. Phys. Chem. B* **112**, 1267 (2008).
- ¹⁷J. Wang, R. M. Wolf, J. W. Caldwell, P. A. Kollman, and D. A. Case, *J. Comput. Chem.* **25**, 1157 (2004).
- ¹⁸W. D. Cornell, P. Cieplak, C. I. Bayly, I. R. Gould, K. M. Merz, D. M. Ferguson, D. C. Spellmeyer, T. Fox, J. W. Caldwell, and P. A. Kollman, *J. Am. Chem. Soc.* **117**, 5179 (1995).
- ¹⁹S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. I. J. Probert, K. Refson, and M. C. Payne, *Z. Kristallogr.* **220**, 567 (2005).
- ²⁰D. Case, T. Darden, T. E. Cheatham III, C. Simmerling, J. Wang, R. Duke, R. Luo, K. Merz, D. Pearlman, M. Crowley, R. Walker, W. Zhang, B. Wang, S. Hayik, A. Roitberg, G. Seabra, K. Wong, F. Paesani, X. Wu, S. Brozell, V. Tsui, H. Gohlke, L. Yang, C. Tan, J. Mongan, V. Hornak, G. Cui, P. Beroza, D. Mathews, C. Schafmeister, W. Ross, and P. Kollman, *AMBER 9* (University of California, San Francisco, 2006).
- ²¹M. S. Lehmann, T. F. Koetzle, and W. C. Hamilton, *J. Am. Chem. Soc.* **94**, 2657 (1972).
- ²²D. C. Chatfield and S. E. Wong, *J. Phys. Chem. B* **104**, 11342 (2000).
- ²³D. Vanderbilt, *Phys. Rev. B* **41**, 7892 (1990).
- ²⁴J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).
- ²⁵H. J. Monkhorst and J. D. Pack, *Phys. Rev. B* **13**, 5188 (1976).
- ²⁶M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Oxford University Press, New York, 1989).
- ²⁷D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Academic, New York, 2002).
- ²⁸D. Quigley and M. I. J. Probert, *J. Chem. Phys.* **120**, 11432 (2004).
- ²⁹R. Kubo, *Rep. Prog. Phys.* **29**, 255 (1966).
- ³⁰C. Ye, R. Fu, J. Hu, L. Hou, and S. Ding, *Magn. Reson. Chem.* **31**, 699 (1993).
- ³¹I. Schnell, A. Lupulescu, S. Hafner, D. E. Demco, and H. W. Spiess, *J. Magn. Reson.* **133**, 61 (1998).
- ³²Y.-G. Yoon, B. G. Pfommer, S. G. Louie, and A. Canning, *Solid State Commun.* **131**, 15 (2004).
- ³³P. Tangney and S. Scandolo, *J. Chem. Phys.* **117**, 8898 (2002).
- ³⁴W. L. Jorgensen, D. S. Maxwell, and J. Tirado-Rives, *J. Am. Chem. Soc.* **118**, 11225 (1996).