

Uncoupling System and Environment Simulation Cells for Fast-Scaling Modeling of Complex Continuum Embeddings

G. Medrano¹, E. Bainglass², and O. Andreussi³

¹Department of Physics, University of North Texas, Denton, Tx, 76203, USA

²Materials Software and Data Group, Paul Scherrer Institut, 5232 Villigen PSI, Switzerland

³Department of Chemistry and Biochemistry, Boise State University, Boise, Id, 83725, USA

***Author to whom correspondence should be addressed:**

olivieroandreuss@boisestate.edu

March 2, 2024

Abstract

Continuum solvation models are becoming increasingly relevant in condensed matter simulations, allowing to characterize materials interfaces in the presence of wet electrified environments at a reduced computational cost with respect to all atomistic simulations. However, some challenges with the implementation of these models in plane-wave simulation packages still persists, especially when the goal is to simulate complex and heterogeneous environments. Among these challenges is the computational cost associated with large heterogeneous environments, which in plane-wave simulations has a direct effect on the basis-set size and, as a result, on the cost of the electronic structure calculation. Moreover, the use of periodic simulation cells are not well-suited for modeling systems embedded in semi-infinite media, which is often the case in continuum solvation models. To address these challenges, we present the implementation of a double-cell formalism, in which the simulation cell used for the continuum environment is uncoupled from the one used for the electronic-structure simulation of the quantum-mechanical system. This allows for a larger simulation cell to be used for the environment, without significantly increasing computational time. In this work, we show how the double-cell formalism can be used as an effective periodic boundary conditions (PBC) correction scheme for non-periodic and partially periodic systems. The accuracy of the double-cell formalism is tested using representative examples with different dimensionalities, both in vacuum and in a homogeneous continuum dielectric environment. Fast convergence and good speedups are observed for all the simulation setups, provided the quantum-mechanical simulation cell is chosen to completely fit the electronic density of the system.

1 Introduction

Materials simulations have seen an incredible growth in recent years, thanks to the increasing power of computer hardware and simulation software, together with the development of new computational infrastructures to handle high-throughput simulations [1, 2, 3, 4, 5]. Unbiased first-principles simulations based on density functional theory (DFT) represent the work-horse of most computational approaches to materials design [6], with many research groups pushing to improve the accuracy of computational predictions and the extend the scope of first-principles simulations to more challenging systems and properties.

Traditional DFT simulations of materials focus on bulk periodic crystalline structures. For this reason, most condensed-matter and materials simulation packages rely on basis sets composed of periodic functions (e.g. plane-waves, PW) which uniformly occupy the simulation cell. This approach provides a systematically improvable basis set, whose accuracy can be tuned with a single parameter, usually associated with the kinetic energy of the corresponding plane-wave. PWs also provide easy

access to the solution of the core electrostatic equations involved in first-principles simulations, with the Poisson equation in vacuum being readily solved in reciprocal space. Fast Fourier transforms (FFTs)[\[7\]](#) have proven instrumental to the wide-spread success and fast scaling of PW-based DFT packages, by providing real-to-reciprocal space mapping, and vice-versa, at a cost proportional to $N_{pw} \log(N_{pw})$, with N_{pw} being the number of PWs in the simulation. However, what is advantageous for bulk crystalline materials can be a source of artifacts when non-periodic or partially periodic systems need to be characterized. Traditionally, materials interfaces are simulated using the slab approach, in which a few-atoms thick slab of the material is described in a 3-dimensional periodically repeated simulation cell. By increasing the thickness of the slab and by increasing the cell size in the direction orthogonal to the material interface, the periodic replicas of the slab are effectively uncoupled from each other and the results are considered to be representative of the surface of a semi-infinite bulk material. A similar strategy can be devised for systems that are only periodic in one dimension, such as nano-wires or nanotubes, or that are non-periodic. While adjusting the cell size can provide an easy solution to periodic boundary conditions (PBC) artifacts, the number of PWs grows linearly with the cell volume, leading to a polynomial increase in the cost of the corresponding DFT calculation. For this reason, multiple alternative strategies to correct for PBC artifacts have been proposed in the past years, including real-space and reciprocal-space approaches. Most of these strategies are focused on simulations of two-dimensional and zero-dimensional (isolated) systems, while one-dimensional correction schemes are usually less wide-spread.

Unfortunately, as the typical scaling of DFT simulations varies as N^3 , where N is the number of electrons in the system, complex and large systems with more than a few hundreds of electrons are still out of reach for systematic DFT studies. To overcome this limitation, a variety of hierarchical approaches [\[8, 9\]](#) and/or divide and conquer[\[10, 11, 12\]](#) strategies have been developed in the literature. Without sacrificing the important atomistic characteristics of the relevant part of the process under investigation (the system), these methods take advantage of different, possibly simplified, models to handle the more macroscopic or complex part of the process (the environment), thus allowing a significant reduction of the computational cost.

In particular, implicit solvation models based on continuum embedding media have proven to provide a reasonable qualitative accuracy, while significantly reducing the need for statistical sampling of disordered configurations [\[13, 14, 15, 16\]](#). Indeed, removing the atomistic details of solvent molecules not only reduces the number of electrons involved in the DFT simulation, but it allows to avoid the use of molecular dynamics (MD) to sample the configurations of the molecules of the liquid. Continuum methods in computational chemistry usually involve more or less empirical definitions of the different contributions to solvation free energy, which are then parameterized on experimental databases generated from solubility data. Methods such as the Polarizable Continuum Model (PCM)[\[17\]](#) of Tomasi and co-workers have shown a good success in predicting ground state and response properties of molecules in solution. More recently, similar approaches have been translated into condensed-matter simulation packages in order to study wet, possibly electrified, interfaces of materials[\[8\]](#). When compared to the more widespread continuum solvation models in the computational chemistry community, the models implemented in periodic PW-based simulation packages rely on a smoothly-varying boundary between the quantum-mechanical system and the continuum environment, instead of a sharp two-dimensional interface. Definition of the interface based off the electronic density of the solute[\[18\]](#) or its atomic positions[\[19\]](#) have been developed and show similar accuracy, when properly parameterized. While computationally more demanding, the use of smooth interfaces allows continuum models in condensed matter to more seamlessly introduce non-local corrections [\[20, 21\]](#) to avoid some of the typical artifacts of these solvation methods.

Beyond the use of continuum solvation to model bulk neutral solutions, extensions to handle liquid crystals and diluted ionic solutions were implemented exploiting the integral equation formalism (IEF) of PCM[\[22\]](#). Still following the same strategy, PCM was extended to the study of a single two-dimensional interface between different dielectric media (liquid-air, liquid-liquid)[\[23\]](#). Corni et al. extended the use of the polarizable dielectric model to handle the electrostatic effects of metal surfaces and nanoparticles on nearby molecular dyes, thus allowing a more quantum-mechanical characterization of surface-enhanced spectroscopies[\[24, 25, 26\]](#).

Similar to the PCM literature, more complex and heterogeneous environments can be introduced in condensed-matter simulations. In particular, a significant focus has been devoted to study electrolyte solutions [\[23, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36\]](#), so as to unlock the use of these models for simulations

of electrified interfaces. Going beyond neutral solutions and electrolyte distributions, Campbell et al. [37, 38] exploited a continuum approach to model the charge reorganization in a semiconductor substrate.

As the definition of the continuum environment in smooth-interface solvation models follows a function defined everywhere in the simulation cell, it is relatively straightforward to introduce more complex and heterogeneous environments, e.g. by introducing multiple continuum media, each with its own boundaries and physical properties (e.g. dielectric constant or surface tension). These media can be used to model substrate effects on the properties of overlaying materials or molecular compounds. Following this idea, Bononi and collaborators[39, 40] exploited the possibility to introduce flat two-dimensional dielectric regions in the simulation cell in order to characterize in a computationally inexpensive way the effect of an ice substrate on the absorption properties and photo-degradation of molecular dyes.

Despite the flexibility of continuum embedding approaches in condensed matter environments, PW-based simulations can pose significant challenges to these class of multiscale methods. The use of periodic simulation cells is intrinsically unfit to model systems that are embedded in a semi-infinite medium, i.e. most solvated systems would typically require non-periodic simulation cells in the directions where the continuum solvent resides. Moreover, the fact that the simulation cell size is directly related to the basis-set size N_{pw} makes simulating large heterogeneous environments computationally more challenging. Extensions of both real-space and reciprocal-space PBC correction schemes have been developed to account for continuum dielectric media in simulations of isolated (0D) and slab (2D) systems[41]. In both continuum electrolyte and semiconductor models, long-range charge reorganizations are handled implicitly, by relying on analytical solutions to the electrostatic problems and/or appropriate boundary conditions. However, simulating the effects of nano-sized environments, such as plasmonic nanoparticles, micelles, liquid nanodroplets, would require to use simulations cells, and thus basis-sets, that are beyond reach of standard computational resources.

The straightforward solution to the aforementioned problem is to uncouple the simulation cell exploited for the continuum environment from the one used for the DFT simulation of the quantum-mechanical system. The use of a large simulation cell for the environment would not cause significant limitations to the simulation time: the time-consuming part of the continuum environment calculation is related to computing the electrostatic potentials, by solving modified and generalized forms of the Poisson or Poisson-Boltzmann equations. Most of the algorithms implemented rely on iterative strategies in which each iteration involve the solution of a simple Poisson equation, which can be performed via FFTs in a time almost linear in the cell size. By keeping the simulation cell of the DFT calculation small, the largest overhead in computational time is avoided. While ideal in theory, in practice this double-cell formalism requires to fully uncouple the two simulation methods and to carefully design the mapping between the two simulation cells.

While designed to account for large heterogenous environments, the double-cell formalism can also be used for non-periodic and partially periodic systems as an effective PBC correction scheme. By increasing the length of the lattice vectors going along the non-periodic directions in the environment cell, electrostatics can be fully converged without affecting the basis-set size and cost of the quantum-mechanical calculation. This general strategy works equally well for isolated (0D), slab (2D), but also one-dimensional (1D) systems. In this work, the details of the double-cell methodology are presented and the approach is tested as a PBC correction scheme for some representative examples with different dimensionalities, both in vacuum and in a continuum dielectric environment. Along with this new methodology and in order to test the accuracy and the speed-up of the double-cell formalism for PBC corrections, the auxiliary-function countercharge corrections library (LIBAFCC) of Dabo et al.[42] was coupled to the continuum embedding software, giving access to an alternative exact correction scheme for all types of partially periodic systems.

The paper is organized as follows: in Section 2 we review the main details of the implementation of the double cell formalism, as well as the coupling of the auxiliary-function countercharge corrections library (LIBAFCC)[42]; in Section 3 the computational details of the DFT calculations as well as how the double cell calculations are reported. Eventually, in Section 4 we report a detailed benchmark of the accuracy and performance of the double cell approach for modelling partially periodic systems, including 0D, 1D, and 2D systems, in vacuum and in a uniform dielectric medium.

2 Methods

2.1 Double-Cell Formalism




Figure 1: A schematic view of the double cell formalism with an isolated water molecule. The green cell is the system cell where all the PWs are introduced while the blue cell is the environment cell that are used for PBC correction.

The design of a double-cell algorithm is strongly dependent on the numerical details that are associated with the calling program (e.g. the PW DFT simulation package), the continuum embedding library, and with their coupling. In particular, the key aspects that affect the implementation are:

1. which scalar fields need to be passed from one cell to the other
2. how the scalar fields are stored and parallelized in the two cells
3. the way the gridpoints of the two cells are aligned

In the following we will refer to the specific implementation involving Quantum Espresso as the DFT driver and Environ as the continuum embedding library. A schematic representation of the steps involved in the algorithm is visualized in Figure 1, while a more detailed example in a simplified two-dimensional orthorhombic cell is reported in Figure 2. In order to generate the continuum boundary in Environ, information on the electronic density of the system and/or its atomic positions are required. The same information is also needed for the calculation of the electrostatic potential, i.e. the position and valence of the ions as well as the full electronic density of the system are required inputs for an Environ calculation. In alternative polarizable dielectric models, such as PCM, the electrostatic potential of the system in vacuum is passed to the continuum embedding module, to be used to solve for the polarization charge on the environment. However, passing the full information of the charge of the solute allows to have a consistent definition of the electrostatic potential, e.g. when smeared ionic charges are used to account for core-electrons. In the case of the double-cell formalism, passing the full charge density of the solute also allows to seamlessly adapt the calculation to a different simulation cell. In the presented implementation, the key scalar quantity that needs to be passed from the system cell into the environment cell is thus the electronic density of the system.

The results of the continuum embedding calculations are usually reported in terms of corrections to the corresponding quantities computed in vacuum and periodic boundary conditions by the host DFT program. In particular, corrections to the total energy, interatomic forces, and Kohn-Sham potential are the key outputs of Environ calculations. Of these quantities, only the latter needs to be mapped from the environment cell into the system cell.

Scalar fields in a 3D simulation cell are discretized in PW simulation packages in terms of their values on a 3D structured grid, whose spacing is inversely related to the PW cutoff specified in the simulation input. However, the values on the 3D grid are stored by Quantum-Espresso as one-dimensional arrays, by scanning grid-points sequentially along the three cell axis. As part of the hierarchy of parallelization schemes adopted by Quantum-Espresso, parallelization on the real-space simulation cell is implemented by dividing it into slices along the third axis or into sticks along the second and third




Figure 2: Visualization of the mapping algorithm in a simplified two-dimensional rectangular grid. Left panels: a single water molecule and its electronic density (dark green circle) are mapped on a periodically repeated system cell (black rectangle), with a grid size of 10 times 7 points along the horizontal (x) and vertical (y) axes, respectively. A generic gridpoint (the pink diamond) can be identified by its pair of indexes (ix and iy) or by a combined single index (i). A scalar field in the cell is thus represented as a one-dimensional array with index i ranging from 1 to $nx \times ny$. Central panels: a re-centering of the system cell on the center of mass of the system is required to ensure that the electronic density of the system is fully connected when the cell is expanded into the environment cell. The position of the center of mass with respect to the center of the system cell (black arrow) is used to determine the shift (green arrow) of the origin of the re-centered cell (green rectangle). To ensure that gridpoints of the two cells are at the same physical positions, integer operations (NINT and FLOOR) are adopted in the algorithm, resulting in an origin shift that may be slightly different from the real-valued shift of the center of mass of the system: the green and black arrows are meant to be different. While the number of gridpoints in the centered cell is the same as the system cell, their indexes will be different, possibly including the fact that a gridpoint from a periodic image of the system cell needs to be selected. In the example, the original diamond is now replaced by one of its periodic images, its new ix and iy indexes give rise to a new global index i , with the result that each entry in the original scalar field array is now mapped to a different entry of the re-centered array in a one-to-one fashion. Right panels: the centered cell is expanded in the two directions by one replica per side, for a total environment cell nine times larger than the original system cell. The new gridpoint indexes need to account for the additional gridpoints in the replicas and the global gridpoint index is computed using the expanded grid size. While each gridpoint index in the system cell has a corresponding location in the environment cell, the vice-versa is not true, as most of the gridpoints in the environment cell are not meant to be mapped back to the system cell.

axes, depending on the total number of available processors. This allows to perform operations that only depend on local values of the scalar field in a fully parallelized way, while still requiring gathering and scattering of the whole simulation data for real to reciprocal-space transformations (FFTs and

inverse FFTs). Owing to the origin of Environ as an internal plugin of Quantum Espresso, the same strategies and the same numerical libraries are exploited by Environ for its internal calculations. For sake of simplicity in the equations, the following discussion will only focus on the serial algorithm, but the developed strategy is seamlessly extended to parallel implementations.

In practice, given a simulation cell with cell axes

$$\vec{h}_1 = \begin{pmatrix} h_{1x} \\ h_{1y} \\ h_{1z} \end{pmatrix}, \quad \vec{h}_2 = \begin{pmatrix} h_{2x} \\ h_{2y} \\ h_{2z} \end{pmatrix}, \quad \vec{h}_3 = \begin{pmatrix} h_{3x} \\ h_{3y} \\ h_{3z} \end{pmatrix}, \quad (1)$$

and a corresponding cell matrix

$$\hat{h} = (\vec{h}_1, \vec{h}_2, \vec{h}_3) = \begin{pmatrix} h_{1x} & h_{2x} & h_{3x} \\ h_{1y} & h_{2y} & h_{3y} \\ h_{1z} & h_{2z} & h_{3z} \end{pmatrix} \quad (2)$$

the choice of the density cutoff in the simulation input determines the number of the gridpoints along the three axes, $\mathbf{n}^s = (n_1^s, n_2^s, n_3^s)$. A generic scalar property (e.g., the electronic density or the electrostatic potential) in the simulation cell, $s(\vec{r})$, is thus discretized into an array, $s(\mathbf{i}^s) \equiv s(i_1^s, i_2^s, i_3^s) \equiv s(i^s)$ that contains the values of the property at each gridpoint's position, $\vec{r}(\mathbf{i}^s)$, with

$$\vec{r}(\mathbf{i}^s) = \hat{h} \cdot ((\mathbf{i}^s)/\mathbf{n}^s) \equiv \begin{pmatrix} h_{1x} & h_{2x} & h_{3x} \\ h_{1y} & h_{2y} & h_{3y} \\ h_{1z} & h_{2z} & h_{3z} \end{pmatrix} \cdot \begin{pmatrix} i_1^s/n_1^s \\ i_2^s/n_2^s \\ i_3^s/n_3^s \end{pmatrix}, \quad (3)$$

where the integer gridpoint indexes i_{1-3}^s range from 0 to $n_{1-3}^s - 1$. A one-dimensional version of the array, with $i^s = 0, \dots, n_1^s n_2^s n_3^s - 1$, is obtained by combining the three indexes, e.g., by using the expression

$$i^s = 1 + i_1^s + i_2^s n_1^s + i_3^s n_1^s n_2^s, \quad (4)$$

which assumes that the scanning of the gridpoints follows the order of the axes, i.e., faster along the first axis, than along the second, and slower on the third. Vice-versa, given the global index i^s of a gridpoint, its three-dimensional indexes can be reconstructed as

$$i_3^s = \frac{i^s - 1}{n_1^s n_2^s}, \quad (5)$$

$$i_2^s = \frac{i^s - 1 - i_3^s n_1^s n_2^s}{n_1^s}, \quad (6)$$

$$i_1^s = i^s - 1 - i_3^s n_1^s n_2^s - i_2^s n_1^s, \quad (7)$$

where the fractions are intended to represent integer division operations.

In designing the coupling between system and environment simulation cell, it is crucial to realize that the quantum-mechanical system lives in a fully periodic simulation cell: in the DFT calculation the system energy and properties are not affected by arbitrary translations of the system's degrees of freedom. In particular, it is possible for a system to have some of its atoms or part of its electronic density on opposite sides of the simulation cell (e.g., as schematized in the left panel of Figure 2), as the minimum image convention would still give rise to a well-defined fully connected system. However, when expanding the system cell into the environment cell, care must be taken to ensure that a fully connected system is passed to the continuum embedding module and that the environment is properly defined with respect to a meaningful center of the system. Thus, for the mapping of scalar fields between the two simulation cells, a two step process was implemented as visualized in Figures 1 and 2: first a centering of the system and second an expansion of the environment cell along the designated axes.

The first step is accomplished by computing the center of mass, \vec{r}_{cm} , of the system within its original simulation cell, using the minimum image convention to account for potential splits in neighbouring cells. We then introduce a new cell with the same cell parameters and the same grid spacing as the system cell, but with an origin shifted from the Cartesian origin by a vector \vec{o} , so as to place the system's center of mass near its central gridpoint. In practice, we set the indexes of a central gridpoint as $\mathbf{c} = \text{NINT}(\mathbf{n}^s/2)$, with the NINT() operation selecting the nearest integer to the result of the

division, in case of grids with odd numbers of gridpoints along an axis. We then select a gridpoint near the center of mass of the system as

$$\mathbf{m} = \text{NINT}(\hat{h}^{-1} \cdot \vec{r}_{cm} \cdot \mathbf{n}^s). \quad (8)$$

The integer shift between the indexes of the center of mass and the indexes of the center of the cell, $\mathbf{d} = \mathbf{c} - \mathbf{m}$, provides the origin vector for the new re-centered cell,

$$\vec{o} = -\hat{h} \cdot (\mathbf{d}/\mathbf{n}^s) = - \begin{pmatrix} h_{1x} & h_{2x} & h_{3x} \\ h_{1y} & h_{2y} & h_{3y} \\ h_{1z} & h_{2z} & h_{3z} \end{pmatrix} \cdot \begin{pmatrix} d_1/n_1^s \\ d_2/n_2^s \\ d_3/n_3^s \end{pmatrix}, \quad (9)$$

as well as the shift used to remap the integer gridpoint indexes, i_{1-3}^c , needed to compute the single-index representation in Eq. 4. However, in applying the shift to the gridpoint indexes we need to make sure that gridpoints in periodic images are mapped to the corresponding points in the re-centered cell (e.g. as visualized by the pink star in the central panel of Figure 2). In the algorithm this is enforced as follows

$$i_{1-3}^c = (i_{1-3}^s + d_{1-3}) - \text{FLOOR}((i_{1-3}^s + d_{1-3})/n_{1-3}^s) \times n_{1-3}^s, \quad (10)$$

where the FLOOR() integer operation returns the greatest integer less than or equal to the argument. As defined above, the re-centering of the cell corresponds to a one-to-one mapping of gridpoints, which is univocally defined by the position of the center of mass and the grid dimensions. Such a mapping does not involve any interpolation and can be computed on the fly without loss of information and without affecting the real-space parallelization scheme. In particular, given an initial gridpoint in the system cell with global index i^s , its corresponding one-dimensional index in the re-centered cell is given by

$$i^c = 1 + i_1^c + i_2^c n_1^s + i_3^c n_1^s n_2^s, \quad (11)$$

where i_{1-3}^c are computed by applying Eq. 10 to the indexes obtained from Eq. 5.

The second step of the double-cell algorithm involves the expansion of the system cell in the different directions in order to generate a larger environment cell. While arbitrary expansions could in principle be considered, they would usually result in an environment grid spacing and gridpoint positions that would not match the initial ones on the system cell. This numerical mismatch could be easily addressed by interpolation algorithms, which would add an additional layer in the mapping between the two cells. However, in order to keep the process as simple and robust as possible, the current implementation of the double-cell formalism restricts the cell expansion to integer multiples of the initial cell. An expansion vector $\mathbf{n} = (n_1, n_2, n_3)$ is defined in the input of the environment module, with n_i representing the number of fictitious replicas added on both sides of the re-centered system cell along its i -th axis. The combination of these two steps creates the desired environment cell, with lattice vectors 2 \mathbf{n} + 1 times as large as the system cell, with the center of mass of the system located near its central gridpoint. In practice, the gridpoint indexes in the environment cell, i_{1-3}^e need to account for the additional replicas added before the cell:

$$i_{1-3}^e = i_{1-3}^c + n_{1-3} n_{1-3}^s. \quad (12)$$

The corresponding global gridpoint index can thus be computed as

$$i^e = 1 + i_1^e + i_2^e n_1^e + i_3^e n_1^e n_2^e, \quad (13)$$

with $n_{1-3}^e = 2(n_{1-3} + 1)n_{1-3}^s$.

It is important to stress that in the above algorithm each gridpoint of the system cell has a corresponding gridpoint of the environment cell at the same position: the system-to-environment mapping of a smooth charge density only involves a re-indexing of the entries of the array in the system cell and a zeroing of the elements of the environment cell that lie outside of the system cell. Vice-versa, the environment-to-system mapping of the electrostatic potential involves the re-indexing of the entries of the array in the environment cell that only correspond to gridpoints of the system cell, while the information stored in the other gridpoints is discarded. This corresponds to a mapped electrostatic potential that, in general, will not possess the periodicity of the system cell and will likely present discontinuities at the cell borders corresponding to the axes that were expanded in the environment cell. This is consistent with other real-space PBC correction schemes, such as the

parabolic/dipole correction for isolated and slab systems [41]. The presence of such a discontinuity will not affect the energy and convergence of the simulated system, provided that the atomic and electronic charges of the system are not present at the borders of the cell. While the re-centering step of the algorithm enforces that the borders are as far as possible from the system's center of mass, for small system cells it is possible that part of the charges of the system still overlaps or gets close to the borders of the cell. Convergence of the double-cell algorithm with system-cell size can provide the necessary validation of the approach, as shown in the Results sections.

2.2 Auxiliary-Function Correction

The long-range nature of Coulomb interactions leads to a non-convergent or ill-convergent result for charged or dipolar systems, when summing over all the periodic replicas of a infinite 3D periodic system[43]. The slow non-integrable behavior of the Coulomb potential at long distances has a mirror counterpart in its singularity at vanishing distances, which affects the calculation of electron interactions and exact exchange terms in periodic DFT simulations [44, 45]. These singularities occur when the reciprocal-space vectors approach the Brillouin zone boundaries of partially periodic systems or when the interaction potentials exhibit rapid variations. Most simulation softwares set these divergences to zero, introducing negligible errors that approach zero as the size of the computational supercell increases [43]. Meanwhile others employ special techniques to mitigate these divergences, such as implementations of real-space countercharge corrections schemes[46, 47, 48, 49, 50, 51]. Even though the countercharge correction schemes were created as a method for charged systems, they can still be used with neutral systems for a more comprehensive treatment of long-range electrostatic interactions. Example calculations, both neutral and charged, can be found in Ref. [51]. As an alternative, real-space electrostatic corrections have been introduced that can remove these singularities although they can become costly due to these corrections requiring different interaction potentials per self-consistent iteration.

Because of this additional computational cost, reciprocal-space countercharge corrections have been studied in the form of auxiliary-function techniques which are the basis of the LIBAFCC[42] library. These corrections involve adding the auxiliary-function to the point-charge electrostatic kernel which allows for efficient removal of the singularities associated with the plane-wave representation. The LIBAFCC library implements exact point-charge reciprocal-space auxiliary functions for all types of partially periodic systems with the 0D correction giving the same formulas as the Martyna-Tuckerman[52] (MT) correction scheme.

Following an approach similar to the MT implementation, we coupled the library of Li and Dabo to a development version of Environ. The point-charge correction, as computed by the library, is mapped in the real-space grid, centered on the origin. Fast Fourier transforms allow to convert the correction to reciprocal space and to store it, so that it can be used to complement the standard vacuum FFT Poisson solver at no extra cost. More detailed information can be found in the LIBAFCC[42] reference.

3 Computational Details

All calculations were performed using Quantum Espresso (QE) v7.1[53, 54] compiled with the Environ library for continuum embedding effects [18, 53, 55]. While the double-cell implementation has been officially released in Environ 3.0, for the results reported in the following we used a local development version of the library that includes the coupling with the AFCC library for periodic boundary corrections.

All reported benchmarks, in vacuum and in dielectric environments, involve a single self-consistent field (SCF) optimization of the electronic density, starting from a random initial guess. In an effort to keep benchmarks as consistent as possible, we did not include geometry optimization calculations in our results. However, convergence of the inter-atomic forces with the newly implemented algorithms was thoroughly tested, showing an accuracy consistent with fully periodic simulations and with other PBC correction schemes.

The 0D and 2D simulations are based off the Environ's examples distributed with the library. In particular, the isolated system considered in this study is an acetamine cation, while a thin Pt (111) slab with an adsorbed carbon monoxide molecule was used for testing 2D systems. A BN nanoribbon was chosen as benchmark for the 1D systems. The selected systems are known to display more pronounced




Figure 3: Graphic representation of the two sets of calculations adopted to investigate convergence of the developed formalism. The double-cell calculations involve a small system cell (green box) embedded in an expanded environment cell (blue box). Reference calculations adopt the same expanded cell for both the system (DFT) and the environment (continuum electrostatic) calculations.

PBC artifacts. However, it is important to stress that the reported analysis is meant to focus on the consistency of the results with alternative PBC correction schemes and on the overall computational performance. For this reason, the physical approximations and the simulation details that are not relevant for the convergence of the electrostatic energy have been selected according to the most widespread choices in the literature. All the calculations were performed using the Perdue-Burke-Ernzerhof generalized gradient approximation (GGA) density functional[56]. For the 0D and 2D systems, the standard ultra-soft pseudopotentials distributed on the QE website were adopted. For the 1D system, the pseudopotentials from the SSSP efficiency library were selected[57, 58]. The Simulations were performed by only sampling the gamma point sampling of the Brillouin zone.

A dielectric embedding environment was modelled using the self-consistent interface function of Andreussi et al[18]. A homogeneous static dielectric permittivity of 100 was selected for the bulk of the environment. The preconditioned conjugate gradient (PCG) approach of Fisicaro et al [59] was used for the convergence of the generalized Poisson equation, with a convergence threshold on the computed electrostatic potential of $5.0e - 13$ Ry [18, 53].

For the acetamine cation (0D system) the wave function cutoff and density cutoff were chosen to be 30 Ry and 300 Ry, respectively, with a estimated convergence threshold on the total energy of $5.0e - 6$ Ry. A simple cubic simulation cell was adopted, in order to allow the use of a parabolic correction to PBCs [41, 51]. For this isolated system, reference simulations were also performed using the Martyna-Tuckerman reciprocal space correction [52], both in vacuum and in the dielectric medium.

For the 2D system, an orthorhombic cell was used with the cell parameter in the z-direction being varied to evaluate the interaction with periodic replicas. The wave function cutoff and density cutoff were set to 35 Ry and 300 Ry, respectively, with an estimated convergence threshold on the total energy of $1.0e - 6$ Ry. Marzari-Vanderbilt[60] smearing of the band occupations was adopted with a spread value of 0.03 Ry. A parabolic correction scheme, as implemented in the Environ library, was adopted for simulations in vacuum and in the dielectric medium.

Finally, for the 1D system, an ideal boron-nitride (BN) 2D surface, as obtained from the aflowlib[4] repository, was used to create a 1D ribbon. The wave function and density cutoffs were set to 60 Ry and 480 Ry, respectively, with an estimated convergence threshold on the total energy of $1.0e - 7$ Ry. A Gaussian smearing with a spread value of 0.02 Ry was adopted, together with a local Tomas-Fermi mixing and a mixing parameter of 0.7. Although the parabolic correction scheme implemented in

Environ is only compatible with 0D and 2D systems, a 2D correction with a large simulation cell in the direction along the plane of the ribbon was also exploited as reference for the 1D system.

In order to demonstrate the accuracy and performance of the double-cell formalism, two DFT calculations were done for each unit cell size, as schematized in Figure 3. The first calculation is performed using double-cell approach, i.e. having the environment cell defined as an integer expansion of the system cell in at least one direction. A reference calculation is then performed by imposing the system and environment cells to be of the same size and to match the environment cell of the first calculation. As an example, for a simulation of an isolated (0D) system with a cubic cell of size 10 Bohr, a double-cell calculation with one replica in every direction, $\mathbf{n} = (1, 1, 1)$, would involve a cubic environment cell with a 30 Bohr side length. The reference calculation would be performed with identical system and environment cells, both set to a size of 30 Bohr. In the following section, the energy difference between the double-cell calculation and the reference are reported together with the speed up in calculation time.

4 Results

4.1 Isolated Systems (0D)

As discussed above, one of the uses of the double-cell formalism is to reduce PBC artifacts in isolated systems, thus overcoming one of the main sources of artifacts of plane-wave based codes. The system selected for testing the accuracy and performance of the double-cell strategy is an acetamine cation, visualized in Figure 4. To simplify the handling of PBC artifacts and to allow the use of the parabolic point-countercharge correction, a cubic cell is used to study this system.

In the top panels of Figure 5 the changes in total energy with respect to the reference are reported, for both simulations in vacuum and in a dielectric environment. The considered molecule has overall size of about 8.0 Bohr. However, from the convergence of the double-cell calculations with respect to the reference it is clear that system cells smaller than 15 Bohr present artifacts with respect to the same calculation in the reference cell. As the electrostatic calculations in the double-cell and the reference are the same, these small difference must be due to the fact that the double-cell mapping is affecting the DFT convergence. In fact, inspection of the density of the molecule reveals that the difference is due to the non-negligible spilling of the electronic density of the molecule beyond the cell boundaries. While in a periodic cell this still corresponds to a connected smooth density, the mapping step of the double-cell formalism breaks the spilling density apart. This result suggests a general rule for the identification of the minimal system cell in a double-cell calculation, which needs to be about 10.0 Bohr larger than the system size.

When looking at the behaviour of the double-cell energy as a function of system size, we notice that the results are within 0.1 Ry from the fully converged energy for all simulation cells larger than the identified minimum size. As the simulated system is a molecule with a non-zero total net charge, its electrostatic energy in PBC is expected to vary as $1/L$ with respect to the cubic simulation cell size, L (red curves in Figure 5). With the parabolic correction scheme, such a dependence is corrected into a $1/L^5$, with energies virtually converged for cell sizes as large as 15 Bohr. Reciprocal-space corrections fully correct such a cell-size dependence, provided cell sizes twice as large as the system size are chosen. In the double-cell algorithm, the $1/L$ behavior is instead still present, although significantly smaller than the small-cell counterpart, as it corresponds to significantly larger periodic cell sizes. The same behavior is less evident for the simulations in a dielectric environment, as a significant fraction of the monopole-monopole interaction between periodic cells is screened by the embedding medium. In the bottom panels of Figure 5 we report the speedup per SCF cycle of the double-cell formalism when compared with the reference calculations. This speedup accounts for the relative weight of the




Figure 4: The acetamine cation used for testing the double-cell formalism in isolated systems.




Figure 5: Plots a) and b) show the energy difference, in Ry, between an SCF calculation using the double cell formalism and the corresponding reference calculation (as visualized in Figure 2). For plots c) and d), the convergence of different PBC schemes as a function of DFT system cell size is reported. Plots e) and f) show the time speedup between these simulations. All reported tests were performed for systems in vacuum (left column) and in a continuum dielectric medium (right column).

FFT-based electrostatic calculation and the cost of inverting the DFT Hamiltonian in a single SCF step. The reported results show an average five-fold decrease in computational time when the system cell is allowed to be smaller than the environment cell. In the dielectric embedding, the electrostatic problem needs to be solved iteratively (either using a fixed-point method [18] or using a preconditioned conjugate gradient approach [59]) where each iteration requires to solve a Poisson equation in vacuum. Given the multiple number of times the Poisson equation needs to be solved at each SCF step, the overhead in mapping and remapping scalar fields may become more evident in dielectric environments. In addition, given the fact that the dielectric medium is defined on the electronic density, for small simulation sizes the convergence of the dielectric medium effects may suffer for the truncation of the electronic density at the cell boundaries.

While the results in this section aim to provide a comparison between alternative PBC schemes, it is important to stress that the double-cell formalism can also be used in combination with any other PBC corrections scheme reported above: the electrostatic calculation in the expanded environment cell can be corrected with a real-space term (parabolic correction) or with any reciprocal-space scheme (Martyna-Tuckerman or AFC) already implemented in Environ. Those results are not reported in this section because the simulations are already fully converged without the need for additional corrections.

4.2 Slab (2D) System

To study 2D systems, systems that should be periodic in two dimensions, we used a bi-layer of platinum atoms in the (111) orientation and in the presence a CO molecule in the atop position, as visualized in Figure 6. For this system, the slab is included in an orthorhombic cell, oriented in the x-y plane and with the non-periodic axis along third direction. For the double-cell calculations, the system cell is expanded only in the z direction, while the system and environment cells are the same in the in-plane directions. Convergence tests are reported in Figure 7 as a function of the size of the system cell along the z direction.

The considered system has an extension along the vertical axis of about 11.0 Bohr. When looking at the difference in energy between the double-cell and the reference calculations, it appears that minimal artifacts are present at the smallest cell sizes (22-23 Bohr). This suggests that no significant spills in the electronic density are present for these DFT cells, consistent with our previous observations for the isolated system. As observed in the 0D case, double-cell energies are fully converged with respect to the DFT system cell already for the smallest cells considered. The results are in line with the convergence observed for a parabolic (also known as point-counter-charge or dipole correction) scheme, while AFC90 shows a slower convergence, with exact results for cells larger than 35 Bohr.

Eventually, in the bottom panels of Figure 7 we report the computational speedup of using a double cell approach, compared to using the same expanded cell for both electrostatic and DFT simulations and inverting the DFT Hamiltonian for this small system provide a speedup of about 2.5-2.8 in vacuum, that drops to about 1.5 in a dielectric medium. While the scaling of the FFT-based electrostatic is expected to be better than the DFT one as a function of the cell size, speedups appear to be independent of cell size for the considered system. This is probably due to the relatively small number of electrons in the system and to the other computational overheads associated with larger FFT grids.




Figure 7: Plots a) and b) show the energy difference, in Ry, between an SCF calculation using the double cell formalism and the corresponding reference calculation (as visualized in Figure 2). For plots c) and d), the convergence of different PBC schemes as a function of DFT system cell size is reported. Plots e) and f) show the time speedup between these simulations. All reported tests were performed for systems in vacuum (left column) and in a continuum dielectric medium (right column).




Figure 6: Visualization of the system used for testing the double-cell algorithm in two dimensions: a Pt (111) surface with two layers of Pt atoms in the presence of a CO molecule adsorbed in the atop position.

Figure 7: Plots a) and b) show the energy difference, in Ry, between an SCF calculation using the double cell formalism and the corresponding reference calculation (as visualized in Figure 2). For plots c) and d), the convergence of different PBC schemes as a function of DFT system cell size is reported. Plots e) and f) show the time speedup between these simulations. All reported tests were performed for systems in vacuum (left column) and in a continuum dielectric medium (right column).

4.3 Wire (1D) System

For the 1D systems, we studied the boron nitride (BN) nanoribbon, as shown in Figure 8. An orthorhombic cell is adopted, with the ribbon oriented along the x-axis, making the y and z axes the two non-periodic directions. Also in this case, only these two directions will be considered in the expansion of the system cell into the environment cell, while the system and environment cells have the same size along the x direction. For a sake of simplifying the discussion, the sizes of the cell axes in the non-periodic directions are kept identical, allowing to focus our analysis on a single parameter. Convergence tests are reported in Figure 9 as a function of the such a parameter.

The BN nanoribbon has a width along the y direction of 13 Bohr, with a one-atom thickness of less than 1 Bohr in the z direction. Convergence of the double-cell calculations with respect to the reference setup show minor artifact at the smallest cells considered (10 Bohr). This is consistent with our understanding of these small-cell artifacts reported for the 0D case and with the size of the




Figure 8: Visualization of the BN nanoribbon (boron atoms in pink, nitrogen atoms in blue) investigated to test the double-cell formalism for one-dimensional systems.




Figure 9: Plots a) and b) show the energy difference, in Ry, between an SCF calculation using the double cell formalism and the corresponding reference calculation (as visualized in Figure 2). For plots c) and d), the convergence of different PBC schemes as a function of DFT system cell size is reported. Plots e) and f) show the time speedup between these simulations. All reported tests were performed for systems in vacuum (left column) and in a continuum dielectric medium (right column).

studied system. While a small drift in convergence is observed at the larger cell sizes, the associated errors are within the convergence threshold of the performed SCF calculations.

Similarly to the results reported for 2D systems, double-cell energies appear to be well converged at all the considered system-cell sizes. AFC90 results show a slightly slower convergence, with some fluctuations for cell sizes smaller than 20 Bohr. While the parabolic (point-counter-charge) correction is not implemented for 1D systems, in Figure 9 we report results of using a 2D scheme applied considering the y axis as the only non-periodic direction of the system. As the system is very homogeneous along the z direction, we expect this scheme to capture most of the interaction energy between periodic replicas of the nanoribbon. Indeed, our simulations show that this parabolic correction is close to the double-cell results, although some small deviations are still present for cell sizes as large as 20 Bohr.

For this application, more substantial computational speedups are observed in using the double-cell formalism. Moreover, a linear increase in the speedup is observed as a function of the cell size, consistent with the expected difference in computational scaling between the electrostatic and the DFT components of the calculations. As for the 0D and 2D cases, the speedup is significantly reduced for the simulations in a dielectric medium, as the iterations involved in the solution of the more complex Generalized Poisson problem increase the weight of the electrostatic calculation.

5 Conclusion

We report the development and implementation of a double-cell numerical algorithm to decouple electrostatic calculations and environment effects from the underlying DFT simulation. The proposed formalism was tested as a tool to remove PBC artifacts for partially periodic and non-periodic systems. By comparing double-cell calculations with a benchmark reference we identified a source of potential artifacts for very small system cell sizes. Indeed, if the DFT cell is so small that the electronic density of the system spills across the cell boundaries, the mapping of the DFT cell into an expanded environment cell introduces artificial cuts in the electronic density, thus forbidding the use of the double cell for minimal DFT cell sizes. However, the presented results should full convergence with cell size for most simulation setups and system dimensionalities. This fast convergence can be further improved by the additional use of PBC corrections on the environment cell. While a general speedup in simulation is observed by allowing the DFT cell to be smaller than the environment cell, the most impressive result were observed the 1D system, with speedups of up to 30 times with respect to corresponding single-cell simulations. The effectiveness and accuracy of the presented approach give us confidence in the future use of this formalism for simulations with heterogeneous and nano-scaled continuum environments.

6 Supplementary Material

The supplementary material includes information on the the effect of the spread value when using LIBAFCC for different system dimentionalities as well as how the double-cell formalism affects the calculated forces.

References

- [1] Sebastiaan P. Huber, Spyros Zoupanos, Martin Uhrin, Leopold Talirz, Leonid Kahle, Rico Häuselmann, Dominik Gresch, Tiziano Müller, Aliaksandr V. Yakutovich, Casper W. Andersen, and et al. Aiida 1.0, a scalable computational infrastructure for automated reproducible workflows and data provenance. *Scientific Data*, 7(1), 2020.
- [2] Martin Uhrin, Sebastiaan P. Huber, Jusong Yu, Nicola Marzari, and Giovanni Pizzi. Workflows in aiida: Engineering a high-throughput, event-based engine for robust and modular computational workflows. *Computational Materials Science*, 187:110086, 2021.
- [3] Shyue Ping Ong, William Davidson Richards, Anubhav Jain, Geoffroy Hautier, Michael Kocher, Shreyas Cholia, Dan Gunter, Vincent L. Chevrier, Kristin A. Persson, Gerbrand Ceder, and et al. Python materials genomics (pymatgen): A robust, open-source python library for materials analysis. *Computational Materials Science*, 68:314–319, 2013.

- [4] Stefano Curtarolo, Wahyu Setyawan, Gus L.W. Hart, Michal Jahnatek, Roman V. Chepulskii, Richard H. Taylor, Shidong Wang, Junkai Xue, Kesong Yang, Ohad Levy, and et al. Aflow: An automatic framework for high-throughput materials discovery. *Computational Materials Science*, 58:218–226, 2012.
- [5] Claudia Draxl and Matthias Scheffler. The nomad laboratory: From data sharing to artificial intelligence. *Journal of Physics: Materials*, 2(3):036001, 2019.
- [6] Nicola Marzari, Andrea Ferretti, and Chris Wolverton. Electronic-structure methods for materials design. *Nature Materials*, 20(6):736–749, 2021.
- [7] David Padua. Fftw. *Encyclopedia of Parallel Computing*, page 671–671, 2011.
- [8] Oliviero Andreussi and Giuseppe Fisicaro. Continuum embeddings in condensed-matter simulations. *International Journal of Quantum Chemistry*, 119(1), 2018.
- [9] Hai Lin and Donald G. Truhlar. Qm/mm: What have we learned, where are we, and where do we go from here? *ChemInform*, 38(22), 2007.
- [10] Wenhui Mi, Xuecheng Shao, Alessandro Genova, Davide Ceresoli, and Michele Pavanello. Eqe 2.0: Subsystem dft beyond gga functionals. *Computer Physics Communications*, 269:108122, 2021.
- [11] Christoph R. Jacob and Johannes Neugebauer. Subsystem density-functional theory. *Wiley Interdisciplinary Reviews: Computational Molecular Science*, 4(4):325–362, 2014.
- [12] Xuecheng Shao, Wenhui Mi, and Michele Pavanello. Gga-level subsystem dft achieves sub-kcal/mol accuracy intermolecular interactions by mimicking nonlocal functionals. *Journal of Chemical Theory and Computation*, 17(6):3455–3461, 2021.
- [13] Modesto Orozco and F. Javier Luque. Theoretical methods for the description of the solvent effect in biomolecular systems. *Chemical Reviews*, 101(1):203–204, 2001. PMID: 11749364.
- [14] Christopher J. Cramer and Donald G. Truhlar. Implicit solvation models: Equilibria, structure, spectra, and dynamics. *Chemical Reviews*, 99(8):009–2665, 1999.
- [15] Jacopo Tomasi and Maurizio Persico. Molecular interactions in solution: An overview of methods based on continuous distributions of the solvent. *Chemical Reviews*, 94(7):2027–2094, 1994.
- [16] Jacopo Tomasi, Benedetta Mennucci, and Roberto Cammi. Quantum mechanical continuum solvation models. *Chemical Reviews*, 105(8):2999–3094, 2005. PMID: 16092826.
- [17] Jacopo Tomasi and Maurizio Persico. Molecular interactions in solution: An overview of methods based on continuous distributions of the solvent. *Chemical Reviews*, 94(7):2027–2094, 1994.
- [18] Oliviero Andreussi, Ismaila Dabo, and Nicola Marzari. Revised self-consistent continuum solvation in electronic-structure calculations. *The Journal of Chemical Physics*, 136(6):064102, 2012.
- [19] Giuseppe Fisicaro, Luigi Genovese, Oliviero Andreussi, Sagarmoy Mandal, Nisanth N. Nair, Nicola Marzari, and Stefan Goedecker. Soft-sphere continuum solvation in electronic-structure calculations. *Journal of Chemical Theory and Computation*, 13(8):3829–3845, 2017. PMID: 28628316.
- [20] Oliviero Andreussi, Nicolas Georg Hörmann, Francesco Nattino, Giuseppe Fisicaro, Stefan Goedecker, and Nicola Marzari. Solvent-aware interfaces in continuum solvation. *Journal of Chemical Theory and Computation*, 15(3):1996–2009, 2019.
- [21] Matthew Truscott and Oliviero Andreussi. Field-aware interfaces in continuum solvation. *The Journal of Physical Chemistry B*, 123(16):3513–3524, 2019.
- [22] E. Cancès, B. Mennucci, and J. Tomasi. A new integral equation formalism for the polarizable continuum model: Theoretical background and applications to isotropic and anisotropic dielectrics. *The Journal of Chemical Physics*, 107(8):3032–3041, 1997.

- [23] Luca Frediani, Roberto Cammi, Stefano Corni, and Jacopo Tomasi. A polarizable continuum model for molecules at diffuse interfaces. *The Journal of Chemical Physics*, 120(8):3893–3907, 2004.
- [24] S. Corni and J. Tomasi. Enhanced response properties of a chromophore physisorbed on a metal particle. *The Journal of Chemical Physics*, 114(8):3739–3751, 2001.
- [25] Oliviero Andreussi, Stefano Corni, Benedetta Mennucci, and Jacopo Tomasi. Radiative and nonradiative decay rates of a molecule close to a metal particle of complex shape. *The Journal of Chemical Physics*, 121(20):10190–10202, 2004.
- [26] Oliviero Andreussi, Alessandro Biancardi, Stefano Corni, and Benedetta Mennucci. Plasmon-controlled light-harvesting: Design rules for biohybrid devices via multiscale modeling. *Nano Letters*, 13(9):4475–4484, 2013.
- [27] Ravishankar Sundararaman, William A. Goddard, and Tomas A. Arias. Grand canonical electronic density-functional theory: Algorithms and applications to electrochemistry. *The Journal of Chemical Physics*, 146(11):114104, 2017.
- [28] Ravishankar Sundararaman and Kathleen Schwarz. Evaluating continuum solvation models for the electrode-electrolyte interface: Challenges and strategies for improvement. *The Journal of Chemical Physics*, 146(8):084111, 2017.
- [29] Ryosuke Jinnouchi and Alfred B. Anderson. Electronic structure calculations of liquid-solid interfaces: Combination of density functional theory and modified poisson-boltzmann theory. *Physical Review B*, 77(24), 2008.
- [30] Francesco Nattino, Matthew Truscott, Nicola Marzari, and Oliviero Andreussi. Continuum models of the electrochemical diffuse layer in electronic-structure calculations. *The Journal of Chemical Physics*, 150(4):041722, 2019.
- [31] Ravishankar Sundararaman, Kendra Letchworth-Weaver, Kathleen A. Schwarz, Deniz Gunceler, Yalcin Ozhabes, and T.A. Arias. Jdftx: Software for joint density-functional theory. *SoftwareX*, 6:278–284, 2017.
- [32] Joseph A. Gauthier, Stefan Ringe, Colin F. Dickens, Alejandro J. Garza, Alexis T. Bell, Martin Head-Gordon, Jens K. Nørskov, and Karen Chan. Challenges in modeling electrochemical reaction energetics with polarizable continuum models. *ACS Catalysis*, 9(2):920–931, 2018.
- [33] Joseph A. Gauthier, Colin F. Dickens, Stefan Ringe, and Karen Chan. Practical considerations for continuum models applied to surface electrochemistry. *ChemPhysChem*, 20(22):3074–3080, 2019.
- [34] Stefan Ringe, Harald Oberhofer, Christoph Hille, Sebastian Matera, and Karsten Reuter. Function-space-based solution scheme for the size-modified poisson–boltzmann equation in full-potential dft. *Journal of Chemical Theory and Computation*, 12(8):4052–4066, 2016.
- [35] Stefan Ringe, Harald Oberhofer, and Karsten Reuter. Transferable ionic parameters for first-principles poisson-boltzmann solvation calculations: Neutral solutes in aqueous monovalent salt solutions. *The Journal of Chemical Physics*, 146(13):134103, 2017.
- [36] Stefan Ringe, Nicolas G. Hörmann, Harald Oberhofer, and Karsten Reuter. Implicit solvation methods for catalysis at electrified interfaces. *Chemical Reviews*, 122(12):10777–10820, 2021.
- [37] Quinn Campbell and Ismaila Dabo. Quantum-continuum calculation of the surface states and electrical response of silicon in solution. *Physical Review B*, 95(20), 2017.
- [38] Quinn Campbell and Ismaila Dabo. Erratum: Quantum-continuum calculation of the surface states and electrical response of silicon in solution [phys. rev. b 95, 205308 (2017)]. *Physical Review B*, 96(3), 2017.

[39] Ted Hullar, Fernanda C. Bononi, Zekun Chen, Danielle Magadia, Oliver Palmer, Theo Tran, Dario Rocca, Oliviero Andreussi, Davide Donadio, Cort Anastasio, and et al. Photodecay of guaiacol is faster in ice, and even more rapid on ice, than in aqueous solution. *Environmental Science: Processes & Impacts*, 22(8):1666–1677, 2020.

[40] Fernanda C. Bononi, Zekun Chen, Dario Rocca, Oliviero Andreussi, Ted Hullar, Cort Anastasio, and Davide Donadio. Bathochromic shift in the uv-visible absorption spectra of phenols at ice surfaces: Insights from first-principles calculations. *The Journal of Physical Chemistry A*, 124(44):9288–9298, 2020.

[41] Oliviero Andreussi and Nicola Marzari. Electrostatics of solvated systems in periodic boundary conditions. *Physical Review B*, 90(24), 2014.

[42] Yanli Li and Ismaila Dabo. Electronic levels and electrical response of periodic molecular structures from plane-wave orbital-dependent calculations. *Phys. Rev. B*, 84:155127, Oct 2011.

[43] G. Makov and M. C. Payne. Periodic boundary conditions in ab initio calculations. *Physical Review B*, 51(7):4014–4022, 1995.

[44] Peter Broqvist, Audrius Alkauskas, and Alfredo Pasquarello. Hybrid-functional calculations with plane-wave basis sets: Effect of singularity correction on total energies, energy eigenvalues, and defect energy levels. *Physical Review B*, 80(8), 2009.

[45] F. Gygi and A. Baldereschi. Self-consistent hartree-fock and screened-exchange calculations in solids: Application to silicon. *Physical Review B*, 34(6):4405–4408, 1986.

[46] P. E. Blöchl. Electrostatic decoupling of periodic images of plane-wave-expanded densities and derived atomic point charges. *The Journal of Chemical Physics*, 103(17):7422–7428, 1995.

[47] Ikutaro Hamada, Minoru Otani, Osamu Sugino, and Yoshitada Morikawa. Green’s function method for elimination of the spurious multipole interaction in the surface/interface slab model. *Physical Review B*, 80(16), 2009.

[48] N. D. Hine, K. Frensch, W. M. Foulkes, and M. W. Finnis. Supercell size scaling of density functional theory formation energies of charged defects. *Physical Review B*, 79(2), 2009.

[49] Peter A. Schultz. Charged local defects in extended systems. *Physical Review Letters*, 84(9):1942–1945, 2000.

[50] Peter A. Schultz. Theory of defect levels and the “band gap problem” in silicon. *Physical Review Letters*, 96(24), 2006.

[51] Ismaila Dabo, Boris Kozinsky, Nicholas E. Singh-Miller, and Nicola Marzari. Electrostatics in periodic boundary conditions and real-space corrections. *Physical Review B*, 77(11), 2008.

[52] Glenn J. Martyna and Mark E. Tuckerman. A reciprocal space based method for treating long range interactions in ab initio and force-field-based calculations in clusters. *The Journal of Chemical Physics*, 110(6):2810–2821, 1999.

[53] P Giannozzi, O Andreussi, T Brumme, O Bunau, M Buongiorno Nardelli, M Calandra, R Car, C Cavazzoni, D Ceresoli, M Cococcioni, and et al. Advanced capabilities for materials modelling with quantum espresso. *Journal of Physics: Condensed Matter*, 29(46):465901, 2017.

[54] Paolo Giannozzi, Stefano Baroni, Nicola Bonini, Matteo Calandra, Roberto Car, Carlo Cavazzoni, Davide Ceresoli, Guido L Chiarotti, Matteo Cococcioni, Ismaila Dabo, and et al. Quantum espresso: A modular and open-source software project for quantum simulations of materials. *Journal of Physics: Condensed Matter*, 21(39):395502, 2009.

[55] Edan Bainglass, Medrano Gabriel, Matthew Truscott, Quinn Campbell, Ismaila Dabo, Iurii Timrov, Nicola Marzari, and Oliviero Andreussi. Environ 3.0: A continuum solvation package [unpublished manuscript].

- [56] John P. Perdew, Kieron Burke, and Matthias Ernzerhof. Generalized gradient approximation made simple. *Physical Review Letters*, 77(18):3865–3868, 1996.
- [57] Gianluca Prandini, Antimo Marrazzo, Ivano E. Castelli, Nicolas Mounet, and Nicola Marzari. Precision and efficiency in solid-state pseudopotential calculations. *npj Computational Materials*, 4(1), 2018.
- [58] Kurt et al. Lejaeghere. Reproducibility in density functional theory calculations of solids. *Nature*, 351(6280), 2016.
- [59] G. Fisicaro, L. Genovese, O. Andreussi, N. Marzari, and S. Goedecker. A generalized poisson and poisson-boltzmann solver for electrostatic environments. *The Journal of Chemical Physics*, 144(1):014103, 2016.
- [60] Nicola Marzari and David Vanderbilt. Maximally localized generalized wannier functions for composite energy bands. *Physical Review B*, 56(20):12847–12865, 1997.