Realspace local pseudopotential in ONETEP

School of Chemistry, University of Southampton,
Highfield, Southampton SO17 1BJ, UK

Written by Jacek Dziedzic
based on notes by Chris-Kriton Skylaris

September 28, 2010

1 Motivation

In standard ONETEP the local pseudopotential is obtained in reciprocal space by a discrete Fourier transform, by assuming the cell is periodically repeated in space. However, there are certain use-cases, where one is interested in the properties of an isolated (not periodically repeated) system. This is especially true if other energy terms, such the Hartree energy or the ion-ion energy are already calculated with open boundary conditions, which is the case, e.g., for implicit solvent calculations in ONETEP.

2 Theory

Assume that $v_{loc}(\vec{r})$ is located on an atom A at a position $\vec{R}_A$ and we want to determine the contribution to the local pseudopotential coming from this atom. Owing to the spherical symmetry of the potential, we have

$$v_{loc,A}(\vec{r}) = v_{loc}\left(\vec{r} - \vec{R}_A\right) = v_{loc}\left(|\vec{r} - \vec{R}_A|\right). \tag{1}$$

The local pseudopotential is given to us in terms of its continuous Fourier coefficients, $\tilde{v}_{loc}(|\vec{g}|)$, read from a recpot file. To generate the pseudopotential at a point $\vec{r}$ in real space, we use the continuous Fourier transform:

$$v_{loc}\left(\vec{r} - \vec{R}_A\right) = \frac{1}{(2\pi)^3}\int \tilde{v}_{loc}(\vec{g})\, e^{i\vec{g}\cdot\left(\vec{r}-\vec{R}_A\right)} d\vec{g} = \int \tilde{v}_{loc}(\vec{g})\, e^{i\vec{g}\cdot\vec{x}} d\vec{g}, \tag{2}$$

where we have set $\vec{x} = \vec{r} - \vec{R}_A$. Expanding the plane wave $e^{i\vec{g}\cdot\vec{x}}$ in terms of localised functions, we get

$$v_{loc,A}(\vec{r}) = \frac{1}{(2\pi)^3}\int \tilde{v}_{loc}(\vec{g}) \cdot \left[4\pi \sum_{l=0}^{\infty}\sum_{m=-l}^{l} i^l j_l(gx)\, Z_{lm}\left(\Omega_{\vec{g}}\right) Z_{lm}\left(\Omega_{\vec{x}}\right) d\vec{g}\right], \tag{3}$$

$$v_{loc,A}(\vec{r}) = \frac{1}{(2\pi)^3} 4\pi \sum_{l=0}^{\infty}\sum_{m=-l}^{l} i^l Z_{lm}\left(\Omega_{\vec{x}}\right) \underbrace{\int \tilde{v}_{loc}(\vec{g})\, j_l(gx)\, Z_{lm}\left(\Omega_{\vec{g}}\right) d\vec{g}}_{I_1}. \tag{4}$$

The orthogonality of harmonics means that all of the terms, except for that of $l = m = 0$, disappear and, after a change of coordinates ($g^2 \sin\theta$ being the Jacobian), we obtain a new expression for the integral in (4):

$$I_1 = \int_0^{2\pi} \int_0^{\pi} Z_{lm}\left(\Omega_{\vec{g}}\right) Z_{00} \sin\theta \, d\theta \, d\varphi \cdot \int_0^{\infty} \tilde{v}_{loc}\left(g\right) j_l\left(gx\right) g^2 dg. \tag{5}$$

With $Z_{00} = 1/\sqrt{4\pi}$, the double integral simplifies to 1 and we obtain, after realizing that all terms except for $l = 0$ disappear,

$$v_{loc,A}\left(\vec{r}\right) = \frac{1}{(2\pi)^3} 4\pi \int \tilde{v}_{loc}\left(g\right) j_0(gx)\, g^2 dg = \frac{1}{(2\pi)^3} 4\pi \int \tilde{v}_{loc}\left(g\right) \frac{\sin\left(gx\right)}{gx} g^2 \, dg. \tag{6}$$

ONETEP uses a convention where an additional factor of 4π is needed when transforming between real and reciprocal space. Thus the final formula for the local pseudopotential at a distance of x from an atom of species s becomes

$$v^s_{loc}\left(x\right) = \frac{2}{\pi} \int_0^{\infty} \tilde{v}^s_{loc}\left(g\right) \frac{\sin\left(gx\right)}{x} g \, dg. \tag{7}$$

3 Implementation

In practice, however, it is not possible to evaluate the integral (7) with ∞ as the upper limit, because $\tilde{v}^s_{loc}\left(g\right)$ is defined in the recpot file only up to a g_{max} of 100 Å^{-1}. Furthermore, to ensure the results are consistent with standard ONETEP, we must lower this limit even more, to prevent aliasing, as high g's will not be representable on our reciprocal space grid. Thus, in practice we evaluate

$$v^s_{loc}\left(x\right) = \frac{2}{\pi} \int_0^{g_{cut}} \tilde{v}^s_{loc}\left(g\right) \frac{\sin\left(gx\right)}{x} g \, dg, \tag{8}$$

where $g_{cut} = 2\pi \max\left(d_1, d_2, d_3\right)$ (d_i being the grid spacings of `pub_cell`) and will usually be in the order of 20-30 a_0^{-1}.

The integral is evaluated for x's on a fine radial grid running from 0 to the maximum possible distance, which is the magnitude of the cell diagonal. The calculation is distributed across nodes (each node deals with a portion of the fine radial grid). The total pseudopotential for any point on the real space fine grid is evaluated by interpolation from the fine radial grid and by summing over all atoms. This calculation is distributed across nodes as well (each node deals with its own slabs of the real space fine grid). The default number of points in the radial grid is 100000 and can be changed with the directive `openbc_pspot_finetune_nptsx`.

The integral (8) is difficult to evaluate numerically. One source of difficulties is the oscillatory nature of $\sin\left(gx\right)$. For larger cells, where the maximum interesting x is in the order of $100\, a_0$, this oscillates so rapidly that the resolution of the recpot file (0.05 Å^{-1}) is not enough and it becomes necessary to interpolate $\tilde{v}^s_{loc}\left(g\right)$, and the whole integrand, between the g-points specified in the recpot file. The result of the interpolation is stored on a fine radial g-grid, which is f times as fine as the original radial g-grid of the recpot file. f is determined automatically so that every full period of $\sin\left(gx\right)$ is sampled by at least 50 points. For typical cells, this yields f in the order of 5-50, depending on the cell size. Alternatively, f may be specified manually by the `openbc_pspot_finetune_f` directive.

Another difficulty is caused by the singularity in $\tilde{v}^s_{loc}(g)$ as $g \to 0$, where the behaviour of $\tilde{v}^s_{loc}(g)$ approaches that of $-Z_s/g^2$. Although the integral is convergent, this singularity cannot be numerically integrated in an accurate fashion. The singularity also presents problems when interpolating between the g points – the usual cubic interpolation of `services_1d_interpolation` becomes inaccurate at low g's. The second problem is solved by subtracting the Coulombic potential, $-Z_s/g^2$, before interpolation to the fine radial g-grid and then adding it back. The first problem is difficult to treat. An approach where at low g's $\tilde{v}^s_{loc}(g)$ is assumed to be exactly equal to $-Z_s/g^2$ (which allows the low-g part of integral (8) to be evaluated analytically) gives better results than attempting to numerically integrate the singularity, but is not accurate enough, leading to errors in the order of $50-100\,\mu$Ha in the energy for a hydrogen atom test-case (with a total energy of ca. 0.477 Ha. Attempting to fit $A/g^2 + B/g + C$ (which also allows analytical integration at low g's) gives similar results. The numerical inaccuracy presents itself as a near-constant shift of the obtained pseudopotential and clearly affects total energy.

To solve this problem, we observe that the local pseudopotential can be split into a long-range part and a short-range part:

$$v^s_{loc}(x) = v^{s(long)}_{loc}(x) + v^{s(short)}_{loc}(x)\,, \tag{9}$$

$$\tilde{v}^s_{loc}(g) = \tilde{v}^{s(long)}_{loc}(g) + \tilde{v}^{s(short)}_{loc}(g)\,. \tag{10}$$

Following [1], we observe that $\tilde{v}^{s(long)}_{loc}(g) = \frac{4\pi}{g^2}\exp\left(\frac{-g^2}{4\alpha^2}\right)$ (where α is an adjustable parameter, controllable with `openbc_pspot_finetune_alpha`) which easily transforms to real space to give $v^{s(long)}_{loc}(x) = -\frac{\mathrm{erf}(\alpha x)}{x}$ and is conveniently calculated in real space. The short-range part (corresponding to high g's) is $\tilde{v}^{s(long)}_{loc}(g) = \tilde{v}^s_{loc}(g)\cdot\left[1-\exp\left(\frac{-g^2}{4\alpha^2}\right)\right]$. In this way, the integral (8) can be rewritten as

$$v^s_{loc}(x) = -\frac{\mathrm{erf}(\alpha x)}{x} + \frac{2}{\pi}\underbrace{\int_0^{g_{cut}} \tilde{v}^s_{loc}(g)\cdot\left[1-\exp\left(\frac{-g^2}{4\alpha^2}\right)\right]\cdot\frac{\sin(gx)}{x}g\,dg}_{I_s(x)}\,. \tag{11}$$

Owing to the $\left[1-\exp\left(\frac{-g^2}{4\alpha^2}\right)\right]$ factor, the integral $I_s(x)$ is no longer singular at $g=0$ and can be accurately evaluated numerically, if α is large enough. Small values of α make the numerical integration more difficult (requiring larger values for f), because the oscillations at low g's are large in magnitude. Larger values of α allow for easy integration, but they cause the long-range behaviour to "kick in" earlier. As this long-range behaviour is calculated in real space, it lacks the oscillations that are present in standard ONETEP because of a finite value for g_{cut}. Even though these oscillations are an artifact, obtaining a long-range behaviour that is physically more correct, but without the oscillations, leads to aliasing in reciprocal space and to a departure from the results of standard ONETEP. For this reason we want α to be as small as possible, without negatively impacting the numerical integration. The accuracy of the obtained method can be judged by comparing the real space tail of the obtained pseudopotential with the Coulombic potential. Since we expect the obtained pseudopotential to oscillate slightly around $-Z_s/x$, a good measure of accuracy, which we will call b, is the average value of $\frac{v^s_{loc}(x)-(-Z_s/x)}{-Z_s/x}$ over the tail of the pseudopotential, from, say, $5\,a_0$ to the maximum x for which $v^s_{loc}(x)$ is evaluated. Ideally, b should be zero. Numerical inaccuracies will cause a shift in $v^s_{loc}(x)$ which will present itself as a finite, non-zero value of b. Naïve numerical integration by a direct calculation of (8) led, for our test-case, to b in the order of 0.01, which can be reduced by an order of magnitude

by using a very fine radial g-grid (high value of f). Subtracting out the Coulombic potential and integrating only the difference between $\tilde{v}^s_{loc}(g)$ and the Coulombic potential numerically, while integrating the remaining part analytically reduced b to about 0.0005. Application of the proposed formula (11) yielded $b = 5 \cdot 10^{-8}$ for $\alpha = 0.5/l$ and $b = 3 \cdot 10^{-9}$ for $\alpha = 0.1/l$ with a suitably large f to ease the numerical integration at low g (l is the box length). With the default value for f, the total energy is not sensitive (to more than 0.0001%) to the choice of α, provided it is in a resonable range of $0.1/l - 2/l$. The value of $0.3/l$ was chosen as a default.

The calculation of the realspace local pseudo is implemented in `norm_conserv_pseudo.F90` in the subroutine `pseudo_local_on_grid_openbc` and its internal subroutine `internal_Is_of_x`, which evaluates $I_s(x)$. A typical calculation would use default values for all the parameters. The realspace local pseudo is off by default and is turned on automatically when smeared ions or implicit solvent is in use. It can also be forced to be on (for development tests) by using `openbc_pspot T`.

Table 1: Directives controlling the calculation of the realspace local pseudo

Directive	Action	Rationale for use
`openbc_pspot T`	Forces the realspace pseudo to be used	Normally not needed, the realspace pseudo will be turned on when necessary. This directive allows turning it on even though the Hartree potential calculation and Ewald calculation proceed in reciprocal space, which might be useful for certain test calculations. A related directive, `openbc_ion_ion T` may be used in conjuction, to replace Ewald with a direct Coulombic sum.
`openbc_pspot_finetune_f` *value* *value* is an integer.	Sets the fineness parameter, f, to *value*.	Default value of -1 causes f to be determined automatically. Positive values can be used to increase f to obtain extra accuracy. Decreasing f will reduce accuracy and is not recommended.
`openbc_pspot_finetune_nptsx` *value* *value* is an integer.	Sets the number of radial grid points (distinct values of x) to *value*.	The default of 100000 should be enough, unless huge boxes are used, where it might make sense to increase it. Decreasing this value is not recommended, as it will impact accuracy.
`openbc_pspot_finetune_alpha` *value* *value* is a real.	Sets the short-range-long-range crossover parameter α to *value*$/l$, where l is the maximum dimension of the cell.	A default value of 0.3 should be OK for most applications. Increasing α will reduce the numerical inaccuracy in $I_s(x)$, but will cause the long-range behaviour to lack the oscillations of usual ONETEP and thus increase aliasing. Decreasing α will make $I_s(x)$ inaccurate, which can be helped, to a certain extent, by increasing f.

References

[1] G. J. Martyna and M. E. Tuckerman *J. Chem. Phys.* **110** (1999).