

Computer Programs in Physics



Berry: A code for the differentiation of Bloch wavefunctions from DFT calculations ☆☆☆

Leander Reascos^a, Fábio Carneiro^a, André Pereira^d, Nuno Filipe Castro^c,
Ricardo Mendes Ribeiro^{a,b,*}

^a Department and Centre of Physics, University of Minho, Campus of Gualtar, 4710-057, Braga, Portugal

^b International Iberian Nanotechnology Laboratory (INL), Av. Mestre Jose Veiga, 4715-330, Braga, Portugal

^c Department of Physics and LIP, University of Minho, Campus of Gualtar, 4710-057, Braga, Portugal

^d HasLab, INESC TEC & University of Minho, Campus of Gualtar, 4710-057, Braga, Portugal

ARTICLE INFO

Keywords:

Electronic structure methods
Berry geometries
Second harmonic generation

ABSTRACT

Density functional calculation of electronic structures of materials is one of the most used techniques in theoretical solid state physics. These calculations retrieve single electron wavefunctions and their eigenenergies. The berry suite of programs amplifies the usefulness of DFT by ordering the eigenstates in analytic bands, allowing the differentiation of the wavefunctions in reciprocal space. It can then calculate Berry connections and curvatures and the second harmonic generation conductivity. The berry software is implemented for two dimensional materials and was tested in hBN and InSe. In the near future, more properties and functionalities are expected to be added.

Program summary

Program Title: berry

CPC Library link to program files: <https://doi.org/10.17632/mpbbksz2t7.1>

Developer's repository link: <https://github.com/ricardoribeiro-2020/berry>

Licensing provisions: MIT

Programming language: Python3

Nature of problem: Differentiation of Bloch wavefunctions in reciprocal space, numerically obtained from a DFT software, applied to two dimensional materials. This enables the numeric calculation of material's properties such as Berry geometries and Second Harmonic conductivity.

Solution method: Extracts Kohn-Sham functions from a DFT calculation, orders them by analytic bands using graph and AI methods and calculates the gradient of the wavefunctions along an electronic band.

Additional comments including restrictions and unusual features: Applies only to two dimensional materials, and only imports Kohn-Sham functions from Quantum Espresso package.

1. Introduction

Ventura et al. [1] derived an iterative formula to calculate the optical conductivity of crystals to any order based on the knowledge of the Berry connections of the material. Although the calculation of the linear optical conductivity of materials is straightforward and actually implemented in most density functional calculation (DFT) software packages,

the same cannot be said of the higher orders. Given the technological importance of second order processes like second harmonic generation and photoconductivity, or higher order like third harmonic generation, etc., the implementation of Ventura's formula departing from first principles calculations like DFT seems to be very interesting.

But this endeavor faces a major hurdle. The Berry connections in a crystal can be obtained by [2,3]:

☆ The review of this paper was arranged by Prof. Weigel Martin.

☆☆ This paper and its associated computer program are available via the Computer Physics Communications homepage on ScienceDirect (<http://www.sciencedirect.com/science/journal/00104655>).

* Corresponding author at: Department and Centre of Physics, University of Minho, Campus of Gualtar, 4710-057, Braga, Portugal.
E-mail address: ricardo@fisica.uminho.pt (R.M. Ribeiro).

$$\xi_{ks,s'} = i \langle u_{ks} | \nabla_{\mathbf{k}} u_{ks'} \rangle = \frac{i}{v_{uc}} \int_{uc} d^3 \mathbf{r} u_{ks}^*(\mathbf{r}) \nabla_{\mathbf{k}} u_{ks'}(\mathbf{r}) \quad (1)$$

which includes the gradient of the Bloch factor in reciprocal space. In this equation, s and s' are electronic bands, u_{ks} is the Bloch factor of the Bloch wavefunction at reciprocal space point \mathbf{k} and band s , v_{uc} is the volume of the unit cell and the integral is over the unit cell. This means that we need the wavefunctions classified in bands where analyticity applies, that is, are numerically differentiable in reciprocal space. Strinati [4] has already noted that and proposed that bands should be labeled in such a way as to guarantee the analyticity of the energy eigenvalues.

There are known problems in achieving this [4–6]. One is that the DFT programs calculate the Kohn-Sham equations for each \mathbf{k} -point independently of the others. Since the solutions of the differential equation come with an overall random phase, that phase will be different from point to point, precluding analyticity. The way this is dealt in the BERRY package is explained in subsection 2.2. Another major obstacle is the degeneracy of the electronic bands at some points, i.e. bands are not isolated but cross. This problem also appears when trying to generate Wannier functions from DFT numerical calculations, and it is dealt by isolating a set of bands and performing a basis rotation on the set [7,5]. Here the adopted procedure is similar, as can be seen in subsection 2.4.

Even surpassing these obstacles, the classification of the Bloch wavefunctions in analytic bands in an ordered way is not simple. BERRY uses techniques of Graph Theory (GT) and Unsupervised Machine Learning (UML) iteratively to achieve the final result. It also includes some visualization tools to help discern the quality of the calculation.

Berry connections and curvatures are then calculated and used to obtain the linear and second harmonic generation optical conductivity. Of course, once having the Berry connections, other non-linear optical properties can be implemented, as well as other properties that can be calculated using the organization of wavefunctions in analytic bands. We expect to implement them in the near future.

The paper is organized as follows. In Sec. 2 the algorithm and the implementation are described. Then, in Sec. 3 the installation and usage of the software is explained. In Sec. 4 the software is applied to two materials: single layer hexagonal boron nitride and single layer indium selenide, with some convergence studies. For these two materials, we calculated linear and second harmonic conductivity using the length gauge description of the reduced density matrix, as described below. Other works on linear and non linear optical response of hBN using the length gauge description can be found in [8] and [9]. Some considerations regarding the performance of BERRY are exposed in Sec. 5. We comment on future perspectives of the software in Sec. 6 and we draw some conclusions in Sec. 7.

2. Algorithm and implementation

2.1. Preprocessing

BERRY departs from first principles calculations. It aims at transforming the results obtained from a DFT software (only QUANTUM ESPRESSO is currently supported, but we expect to extend the possibility to other DFT software) in such a way as to be able to use in other calculations.

So the first step is to obtain the results from a DFT calculation. This is accomplished by a package that, given a self consistent field (scf) input file of QUANTUM ESPRESSO (QE), with the pseudopotentials for the DFT run, and a small input file with the area and density of points in the reciprocal space wanted, automatically runs the DFT scf calculation, generates the non-scf (nscf) input file for the \mathbf{k} -points wanted and runs it. This procedure assures that the wavefunctions are well ordered and with their position identified in reciprocal space.

The package also reads all relevant data from the DFT run and saves them in a manner useful for the other packages of the software. It works like an interface between the DFT software and the rest of the packages, and so is adaptable for use with other DFT packages other than QE. (See flux chart in Fig. 1.)

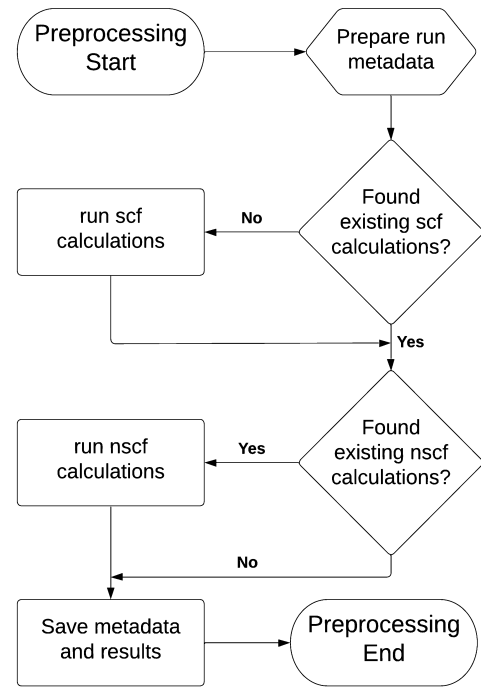


Fig. 1. Flux chart for the preprocessing implementation. This part of the code runs the DFT calculations and extracts the data for the run, saving it for later use.

2.2. Eliminating random phases

The next step is to read the wavefunctions obtained from the DFT calculation, make them coherent by multiplying by a properly chosen phase, and save them in a format suitable for the next steps.

Wavefunctions in a crystalline material can be written as:

$$\psi_{ks}(\mathbf{r}) = e^{i\mathbf{k}\cdot\mathbf{r}} u_{ks}(\mathbf{r}) \quad (2)$$

with the Bloch factor $u_{ks}(\mathbf{r}) = u_{ks}(\mathbf{r} + \mathbf{R})$ for \mathbf{R} any lattice vector.

A procedure to solve this problem is hinted, for instance, in [10] and [11]. The idea is that all wavefunctions have to have a common reference. We choose that the phases of all wavefunctions of the system should be synchronized in a chosen point \mathbf{r}_0 of real space. For instance, we can set the phase to zero at that chosen point by multiplying each wavefunction $\psi_{ks}(\mathbf{r})$ by the reciprocal of the phase $e^{i\phi_0}$ it has in that point: $\psi_{ks}(\mathbf{r}) \rightarrow e^{-i\phi_0} \psi_{ks}(\mathbf{r})$, where ϕ_0 is the argument of wavefunction $\psi_{ks}(\mathbf{r})$ at $\mathbf{r} = \mathbf{r}_0$. The choice of the point in real space is not relevant (as long as all wavefunctions are different from zero at that point) since choosing another point would simply add a global phase to all wavefunctions, that would cancel in equation (1). A low symmetry point in real space is chosen so that wavefunctions at that point in real space are different from zero, a condition that has to be verified. This results in a totally coherent set of wavefunctions, for which we can apply the usual operations.

Wavefunctions are saved in individual files, labeled by \mathbf{k} -point number and band, and are written as a function of position, already corrected for the phase.

2.3. Establishing the bands: solver algorithm

This is the main part of the project, by which each eigenstate is put in a set that form a continuous and analytic band. By analytic band we mean that the gradient of the Bloch factors $\nabla_{\mathbf{k}} u_{ks}$ of the set s that we call a band may be calculated.

There are two main data that can be used to establish to which band an eigenstate belongs. One is the eigenenergy $E_{k,s}$ which has to be continuous in \mathbf{k} in every band s . This criterion is necessary but clearly

insufficient, since there are degenerate points (i.e. with the same energy, although belonging to different bands) and even if the states are not degenerate they can be sufficiently close in energy to be difficult to discern numerically if there is continuity or not.

A more natural criterium is the continuity of the Bloch factors u_{k_s} along the reciprocal space, which are the ones used to calculate the gradient. Continuity means that u_{k_s} should change little from one k to another very close k' . Therefore the dot product of two neighboring Bloch factors u_{k_s} and $u_{k'_s}$ should be close to 1 (they are normalized) if there is continuity between the factors in the same band s . Otherwise, the dot product should be close to zero, since Bloch factors of eigenstates in the same k -point are orthogonal.

To verify the continuity of the Bloch factors we calculated the inner product $S_{kk's's'} = \langle u_{k_s}(\mathbf{r}) | u_{k'_s}(\mathbf{r}) \rangle$ between Bloch factors of neighboring k and $k' = k + \Delta k$ points.

For sufficiently small Δk the modulus of the inner product of the normalized Bloch factors $|S|$ is either close to one for states that vary continuously and belong to the same band or much smaller than one if the states belong to different bands.

This way we were able to classify, for each k , which eigenstate corresponds (in the sense that it is continuous) to which eigenstate of neighbor $k + \Delta k$ point. We then propagate this procedure through all k points of the Brillouin Zone (BZ) and all eigenstates and make sets of eigenstates that are linked by continuity. This gives a straightforward way of determining which solutions $|u_{k_s}\rangle$ should go together to form a band.

Unfortunately, although this works for most of the cases, there are some for which it doesn't. One example where it doesn't work is where we have the degenerate states, since any linear combination of degenerate states is also a solution to the eigenvalue problem; thus the orthogonality criterion fails. Quasi-degeneracies are also problematic, but the algorithm identifies them since they join two different bands that are superposed. The algorithm deals with these pairs of eigenstates by assigning each one to a band.

To address the cases where these criteria fail, we resort to Graph Theory (GT) and Unsupervised Machine Learning (UML). Our approach involves two steps, that are repeated iteratively, and is called *Solver Algorithm*. The first step is a component identification based on GT and the second is a samples assignment using UML methods.

The *vertices* or *nodes* of GT are points of the form $p = (k, E_{k_s})$ and a *connection* or *edge* between points is formed whenever $|S_{kk's's'}| > \text{TOL}$, where TOL is a tolerance the user can define. For simplicity, this graph is undirected and does not have weights on its edges.

Component detection algorithms based on Graph Theory are used to identify groups of points for which the dot product of the Bloch factors $S_{kk's's'}$ is close to 1. These groups are referred to as *components* [12]. Specifically, we use the algorithm implemented by the "Networkx" Python library [13]. Once identified, these components correspond to patches of a continuous band. These patches have to be "glued" together consistently to form the full band.

This prepares an initial data structure in the form of components, convenient for an Unsupervised Machine-Learning clustering algorithm. The purpose of UML algorithms is to identify patterns within unclassified data by analyzing the differences or similarities between data points using specific metrics. One of the most commonly used UML algorithms is the K-means clustering algorithm. This algorithm classifies the data points into a predefined number of *clusters* referred to a collection of data points that exhibit similarities according to a specific metric [14,15].

The implemented UML algorithm in this program is based on the K-means approach, where the component detection method is used to identify the initial K clusters and data points to be classified. However, the current implementation does not rely on pre-existing algorithms, as some data points may require adjustments to be correctly assigned.

The implementation of the algorithm offers an added advantage, as it provides the flexibility to adjust and correctly assign data points for more accurate classification.

UML then forms groups of components called *clusters* (which are patches of bands glued together). Each data point belongs exclusively to a single cluster. Ideally, in the end, each cluster should contain all points of a continuous (in the analytic sense) band and only those. Yet this usually is not enough to build complete bands, and so the process has to be iterated, as shown in Fig. 2. Each iteration uses as a starting point the best result of previous iterations. One iteration involves identifying the components, given the initial state in the form of a graph. Clustering these components gives larger components, that will form a new state that will be used as an initial state of the next iteration. The optimization process ends when the result does not change between iterations, reaching a stationary solution, or the relaxation parameter reaches the minimum value.

The final result is evaluated to ascertain which bands are correctly built and for the identification of problems such as degenerate states, that have to be dealt afterwards. A table is then produced with this information and the number of useful bands generated is given, for the next packages of the software.

2.4. Dealing with degeneracy

Degenerate points are easy to spot, since they have the same energy, at the same k -point.

Yet, there is no guarantee that the Bloch factors of these degenerate states are continuous to the adjacent Bloch factors, that are already attributed to specific bands, since a linear combination of two degenerate eigenstates is also an eigenstate.

The solution is to make a local basis rotation of these two states, by applying an unitary transformation in such a way as to maximize the continuity with the two bands that cross. This is what is done for similar problems when constructing Wannier functions. It can be applied after the previous step of establishing the bands, to the identified degenerate points.

Consider a point in k -space and two Bloch factors $|u_{k_s}\rangle$ and $|u_{k'_s}\rangle$ that have been signaled as not continuous to the neighboring Bloch factors $|u_{k'_A}\rangle$ and $|u_{k'_B}\rangle$ that belong to band A and band B , respectively.

We will apply a unitary transformation

$$|u'_{kA}\rangle = a_1 |u_{k_s}\rangle + a_2 |u_{k'_s}\rangle \quad (3)$$

$$|u'_{kB}\rangle = b_1 |u_{k_s}\rangle + b_2 |u_{k'_s}\rangle \quad (4)$$

where we have the restrictions due to orthonormalization:

$$a_1 a_1^* + a_2 a_2^* = 1 \quad (5)$$

$$b_1 b_1^* + b_2 b_2^* = 1 \quad (6)$$

$$a_1^* b_1 + a_2^* b_2 = 0 \quad (7)$$

Then we want this transformation to be such that maximizes continuity to bands A and B , so that the dot products $\langle u'_{kA} | u'_{kA} \rangle$ and $\langle u'_{kB} | u'_{kB} \rangle$ are close to 1, which is the maximum they can be.

We look for the set a_1, a_2, b_1, b_2 that give:

$$\max \langle u'_{kA} | u'_{kA} \rangle = \max (\langle u_{k'_A} | u_{k_s} \rangle a_1 + \langle u_{k'_A} | u_{k'_s} \rangle a_2) \quad (8)$$

$$\max \langle u'_{kB} | u'_{kB} \rangle = \max (\langle u_{k'_B} | u_{k_s} \rangle b_1 + \langle u_{k'_B} | u_{k'_s} \rangle b_2) \quad (9)$$

Using the Bloch factors in the rotated basis assures the analyticity at the degenerate point.

2.5. Calculating Berry connections

Having obtained the set of states that belong to each band of interest, ordered in reciprocal space and covering all BZ, we can then determine

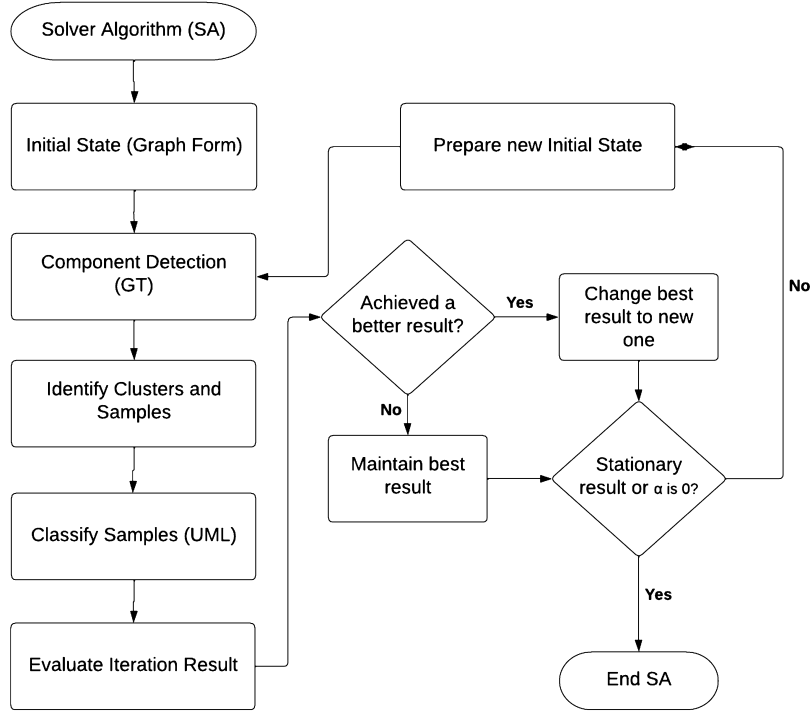


Fig. 2. The flux chart represents a simplified Solver Algorithm version. It starts converting the problem into a graph form using the dot-product data. The graph's components are detected by Graph Theory (GT) methods. Then, the clusters and samples are identified to be classified with Unsupervised Machine Learning (UML) techniques. A result evaluation identifies the initial state for the next iteration. The algorithm stops when it reaches a stationary result, or the relaxation parameter α is zero.

the Berry connections between every pair of bands s and s' at each \mathbf{k} point of the BZ.

To calculate the gradient in \mathbf{k} -space, $\nabla_{\mathbf{k}} u_{\mathbf{k}s}$, we have to have the Bloch factors written as a function of \mathbf{r} for each \mathbf{k} -point and not as a function of \mathbf{k} for each real space point \mathbf{r} . So a conversion is done before differentiation.

Gradients are then calculated using the *Gradient* function of the python package library *findiff*, using a centered two point formula to 2nd order and saved in files. Berry connections are calculated for each pair of bands using Equation (1), and saved in files.

Berry curvatures can also be calculated using Equation (10) [16]

$$\Omega_{s,s'} = i \int_{uc} d^3 \mathbf{r} (\nabla_{\mathbf{k}} u_{\mathbf{k}s}^* \times \nabla_{\mathbf{k}} u_{\mathbf{k}s'}) \quad (10)$$

2.6. Optical conductivity

Now it is possible to proceed to the calculation of the linear optical conductivity and the second harmonic generation. For this, we use the length gauge description of the reduced density matrix [17–19], as described in reference [19].

For the case of a cold 2 dimensional semiconductor, the linear conductivity can be expressed in terms of a single, interband, contribution, given by Equation (11)

$$\sigma^{\beta\alpha}(\omega) = \frac{ie^2}{\hbar} \sum_{s's} \int_{BZ} \frac{d^2 \mathbf{k}}{(2\pi)^2} \frac{\Delta \epsilon_{\mathbf{k}s's} \xi_{\mathbf{k}s's}^{\beta} \xi_{\mathbf{k}s's}^{\alpha}}{(\hbar\omega + i\Gamma) - \Delta \epsilon_{\mathbf{k}s's}} \Delta f_{\mathbf{k}s's'}, \quad (11)$$

where $\Delta \epsilon_{\mathbf{k}s's} = \epsilon_{\mathbf{k}s'} - \epsilon_{\mathbf{k}s}$, $\Delta f_{\mathbf{k}s's'} = f_{\mathbf{k}s} - f_{\mathbf{k}s'}$, $\epsilon_{\mathbf{k}s}$ is the eigenvalue of band s at \mathbf{k} , and $f_{\mathbf{k}s} = f(\epsilon_{\mathbf{k}s})$ is the Fermi-Dirac distribution function. Note that with the knowledge of the band energies, $\epsilon_{\mathbf{k}s}$, and the Berry connections, $\xi_{\mathbf{k}s's'}^{\beta}$, both of which can be determined by a DFT calculation, we can compute the optical response of a crystal. The indexes α and β run through the spatial coordinates x and y as we are dealing with a two dimensional material, and the conductivity tensor has,

in general, four relevant components. The results from this calculation (see Section 4) are the same as the ones obtained on the same system calculated using the dipole approximation within the random phase approximation (RPA) [20], that usually is available in DFT packages. This very good agreement validates the model: the same physical quantity calculated in two totally different models gives the same result.

The second order response can be calculated for the case of the second harmonic generation. Following the procedure described in ref. [19] we obtain, for a diagonal component of the SHG, $\sigma^{\beta\beta\beta}(\omega, \omega)$, Eq. (12).

$$\begin{aligned} \sigma^{\beta\beta\beta}(\omega, \omega) = & -\frac{e^3}{\hbar} \int_{BZ} \frac{d^2 \mathbf{k}}{(2\pi)^2} \left[\sum_{s' \neq s} \frac{1}{2} \frac{\Delta \epsilon_{\mathbf{k}s's'}}{g_{\omega \mathbf{k}s's'}} \left\{ \left[\xi_{\mathbf{k}s's'}^{\beta} (\xi_{\mathbf{k}s's'}^{\beta})_{;\beta} \right] \frac{\Delta f_{\mathbf{k}s's'}}{h_{\omega \mathbf{k}s's'}} \right\} \right. \\ & \left. - \frac{i}{2} \sum_{s' \neq s \neq r} \frac{\Delta \epsilon_{\mathbf{k}s's'}}{g_{\omega \mathbf{k}s's'}} \left\{ \frac{\Delta f_{\mathbf{k}s'r}}{h_{\omega \mathbf{k}rs'}} - \frac{\Delta f_{\mathbf{k}rs}}{h_{\omega \mathbf{k}rs}} \right\} \left(\xi_{\mathbf{k}'s'}^{\beta} \xi_{\mathbf{k}s'r}^{\beta} \xi_{\mathbf{k}rs'}^{\beta} + \xi_{\mathbf{k}'r}^{\beta} \xi_{\mathbf{k}rs}^{\beta} \xi_{\mathbf{k}s's'}^{\beta} \right) \right] \quad (12) \end{aligned}$$

where

$$g_{\omega \mathbf{k}s's'} = 2(\hbar\omega + i\Gamma) - \Delta \epsilon_{\mathbf{k}s's'} \quad (13)$$

$$h_{\omega \mathbf{k}s's'} = (\hbar\omega + i\Gamma) - \Delta \epsilon_{\mathbf{k}s's'} \quad (14)$$

$$\left[\xi_{\mathbf{k}'s'}^{\beta} (\xi_{\mathbf{k}s's'}^{\beta})_{;\beta} \right] = \xi_{\mathbf{k}'s'}^{\beta} (\xi_{\mathbf{k}s's'}^{\beta})_{;\beta} - (\xi_{\mathbf{k}s's'}^{\beta})_{;\beta} \xi_{\mathbf{k}'s'}^{\beta} \quad (15)$$

$$(\xi_{\mathbf{k}s's'}^{\beta})_{;\beta} = \left(\nabla^{\beta} \xi_{\mathbf{k}s's'}^{\beta} \right) - i \left(\xi_{\mathbf{k}ss}^{\beta} - \xi_{\mathbf{k}s's'}^{\beta} \right) \xi_{\mathbf{k}s's'}^{\beta} \quad (16)$$

3. Installation and usage

3.1. Installation

The BERRY software can be installed using pip package manager with the command

```
pip install berry-suite
```

and all the dependencies will be installed automatically. To install the software with the development dependencies, use:

Table 1
List of packages included in the software.

Package	Description
<code>preprocess</code>	Runs DFT calculations and prepares for the following runs
<code>wfcgen</code>	Reads wavefunctions from DFT results and saves them
<code>dot</code>	Calculates the dot product of the Bloch factor
<code>cluster</code>	Finds the bands
<code>basis</code>	Rotates basis of degenerate states
<code>r2k</code>	Converts from r -space to k -space
<code>geometry</code>	Calculates Berry connections and curvatures
<code>conductivity</code>	Calculates linear optical conductivity
<code>shg</code>	Calculates SHG optical conductivity

```
pip install berry-suite[dev]
```

Alternatively, the software can be downloaded from the GitHub repository, and installed using the pip command

```
pip install -e .
```

A working QUANTUM ESPRESSO installation, version v.6.6 or higher, is required. BERRY was tested with python 3.8 and 3.10, but likely works with other versions of python3.

3.2. Usage

The BERRY software has a Command Line Interface (CLI), that is used to run the packages.

```
berry [package options] package parameter
[package options]
```

where the possible packages are shown in Table 1 and the parameter is dependent on which package is being run. Simply writing `berry` will output the help message. From there, the available options can be seen.

The `preprocess` package is the first one to be run. A working directory has to be created, with a directory called `dft` in it, with the files needed for running QE for the desired system. It needs an input file (given as a parameter) with the number of k -points in each direction and the area in reciprocal space that is covered. It can include some change to the default values of the run.

The rest of the packages should be run in the order given in Table 1 and no input is needed except after the `cluster` package, where the number of the highest band has to be given. This is the highest analytic band that was possible to recover, and can be read from the output of the `cluster` package.

All packages give information about the run in a file named after the name of the package with a `.log` extension. Also, help can be obtained adding `-h` as an option to the package.

3.3. Parallelization

BERRY supports parallel execution of QE tasks on a single or multiple interconnected multicore servers. It relies on QE's MPI-based multi-process parallelization [21] through the configuration of specific tasks based on user input. The number of processes can be specified when running any of BERRY's packages through an input file in the case of `preprocess` and `wfcgen`, or by the `-np` CLI argument for the rest of the packages.

While this is initially the best approach to parallelize BERRY, since it requires a smaller effort for significant gains in performance and ensures that it benefits from future improvements in QE, it may be limiting in some edge cases. Section 5 analyzes the scalability limitations of its parallelization, which could be further improved in the future.

Table 2
List of packages for visualization.

Package	Description
<code>debug</code>	Prints data for debugging
<code>geometry</code>	Draws Berry connection and curvature
<code>wave</code>	Shows the electronic band structure

3.4. Output visualization

A set of packages was created for a quick visualization of the results and debugging. To run them, one has to use the command:

```
berry-vis [-h] package options
```

where the package can be `geometry`, `debug` or `wave`. Table 2 gives a summary of the possibilities.

Each program contains a set of sub-commands that can be used to visualize different aspects of the results. The help message of each package gives a summary of the available options.

4. Examples

The software was tested for two materials, namely single layer hexagonal boron nitride (hBN) and InSe. Fig. 3 shows the eight lower bands of hBN obtained from this software. For both materials, the software was run using different meshes of points in reciprocal space and different energy cutoffs for the DFT runs.

The higher the energy cutoff, the higher the numerical precision of the calculation, since a larger energy cutoff implies a larger base of plane waves to describe the wavefunctions. The numerical precision increases monotonically with the size of the plane wave basis, but also the size of the files that store the wavefunctions.

A higher density of points in reciprocal space allows for a better construction of the analytic band, and a better precision in the calculation of the gradient of the wavefunctions. But the number of wavefunctions is proportional to the number of k -points, and so the amount of calculus needed. Convergence tests are then convenient.

The DFT calculations were performed using QUANTUM ESPRESSO [22,23]. The exchange-correlation functional used was the generalized gradient approximation of Perdew-Burke-Ernzerhof (GGA-PBE) [24] and the integration over the Brillouin-zone was performed using the scheme proposed by Monkhorst-Pack [25] with a grid of $24 \times 24 \times 1$ k -points. A vacuum size of 30 bohr was used to avoid interactions between the periodic images. Γ , the relaxation parameter, is fixed to $\Gamma = 0.01$ Ry = 0.136 eV in all calculations.

We also compared the optical conductivity obtained using the dipole approximation in the QUANTUM ESPRESSO implementation with the one obtained using our software, and the result is the same, as shown in Fig. 4.

4.1. hBN

For hBN scalar-relativistic norm-conserving pseudopotentials [26, 27] were used. The energy cutoff and the density of k -points were varied (the area covered in reciprocal space is always the Brillouin Zone (BZ) or equivalent) with different results for the number of correct bands obtained, as shown in Table 3. It is clear that the density of k -points is important in order to achieve a larger number of successful reconstructed bands. On the other hand, there is no such improvement with energy cutoff, perhaps because the values used were already very high. Also, there is some inconsistency for the lower number of k -points where increasing the energy cutoff doesn't necessarily leads to improvement. Nevertheless, detailed examination of the output of the `cluster` package shows that for $E_{cut} = 120$ Ry the 8th band is correct and there is only one k -point of the 7th band that is estimated to be wrong, and

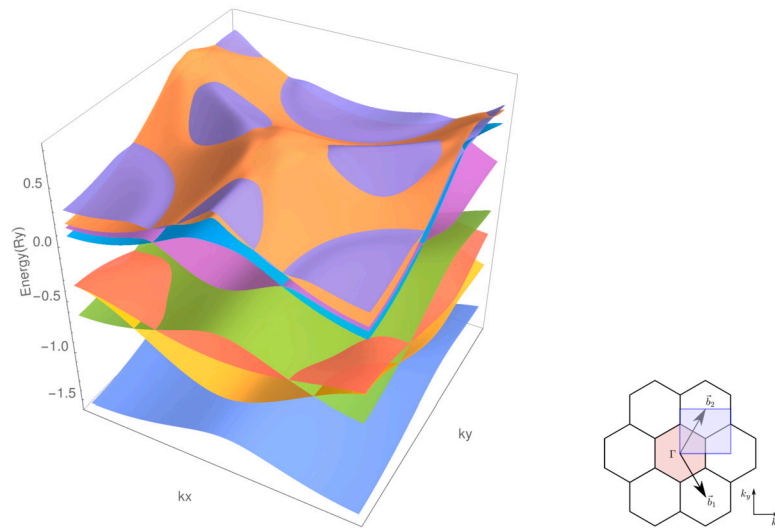


Fig. 3. The first eight bands of hBN (left). k_x runs from Γ -M- Γ and k_y runs from Γ -K-M as shown in the representation of the reciprocal space (right) by the rectangle.




Fig. 4. Real part of the linear optical conductivity along the armchair direction calculated using the Berry software (orange) and using the dipole approximation within the random phase approximation (blue), for hBN. The units of conductivity are Rydberg atomic units.




Fig. 5. Linear optical conductivity of hBN for several energy cutoffs, and 47×41 k -points.

Table 3

Number of correct bands achieved for different values of the energy cutoff E_{cut} and density of k -points. In parentheses is the number actually used for optical conductivity calculation.

k -points	E_{cut} (Ry)	Bands
47×41	80	6
47×41	100	8
47×41	120	6 (8)
47×41	150	8
60×52	80	8 (10)
60×52	100	10
60×52	120	10
60×52	150	10

so we decided to proceed with eight bands. The same happened for the 9th and 10th bands of $E_{cut} = 80$ Ry for the larger density of k -points, and so we proceeded with 10 bands.

Fig. 5 shows the real and imaginary parts of the optical conductivity for 47×41 k -points, while Fig. 6 shows the same for 60×52 k -points. In both cases, the dependence in energy cutoff is minimal. The greatest difference is for $E_{cut} = 80$ Ry and 47×41 k -points but that is because we are using just 6 bands, while the others use 8, as shown in Table 3.




Fig. 6. Linear optical conductivity of hBN for several energy cutoffs, and 60×52 k -points.

In this case, the number of bands affects the height and ratio of the main peaks, but that is natural, and is the only relevant difference between the results shown in Fig. 5 and Fig. 6.

The Second Harmonic Generation optical conductivity is shown in Figs. 7 and 8. The same behavior as for optical conductivity can be seen: the number of bands that were used is what actually makes a difference.




Fig. 7. Second harmonic generation optical conductivity of hBN for several energy cutoffs, and 47×41 k -points.




Fig. 8. Second harmonic generation optical conductivity of hBN for several energy cutoffs, and 60×52 k -points.

The dependence on energy cutoff is small (only noticeable for the lowest value of 80 Ry) and the number of k -points used seems to affect the result indirectly via the number of bands that can be solved rather than directly changing the result. The position of the peaks does not change but their intensity increases with the number of bands, and for higher energies more peaks appear, as expected, since there are more bands with higher energy.

4.2. InSe

For InSe full relativistic norm-conserving pseudopotentials with non-linear core-correction and spin-orbit information were used. Table 4 shows the number of correct bands achieved for two different sets of k -points and two energies. There are 18 electrons in InSe, so 9 bands are full. For 47×41 k -points only two conduction bands were obtained which is a low number. Increasing the number of k -points to 59×51 improved the conduction bands achieved to seven. Yet, energy cutoff lower than 100 Ry lead to less than 9 solved bands, which is useless for the optical conductivity calculations.

Fig. 9 shows the linear optical conductivity of InSe for the five cases shown in Table 4. Naturally, there is a notable difference between the cases where only 11 bands were considered and the ones where 16 bands were used. With 16 bands there are more possibilities of transitions at higher energies, and so more peaks show up. It is important to note the consistency and resilience of the results for the linear optical conductivity: no changes for different energy cutoffs.

Table 4

Number of correct InSe bands achieved for different values of the energy cutoff E_{cut} and density of k -points.

k -points	E_{cut} (Ry)	Bands
47×41	100	11
47×41	120	11
59×51	100	16
59×51	120	16
59×51	150	16




Fig. 9. Linear optical conductivity of InSe for different cutoff energies and density of k -points. Lines are for 47×41 k -points and dots are for 59×51 k -points.




Fig. 10. Second Harmonic Generation optical conductivity for InSe. Lines are for 47×41 k -points and dots are for 59×51 k -points.

Fig. 10 shows the Second Harmonic Generation optical conductivity for InSe for the five cases shown in Table 4.

Again, there is a large difference between having more or less bands, although the position of the main peaks is fairly accurate. This time there are differences at low energies for different energy cutoffs.

5. Computational profile

The computational performance and limitations of BERRY were assessed using the single layer hexagonal boron nitride (hBN) material as a representative case study. A variation of the analysis parameters, such as the number of points in the reciprocal space meshes and the number of bands, allows to manipulate the problem size, i.e. the amount of





Fig. 11. Execution time (left) and speedup (right) of BERRY's packages using up to 64 processes. Speedup of n processes is measure relative to the sequential run.

Table 5

Bandwidth of the different storage devices on the test server.

	Read		Write	
	Bandwidth (MiB/s)	Latency (μ s)	Bandwidth (MiB/s)	Latency (μ s)
NAS	347.2	3771	158.2	4182
SSD	415.3	689	470.4	892
NVMe	3559.2	295	3317.8	348

computation and/or data processed, to evaluate BERRY's subroutines on different scenarios.

BERRY was evaluated in a dual-socket server with two 16-core Intel Xeon E5-2683v4 Broadwell devices @2.1 GHz with 2-way simultaneous multithreading and 256 GiB RAM, connected to three types of storage: a high-performance Network-Attached Storage (NAS) with HDDs and a SSD-based caching system over 10 Gbit ethernet, common in cluster environments; a local Crucial 1050 MX SSD; and a local Samsung 970 Evo Plus NVMe SSD. Table 5 summarizes the relevant performance metrics of the storage devices tested. Bandwidth is more relevant for I/O interaction with larger amounts of data, while latency affects the I/O of random accesses to small amounts of data. The software was tested using CentOS 7.9.2 operating system, with Python 3.10 and Quantum Espresso 7.1. A k -best measurement heuristic was used to ensure that the results can be replicated, with $k = 3$ with a 5% tolerance and a minimum/maximum of 10/20 measurements.

5.1. Performance assessment

Fig. 11 presents the individual performance of BERRY's packages on the test server using the NAS storage over Ethernet, varying the number of parallel processes used to assess its scalability. The hBN material is used with 24×24 k -points, 20 bands, and an energy cutoff of 100 Ry. The *wfc* and *shg* aggregate processing accounts between 76% and 89% of the overall execution time, for 1 and 32 processes respectively. This relative increase is due to *wfc* performance not being affected by the increased amount of parallel processes used, since this package is limited by the storage I/O performance (it performs a small amount of computation per byte read/written). Similar behavior is shown by *cluster* and *basis*, but with a less significant impact on the overall performance due to their faster execution times, accounting at most for 1.2% of BERRY's time. An overview of the relative execution time of each of BERRY's packages using 1 and 32 processes is presented in Table 6.

shg and *conductivity* are the best scaling packages, improving the performance over the sequential code by 28 \times and 17 \times using 32 processes. This would be expected as they are the most computational intensive packages of BERRY. *r2k*, *preprocess*, *geometry*, and *dot* exhibit less significant scalability, with peak speedups around 8.5 \times ,

Table 6

Relative execution time of each of BERRY's packages.

	preprocess	wfc	dot	cluster	basis
sequential	6.31%	21.07%	2.21%	0.39%	0.01%
32 processes	2.29%	81.13%	1.18%	1.20%	0.02%
	r2k	geometry	conductivity	shg	
sequential	5.54%	9.53%	0.24%	54.70%	
32 processes	2.42%	3.83%	0.05%	7.89%	

10.8 \times , 9.2 \times , and 6.9 \times , respectively, for 24 and 32 processes. However, none of these packages effectively use the extra computing resources, with parallelization efficiencies below 50%, with the exception of *shg*. *cluster*, and *wfc* do not scale with the number of processes as they are significantly bound either memory (latency or bandwidth) or I/O performance, thus not benefiting by the extra processes used. None of the packages benefit from using more processes than the available physical cores in the server.

Fig. 12 shows the impact of using a faster SSD or NVMe storage on the performance of BERRY's packages. *r2k*, *cluster*, and *conductivity* are omitted as they did not significantly benefit from the improved storage, as they are compute-bound and have small execution times. The remaining packages benefit from using an SSD over NAS mostly due to the lack of performance of the traditional HDDs in the NAS, with improvements generally between 10% and 30%. *shg* benefit less from using a NVMe over a SSD as the bump in performance provided by the SSD ensures they are mostly compute-bound. *preprocess* and *wfc* greatly benefit from both SSD and NVMe devices, due to their bandwidth intensive operations (large amounts of sequential data being read and written), especially when using a higher amount of processes. *dot* and *geometry* benefit the most from the NVMe due to its latency, as these packages have more random memory access patterns with less data being required per access.

The overall efficiency of BERRY's parallelization is relatively low since it is mostly limited by *wfc*, whose performance does not scale with the number of processes, making it memory-bound when considering all packages. Significant performance improvements may be possible by tackling the I/O bottlenecks limiting this package, especially when using more than 4 processes where BERRY's parallelization efficiency drops below 50%. The use of faster storage also provided improvements for compute-bound packages due to the contention when writing the output data using large number of processes, which may be especially problematic when using HDD-based NAS.





Fig. 12. Speedup of the SSD vs NAS (left) and the NVMe vs SSD (right) storage devices for each of BERRY's packages using up to 64 processes.


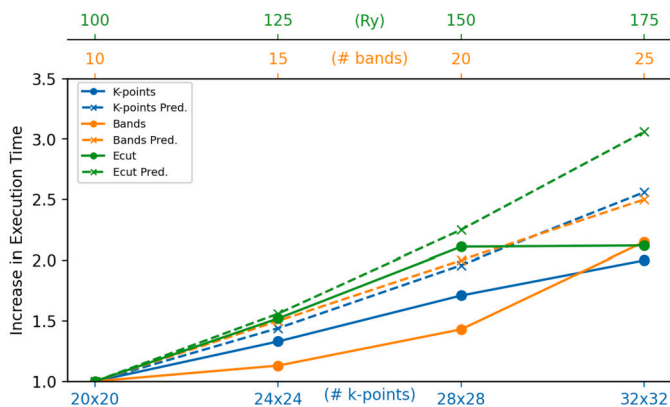


Fig. 13. BERRY's predicted (*Pred*) and measured execution time (left) and output data size (right) for different problem sizes.

5.2. Performance scalability with problem size

This subsection focuses on the analysis of the impact that different input data size and configuration has on the execution time and size of the output data of BERRY, using 32 processes in the test server and SSD storage.

Fig. 13 presents predictions and measurements of the increase in BERRY's execution time, compared to a reference configuration, and storage requirements for three different problem sizes using the hBN material: (i) varying the number of k -points, which has a linear impact on computation; (ii) varying the number of bands to be considered, which also has a linear impact on computation; and (iii) the energy cutoff for the electrons, which should quadratically increase the computing and storage of its data.

Duplicating the number of bands only increases the execution time by 40%, which is significantly below the predicted 2x increase. This behavior is expected since the amount of bands affects especially the `cluster` package, which is compute-bound but only accounts for roughly 1% of BERRY's execution time, and other memory-bound packages to a lower degree, thus mitigating the extra computing required. However, the required storage increases almost as expected. The execution time for the increase in k -points also falls short of the prediction even though the amount of computation required is almost doubled, which is also due to the memory-bound nature of BERRY. The output size increase is almost inline with expectations (7% short of the prediction for 32×32 k -points). A 50% increase in the energy cutoff, from 100 to 150 Ry, results in a quadratic increase in computing time and storage, falling short of the predicted cube increase in both metrics. The shorter increase in BERRY's execution time is again due to it being memory-bound that mitigates the increase in computations. In terms of storage, there is a cubic increase in the data output by `wfc`, whose ef-

fect is diluted by other non-scaling data outputs. Between 150 and 175 Ry, the execution time barely changes; this is because increasing the energy cutoff increases the numerical accuracy of the DFT calculations, and this results in a faster convergence in some programs, specially the `cluster` package. This compensates the increase in time due to the larger amount of calculations in other packages.

These analyses of BERRY's performance can be valuable from the perspective of a user, where the improvements in execution time do not necessarily mean that the time-to-solution will decrease, but rather allow for more complex and/or larger problems to be tackled [28,29]. Assuming the extrapolations of the data in Fig. 13 and considering the best performing configuration presented in subsection 5.1, BERRY is capable of processing either 16×16 k -points, 20 bands more, or an energy cutoff twice higher when using 32 processes within the time frame of a sequential run.

6. Future perspectives

We recognize that BERRY can be improved in both functionality and performance.

The inefficient parallelization of specific packages can be improved, especially targeting I/O- and memory-bound tasks, as it is not currently feasible to analyze 3 dimensional materials in a reasonable time frame. Such improvements should be detached from QUANTUM ESPRESSO to ensure compatibility with previous and future versions of this software. There will also be improvements in the determination of bands.

Another important feature is to be able to use noncolinear DFT calculations. This will allow a large expansion on the number of materials that can be dealt with, in a rigorous manner.

Other properties will be implemented, like excitons and Hall conductivity, and higher order non-linear optical properties.

7. Conclusions

We developed a software that successfully generates sets of eigenstates that are differentiable in reciprocal space, departing from purely numerical density functional calculations.

Several hurdles had to be overcome to achieve this, namely establishing a procedure to solve the problem of random phases in numerical calculations of electronic structure due to the independence of basis at different \mathbf{k} points, and thus making all wavefunctions coherent; finding degenerate points and correcting the basis to achieve continuity; and sorting out the eigenstates in analytic sets. This last step involved the use of Graph Theory and Unsupervised Machine Learning in an iterative process.

The solution allowed us to calculate the gradients of the Bloch factor in reciprocal space, and the Berry connections and curvatures from first principles. We then used this knowledge for the calculation of optical response in linear order and the second harmonic generation in hBN and InSe. We were able to identify the density of \mathbf{k} -points sampled to be the most important factor in achieving a larger number of corrected bands, while the energy cutoff of the DFT calculations is less important.

This software opens huge possibilities in the study of responses in 2D crystals as well as other properties that may be calculated from the eigenstates along bands in reciprocal space.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

We acknowledge support by the Portuguese Foundation for Science and Technology (FCT) in the framework of the Strategic Fundings UIDB/04650/2020, UIDB/50007/2020 and Associate Laboratory LA/P/0016/2020 and COMPETE 2020, PORTUGAL 2020, FEDER and FCT through project PTDC/FIS-MAC/2045/2021 and from the European Commission through the project Graphene Driven Revolutions in ICT and Beyond (Ref. No. 881603, CORE 3).

References

- [1] G.B. Ventura, D. Passos, J.M.V.P. Lopes, J.M.B.L. dos Santos, A study of the nonlinear optical response of the plain graphene and gapped graphene monolayers beyond the Dirac approximation, *J. Phys. Condens. Matter* (dec 2019), <https://doi.org/10.1088/1361-648x/ab262bc>.
- [2] R. Resta, *Theory of the electric polarization in crystals*, *Ferroelectrics* 136 (1) (1992) 51–55.
- [3] J. Zak, Berry's phase for energy bands in solids, *Phys. Rev. Lett.* 62 (1989) 2747–2750, <https://doi.org/10.1103/PhysRevLett.62.2747>.
- [4] G. Strinati, Multipole wave functions for photoelectrons in crystals. iii. the role of singular points in the band structure and the tails of the Wannier functions, *Phys. Rev. B* 18 (1978) 4104–4119, <https://doi.org/10.1103/PhysRevB.18.4104>.
- [5] K.S. Virk, J.E. Sipe, Semiconductor optics in length gauge: a general numerical approach, *Phys. Rev. B* 76 (2007) 035213, <https://doi.org/10.1103/PhysRevB.76.035213>.
- [6] T. Fukui, Y. Hatsugai, H. Suzuki, Chern numbers in discretized Brillouin zone: efficient method of computing (spin) Hall conductances, *J. Phys. Soc. Jpn.* 74 (6) (2005) 1674–1677, <https://doi.org/10.1143/JPSJ.74.1674>.
- [7] N. Marzari, D. Vanderbilt, Maximally localized generalized Wannier functions for composite energy bands, *Phys. Rev. B* 56 (1997) 12847–12865, <https://doi.org/10.1103/PhysRevB.56.12847>.
- [8] F. Hipolito, T.G. Pedersen, V.M. Pereira, Nonlinear photocurrents in two-dimensional systems based on graphene and boron nitride, *Phys. Rev. B* 94 (4) (2016) 045434, <https://doi.org/10.1103/PhysRevB.94.045434>.
- [9] A. Taghizadeh, F. Hipolito, T.G. Pedersen, Linear and nonlinear optical response of crystals using length and velocity gauges: effect of basis truncation, *Phys. Rev. B* 96 (2017) 195413, <https://doi.org/10.1103/PhysRevB.96.195413>.
- [10] L.G. Ferreira, N.J. Parada, Wannier functions and the phases of the Bloch functions, *Phys. Rev. B* 2 (6) (1970) 1614–1618, <https://doi.org/10.1103/PhysRevB.2.1614>.
- [11] F.B. Pedersen, G.T. Einevoll, P.C. Hemmer, Wannier functions for the Kronig-Penney model, *Phys. Rev. B* 44 (1991) 5470–5475, <https://doi.org/10.1103/PhysRevB.44.5470>.
- [12] M. Newman, *Networks*, Oxford University Press, 2018.
- [13] A.A. Hagberg, D.A. Schult, P.J. Swart, Exploring network structure, dynamics, and function using networkx, in: G. Varoquaux, T. Vaught, J. Millman (Eds.), *Proceedings of the 7th Python in Science Conference, Pasadena, CA USA, 2008*, pp. 11–15.
- [14] M. Alloghani, D. Al-Jumeily, J. Mustafina, A. Hussain, A.J. Aljaaf, *A Systematic Review on Supervised and Unsupervised Machine Learning Algorithms for Data Science*, Springer International Publishing, 2020.
- [15] K.P. Sinaga, M.-S. Yang, Unsupervised k-means clustering algorithm, *IEEE Access* 8 (2020) 80716–80727, <https://doi.org/10.1109/ACCESS.2020.2988796>.
- [16] M. Gradhand, D.V. Fedorov, F. Pientka, P. Zahn, I. Mertig, B.L. Györfly, First-principle calculations of the Berry curvature of Bloch states for charge and spin transport of electrons, *J. Phys. Condens. Matter* 24 (21) (2012) 213202, <https://doi.org/10.1088/0953-8984/24/21/213202>.
- [17] C. Aversa, J.E. Sipe, Nonlinear optical susceptibilities of semiconductors: results with a length-gauge analysis, *Phys. Rev. B* 52 (1995) 14636–14645, <https://doi.org/10.1103/PhysRevB.52.14636>.
- [18] J.E. Sipe, A.I. Shkrebti, Second-order optical response in semiconductors, *Phys. Rev. B* 61 (2000) 5337–5352, <https://doi.org/10.1103/PhysRevB.61.5337>.
- [19] G.B. Ventura, D.J. Passos, J.M.B. Lopes dos Santos, J.M. Viana Parente Lopes, N.M.R. Peres, Gauge covariances and nonlinear optical responses, *Phys. Rev. B* 96 (2017) 035431, <https://doi.org/10.1103/PhysRevB.96.035431>.
- [20] N.E. Brener, Random-phase-approximation dielectric function for diamond, with local field effects included, *Phys. Rev. B* 12 (4) (1975) 1487–1492, <https://doi.org/10.1103/PhysRevB.12.1487>.
- [21] P. Giannozzi, O. Andreussi, T. Brumme, O. Bunau, M.B. Nardelli, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, M. Cococcioni, N. Colonna, I. Carnimeo, A.D. Corso, S. de Gironcoli, P. Delugas, R.A. DiStasio, A. Ferretti, A. Floris, G. Fratesi, G. Fugallo, R. Gebauer, U. Gerstmann, F. Giustino, T. Gorni, J. Jia, M. Kawamura, H.-Y. Ko, A. Kokalj, E. Küçükbenli, M. Lazzeri, M. Marsili, N. Marzari, F. Mauri, N.L. Nguyen, H.-V. Nguyen, A.O. de-la Roza, L. Paulatto, S. Poncè, D. Rocca, R. Sabatini, B. Santra, M. Schlipf, A.P. Seitsonen, A. Smogunov, I. Timrov, T. Thonhauser, P. Umari, N. Vast, X. Wu, S. Baroni, Advanced capabilities for materials modelling with quantum espresso, *J. Phys. Condens. Matter* 29 (46) (oct 2017), <https://doi.org/10.1088/1361-648X/aa8f79>.
- [22] P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G.L. Chiarotti, M. Cococcioni, I. Dabo, et al., Quantum espresso: a modular and open-source software project for quantum simulations of materials, *J. Phys. Condens. Matter* 21 (39) (2009) 395502, <https://doi.org/10.1088/0953-8984/21/39/395502>.
- [23] P. Giannozzi, O. Andreussi, T. Brumme, O. Bunau, M. Buongiorno Nardelli, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, M. Cococcioni, et al., Advanced capabilities for materials modelling with quantum espresso, *J. Phys. Condens. Matter* 29 (46) (2017) 465901, <https://doi.org/10.1088/1361-648X/aa8f79>.
- [24] J.P. Perdew, K. Burke, M. Ernzerhof, Generalized gradient approximation made simple, *Phys. Rev. Lett.* 77 (18) (1996) 3865–3868.
- [25] H.J. Monkhorst, J.D. Pack, Special points for Brillouin-zone integrations, *Phys. Rev. B* 13 (12) (1976) 5188.
- [26] D.R. Hamann, Optimized norm-conserving Vanderbilt pseudopotentials, *Phys. Rev. B* 88 (2013) 085117, <https://doi.org/10.1103/PhysRevB.88.085117>.
- [27] M. Schlipf, F. Gygi, Optimization algorithm for the generation of oncv pseudopotentials, *Comput. Phys. Commun.* 196 (2015) 36–44, <https://doi.org/10.1016/j.cpc.2015.05.011>.
- [28] J.L. Gustafson, Reevaluating Amdahl's law, *Commun. ACM* 31 (1988) 532–533, <https://doi.org/10.1145/424111.424115>.
- [29] M.D. Hill, M.R. Marty, Amdahl's law in the multicore era, *Computer* 41 (7) (2008) 33–38, <https://doi.org/10.1109/MC.2008.209>.