

Large scale CASTEP calculations to interpret solid-state NMR and Vibrational Spectroscopy experiments

Bi-Ching Shih,[†] Dominik B. Jochym,[‡] Keith Refson,^{¶,§} and Jonathan R. Yates^{*,†}

[†]*Department of Materials, University of Oxford, Oxford, OX1 3PH, United Kingdom*

[‡]*Scientific Computing Department, Science and Technology Facilities Council, Rutherford Appleton Laboratory, Didcot OX11 0QX, United Kingdom*

[¶]*Department of Physics, Royal Holloway, University of London, Egham TW20 0EX, United Kingdom,*

[§]*ISIS Facility, Science and Technology Facilities Council, Rutherford Appleton Laboratory, Didcot OX11 0QX, United Kingdom*

E-mail: jonathan.yates@materials.ox.ac.uk

Abstract

In this project we have significantly improved the parallel scaling of the spectroscopic modules in CASTEP (specifically the Nuclear Magnetic Resonance (NMR) and vibrational spectroscopy modules) by introducing an additional layer parallelism into the code; parallelism over perturbations. Many spectroscopic parameters can be considered as the response of the electronic structure to a perturbation. For example the NMR J-coupling is the response of a system to the magnetic field generated by excitation of a nuclear spin; the NMR magnetic shielding is the response of a system to an applied external magnetic field; lattice dynamics is the response of the system to the movement of the ionic cores. Utilising this new level of parallelism means that

for large core counts the parallel FFT is only computed within a small subset of the participating cores dramatically increasing the scaling performance of the code.

Introduction

In recent years it has become clear that in the broad discipline of Materials Science the power of modern experimental spectroscopic techniques is inextricably linked to the availability of advanced computational tools and the infrastructure to run such simulations. This is particularly true in the fields of solid-state Nuclear Magnetic Resonance (NMR) and vibrational spectroscopy (neutron and X-ray scattering, Raman and IR spectroscopy), where deployment of first-principles methods to predict experimental observables has enabled many research groups to employ a combined experimental and computational approach.

The key quantity of Materials Modelling approaches based on Density Functional Theory (DFT) is the total energy of the system.¹ This is obtained by solving the Kohn-Sham² eigenvalue equations corresponding to different sampling points ('k-points') of the reciprocal-space Brillouin zone of a periodic simulation cell,

$$\hat{H}_k \psi_{bk}(\vec{r}) = E_{bk} \psi_{bk}(\vec{r}) \quad (1)$$

where \hat{H}_k is the Hamiltonian at k-point k , $\{\psi_{bk}(\vec{r})\}$ are the eigenstates, known as *bands* and $\{E_{bk}\}$ are the corresponding band-energies (i.e. the eigenvalues).³ The Hamiltonian is not a fixed operator but contains terms that depend on the electronic charge density n , which is given by

$$n(\vec{r}) = \sum_{bk} f_{bk} \int |\psi_{bk}(\vec{r})|^2 d^3k \quad (2)$$

The total energy can be used to rank the stability of structures - however, many experimental observables can be obtained as changes in total energy in response to a perturbation to the system. A simple example is the change in total energy on displacing an atom -

which gives rise to the force acting on the atom. In the context of NMR spectroscopy the key observables can be formulated as changes in the total energy. For example the magnetic shielding is the 2nd order change in the total energy with respect to an external magnetic field and a nuclear magnetic moment.

NMR Spectroscopy

Magnetic Shielding

The Zeeman interaction between a magnetic field, \mathbf{B} , and a set of spin 1/2 nuclei K , is given by

$$H = -\sum_K \gamma_K \mathbf{I}_K \cdot \mathbf{B}, \quad (3)$$

where $\hbar I_K$ is the spin angular momentum and γ_K the gyromagnetic ratio of nucleus K . If we consider \mathbf{B} as the field at the nucleus due to the presence of an externally applied field B_{ext} we can express Eqn. 3 as

$$H = -\sum_K \gamma_K \mathbf{I}_K (\mathbf{1} - \bar{\sigma}) \mathbf{B} \quad (4)$$

The first term is the interaction of the bare nucleus with the applied field while the second accounts for the response of the electrons to the field. This electronic response is characterized by the absolute magnetic the shielding tensor is defined as the ratio between this induced field, and the external applied field

$$\mathbf{B}_{in}(\mathbf{r}) = -\bar{\sigma}(\mathbf{r}) \mathbf{B}_{ext}. \quad (5)$$

In a diamagnetic insulator this arises from orbital currents induced by an external magnetic field.^{4,5} This current $\mathbf{j}(\mathbf{r})$, produces a non-uniform induced magnetic field in the material,

which is given by the Biot-Savart law as

$$\mathbf{B}_{\text{in}}(\mathbf{r}) = \frac{1}{c} \int d^3r' \mathbf{j}(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3}. \quad (6)$$

The shielding tensor can equivalently be written as a second derivative of the electronic energy of the system

$$\bar{\sigma} = \frac{\partial^2 E}{\partial m_K \partial B_{\text{ext}}}. \quad (7)$$

J-coupling

In the previous section we considered the effect of the magnetic field at a nucleus resulting from an externally applied field. However, there may also be a contribution to the magnetic field at a nucleus arising from the magnetic moments of the other nuclei in the system. In an effective spin Hamiltonian we may associate this spin-spin coupling with a term of the form

$$H = \sigma_{K < L} \mathbf{I}_K (\bar{D}_{KL} + \bar{J}_{KL}) \mathbf{I}_L \quad (8)$$

\bar{D}_{KL} is the direct dipolar coupling between the two nuclei and is a function of only the nuclear constants and the inter-nuclear distance. The J-coupling, \bar{J}_{KL} is a small perturbation to the electronic ground-state of the system and we can identify it as a derivative of the total energy E , of the system

$$\mathbf{J}_{KL} = \frac{\hbar \gamma_K \gamma_L}{2\pi} \frac{\partial^2 E}{\partial \mathbf{m}_K \partial \mathbf{m}_L} \quad (9)$$

An equivalent expression arises from considering one nuclear spin (L) as perturbation which creates a magnetic field at a second (receiving) nucleus (K)

$$\mathbf{B}_{\text{in}}^{(1)}(\mathbf{R}_K) = \frac{2\pi}{\hbar \gamma_K \gamma_L} \mathbf{J}_{KL} \cdot \mathbf{m}_L. \quad (10)$$

Eqn. 10 tells us that the question of computing J is essentially that of computing the magnetic field induced indirectly by a nuclear magnetic moment.⁶ When spin-orbit coupling is

neglected we can consider the field as arising from two, essentially independent, mechanisms. Firstly, the magnetic moment can interact with electronic charge inducing an orbital current $\mathbf{j}(\mathbf{r})$, which in turn creates a magnetic field at the other nuclei in the system. This mechanism is similar to the case of magnetic shielding in insulators. The second mechanism arises from the interaction of the magnetic moment with the electronic spin, causing an electronic spin polarisation. By working to first order in these quantities we can write the magnetic field at atom K induced by the magnetic moment of atom L as

$$\begin{aligned} \mathbf{B}_{\text{in}}^{(1)}(\mathbf{R}_K) &= \frac{\mu_0}{4\pi} \int \mathbf{m}^{(1)}(\mathbf{r}) \cdot \left[\frac{3\mathbf{r}_K\mathbf{r}_K - |\mathbf{r}_K|^2}{|\mathbf{r}_K|^5} \right] d^3\mathbf{r} \\ &+ \frac{\mu_0}{4\pi} \frac{8\pi}{3} \int \mathbf{m}^{(1)}(\mathbf{r}) \delta(\mathbf{r}_K) d^3\mathbf{r} \\ &+ \frac{\mu_0}{4\pi} \int \mathbf{j}^{(1)}(\mathbf{r}) \times \frac{\mathbf{r}_K}{|\mathbf{r}_K|^3} d^3\mathbf{r}. \end{aligned} \quad (11)$$

The first two terms arise from the induced spin density and are respectively termed the dipolar (DIP) and Fermi-contact (FC) contributions. The final term is due to the induced orbital current and can be further decomposed into a paramagnetic (PARA) and diamagnetic (DIA) contributions. The isotropic contributions to the J-coupling tensor can hence be written as^{7,8}

$$J_{iso} = J_{iso}^{SD} + J_{iso}^{FC} + J_{iso}^{DIA} + J_{iso}^{PARA} \quad (12)$$

Summary

The magnetic shielding tensors for a system can be obtain from the response to a single external perturbation - i.e. the current induced by an external magnetic field. However, for a system of N atoms there will be N^2 J-coupling tensors. A single perturbation gives the current (or spin density) induced by a single atomic movement. N perturbations will be required to obtain the full set of J-coupling tensors. In some experiments only a subset of the J-couplings will be required (e.g. those involving particular pairs of NMR active nuclei), however, in other materials, such as glasses, all possible J-couplings will be required. It can

hence be seen that while the magnetic shieldings cannot take advantage from perturbation parallelism, the J-coupling presents an ideal case.

Prior parallel performance

Response calculations such as J-couplings made use of two levels of parallelism within CASTEP. The calculation of contributions from different k-points are almost independent of each other and so it is natural to distribute the data over the k-points. Calculations scale almost perfectly with this distribution of data. However, as the size of the simulation cell grows, the size of the Brillouin Zone correspondingly reduces. This means that for larger calculations often a single k-point is required. This means that k-point parallelism is of limited use in a HPC environment. The second strategy is to distribute the data by the G-vectors. Since the number of plane-waves is often large, and grows approximately linearly with simulation cell size, this enables efficient data distribution and load-balancing. During the application of the Hamiltonian the data must be transferred from real to reciprocal space using a 3D FFT. This operation involves a transpose of the data and all-to-all communications. This limits the number of cores that can be used effectively. Figure 1 shows the scaling of a typical NMR calculation with number of processing elements.

Perturbation parallelism

The implementation made use of the MPI communicators previously introduced to handle task farming parallelism. CASTEP's modular structure means that all mpi calls are contained within the comms module - and so the additional communication routines only needed to be added to this file. The J-coupling code was original written to be called only a single time in each run. Careful testing was needed to make sure the routines were re-entrant and numerically identical results were obtain no matter which order the perturbations were computed. To enable the user to easily select which perturbations to include in the calculation

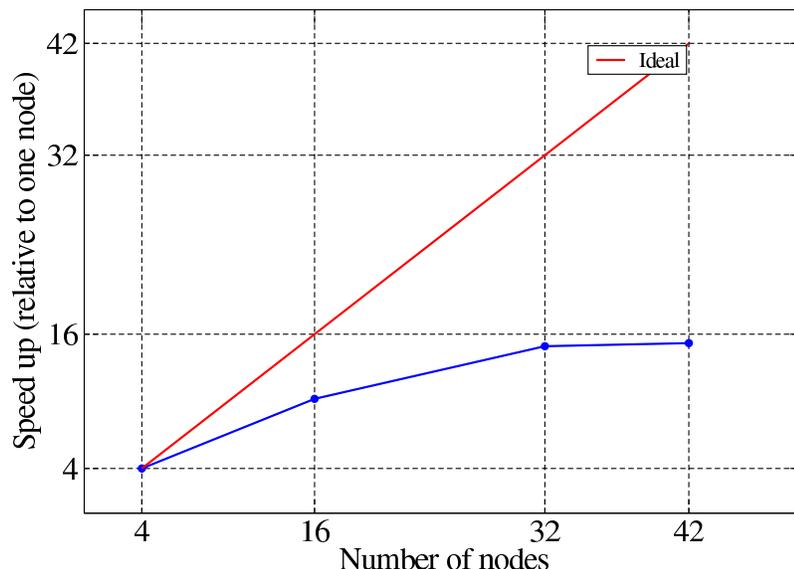


Figure 1: Performance of J-coupling code prior to e-CSE work. Scaling is limited by the 3D FFT and the all-to-all communication.

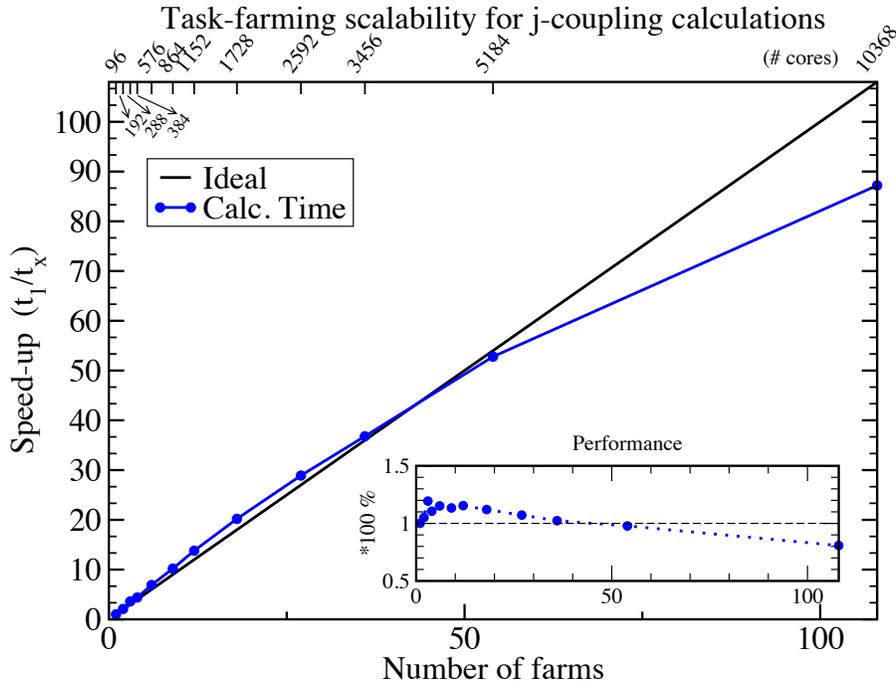
the following keyword was added to the CASTEP input file (via the “cell” file)

```

jcoupling_site : all ! All atomic sites
jcoupling_site : C N ! All sites belonging to the listed species
jcoupling_site : C 1 !The site C 1
jcoupling_site : C{1,3,5-11} ! The sites C1,C3,C5,C6,C7,C8,C9,C10,C11

```

The output of the calculation (the J-coupling tensors) is written to a single file in the magres file format.⁹ Tools for the visualization and processing of magres files have recently been developed.⁹ Initial very poor parallel performance was tracked down to i/o issues (namely multiple cores trying to write data to the same, admittedly small, file). Once this was corrected we found very good scaling up to large number of cores (10k+). Figure2 shows the scaling behaviour for a 108 atom model of a silicate glass. The calculations reach 80% speed up at the largest core count (approx 10k) when each task farm is assigned a single perturbation. On further investigation we believe this reduction from ideality is due to the



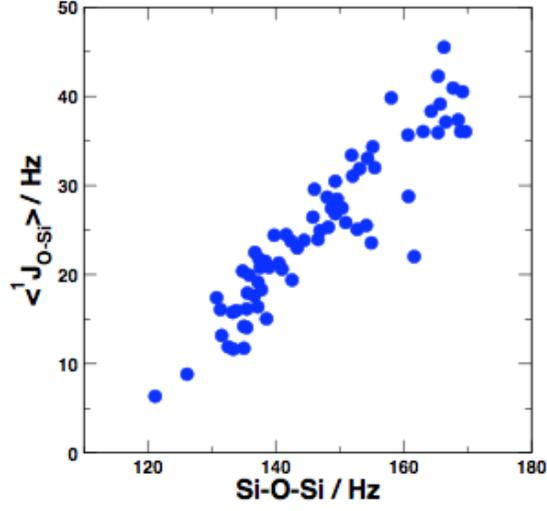


Figure 3: Distribution of ${}^1J_{O-Si}$ with Si-O-Si bond angle for a model of a silicate glass.

Vibrational Spectroscopy

The theory of lattice dynamics begins with an expansion of the total energy of the system in terms of the atomic coordinates,¹⁰

$$E = E_0 + \sum_{\kappa, \alpha, a} \frac{\partial E}{\partial \mathbf{u}_{\kappa, \alpha, a}} + \frac{1}{2} \sum_{\kappa, \alpha, \kappa', \alpha', a} \mathbf{u}_{\kappa, \alpha, a} \Phi_{\alpha, \alpha'}^{\kappa, \kappa'} \mathbf{u}_{\kappa', \alpha', a} + \dots \quad (13)$$

where $\mathbf{u}_{\kappa, \alpha, a}$ is the vector of atomic displacements from equilibrium. $\Phi_{\alpha, \alpha'}^{\kappa, \kappa'}$ is the matrix of force constants,

$$\Phi_{\alpha, \alpha'}^{\kappa, \kappa'}(a) = \frac{\partial^2 E}{\partial \mathbf{u}_{\kappa, \alpha} \partial \mathbf{u}_{\kappa', \alpha'}} \quad (14)$$

At equilibrium the first order term in the expansion is zero. The vibrational modes of the system are obtained by solving the eigenvalue equation,

$$D_{\alpha, \alpha'}^{\kappa, \kappa'}(\mathbf{q}) \varepsilon_{m\kappa, \alpha\mathbf{q}} = \omega_{m, \mathbf{q}}^2 \varepsilon_{m\kappa, \alpha\mathbf{q}} \quad (15)$$

Where the dynamical matrix, $D_{\alpha,\alpha'}^{\kappa,\kappa'}(\mathbf{q})$ is the mass weighted fourier transform of the force constant matrix,¹¹ $\Phi_{\alpha,\alpha'}^{\kappa,\kappa'}$

$$D_{\alpha,\alpha'}^{\kappa,\kappa'}(\mathbf{q}) = \frac{\mathbf{1}}{\sqrt{\mathbf{M}_{\kappa}\mathbf{M}_{\kappa'}}} \sum_{\mathbf{a}} \Phi_{\alpha,\alpha'}^{\kappa,\kappa'}(\mathbf{a}) e^{-i\mathbf{q}\cdot\mathbf{R}_{\mathbf{a}}} \quad (16)$$

It can be seen that the task of computing the vibrational frequencies, $\omega_{m,\mathbf{q}}^2$, of the system is essentially that of computing the force constant matrix - the response of the system to the movement of the ionic cores. In general there will be $3N$ perturbations, however in most situations the symmetry of the crystal will reduce the number of unique perturbations.

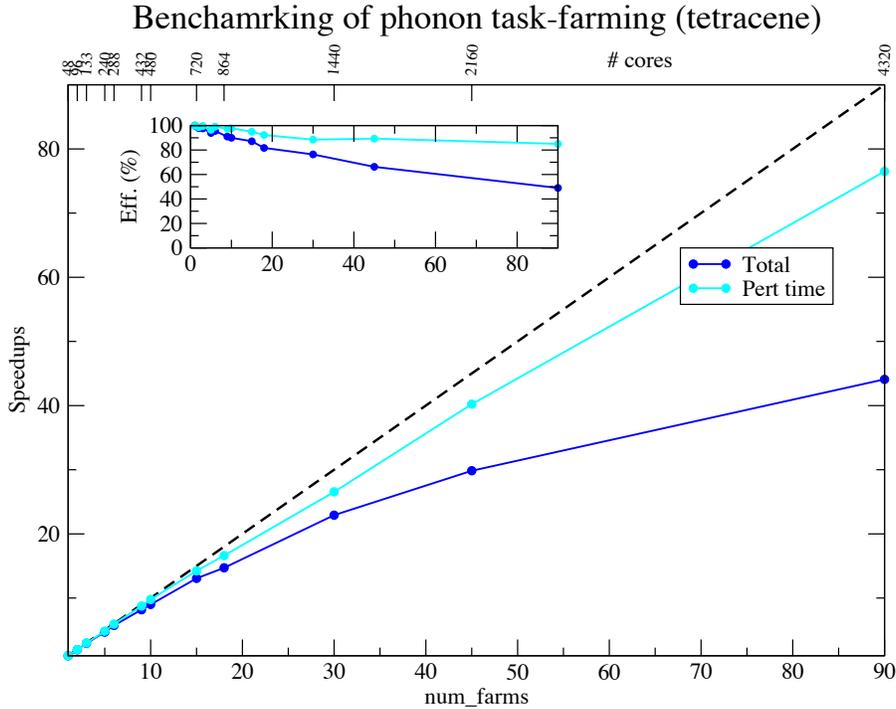


Figure 4: Parallel performance post-eCSE of the calculation of phonon modes in the case of a single k-point to sample the Brillouin Zone

The distribution of the perturbations between the task farms made use of the same communicators as the J-coupling in the previous section. However, there are two additional challenges. Firstly, the calculations of the force constant matrix includes an element of work that is not distributed over the task farms. In the following graphs the timings with and

without this contribution are included. The solution to this is to perform an initial non-task-parallel calculation to pre-compute the necessary quantities. As a prerequisite to a phonon calculation is a optimization of the geometry, these contributions could be calculated at the end of the geometry optimization - thus avoiding the need for an additional job. A more significant challenge is load balancing between the farms. For calculations on an isolated molecule, or a very large crystal it is sufficient to sample the electronic Brillouin zone with a single k-point. In this case the time for each perturbation is almost identical and there is no load balancing issue. However, for small and even moderately large crystals it is necessary to sample the Brillouin Zone with multiple k-points. Depending on the symmetry of a particular perturbation the number of unique k-points will be different. The calculation time is well

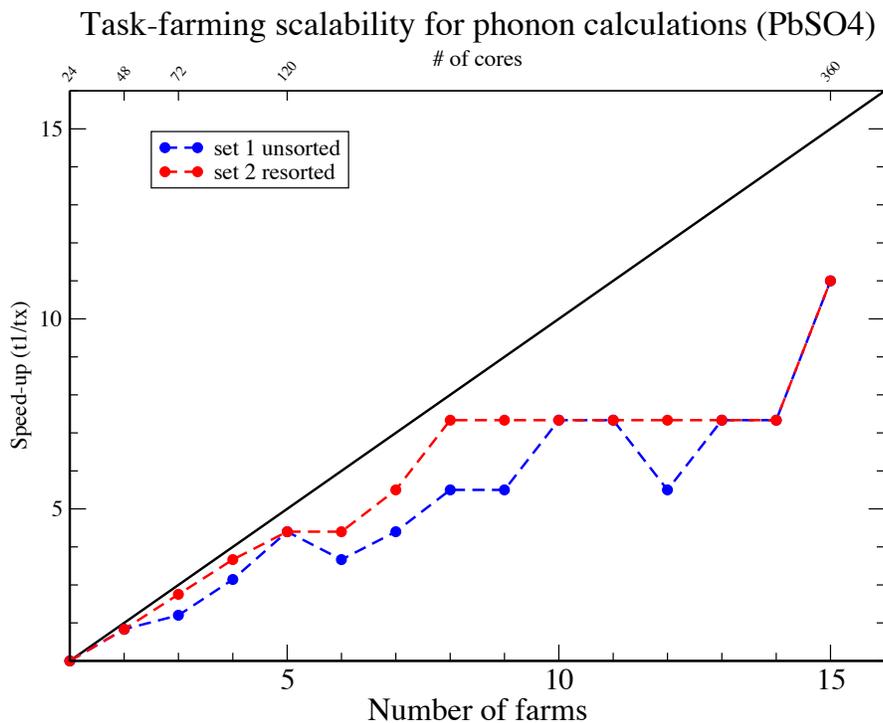


Figure 5: Parallel performance post-eCSE of the calculation of phonon modes in the case where multiple k-points are needed to sample the Brillouin Zone.

correlated to the number of kpoints - and so the time to calculate each perturbation will vary significantly. One approach to addressing this would be to assign variable numbers of processors to each farm. It was felt that this would introduce significant complexity into the

code - and would also be problematic in the case of farms dealing with multiple perturbations (i.e. communicators would need to be dynamically re-assigned during the code run). Instead a fixed farm size was chosen. This means that in the limit that each farm calculates one perturbation the efficiency will be limited by the most expensive perturbation. This can be seen in Figure 5 where using the maximum number of farms (15) gives only an 11x speed up. However, when each farm can deal with multiple perturbations the perturbations can be assigned so as to choose the optimal combination of perturbations to assign to each farm. This is done by pre-computing the number of k-point required for each perturbation and estimating the workload required. Figure 5 shows the effect of the calculated assignment of perturbations in comparison to a random assignment. For 8 farms there is an almost perfect scaling. It is clearly important for a user to be able to determine the optimal number of farms for a given system. A simple command line tool has been developed to estimate the load balancing for different numbers of farms, and hence allow the user to identify the most efficient.

Acknowledgement

This work was funded under the embedded CSE programme of the ARCHER UK National Supercomputing Service (<http://www.archer.ac.uk>)

References

- (1) Hohenberg, P.; Kohn, W. *Phys. Rev.* **1964**, *136*, B864–B871.
- (2) Kohn, W.; Sham, L. J. *Phys. Rev.* **1965**, *140*, A1133–A1138.
- (3) Payne, M. C.; Teter, M. P.; Allan, D. C.; Arias, T.; Joannopoulos, J. D. *Rev. Mod. Phys.* **1992**, *64*, 1045–1097.
- (4) Pickard, C.; Mauri, F. *Phys. Rev. B* **2001**, *63*, 245101.

- (5) Yates, J.; Pickard, C.; Mauri, F. *Phys. Rev. B* **2007**, *76*, 024401.
- (6) Bonhomme, C.; Gervais, C.; Babonneau, F.; Coelho, C.; Pourpoint, F.; Azais, T.; Ashbrook, S. E.; Griffin, J. M.; Yates, J. R.; Mauri, F.; Pickard, C. J. *Chem. Rev.* **2012**, *112*, 5733.
- (7) Joyce, S.; Yates, J.; Pickard, C.; Mauri, F. *J. Chem. Phys.* **2007**, *127*, 204107.
- (8) Green, T.; Yates, J. *J. Chem. Phys.* **2014**, *140*, 234106.
- (9) Sturniolo, S.; Green, T. F.; Hanson, R. M.; Zilka, M.; Refson, K.; Hodgkinson, P.; Brown, S. P.; Yates, J. R. *Solid State Nuclear Magnetic Resonance* **2016**, *78*, 64 – 70.
- (10) Refson, K.; Tulip, P. R.; Clark, S. J. *Phys. Rev. B* **2006**, *73*, 155114.
- (11) Xavier Gonze, J.-C. C.; Teter, M. P. *Phys. Rev. B* **1994**, *50*, 13035–13038.