

# ELSI: A Unified Software Interface for Kohn-Sham Electronic Structure Solvers

Victor Wen-zhe Yu<sup>a</sup>, Fabiano Corsetti<sup>b</sup>, Alberto García<sup>c</sup>, William P. Huhn<sup>a</sup>, Mathias Jacquelin<sup>d</sup>, Weile Jia<sup>d,e</sup>, Björn Lange<sup>a</sup>, Lin Lin<sup>d,e</sup>, Jianfeng Lu<sup>f</sup>, Wenhui Mi<sup>a</sup>, Ali Seifitokaldani<sup>a</sup>, Álvaro Vázquez-Mayagoitia<sup>g</sup>, Chao Yang<sup>d</sup>, Haizhao Yang<sup>f</sup>, Volker Blum<sup>a,\*</sup>

<sup>a</sup>*Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC 27707*

<sup>b</sup>*Departments of Materials and Physics, and the Thomas Young Centre for Theory and Simulation of Materials, Imperial College London, London SW7 2AZ, United Kingdom*

<sup>c</sup>*Institut de Ciència de Materials de Barcelona (ICMAB-CSIC), Bellaterra E-08193, Spain*

<sup>d</sup>*Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720*

<sup>e</sup>*Department of Mathematics, University of California, Berkeley, CA 94720*

<sup>f</sup>*Department of Mathematics, Duke University, Durham, NC 27707*

<sup>g</sup>*Argonne Leadership Computing Facility, Argonne National Laboratory, Argonne, IL 60439*

---

## Abstract

Solving the electronic structure from a generalized or standard eigenproblem is often the bottleneck in large scale calculations based on Kohn-Sham density-functional theory. This problem must be addressed by essentially all current electronic structure codes, based on similar matrix expressions, and by high-performance computation. We here present a unified software interface, ELSI, to access different strategies that address the Kohn-Sham eigenvalue problem. Currently supported algorithms include the dense generalized eigensolver library ELPA, the orbital minimization method implemented in libOMM, and the pole expansion and selected inversion (PEXSI) approach with lower computational complexity for semilocal density functionals. The ELSI interface aims to simplify the implementation and optimal use of the

---

\*Corresponding author.

*Email address:* `volker.blum@duke.edu` (Volker Blum )

different strategies, by offering (a) a unified software framework designed for the electronic structure solvers in Kohn-Sham density-functional theory; (b) reasonable default parameters for a chosen solver; (c) automatic conversion between input and internal working matrix formats, and in the future (d) recommendation of the optimal solver depending on the specific problem. Comparative benchmarks are shown for system sizes up to 11,520 atoms (172,800 basis functions) on distributed memory supercomputing architectures.

*Keywords:* Density-Functional Theory, Kohn-Sham eigenvalue problem, Parallel computing

---

## PROGRAM SUMMARY

*Program title:* ELSI Interface

*Licensing provisions:* BSD 3-clause

*Distribution format:* tar.gz

*Programming language:* Fortran 2003, with interface to C/C++

*External routines/libraries:* MPI, BLAS, LAPACK, ScaLAPACK, ELPA, libOMM, PEXSI, ParMETIS, SuperLU-DIST

*Operating system:* Unix-like (Linux, macOS), Windows (not tested)

*Nature of problem:* Solving the electronic structure from a generalized or standard eigenvalue problem in calculations based on Kohn-Sham density functional theory (KS-DFT).

*Solution method:* To connect the KS-DFT codes and the KS electronic structure solvers, ELSI provides a unified software interface with reasonable default parameters, hierarchical control over the interface and the solvers, and automatic conversions between input and internal working matrix formats. Supported solvers are: ELPA (dense generalized eigensolver), libOMM (orbital minimization method), and PEXSI (pole expansion and selected inversion method).

*Restrictions:* The ELSI interface requires complete information of the Hamiltonian matrix.

## 1. Introduction

Molecular and materials simulations based on Kohn-Sham (KS) [1] and generalized Kohn-Sham (gKS) [2, 3] density-functional theory (DFT) are widely used to provide atomic-scale insights, understanding, and predictions

5 across a wide range of disciplines in the sciences and in engineering. The  
6 number of DFT-related publications has grown rapidly over recent decades  
7 [4, 5, 6], exceeding 20,000 in 2016 [6]. In particular, simulations based on  
8 semilocal and hybrid density functionals serve as the production workhorses  
9 for a broad range of applications. Advances in both computational methods  
10 and high-performance computing hardware render it feasible to model large  
11 systems consisting of thousands of atoms, and linear scaling KS-DFT [7, 8, 9]  
12 can reach system sizes of millions of atoms [10, 11]. Higher levels of density  
13 functional approximations, like the Random Phase Approximation (RPA),  
14 can be formulated to scale linearly with system size as well [12, 13].

15 However, approaches for which the computational effort scales lower than  
16  $O(N^3)$ , where  $N$  is some measure of the system size, are, arguably, not yet  
17 fully established as mainstream methods of the field. There are several rea-  
18 sons for this status. Perhaps the simplest reason is that formally  $O(N^3)$   
19 scaling approaches by solving an algebraic eigenvalue problem are generally  
20 applicable to any class of system, and the computational effort associated  
21 with them has a low prefactor, i.e., they are advantageous to use for systems  
22 comprised of up to roughly a few thousands of atoms, which account for the  
23 bulk of KS-DFT applications. In contrast, the transition to lower-scaling so-  
24 lution methods for larger systems is not necessarily simple. Such alternatives  
25 are typically restricted to certain classes of systems or problems. The transi-  
26 tion is, therefore, not trivial to automate, requiring specific intervention and  
27 sometimes specialist knowledge by its users. This creates hurdles both from  
28 a user point of view (complexity of choice) and from a developer point of  
29 view (replication of often complex infrastructure to implement a particular  
30 method efficiently). The KS eigenvalue problem is thus in practice a bottle-  
31 neck of KS-DFT simulations on current HPC architectures and for system  
32 sizes significantly exceeding several thousands of atoms.

33 We here present a software infrastructure, ELSI, that simplifies the ap-  
34 proach to overcome the Kohn-Sham eigenproblem bottleneck as much as  
35 possible for electronic structure users and developers. ELSI provides an inte-  
36 grated and extendable interface to multiple strategies targeting the KS eigen-  
37 problem (referred to as Kohn-Sham electronic structure solvers throughout  
38 this paper). It presently (version: 2017.05) supports three solvers: ELPA  
39 (Eigenvalue solvers for Petaflop-Applications) [14, 15], libOMM (Orbital  
40 Minimization Method) [16], and PEXSI (Pole EXpansion and Selected In-  
41 version) [17, 18, 19]. For the future, ELSI is expressly intended to integrate  
42 further solvers such as the linear-scaling solver CheSS(CHEbyshev Sparse

43 Solvers) [20], the iterative solver SIPs (Shift-and-Invert Parallel Spectral  
44 transformation eigensolver) [21], and others. By design, ELSI is an open  
45 infrastructure, intended to serve a community, and it can and should be flex-  
46 ibly adaptable to new solvers and new electronic structure codes’ needs in  
47 the future. In this paper, we describe the outline and basic principles of  
48 ELSI, as well as a comparative assessment of the three solution strategies  
49 that are already supported in ELSI as of its 2017.05 release. The software  
50 presented here is a structural foundation that is already working in several  
51 electronic structure codes, and we expect it to become a focal point for new  
52 developments and solver cross-comparisons in the future.

## 53 2. Kohn-Sham Density-Functional Theory

54 In KS-DFT [1], the many-electron problem for the Born-Oppenheimer  
55 electronic ground state is reduced to a system of single particle equations  
56 known as the Kohn-Sham equations

$$\hat{h}^{\text{KS}}\psi_l = \epsilon_l\psi_l, \quad (1)$$

57 where  $\psi_l$  and  $\epsilon_l$  are Kohn-Sham orbitals and their associated eigenenergies,  
58 and  $\hat{h}^{\text{KS}}$  denotes the Kohn-Sham Hamiltonian:

$$\hat{h}^{\text{KS}} = \hat{t}_s + \hat{v}_{\text{es}} + \hat{v}_{\text{xc}} + \hat{v}_{\text{ext}}, \quad (2)$$

59 which includes the kinetic energy  $\hat{t}_s$ , the average electrostatic potential of  
60 the electron density and of the nuclei  $\hat{v}_{\text{es}}$  (i.e. the Hartree potential), the  
61 exchange-correlation potential  $\hat{v}_{\text{xc}}$ , and possible additional potential terms  
62  $\hat{v}_{\text{ext}}$  from external electromagnetic fields.

63 In almost all practical approaches,  $N_{\text{basis}}$  basis functions  $\phi_i(\mathbf{r})$  are em-  
64 ployed to approximately expand the Kohn-Sham orbitals:

$$\psi_l(\mathbf{r}) = \sum_{j=1}^{N_{\text{basis}}} c_{jl}\phi_j(\mathbf{r}). \quad (3)$$

65 The choice of basis set is one of the critical decisions in the design of an  
66 electronic structure code [22]. Using non-orthogonal basis functions (e.g.,

67 Gaussian functions [22, 23, 24, 25, 26, 27, 28], Slater functions [29, 30], nu-  
68 meric atom-centered orbitals [31, 32, 33, 34, 35, 36], (linearized) augmented  
69 plane waves [37, 38, 39, 40, 41], finite elements [42]) in Eq. 3 converts Eq. 1  
70 to a generalized eigenvalue problem

$$\sum_j h_{ij} c_{jl} = \epsilon_l \sum_j s_{ij} c_{jl}, \quad (4)$$

71 where  $h_{ij}$  and  $s_{ij}$  are the elements of the Hamiltonian matrix  $\mathbf{H}$  and the  
72 overlap matrix  $\mathbf{S}$ , which can be computed through numerical integrations:

$$\begin{aligned} h_{ij} &= \int d^3r [\phi_i^*(\mathbf{r}) \hat{h}^{\text{KS}} \phi_j(\mathbf{r})], \\ s_{ij} &= \int d^3r [\phi_i^*(\mathbf{r}) \phi_j(\mathbf{r})]. \end{aligned} \quad (5)$$

73 Eq. 4 can thus be expressed in the following matrix form:

$$\mathbf{H}\mathbf{c} = \epsilon\mathbf{S}\mathbf{c}. \quad (6)$$

74 Here, the matrix  $\mathbf{c}$  and diagonal matrix  $\epsilon$  contain the eigenvectors and eigen-  
75 values, respectively, of the eigensystem of the matrices  $\mathbf{H}$  and  $\mathbf{S}$ .

76 When using orthonormal basis sets (e.g., plane waves [9, 43, 44, 45, 46, 47],  
77 multi-resolution wavelets [48, 49, 50], adaptive local basis set [51, 52], grid-  
78 discretization based approaches [53, 54]), the eigenproblem described in Eq.  
79 6 reduces to a standard form where  $s_{ij} = \delta_{ij}$ , or even can be circumvented  
80 completely by solving the KS equations in an integral formulation [22].

81 The explicit solution of Eq. 4 or 6 yields the Kohn-Sham orbitals  $\psi_i$ , from  
82 which the electron density  $n(\mathbf{r})$  can be computed following an orbital-based  
83 method that scales as  $O(N^2)$ :

$$n(\mathbf{r}) = \sum_{j=1}^{N_{\text{basis}}} f_j \psi_j^*(\mathbf{r}) \psi_j(\mathbf{r}), \quad (7)$$

84 where  $f_l$  denotes the occupation number of each orbital. In an actual com-  
85 putation, it is sufficient to perform the summation only for the occupied

86 ( $f_l > 0$ ) orbitals. The ratio of occupied orbitals to the total number of ba-  
 87 sis functions can be below 1% for plane wave basis sets, whereas with some  
 88 localized basis sets, fewer basis functions are required, leading to a larger  
 89 fraction of occupied states typically between 10% and 40%.

90 An alternative method that scales as  $O(N)$  can be employed for localized  
 91 basis functions:

$$n(\mathbf{r}) = \sum_{i,j}^{N_{\text{basis}}} \phi_i^*(\mathbf{r}) n_{ij} \phi_j(\mathbf{r}), \quad (8)$$

92 with  $n_{ij}$  being the elements of the density matrix that need to be computed  
 93 before the density update:

$$n_{ij} = \sum_{l=1}^{N_{\text{basis}}} f_l c_{il} c_{jl}. \quad (9)$$

94 Due to the dependence of  $\mathbf{H}$  on  $\psi_l$  via the density and the potentials,  
 95 Eqs. 4 and 6 are in fact non-linear eigenvalue problems, and therefore must  
 96 be solved in an iterative fashion. The most commonly used method is the  
 97 self-consistent field (SCF) or fixed-point iteration approach. To achieve self-  
 98 consistency, the electron density needs to be updated in every iteration until  
 99 converged to an acceptable level. From a viewpoint of computational com-  
 100 plexity, almost all standard pieces of solving the Kohn-Sham equations can  
 101 be formulated in a linear scaling fashion with respect to the system size. The  
 102 only piece that can not, in all cases and for all semilocal and hybrid func-  
 103 tionals, be easily addressed in an  $O(N)$  fashion is the Kohn-Sham eigenvalue  
 104 problem described in Eq. 4.

### 105 **3. Kohn-Sham Electronic Structure Solvers Supported by ELSI**

#### 106 *3.1. ELPA: Eigenvalue solvers for Petaflop-Applications*

107 The Kohn-Sham eigenvalue problem in Eq. 4 can be explicitly solved  
 108 by traditional (tri)diagonalization [55]. In ELSI, the massively parallel di-  
 109 rect solver ELPA [14, 15] facilitates the solution of symmetric or Hermitian  
 110 eigenproblems on high-performance computers. It was initially designed for

111 distributed memory architectures, then extended to exploit multi-threading  
 112 parallelism, and is subject to ongoing work for GPU acceleration.

113 In ELPA, the generalized eigenproblem in Eq. 6 is first transformed to  
 114 the standard form by Cholesky decomposition of the overlap matrix  $\mathbf{S}$ :

$$\mathbf{S} = \mathbf{L}\mathbf{L}^*, \quad (10)$$

115 where  $\mathbf{L}$  is a lower triangular matrix. Eq. 6 is then transformed by applying  
 116 the Cholesky factor:

$$\tilde{\mathbf{H}}\tilde{\mathbf{c}} = \epsilon\tilde{\mathbf{c}} \quad (11)$$

117 with  $\tilde{\mathbf{H}} = \mathbf{L}^{-1}\mathbf{H}(\mathbf{L}^*)^{-1}$  and  $\tilde{\mathbf{c}} = \mathbf{L}^*\mathbf{c}$ .

118 Then, the standard eigenproblem is either directly reduced to the tridi-  
 119 agonal form

$$\mathbf{T} = \mathbf{Q}\tilde{\mathbf{H}}\mathbf{Q}^*, \quad (12)$$

120 or first reduced to a banded intermediate form, then to the tridiagonal form  
 121 [56]:

$$\begin{aligned} \mathbf{B} &= \mathbf{Q}_1\tilde{\mathbf{H}}\mathbf{Q}_1^*, \\ \mathbf{T} &= \mathbf{Q}_2\mathbf{B}\mathbf{Q}_2^*. \end{aligned} \quad (13)$$

122 In Eqs. 12 and 13,  $\mathbf{Q}$ ,  $\mathbf{Q}_1$ ,  $\mathbf{Q}_2$  are transformation matrices;  $\mathbf{T}$  is a tridiagonal  
 123 matrix;  $\mathbf{B}$  is a banded matrix.

124 The key steps of the two-stage tridiagonalization algorithm implemented  
 125 in ELPA are reviewed in Fig. 1. Steps (1) and (2) correspond to Eq. 13,  
 126 i.e. the transformations to the banded and tridiagonal forms. Step (3) cor-  
 127 responds to the solution of the actual eigenvalue problem by a divide-and-  
 128 conquer approach [14, 57], which can be restricted to compute only a fraction  
 129 of the eigenvectors. Finally, the computed eigenvectors are transformed back  
 130 into the representations corresponding to the banded (step (4)) and standard  
 131 forms (step (5)) of the problem. Compared to the one-step tridiagonaliza-  
 132 tion (Eq. 12), the two-step algorithm introduces two additional steps (steps  
 133 (1) and (5) in Fig. 1). Still, the two-step approach has been shown to

134 enable faster computation and better parallel scalability than the one-step  
 135 approach on present-day computers [15]. Specifically, the matrix-vector op-  
 136 erations (BLAS level-2 routines) in the one-step tridiagonalization can be  
 137 mostly replaced by more efficient matrix-matrix operations (BLAS level-3  
 138 routines) in the two-step version of the algorithm [58]. Since steps 2 and 4  
 139 pertain to forward and back transformations between banded and tridiago-  
 140 nal matrices only, the resulting transformations can be efficiently grouped to  
 141 minimize computational overhead, especially for the back transformation in  
 142 step (4) [14]. The computational workload associated with step (4) is further  
 143 alleviated in KS-DFT calculations if only a small fraction of the eigenvectors  
 144 representing the lowest eigenstates is required, and by architecture-specific  
 145 linear-algebra “kernels” provided with the ELPA library [14, 15].




Figure 1: Five computational steps of the ELPA eigensolver with two-stage tridiagonal-  
 ization. (1) Reduction of the full matrix to a banded form. (2) Reduction of the banded  
 matrix to a tridiagonal form. (3) Solution of the eigenvalues and eigenvectors of the  
 tridiagonal system. (4) Back-transformation of the eigenvectors to the banded form. (5)  
 Back-transformation of the eigenvectors to the original full form. This figure is redesigned  
 based on Fig. 1 in Ref. [15].

146 Since ELPA employs the same 2D block-cyclic matrix distribution as does  
 147 the ScaLAPACK library [59] (by way of the basic linear algebra communi-  
 148 cation subroutines (BLACS) [60]), it can easily be substituted into existing  
 149 codes that already support parallel linear algebra by ScaLAPACK.

### 150 3.2. *libOMM: Orbital Minimization Method*

151 Instead of diagonalizing the  $N_{\text{basis}} \times N_{\text{basis}}$  eigenproblem, the orbital min-  
 152 imization method (OMM) relies on efficient iterative algorithms to directly  
 153 minimize an unconstrained energy functional using a set of auxiliary orbitals  
 154 that are not the Kohn-Sham orbitals  $\phi_i$ . These auxiliary orbitals are then



155 used to obtain the density matrix of the system. Specifically, the OMM em-  
 156 ploys  $N_W = N_{\text{electron}}/2$  non-orthogonal Wannier functions  $\chi_k$  to represent the  
 157 occupied subspace of a system with  $N_{\text{electron}}$  electrons:

$$\chi_k = \sum_{j=1}^{N_{\text{basis}}} W_{kj} \phi_j. \quad (14)$$

158 For non-spinpolarized systems, the index  $k$  runs from 1 to  $N_W$ . Then the  
 159 matrices  $\mathbf{H}$  and  $\mathbf{S}$  in the occupied subspace become

$$\begin{aligned} \mathbf{H}_{\text{omm}} &= \mathbf{W}^* \mathbf{H} \mathbf{W}, \\ \mathbf{S}_{\text{omm}} &= \mathbf{W}^* \mathbf{S} \mathbf{W}, \end{aligned} \quad (15)$$

160 where  $\mathbf{W}$  is the coefficient matrix of the Wannier functions. The size change  
 161 of the Hamiltonian matrix facilitated by Eq. 15 is illustrated in Fig. 2.




Figure 2: Schematic representation of sizes of Hamiltonian matrix before and after applying the Wannier function transformation in the orbital minimization method. Matrix dimensions are shown above the matrices.  $N_W$ : Number of Wannier functions.  $N_{\text{basis}}$ : Number of basis functions.

162 The OMM energy functional is defined as

$$E[\mathbf{W}] = 4Tr[\mathbf{H}_{\text{omm}}] - 2Tr[\mathbf{S}_{\text{omm}}\mathbf{H}_{\text{omm}}]. \quad (16)$$

163 This functional, when minimized with respect to the coefficients of Wannier  
 164 functions  $\mathbf{W}$ , can be shown to be equal to the sum of the lowest  $N_{\text{electron}}/2$

165 eigenvalues of the original KS eigenproblem [61, 62, 63, 64]. Furthermore,  
 166 the Wannier functions are driven towards perfect orthonormality at this min-  
 167 imum. The density matrix is then constructed from the  $\mathbf{W}$  that minimizes  
 168  $E[\mathbf{W}]$ . Although this density matrix is sufficient for the electron density up-  
 169 date following Eq. 8, compared to the density matrix in Eq. 9, it is obvious  
 170 that the occupation numbers are restricted to be integers (1 for occupied; 0  
 171 for unoccupied) in this method. Without knowledge of individual eigenstates,  
 172 the OMM cannot handle systems with fractional occupation numbers result-  
 173 ing, e.g., from a finite electronic temperature, such as is typically required  
 174 for metals.

175 Compared to other minimization methods with the orthonormality con-  
 176 straint of eigenstates [47, 65, 66], the advantage of the OMM is that it only  
 177 requires an unconstrained minimization without an explicit orthonormaliza-  
 178 tion step. This makes the OMM a good candidate for linear scaling DFT;  
 179 indeed, the method was originally developed in this context [61, 62, 63, 64].  
 180 However, in order to do so, it is necessary to spatially confine the Wannier  
 181 functions by imposing a certain sparsity to  $\mathbf{W}$ . This introduces a number of  
 182 technical difficulties which have ultimately required the development of more  
 183 involved algorithms [61, 63, 67]. The properties of the original OMM func-  
 184 tional with unconstrained Wannier functions have nevertheless been found to  
 185 result in an extremely efficient iterative solver with conventional cubic scaling  
 186 but a smaller prefactor than diagonalization. This approach has been taken  
 187 by the new implementation in libOMM [16]. It should be noted that for finite-  
 188 range basis sets in which  $\mathbf{W}$  is formally sparse, this sparsity can be taken  
 189 into account to reduce the scaling of the matrix-matrix product  $\mathbf{HW}$  from  
 190 cubic to quadratic, thus effectively eliminating the most expensive matrix op-  
 191 eration in the algorithm. The minimization of the OMM energy functional  
 192 in Eq. 16 is carried out in libOMM by using the conjugate-gradient (CG)  
 193 method with an efficient preconditioning using the kinetic energy matrix, as  
 194 described in Ref. [16].

### 195 3.3. PEXSI: Pole EXpansion and Selected Inversion

196 The density matrix in Eq. 9 is associated with the Kohn-Sham orbitals  
 197 and their occupation numbers  $f_l$ , which are given by the Fermi-Dirac distri-  
 198 bution function [68]:

$$f_l = \frac{1}{1 + e^{\frac{\epsilon_l - \mu}{k_B T}}}. \quad (17)$$

199 Here  $k_B$  is the Boltzmann constant,  $T$  is the temperature, and  $\mu$  is the  
 200 chemical potential that is determined by the normalization condition

$$\sum_{l=1}^{N_{\text{basis}}} f_l = N_{\text{electron}}. \quad (18)$$

201 The pole expansion and selected inversion (PEXSI) method [17, 18, 19,  
 202 69, 70] provides an alternative way for solving the Kohn-Sham electronic  
 203 structure without diagonalization. As a Fermi operator expansion (FOE)  
 204 based method, PEXSI expands the density matrix in Eq. 9 using a  $P$ -term  
 205 pole expansion:

$$n \approx \sum_{l=1}^P \text{Im} (\omega_l^\rho (\mathbf{H} - (z_l + \mu)\mathbf{S})^{-1}). \quad (19)$$

206 Here the complex shifts  $\{z_l\}$  and weights  $\{\omega_l^\rho\}$  are determined through a  
 207 semi-analytic formula based on contour integration, and take only a negligi-  
 208 ble amount of time to compute. The number of terms of the pole expansion  
 209 is proportional to  $\log(\beta\Delta E)$ , where  $\beta = 1/(k_B T)$  is the inverse of the thermal  
 210 energy and  $\Delta E$  is the spectral radius. The logarithmic scaling makes the pole  
 211 expansion a highly efficient approach to expand the Fermi operator. Typi-  
 212 cally 40  $\sim$  80 poles are sufficient for the result obtained from PEXSI to be  
 213 fully comparable ( $\mu\text{eV}/\text{atom}$  [18, 19]) to that obtained from diagonalization.

214 At first it may seem that the entire Green's function-like object  $(\mathbf{H} - (z_l +$   
 215  $\mu)\mathbf{S})^{-1}$  needs to be computed. However, if targeting at the electron density  
 216  $n(\mathbf{r})$ , in general only the entries corresponding to the non-zero pattern of  $\mathbf{H}$   
 217 and  $\mathbf{S}$  are actually needed. Then a selected inversion algorithm can be used  
 218 to efficiently compute these selected elements of the Green's function object,  
 219 and therefore the electron density.

220 The computational cost of the PEXSI technique scales at most as  $O(N^2)$ .  
 221 The actual complexity depends on the dimensionality of the system:  $O(N)$   
 222 i.e. linear scaling for quasi-1D systems such as nanotubes;  $O(N^{1.5})$  for quasi-  
 223 2D systems such as surfaces and slabs; and  $O(N^2)$  for general 3D bulk sys-  
 224 tems. This favorable scaling hinges on the sparse character of the Hamilto-  
 225 nian and overlap matrices, but not on any fundamental assumption about  
 226 the localization properties of the single particle density matrix. This method

227 is not only applicable to the efficient computation of the electron density,  
228 but also to other physical quantities such as the free energy, atomic forces,  
229 density of states and local density of states, all obtainable without comput-  
230 ing any eigenvalues or eigenvectors [18]. These quantities can be given by  
231 pole expansions with the same complex shifts as those used for computing  
232 the electron density, with different weights.

233 PEXSI allows the usage of a hybrid scheme of density of states estimation  
234 based on Sylvester’s law of inertia [71], and Newton’s method to obtain the  
235 chemical potential [19], hereafter referred to as the PEXSI mu iteration. This  
236 is an efficient and relatively robust approach with respect to the initial guess  
237 of the chemical potential, with or without the presence of gap states. A  
238 reasonable initial guess, e.g. obtained from the previous SCF step, can often  
239 converge the PEXSI mu iteration in one step.

240 The PEXSI method has a two-level parallelism structure and is by design  
241 highly scalable. The recently developed massively parallel PEXSI technique  
242 can make efficient use of 10,000  $\sim$  100,000 processors on high performance  
243 machines.

## 244 4. The ELSI Infrastructure

### 245 4.1. Overview of the ELSI Interface

246 KS-DFT is implemented by a broad, diverse ecosystem of different soft-  
247 ware packages with different specialties and different numerical discretization  
248 strategies (see, e.g., Ref. [4] for a listing of 46 packages). The Kohn-Sham  
249 eigenvalue problem is unavoidable in all these packages. Since the most ef-  
250 ficient way to solve the problem may depend on factors such as system size  
251 and character (insulating or metallic), sparsity of matrices involved, density-  
252 functional employed, etc., from a user’s perspective, a library that can dy-  
253 namically switch between different methods according to the features of the  
254 problem is preferred. As a first step to achieve this goal (the objective of this  
255 paper), a flexible interface to different methods should enable user codes to  
256 actively select the most effective method while imposing only a minimum of  
257 format conversions, parameter tweaking, etc. on the user code.

258 Although each solver library supported in ELSI maintains a limited num-  
259 ber of well-explained Application Programming Interfaces (APIs), integrat-  
260 ing all of them into a KS-DFT code is still a complicated, time-consuming,  
261 and error-prone task. ELSI ships a small set of APIs that are designed for  
262 rapid integration of a variety of KS electronic structure solvers into KS-DFT

263 codes, and at the same time provides the user with hierarchical control over  
264 the interface and the solvers. There are three key steps to use ELSI, denoted  
265 by the red boxes (a), (b) and (c) in Fig. 3: (a) The ELSI interface needs  
266 to be initialized at the beginning of an SCF calculation, and potentially  
267 re-initialized if performing successive SCF cycles, e.g. for different system  
268 geometries during a molecular dynamics simulation or during a geometry  
269 optimization calculation. (b) Within the SCF cycle, ELSI serves as a bridge  
270 between the KS-DFT codes and the KS solver libraries, by taking the Hamil-  
271 tonian matrix (and the overlap matrix if it exists) as input, translating the  
272 eigenproblem into a solver-specific format, invoking the solver to compute the  
273 eigenvalues and eigenvectors, or the density matrix, and finally translating  
274 the results back to the native format of the KS-DFT codes. (c) When ELSI  
275 is no longer needed, it should be finalized to deallocate any arrays internally  
276 allocated by ELSI.

#### 277 *4.2. Matrix Storage and Distribution in ELSI*

278 The first emerging practical consideration when developing a unified soft-  
279 ware interface is the choice of matrix storage and distribution strategy. The  
280 sparsity of matrices in KS-DFT varies dramatically from small to large sys-  
281 tems, and from 1D to 3D systems. In general, when using localized basis  
282 functions, the sparsity of matrices increases as the simulated system be-  
283 comes larger. Lower dimensional systems often generate more sparse ma-  
284 trices. Since the effective information is only represented by the non-zero  
285 matrix elements, storing and operating on all the matrix elements lead to  
286 unnecessary memory consumption and computational complexity for very  
287 sparse matrices.

288 Implementing dense linear algebra operations, ELPA and libOMM han-  
289 dle matrices stored densely and distributed in a 2D block-cyclic distribu-  
290 tion, whereas PEXSI performs sparse linear algebra with matrices stored in  
291 compressed sparse column (CSC, also known as compressed column storage,  
292 CCS) format in a 1D block distribution. These two combinations, hereafter  
293 referred to as BLACS\_DENSE and PEXSI\_CSC formats, respectively, are  
294 chosen as the input/output matrix format of the ELSI interface to bridge  
295 the needs of the solvers and of different KS-DFT codes. The comparison  
296 between dense matrix storage and CSC sparse matrix storage is illustrated  
297 in Fig. 4, using an  $8 \times 8$  matrix as an example. The dense storage keeps  
298 all the matrix elements including zeros and non-zeros. The CSC format, in  
299 contrast, drops the zeros and packs the remaining non-zeros into a 1D array,

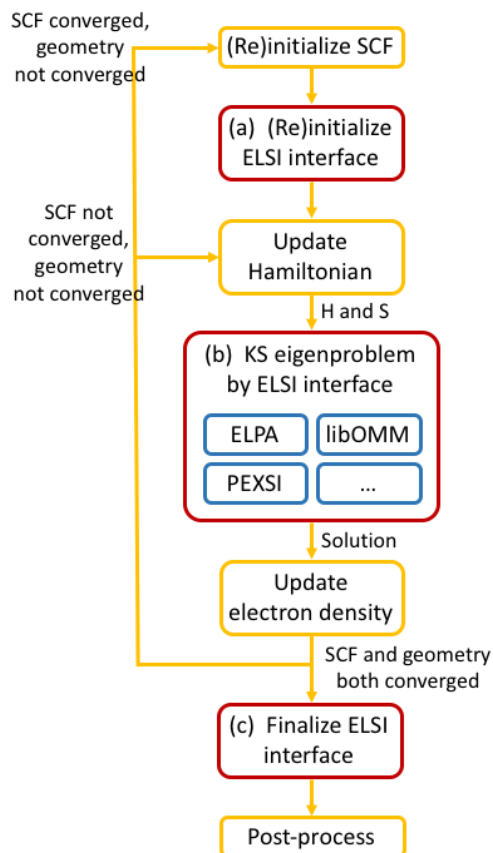


Figure 3: Flow chart describing the key steps in a self-consistent field calculation based on Kohn-Sham Density-Functional Theory. Yellow boxes: Key steps commonly implemented in KS-DFT codes to perform a single SCF cycle or multiple successive SCF cycles with different atomic structures, e.g. for molecular dynamics or for geometry optimizations. Red boxes: Required additions to use the ELSI interface, including (a) initialization of the ELSI interface, (b) computing the eigensolution or the density matrix using the ELSI solvers, and (c) finalization of the ELSI interface.

300 together with the row indices of the non-zero values and the starting points  
 301 of the matrix columns. For a larger matrix with a higher sparsity, the CSC  
 302 format will eventually consume less memory compared to the dense format.

303 To compare the two supported distributions of matrices across multiple  
 304 processors in parallel computations, Fig. 5 shows how the 2D block-cyclic  
 305 and the 1D block distributions are applied to the same  $8 \times 8$  matrix. We note

306 that shown in Fig. 5 are two mathematical matrices, the shapes of which do  
 307 not represent the actual arrays in the computer. The 2D block-cyclic distri-  
 308 bution in Fig. 5 (a) divides the global matrix into several blocks, then maps  
 309 the blocks to the processors in a round-robin fashion in both the row and  
 310 the column directions. The 1D block distribution in Fig. 5 (b) groups con-  
 311 tinuous matrix columns together, then linearly maps the groups of columns  
 312 to the processors. In ELSI, when the input matrices are in a different distri-  
 313 bution from the internally used one, a redistribution of the non-zero matrix  
 314 elements is performed internally, i.e. no unnecessary communication of the  
 315 zero elements. This redistribution is managed by the all-to-all communica-  
 316 tion implemented in the Message Passing Interface (MPI) library. Once the  
 317 matrix is correctly distributed, conversion to various formats is then handled  
 318 concurrently on all the MPI tasks, with each task converting a local matrix  
 319 of the size at most  $N_{\text{basis}}^2/N_{\text{MPI}}$ , where  $N_{\text{MPI}}$  is the number of MPI tasks  
 320 involved.

0	1	0	0	0	0	0	0
0	0	0	0	0	0	8	5
4	0	0	5	0	0	0	7
0	0	0	9	5	0	0	2
0	0	0	1	0	4	0	0
2	0	0	0	0	0	0	0
3	0	0	0	0	0	0	0
0	8	0	0	0	0	0	0

(a)

val	4	2	3	1	8	5	9	1	5	4	8	5	7	2
row_idx	3	6	7	1	8	3	4	5	4	5	2	2	3	4
col_ptr	1	4	6	6	9	10	11	12	15					

(b)

Figure 4: An  $8 \times 8$  matrix stored in (a) dense storage format versus in (b) compressed sparse column (CSC) storage format. In the CSC format, only the values of the non-zero elements, indicated in blue in (a), are stored in the “val” array. The row indexes of the non-zero elements are stored in the “row\_idx” array. The “col\_ptr” array stores the starting points of the matrix columns.

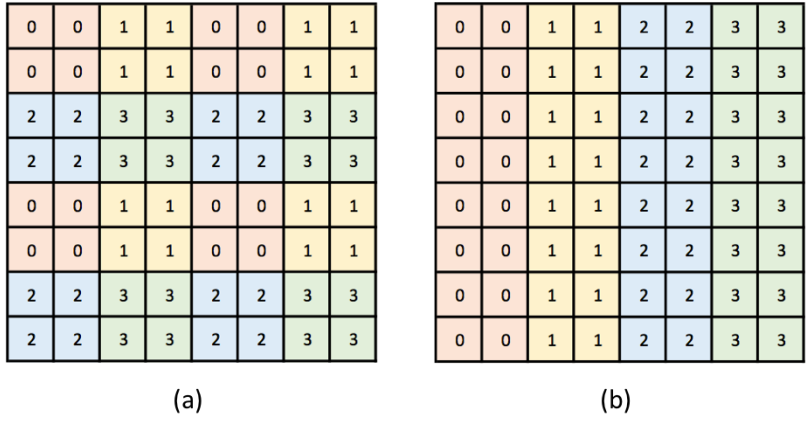


Figure 5: Schematic visualizations of (a) two-dimensional block-cyclic distribution used in the BLACS\_DENSE format, and (b) one-dimensional block distribution used in the PEXSI\_CSC format, of an  $8 \times 8$  matrix on 4 processors. Each unit square represents one matrix element. The integer inside each unit square denotes the index of processor where the element is stored and handled. Different processors are indicated by colors. Shown in the figure are mathematical matrices, not arrays in computers. The actual matrix storage on each processor is arbitrary, e.g. dense storage used by the BLACS\_DENSE format, CSC sparse storage used by the PEXSI\_CSC format.

321 *4.3. Parallelization Strategy and Interaction of ELSI with an Existing KS-*  
 322 *DFT Code*

323 An important distinction in KS-DFT calculations is whether the system  
 324 considered is isolated or is periodically repeated in space. In periodic sys-  
 325 tems, the full problem can be separated into subproblems defined at selected  
 326 k-points in the Brillouin zone, or in a convenient unit cell in reciprocal space.  
 327 The Hamilton and overlap matrices for multiple k-points are block-diagonal,  
 328 such that each block on the diagonal corresponds to an eigenproblem of one  
 329 k-point. These eigenproblems can therefore be solved in a embarrassingly  
 330 parallel fashion side by side. For periodic systems with a small unit cell,  
 331 thousands of k-points or even more can be necessary for an accurate descrip-  
 332 tion of the electronic structure. For a large system, in contrast, the Brillouin  
 333 zone may already be well-represented by the origin of the reciprocal space  
 334 known as the  $\Gamma$  point.

335 Depending on the number of k-points  $N_{\text{kpt}}$  (here defined to be 1 also  
 336 for isolated, non-periodic cases) and the number of MPI tasks  $N_{\text{MPI}}$ , two  
 337 different categories of possible KS-DFT calculations arise, as explained in



338 Fig. 6. Correspondingly, ELSI supports two parallelization strategies that  
339 can be specified by a `parallel_mode` parameter (see also `elsi_init` subroutine  
340 in Section 4.4):

- 341 • `MULTI_PROC` mode, to be used if  $N_{\text{MPI}} \geq N_{\text{kpt}}$ . For instance, there  
342 are 4 k-points in example (a) in Fig. 6, handled by 16 MPI tasks. The  
343 `MULTI_PROC` parallelization divides the 16 MPI tasks into 4 groups.  
344 Each k-point is handled by 4 MPI tasks in the same group, and the  
345 eigenproblems of the 4 k-points are solved simultaneously by the 16  
346 MPI tasks. Since each k-point is solved by multiple MPI tasks (pro-  
347 cesses), this parallelization mode is called `MULTI_PROC`. This mode  
348 should be chosen for isolated systems, periodic systems with only one  
349 k-point, e.g. the  $\Gamma$  point, and periodic systems with  $N_{\text{kpt}}$  k-points  
350 treated by  $N_{\text{MPI}} \geq N_{\text{kpt}}$  MPI tasks.
- 351 • `SINGLE_PROC` mode, to be used if  $N_{\text{MPI}} < N_{\text{kpt}}$ . For instance, there  
352 are 16 k-points in example (b) in Fig. 6, handled by 4 MPI tasks. The  
353 `SINGLE_PROC` parallelization divides the 16 k-points into 4 groups.  
354 Each MPI task handles 4 k-points in the same group, one after another.  
355 Since each k-point is solved by a single MPI task (process), this paral-  
356 lelization mode is called `SINGLE_PROC`. This mode should be chosen  
357 for periodic systems with  $N_{\text{kpt}}$  k-points treated by  $N_{\text{MPI}} < N_{\text{kpt}}$  MPI  
358 tasks.

359 The ELPA eigensolver supported in ELSI is available for both parallel  
360 modes, returning eigensolutions for each k-point. In this case, the KS-DFT  
361 codes can then assemble the pieces of the solutions (eigenvalues and eigenvec-  
362 tors) returned by the solver and construct the electron density. The density  
363 matrix solvers in the 2017.05 release of ELSI do not yet support periodic  
364 calculations with more than one k-point. The ability to return combined  
365 density matrices obtained from ELPA, libOMM, or PEXSI is planned as a  
366 next step of the ELSI interface.

367 Once the matrix storage format and the parallel mode are decided, the  
368 usage of ELSI in KS-DFT codes becomes straightforward. Algorithm 1 sum-  
369 marizes in pseudo-code all the possible use cases of the ELSI interface as of  
370 the 2017.05 release. In Algorithm 1, the main steps are denoted by subrou-  
371 tine names that will be systematically introduced in the following subsec-  
372 tions. Furthermore, the initialization of the SCF calculation, updating the

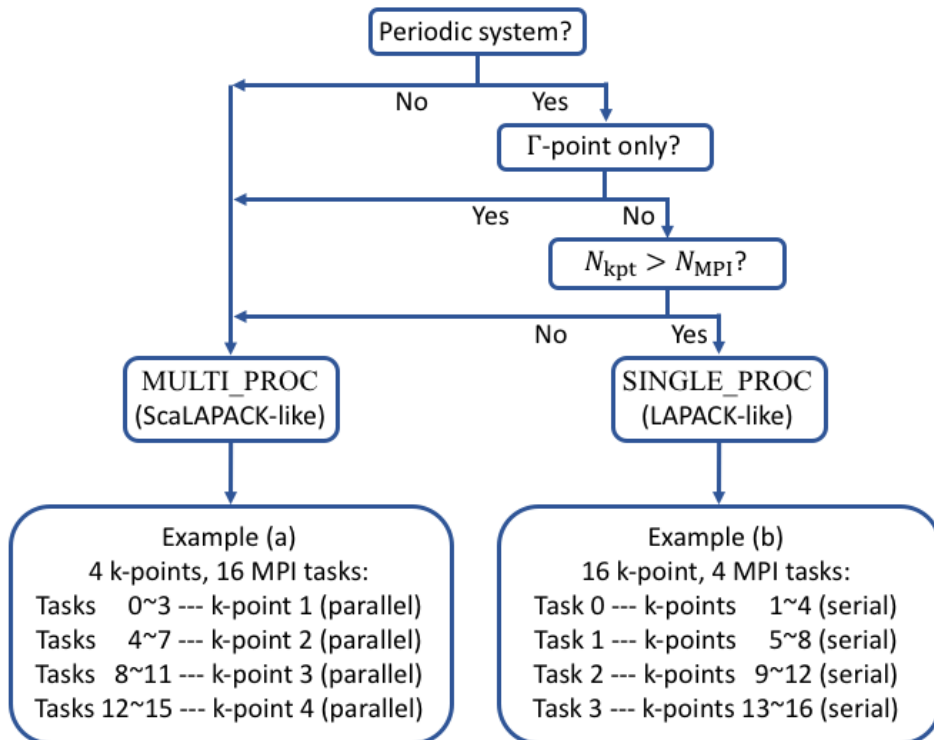


Figure 6: Diagrammatic explanations of the two parallelization strategies supported by ELSI.  $N_{\text{kpt}}$  is the number of k-points.  $N_{\text{MPI}}$  is the number of MPI tasks.

373 Hamiltonian and the electron density, checking the SCF convergence, post-  
 374 processing, and potentially further steps are all tasks that are not handled by  
 375 ELSI but that are instead expected to be executed by the specific KS-DFT  
 376 code calling ELSI.

377 Before showing detailed descriptions of the ELSI API in the next subsections,  
 378 we here first introduce the concept of `elsi_handle`, a Fortran derived  
 379 data type containing all runtime parameters, e.g. the choices of solver, `matrix_storage_format`,  
 380 and `parallel_mode` (see `elsi_init` subroutine in Section 4.4). It is intended to avoid  
 381 global variables in ELSI and to allow concurrent instances of ELSI by passing  
 382 around the handle as arguments. A handle can be initialized with the `elsi_init`  
 383 subroutine, then should be passed to all other ELSI subroutines. The ELSI interface,  
 384 including the `elsi_handle`, is fully interoperable with C and C++ programming  
 385 languages. The `elsi_handle` is

386 defined in C/C++ as an “opaque” pointer, which can be seamlessly con-  
387 nected to a derived data type in Fortran by the `iso_c_binding` feature in  
388 Fortran compilers.

#### 389 *4.4. ELSI Initialization*

390 In this and the following subsections (Sections 4.5, 4.6, 4.7), we provide  
391 details of the capabilities of the ELSI interface as current in the 2017.05  
392 release. Since these capabilities are intimately tied to the actual implemen-  
393 tation, we here explain them grouped by individual subroutines as also shown  
394 in Algorithm 1. In all instances, `elsi_h` denotes the ELSI handle.

395 In the initialization phase, ELSI can be set up to reflect the physical  
396 quantities that usually do not change within an SCF calculation (i.e. fixed  
397 atomic structure), such as the number of basis functions and the number of  
398 electrons in the system. Implementations of SCF typically initialize these  
399 quantities before the SCF cycle begins, then keep reusing them within the  
400 cycle to repeatedly solve KS problems with an updated Hamiltonian matrix  
401 and a fixed overlap matrix. Similarly, the ELSI interface only needs to be (re-  
402 )initialized whenever the SCF cycle is itself (re-)initialized. The subroutines  
403 that are used to initialize ELSI include:

- 404 • **elsi\_init** (`elsi_h`, `solver`, `matrix_storage_format`, `parallel_mode`, `n_basis`,  
405 `n_state`, `n_electron`)  
406 (line 3 in Algorithm 1) – Initializes an ELSI handle with user’s choices of  
407 the solver, the matrix format and distribution, the parallelization strat-  
408 egy, and system information including the number of basis functions,  
409 the number of eigenstates to compute, and the number of electrons.  
  
410 – **elsi\_h** (`type(elsi_handle)`, `output`): An ELSI handle (see Section  
411 4.3 returned by `elsi_init` subroutine. The same handle must be  
412 passed to other ELSI subroutines and be finalized when no longer  
413 needed. Multiple handles can be initialized if needed.  
  
414 – **solver** (`integer`, `input`): The choice of solver. Accepted options  
415 are 1 (ELPA), 2 (libOMM), and 3 (PEXSI).  
  
416 – **matrix\_storage\_format** (`integer`, `input`): Matrix storage and  
417 distribution of the Hamiltonian matrix, the overlap matrix, and  
418 the density matrix or the eigenvectors. Accepted options are 1  
419 (BLACS\_DENSE) and 2 (PEXSI\_CSC) (see Section 4.2). The

---

**Algorithm 1** Usage of ELSI interface in KS-DFT codes. Pseudo-code in line 3-11, line 14-27, and line 31 corresponds to Fig. 3 (a), (b), and (c), respectively.

---

```
1: procedure ELSI
2:   initialize SCF calculation
3:   call elsi_init
4:   if (parallel_mode = MULTI_PROC) then
5:     call elsi_set_mpi
6:     if (matrix_storage_format = BLACS_DENSE) then
7:       call elsi_set_blacs
8:     else if (matrix_storage_format = PEXSI_CSC) then
9:       call elsi_set_csc
10:    end if
11:  end if
12:  while (SCF not converged) do
13:    update Hamiltonian
14:    call elsi_customize
15:    if (desired output: eigensolution) then
16:      if (matrix_storage_format = BLACS_DENSE) then
17:        call elsi_ev_{real|complex}
18:      else if (matrix_storage_format = PEXSI_CSC) then
19:        call elsi_ev_real_sparse
20:      end if
21:    else if (desired output: density matrix) then
22:      if (matrix_storage_format = BLACS_DENSE) then
23:        call elsi_dm_real
24:      else if (matrix_storage_format = PEXSI_CSC) then
25:        call elsi_dm_real_sparse
26:      end if
27:    end if
28:    update electron density
29:    check SCF convergence
30:  end while
31:  call elsi_finalize
32:  post-process
33: end procedure
```

---

420 BLACS\_DENSE format is compatible with ELPA, libOMM, and  
421 PEXSI. If the chosen solver is PEXSI, the input matrices in the  
422 BLACS\_DENSE format are converted to PEXSI\_CSC internally,  
423 and the results in the PEXSI\_CSC format are back-converted  
424 to BLACS\_DENSE. The PEXSI\_CSC format is compatible with  
425 ELPA and PEXSI in the current release. Supporting the PEXSI\_CSC  
426 format with libOMM is on the list of features to be added in the  
427 near future.

- 428 – **parallel\_mode** (integer, input): The choice of parallelization strat-  
429 egy. Accepted options are 1 (MULTI\_PROC) and 2 (SINGLE\_PROC)  
430 (see Section 4.1). In the current release of ELSI, the SINGLE\_PROC  
431 mode is only compatible with ELPA, while the MULTI\_PROC  
432 mode supports all three solvers.
- 433 – **n\_basis** (integer, input): Number of basis functions. This is equal  
434 to the global size of the Hamiltonian matrix, the overlap matrix,  
435 the density matrix, etc.
- 436 – **n\_state** (integer, input): Number of states. For ELPA this is the  
437 number of eigenstates to be solved. For libOMM this must be  
438 the number of occupied states, without any fractional occupation  
439 numbers. PEXSI does not use this information.
- 440 – **n\_electron** (integer, input): Number of electrons.

- 441 • **elsi\_set\_mpi** (elsi.h, mpi\_comm)  
442 (line 5 in Algorithm 1) – Sets the MPI communicator to be used in the  
443 ELSI instance indicated by the handle.
- 444 – **mpi\_comm** (integer, input): An MPI communicator, containing  
445 an ordered group of MPI tasks, is required to use the functionali-  
446 ties implemented in the MPI library. The communicator assigned  
447 to an ELSI calculation can be the default global communicator  
448 of MPI, or a communicator created by the user (e.g. by calling  
449 the MPI subroutine MPI\_Comm\_Split), as long as it is compatible  
450 with the distribution of matrices.
- 451 • **elsi\_set\_blacs** (elsi.h, blacs\_ctxt, block\_size)

452 (line 7 in Algorithm 1) – Sets the BLACS context and the block  
453 size of the 2D block-cyclic distribution to be used in the ELSI in-  
454 stance indicated by the handle. Required before calling `elsi_ev_real`,  
455 `elsi_ev_complex`, and `elsi_dm_real` (see Section 4.5).

456 – **blacs\_ctxt** (integer, input): A BLACS context encloses a group  
457 of processes and arranges them in a particular grid. Processes in  
458 the same context can safely communicate with each other, without  
459 worrying if the operations in one context interfere with operations  
460 in another context [60]. The ELSI interface requires the KS-DFT  
461 code to set up BLACS context(s), by calling BLACS subroutine  
462 `BLACS_Gridinit` or `BLACS_Gridmap`.

463 – **block\_size** (integer, input): The block size parameter of the 2D  
464 block-cyclic distribution. The matrix operations inside ELSI in-  
465 terface, ELPA, and libOMM restrict the block sizes in the row  
466 and column directions to be the same.

467 • **elsi\_set\_csc** (`elsi_h`, `nnz_g`, `nnz_l`, `n_l_cols`, `row_idx`, `col_ptr`)

468 (line 9 in Algorithm 1) – Set the parameters of 1D block distributed  
469 CSC matrix storage (PEXSILCSC) to be used in the ELSI instance  
470 indicated by the handle. Required before calling `elsi_ev_real_sparse`  
471 and `elsi_dm_real_sparse` (see Section 4.5).

472 – **nnz\_g** (integer, input): The global number of non-zero elements  
473 in the sparsity pattern.

474 – **nnz\_l** (integer, input): The local number of non-zero elements in  
475 the sparsity pattern held by an MPI task.

476 – **n\_l\_cols** (integer, input): The local number of matrix columns  
477 held by an MPI task.

478 – **row\_idx** (integer, 1D array, input): The row index array of the  
479 CSC matrix storage format, containing the row index of each non-  
480 zero matrix element. An example is given in Fig. 4.

481 – **col\_ptr** (integer, 1D array, input): The column pointer array of  
482 the CSC matrix storage format, containing the starting point of  
483 each matrix column. An example is given in Fig. 4.

484 The matrices that arise in KS-DFT can be either real or complex-valued.  
 485 ELSI must account for these two possibilities as well. Since the real and com-  
 486 plex arithmetic cases only differ in the data type of input/output matrices,  
 487 they are not distinguishable at the initialization stage.

#### 488 4.5. Tasks during SCF

489 During the SCF cycle, the following tasks may be executed by ELSI solver  
 490 subroutines to compute either the eigensolutions or the density matrix from  
 491 the input Hamiltonian matrix (and overlap matrix, if it is not unity).

- 492 • **elsi\_ev\_real** (elsi\_h, ham, ovlp, eval, evec)  
 493 (line 17 in Algorithm 1) – Computes the eigenvalues and n\_state eigen-  
 494 vectors. Compatible solver: ELPA.
- 495 – **ham** (double precision real, 2D array, input & output): The real-  
 496 valued Hamiltonian matrix in the BLACS\_DENSE format set by  
 497 subroutine elsi\_set\_blacs. This array is used for internal storage  
 498 when solving the eigenproblem, and thus is destroyed on exit.
- 499 – **ovlp** (double precision real, 2D array, input & output): The real-  
 500 valued overlap matrix in the BLACS\_DENSE format set by sub-  
 501 routine elsi\_set\_blacs.

502 A singularity check of the overlap matrix  $\mathbf{S}$  is performed the first  
 503 time `elsi_ev_real` is called. This is because the Cholesky factor-  
 504 ization in Eq. 10 requires  $\mathbf{S}$  to be Hermitian positive-definite.  
 505 While  $\mathbf{S}$  in KS-DFT is guaranteed to be Hermitian by Eq. 5,  
 506 the positive-definite condition can be numerically violated if the  
 507 chosen basis set is large and (near-)singular, i.e. the lowest eigen-  
 508 values of  $\mathbf{S}$  are too close to 0 (although still greater than 0). Using  
 509 a near-singular basis set can lead to completely wrong and unpre-  
 510 dictable numerical results, and thus should be avoided in general.  
 511 In ELSI, this is done by computing all the eigenvalues of  $\mathbf{S}$  and  
 512 comparing them with a user-defined singularity tolerance  $\tau$ . The  
 513 matrix is considered to be singular if it has one or more eigenval-  
 514 ues smaller than  $\tau$ . For a singular overlap matrix, the Cholesky  
 515 decomposition is replaced by an eigendecomposition:

$$\mathbf{S} = (\sqrt{\lambda}\mathbf{x})(\sqrt{\lambda}\mathbf{x})^* = \mathbf{X}\mathbf{X}^*, \quad (20)$$

516 where the matrix  $\mathbf{x}$  and the diagonal matrix  $\boldsymbol{\lambda}$  contain the eigen-  
 517 vectors and eigenvalues of  $\mathbf{S}$ , and the matrix  $\mathbf{X}$  is simply  $\sqrt{\boldsymbol{\lambda}}\mathbf{x}$ . By  
 518 using eigendecomposition, the generalized eigenproblem is again  
 519 transformed to the standard form in Eq. 11, with  $\tilde{\mathbf{H}} = \mathbf{X}^{-1}\mathbf{H}(\mathbf{X}^*)^{-1}$   
 520 and  $\tilde{\mathbf{c}} = \mathbf{X}^*\mathbf{c}$ . In case that only the first  $N_{\text{nonsing}}$  eigenvalues of  
 521  $\mathbf{S}$  are greater than the threshold  $\tau$ ,  $\mathbf{X}$  correspondingly contains  
 522 only the first  $N_{\text{nonsing}}$  eigenvectors by dropping the  $N_{\text{basis}} - N_{\text{nonsing}}$   
 523 eigenvectors associated with small eigenvalues. The eigenproblem  
 524 transformation is still valid, however yields a smaller transformed  
 525  $\tilde{\mathbf{H}}$  ( $N_{\text{nonsing}} \times N_{\text{nonsing}}$ ). The solution of the transformed standard  
 526 eigenproblem must be back-transformed accordingly.

527 On exit, `ovlp` is overwritten by either  $\mathbf{L}$  in Eq. 10 or  $\mathbf{X}$  in  
 528 Eq. 20, depending on which transformation is used. If in the  
 529 `MULTI_PROC` mode, i.e. no MPI task handles more than one k-  
 530 point,  $\mathbf{L}$  or  $\mathbf{X}$  can be stored in `ovlp` and efficiently reused through-  
 531 out the SCF cycle. The Cholesky factorization or the eigendecom-  
 532 position then only needs to be performed once. However, in the  
 533 `SINGLE_PROC` mode, since each MPI task handles a group of k-  
 534 points in serial, memory constraints make it more difficult to reuse  
 535 the matrices  $\mathbf{L}$  or  $\mathbf{X}$ . In this case, the decision to either store  $\mathbf{L}$   
 536 or  $\mathbf{X}$ , or to redo the decomposition in every SCF iteration, is up  
 537 to the KS-DFT code that calls `ELSI`.

538 – **eval** (double precision real, 1D array, output): The eigenvalues in  
 539 ascending order.

540 – **evect** (double precision real, 2D array, output): The real-valued  
 541 eigenvectors in a matrix form in the `BLACS_DENSE` format set  
 542 by subroutine `elsi_set_blacs`.

543 • **elsi\_ev\_complex** (`elsi_h`, `ham`, `ovlp`, `eval`, `evect`)

544 (line 17 in Algorithm 1) – Same as `elsi_ev_real`, except that the Hamil-  
 545 tonian matrix, overlap matrix and eigenvectors are complex-valued.

546 • **elsi\_ev\_real\_sparse** (`elsi_h`, `ham`, `ovlp`, `eval`, `evect`)

547 (line 19 in Algorithm 1) – Computes the eigenvalues and `n_state` eigen-  
 548 vectors. Compatible solver: `ELPA`.

549 – **ham** (double precision real, 1D array, input): The non-zero ele-  
 550 ments of the real-valued Hamiltonian matrix in the `PEXSL_CSC`



551 format set by subroutine `elsi_set_csc`. Inside ELSI, the input Hamil-  
552 tonian matrix is converted to the BLACS\_DENSE format in every  
553 SCF iteration.

- 554 – **ovlp** (double precision real, 1D array, input): The non-zero ele-  
555 ments of the real-valued overlap matrix in the PEXSI\_CSC format  
556 set by subroutine `elsi_set_csc`. Inside ELSI, the input overlap ma-  
557 trix is converted to the BLACS\_DENSE format in the first SCF  
558 iteration. The singularity check of the overlap matrix is performed  
559 as in the `elsi_ev_real` case. Since the sparsity of the eigenproblem  
560 transformation matrix  $\mathbf{L}$  or  $\mathbf{X}$  is not guaranteed, the matrix  $\mathbf{L}$  or  
561  $\mathbf{X}$  is stored internally in the BLACS\_DENSE format for further  
562 reuse throughout the SCF cycle.
- 563 – **eval** (double precision real, 1D array, output): The eigenvalues in  
564 ascending order.
- 565 – **evvec** (double precision real, 2D array, output): The real-valued  
566 eigenvectors in a matrix form in the BLACS\_DENSE format. Note  
567 that the computed eigenvectors are returned in a dense format, for  
568 the reason that they are not in the same sparsity pattern of  $\mathbf{H}$   
569 and  $\mathbf{S}$ , or even not sparse at all.

- 570 • **elsi\_dm\_real** (`elsi_h`, `ham`, `ovlp`, `den_mat`, `energy`)  
571 (line 23 in Algorithm 1) – Computes the density matrix. Compatible  
572 solvers: ELPA, libOMM, PEXSI.
- 573 – **ham** (double precision real, 2D array, input & output): The real-  
574 valued Hamiltonian matrix in the BLACS\_DENSE format set by  
575 subroutine `elsi_set_blacs`. This array is used for internal storage  
576 when computing the density matrix, and thus is destroyed on exit.  
577 If the chosen solver is PEXSI, the input Hamiltonian matrix is  
578 converted to the PEXSI\_CSC format in every SCF iteration.
- 579 – **ovlp** (double precision real, 2D array, input & output): The real-  
580 valued overlap matrix in the BLACS\_DENSE format set by sub-  
581 routine `elsi_set_blacs`. If the chosen solver is PEXSI, the input  
582 overlap matrix is converted to the PEXSI\_CSC format in the first  
583 SCF iteration and reused throughout the SCF cycle. If the chosen  
584 solver is ELPA or libOMM, the singularity check of the overlap

585 matrix is performed as in the `elsi_ev_real` case. The singularity  
586 check is not yet implemented for PEXSI.

587 – **den\_mat** (double precision real, 2D array, output): The density  
588 matrix in the BLACS\_DENSE format set by subroutine `elsi_set_blacs`.  
589 The chemical potential and occupation numbers must be known  
590 when ELPA is chosen to compute the density matrix following  
591 Eq. 9. In ELSI, the chemical potential is found using a bisection  
592 algorithm that starts from an energy interval that includes the  
593 actual solution of the chemical potential. This is often guaran-  
594 teed by using the lowest and highest eigenvalues of the system as  
595 the lower and upper bounds of the interval, and expanding the  
596 interval towards both ends if necessary. In each bisection step  
597 the number of electrons on both bounds and at the middle point  
598 of the interval is computed by Eq. 18 (the summation becomes  
599  $\sum_{i=1}^{N_{\text{kpt}}} \sum_{j=1}^{N_{\text{spin}}} \sum_{k=1}^{N_{\text{basis}}}$  if including k-points and spin channels), to  
600 determine which subinterval the solution lies in. Then the interval  
601 can be repeatedly bisected until the computed number of electrons  
602 on either bound or at the middle point is sufficiently close to the  
603 actual number. During this process, the computation of occupa-  
604 tion numbers  $f_l$  requires a specific broadening scheme, which can  
605 be the Fermi broadening in Eq. 17, or the Gaussian broadening  
606 [72]

$$f_l = 0.5 \cdot [1 - \operatorname{erf}\left(\frac{\epsilon_l - \mu}{k_B T}\right)], \quad (21)$$

607 where `erf` is the Gauss error function:

$$\operatorname{erf} = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt. \quad (22)$$

608 Although the error function is implemented as an intrinsic function  
609 in most programming languages, the error of each single evalua-  
610 tion can accumulate as a consequence of the summation in Eq.  
611 18. This accumulation leads to an error on the order of  $10^{-10}$   
612 in term of the number of electrons, which is small but not neg-  
613 ligible if the desired accuracy is on the same order. During the

614 convergence of the SCF cycle, this small error can become more  
615 noticeable, since fluctuations of the norm of the density matrix  
616 (i.e. the system charge) will have a relatively large electrostatic  
617 effect, and can thus disturb the solution of the nonlinear fixed-  
618 point iteration scheme (e.g. Pulay mixing [73]) that is used to  
619 converge an SCF cycle. Therefore, it is useful and sometimes nec-  
620 essary to avoid charge fluctuations whenever possible, by ensuring  
621 an exact charge norm after the fact. In ELSI, when the accuracy  
622 of electron count can no longer be improved by bisection, then the  
623 remaining discrepancy (surplus of electrons in case of the upper  
624 bisection bound) is successively removed starting from the highest  
625 occupied KS states and proceeding to lower-lying states until the  
626 norm in Eq. 18 is numerically exactly fulfilled.

627 – **energy** (double precision real, output): The energy corresponding  
628 to the occupied eigenstates.

629 • **elsi\_dm\_real\_sparse** (elsi\_h, ham, ovlp, den\_mat, energy)

630 (line 25 in Algorithm 1) – Computes the density matrix. Compatible  
631 solver: PEXSI.

632 – **ham** (double precision real, 1D array, input & output): The  
633 non-zero elements of the real-valued Hamiltonian matrix in the  
634 PEXSI\_CSC format set by subroutine elsi\_set\_csc. This array is  
635 used for internal storage when computing the density matrix, and  
636 thus is destroyed on exit.

637 – **ovlp** (double precision real, 1D array, input & output): The non-  
638 zero elements of the real-valued overlap matrix in the PEXSI\_CSC  
639 format set by subroutine elsi\_set\_csc.

640 – **den\_mat** (double precision real, 1D array, output): The non-zero  
641 elements of the density matrix in the PEXSI\_CSC format set by  
642 subroutine elsi\_set\_csc.

643 – **energy** (double precision real, output): The energy corresponding  
644 to the occupied eigenstates.

645 • **elsi\_collect\_pexsi** (elsi\_h, mu, e\_den\_mat, f\_den\_mat)

646 – Collects additional results computed by PEXSI. Compatible solver:  
647 PEXSI.

- 648       – **mu** (double precision real, output): The chemical potential computed by PEXSI.
- 649
- 650       – **e\_den\_mat** (double precision real, 1D array, output): The non-zero elements of the energy density matrix in the PEXSI\_CSC format set by subroutine `elsi_set_csc`.
- 651
- 652
- 653       – **f\_den\_mat** (double precision real, 1D array, output): The non-zero elements of the free energy density matrix in the PEXSI\_CSC format set by subroutine `elsi_set_csc`.
- 654
- 655

#### 656 4.6. ELSI Customization Options

657       Although ELSI sets reasonable default runtime parameters for each solver whenever possible, no set of parameters can adequately cover all use cases. 658 The `elsi_customize` subroutines allow a user to determine runtime parameters explicitly, thus providing maximum flexibility to control the particulars 659 of ELSI. Designed with the feature of optional arguments in Fortran, the `elsi_customize` subroutines have a general calling syntax: 660

661       call `elsi_customize(elsi_h, keyword=choice)`,

662       where `elsi_h` is the ELSI handle to be customized, “keyword” is the parameter to be customized, and “choice” is the value to overwrite the default value 663 of “keyword”. Calling `elsi_customize` (line 14 in Algorithm 1) only modifies 664 the parameter associated with `elsi_h`, instead of changing the behavior of all 665 handles. 666

- 667       • **elsi\_customize** (`elsi_h, keyword=choice`)

668       The following customizable keywords are particularly important:

- 671       – **overlap\_is\_unit** (logical, input): ELSI by default assumes that 672 the KS eigenproblem is a generalized problem (Eq. 6). Setting the keyword `overlap_is_unit` to true allows the usage of ELSI for 673 a standard eigenproblem, e.g. when using orthonormal basis sets, or the generalized eigenproblem has been transformed to the standard 674 form by the calling code itself. If `overlap_is_unit` is true, the singularity check for the overlap matrix described in Section 4.5 675 will be completely ignored. 676
- 677       – **zero\_threshold** (double precision real, input): Threshold to define “zero” in ELSI matrix format conversions. When converting 678 a dense matrix into a sparse format, any double precision number smaller than this threshold is overwritten by 0. 679
- 680
- 681
- 682

683       – **no\_singularity\_check** (logical, input): The singularity check of  
684       the overlap matrix can be skipped here.

685       – **singularity\_tolerance** (double precision real, input): The toler-  
686       ance of basis singularity  $\tau$  in the singularity check.

687     • **elsi\_customize\_mu** (elsi.h, keyword=choice)

688       Customizes the chemical potential and occupation number computa-  
689       tion in ELSI. Customizable keywords include:

690       – **broadening\_scheme** (integer, input): The broadening scheme to  
691       be used in the determination of occupation numbers and chemical  
692       potential. Accepted options are 1 (Gaussian broadening), 2 (Fermi  
693       broadening), 3 (0<sup>th</sup> order Methfessel-Paxton broadening), and 4  
694       (1<sup>st</sup> order Methfessel-Paxton broadening).

695       – **broadening\_width** (double precision real, input): The broaden-  
696       ing width parameter ( $k_B T$  in Eq. 17 and 21).

697       – **occ\_accuracy** (double precision real, input): Desired accuracy in  
698       terms of the sum of occupation numbers, i.e. the number of elec-  
699       trons, in the determination of occupation numbers and chemical  
700       potential.

701       – **mu\_max\_steps** (integer, input): Maximum steps of the bisection  
702       algorithm (described as a part of subroutine `elsi_dm_real`) to  
703       compute the occupation numbers and chemical potential.

704     • **elsi\_customize\_elpa** (elsi.h, keyword=choice)

705       Customizes the ELPA solver. Customizable keywords include:

706       – **elpa\_solver** (integer, input). The choice of ELPA solvers. Ac-  
707       cepted options are 1 (ELPA 1-stage solver) and 2 (ELPA 2-stage  
708       solver).

709     • **elsi\_customize\_omm** (elsi.h, keyword=choice)

710       Customizes the libOMM solver. Customizable keywords include:

711       – **omm\_flavor** (integer, input): The choice of method to perform  
712       OMM minimization. Accepted options are 0 (the basic flavor that

713 follows Eq. 16 exactly) and 2 (the Cholesky flavor that transforms  
714 the generalized eigenproblem to the standard form using Cholesky  
715 factorization before minimization).

- 716 – **n\_elpa\_steps** (integer, input): ELPA can be employed in the first  
717 n\_elpa\_steps SCF iterations, as these take the longest time to con-  
718 verge with iterative methods. Starting from the (n\_elpa\_steps + 1)<sup>th</sup>  
719 SCF step, the libOMM solver will be used with the eigenvectors  
720 computed by ELPA in the (n\_elpa\_steps)<sup>th</sup> SCF step as the initial  
721 guess for the coefficients of Wannier functions.
- 722 – **omm\_tolerance** (double precision real, input): The stop cri-  
723 terion of the OMM energy functional minimization in Eq. 16.  
724 This minimization is considered to be converged when the rela-  
725 tive energy difference between subsequent line searches given by  
726  $2(E[\mathbf{W}_1] - E[\mathbf{W}_0]) / (E[\mathbf{W}_1] + E[\mathbf{W}_0])$  is smaller than or equal to  
727 this dimensionless value.

728 The convergence rate of the OMM energy functional minimization de-  
729 pends heavily on the minimization method and the initial guess of the  
730 coefficients of the Wannier functions. The effects of omm\_flavor and  
731 n\_elpa\_steps on the performance of OMM are investigated and reported  
732 in Section 5.5.

733 • **elsi\_customize\_pexsi** (elsi\_h, keyword=choice)

734 Customizes the PEXSI solver. Customizable keywords include:

- 735 – **n\_poles** (integer, input): The number of poles in the Fermi oper-  
736 ator expansion, i.e.  $P$  in Eq. 19. The pole expansion is an exact  
737 algorithm if the number of poles is infinitely large. In practice,  
738 40 ~ 80 poles are usually sufficient for the result obtained from  
739 PEXSI to be fully comparable to that obtained from diagonaliza-  
740 tion. Performing a convergence test with increasing number of  
741 poles is a practical approach to estimate the optimal number of  
742 poles for a KS-DFT code.
- 743 – **n\_electron\_accuracy** (double precision real, input): The desired  
744 accuracy in term of the number of electrons out of the density  
745 matrix approximated by Eq. 19.

- 746 – **temperature** (double precision real, input): The physical mean-  
747 ing of the temperature here is the energy  $\beta = K_B T$  in Eq. 17, i.e.  
748 the broadening width.
- 749 – **delta\_e** (double precision real, input): The upper bound for the  
750 spectral radius  $\Delta E$  of  $S^{-1}H$ . This parameter and the  $\beta$  parameter  
751 affect the number of terms of the pole expansion.
- 752 – **max\_iteration** (integer, input): The maximum number of PEXSI  
753 mu iterations to determine the chemical potential.
- 754 – **mu\_0, mu\_min, mu\_max** (double precision real, input): The  
755 initial guess, lower bound, and upper bound for the chemical po-  
756 tential. A good initial guess significantly accelerates the conver-  
757 gence of the PEXSI mu iteration. An estimate of the chemical  
758 potential is available in PEXSI via the inertia counting procedure  
759 based on Sylvester’s law of inertia. Starting from the second SCF  
760 iteration, if the change in chemical potential from the previous  
761 SCF step to the current step is small, ELSI will automatically  
762 skip the inertia counting and use the chemical potential from the  
763 previous step as the initial guess for the current step.
- 764 – **mu\_safeguard** (double precision real, input): A fail-safe approach  
765 designed for the PEXSI mu iteration. If the error in the chemical  
766 potential computed by PEXSI is larger than this safeguard, the  
767 code will exit the mu iteration and re-invoke the inertia counting  
768 to estimate the chemical potential.

#### 769 4.7. ELSI Finalization

- 770 • **elsi\_finalize** (elsi\_h)  
771 (line 31 in Algorithm 1) – Terminates the ELSI instance associated  
772 with the handle. This deallocates any arrays internally allocated by  
773 ELSI.
- 774 – **elsi\_h** (type(elsi\_handle), input & output): On exit, all the pa-  
775 rameters of this handle are reset to “UNSET” or their default  
776 values. To become valid again, the handle must be re-initialized  
777 by `elsi_init`.

778 *4.8. ELSI Software in Practice*

779 The 2017.05 release of the ELSI software package, available on the “ELSI  
780 Interchange” website (<http://elsi-interchange.org>), contains the ELSI inter-  
781 face described in Section 4, as well as redistributed source code of the three  
782 solver libraries ELPA (version 2016.11.001.pre, <http://elpa.mpcdf.mpg.de>),  
783 libOMM (version 0.0.1, <http://esl.cecam.org/LibOMM>), and PEXSI (ver-  
784 sion 0.10.2, <http://pexsi.org>). They are redistributed with ELSI for an op-  
785 tional integrated installation managed by a unified make-based build system  
786 with specific keywords set by the users in “make.sys” files. While we focus  
787 more on the development of a unified interface to connect the KS solvers  
788 and the KS-DFT codes, the ELPA, libOMM, and PEXSI solvers themselves  
789 are being actively developed by their own communities. The three solvers  
790 linked into ELSI can be either the built-in versions shipped with ELSI,  
791 or independently built versions, e.g. pre-installed and optimized versions  
792 available on a given supercomputer. There are two external dependencies  
793 that must be downloaded and installed separately: the ParMETIS library  
794 (<http://glaros.dtc.umn.edu/gkhome/metis/parmetis/overview>) and the Su-  
795 perLU\_DIST library (<http://crd-legacy.lbl.gov/~xiaoye/SuperLU>).

796 ELSI can be integrated directly into relevant pieces of KS-DFT codes  
797 written in Fortran, C, or C++. So far, ELSI has been tested in the DGDFT  
798 [52], FHI-aims [33], NWChem [26] (via Global Arrays Toolkit [74]), and  
799 SIESTA [36] software packages. Detailed instructions on how to obtain,  
800 install, and use the ELSI software are documented in the ELSI User’s Guide  
801 [75].

802 **5. Benchmarks and Discussions**

803 In the final part of this work, we present a comparative study of the three  
804 KS electronic structure solvers ELPA, libOMM, and PEXSI, as currently  
805 supported by ELSI. This study employs a consistent set of systems and  
806 settings, and illustrates the optimal choice of solver strategies in different  
807 scenarios and system size ranges. The Hamiltonian and overlap matrices are  
808 constructed from actual DFT-PBE [76] calculations using the all-electron,  
809 full-potential electronic structure code FHI-aims (Fortran) with a “tier 1”  
810 numeric atom-centered orbital (NAO) basis set [33, 77], and the pseudopo-  
811 tential code DGDFT (C++) with an adaptive local basis (ALB) [51, 52].  
812 Both packages have been demonstrated to perform large-scale DFT calcula-  
813 tions with at least thousands of atoms [15, 52, 78]. Details of the KS-DFT



814 code specific settings are given in Appendix A and Appendix B, respectively.  
 815 As the benchmark systems, we selected 2D graphene supercell models with  
 816 sizes ranging from 1,800 to 11,520 atoms. All calculations reported here  
 817 are  $\Gamma$ -point-only (the ELSI interface is thus in MULTI.PROC mode) and  
 818 real arithmetic. Among the benchmark problems, the graphene  $30 \times 30 \times 1$ ,  
 819  $45 \times 45 \times 1$ , and  $60 \times 60 \times 1$  supercell models have a small band gap of  
 820 about 0.002 meV, since the Dirac cone of graphene, whose coordinates in  
 821 the reciprocal space are  $(1/3, 1/3, 0)$ , is included in the folded images of the  
 822  $\Gamma$  point. The other graphene models have a band gap of  $0.34 \sim 0.51$  eV.  
 823 The dimensions of the models, the number of employed basis functions, and  
 824 the sparsity factor of the corresponding matrices are reported in Table 1.  
 825 The maximum differences in the converged total energies are  $6.3 \mu\text{eV}/\text{atom}$   
 826 between the results obtained with ELPA and libOMM, and  $0.8 \mu\text{eV}/\text{atom}$   
 827 between ELPA and PEXSI. We note that separate benchmarks of ELPA,  
 828 libOMM, and PEXSI applied to insulating/semiconducting, 1D/3D systems  
 829 have been reported in earlier publications [14, 15, 16, 17, 18].

830 We here report, to our knowledge, the first directly comparable bench-  
 831 mark of all three approaches for the same system and using exactly the same  
 832 hardware and software environment. All computations were performed on  
 833 the Cray XC30 supercomputer Edison at National Energy Research Scien-  
 834 tific Computing Center (NERSC). Each node of Edison is equipped with  
 835 two 12-core Intel Ivy Bridge processors. The nodes were fully exploited by  
 836 launching 24 MPI tasks on each node. No multi-threading parallelization  
 837 was employed.

### 838 *5.1. Performance of the ELPA, libOMM, PEXSI Solvers*

839 We first compare the performance of the key computational steps of the  
 840 ELSI solvers that are repeated in every SCF iteration. These repeated steps  
 841 are: transforming the eigenproblem (Eq. 11), solving the standard eigen-  
 842 problem (Fig. 1), and back-transforming the eigenvectors in ELPA; the min-  
 843 imization (CG line search) of OMM energy functional (Eqs. 15 and 16),  
 844 and the construction of density matrix from the final Wannier functions in  
 845 libOMM; the numerical factorization and the selected inversion of the object  
 846  $\mathbf{H} - (z_l + \mu)\mathbf{S}$  (Eq. 19), and the construction of density matrix from the  
 847 poles in PEXSI. There are other computationally expensive steps that only  
 848 occur in the first SCF iteration and have less significant effects on the total  
 849 time of an SCF cycle. The performance of those steps is discussed separately

Table 1: Supercell size, number of atoms  $N_{\text{atom}}$ , number of basis functions  $N_{\text{basis}}$ , and sparsity factor  $N_{\text{zero}}/N_{\text{basis}}^2$  of the graphene systems used in this work.  $N_{\text{zero}}$  is the number of zero elements in the Hamiltonian matrices. FHI-aims models contain 2 carbon atoms in each unit cell, and results are shown in Figs. 7, 8, 9, 10, 11, 12, and A.14. DGDFT models contain 10 graphene layers (20 carbon atoms) in each unit cell, and results are shown in Fig. 13.

Code	Model	Supercell	$N_{\text{atom}}$	$N_{\text{basis}}$	$N_{\text{zero}}/N_{\text{basis}}^2$
FHI-aims	Graphene	$30 \times 30 \times 1$	1800	25200	97.50%
FHI-aims	Graphene	$35 \times 35 \times 1$	2450	34300	98.16%
FHI-aims	Graphene	$40 \times 40 \times 1$	3200	44800	98.58%
FHI-aims	Graphene	$45 \times 45 \times 1$	4050	56700	98.88%
FHI-aims	Graphene	$50 \times 50 \times 1$	5000	70000	99.09%
FHI-aims	Graphene	$55 \times 55 \times 1$	6050	84700	99.25%
FHI-aims	Graphene	$60 \times 60 \times 1$	7200	100800	99.41%
DGDFT	Graphene	$18 \times 18 \times 1$	6480	97200	99.98%
DGDFT	Graphene	$24 \times 24 \times 1$	11520	172800	99.99%

850 in Sections 5.2 and 5.3. For reference, the performance of the remaining com-  
 851 putational steps (in addition to the KS eigenproblem) of standard DFT-PBE  
 852 calculations using FHI-aims code is shown in Appendix A.

853 Fig. 7 shows the wall clock time of the above-mentioned repeated steps  
 854 of the solvers. It is worth noting that, when using the same computational  
 855 resources, the time used by ELPA is theoretically constant during an SCF  
 856 cycle, as the performance of a dense direct eigensolver only depends on the  
 857 size of the matrix to solve. In contrast, the time used by libOMM and PEXSI  
 858 depends on the number of CG line searches and the number of PEXSI mu  
 859 iterations, respectively. Since both the number of CG line searches and the  
 860 number of PEXSI mu iterations can be quickly reduced to 1 as the SCF cycle  
 861 proceeds, shown in Fig. 7 are the timings corresponding to 1 CG line search  
 862 in libOMM using the basic flavor (see Section 5.5 for the effect of flavor  
 863 on the performance of libOMM), and 1 PEXSI mu iteration in PEXSI. In  
 864 future versions of PEXSI, a newly designed algorithm will be used to update  
 865 the chemical potential as the SCF cycle converges, and the number of mu  
 866 iterations will always be 1 in each SCF iteration.

867 In Fig. 7 (a), the scaling of solvers with respect to the basis size is  
 868 shown for DFT-PBE calculations of graphene models consisting of 1,800  
 869 atoms (25,200 basis functions) to 7,200 atoms (100,800 basis functions) using

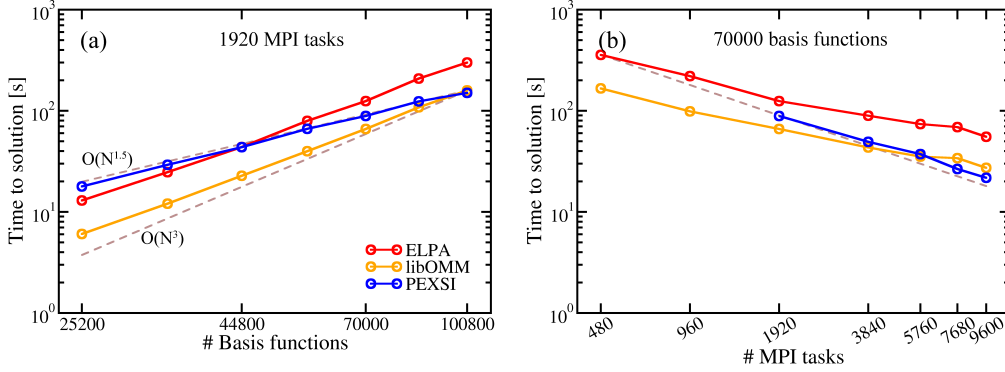


Figure 7: Scaling of the “repeated” steps in ELPA, libOMM, and PEXSI solvers with respect to (a) the number of basis functions and (b) the number of MPI tasks. The number of MPI tasks in (a) is 1,920. The number of basis functions in (b) is 70,000. The “repeated” steps are: transforming the eigenproblem (Eq. 11), solving the standard eigenproblem (Fig. 1), and back-transforming the eigenvectors in ELPA; the minimization of OMM energy functional (Eqs. 15 and 16), and the construction of density matrix from the final Wannier functions in libOMM (the CG line search converges in one step); the numerical factorization and the selected inversion of the object  $\mathbf{H} - (z_l + \mu)\mathbf{S}$  (Eq. 19), and the construction of density matrix from the poles in PEXSI (the PEXSI iteration converges in one step). Ideal scaling is indicated by the dashed lines. PEXSI cannot solve the problem of 70,000 basis functions (5,000 carbon atoms) with 480 or 960 MPI tasks, due to the limited amount of memory assigned to each pole.

870 1,920 MPI tasks. Both ELPA and libOMM exhibit scalings close to  $O(N^3)$ , as  
 871 expected. In this particular set-up, libOMM is consistently faster than ELPA  
 872 by a factor of 2. PEXSI, with a lower computational complexity (theoretically  
 873  $O(N^{1.5})$  for 2D systems), begins to outperform ELPA and libOMM at around  
 874 3,000 atoms and 7,000 atoms, respectively. The benefit of using PEXSI  
 875 should become more significant as we further increase the system size.

876 The strong scaling shown in Fig. 7 (b) demonstrates the scalability of  
 877 the solvers when they are applied to the graphene 5,000-atom model (70,000  
 878 basis functions) using 480 to 9,600 MPI tasks. All three solvers exhibit  
 879 good scalability to 9,600 MPI tasks. In particular, the PEXSI solver scales  
 880 almost ideally up to thousands of MPI tasks. This is attributed to the 2-level  
 881 parallelism employed in PEXSI (Section 3.3). The perfect strong scaling of  
 882 PEXSI can be further extended to at least tens of thousands of MPI tasks  
 883 (this is demonstrated in Section 5.6). However, PEXSI fails to solve the  
 884 problem with 480 or 960 MPI tasks, owing to the limited memory assigned

885 to each pole.

### 886 *5.2. Matrix Redistribution*

887 When using the `elsi_dm_real` subroutine (Section 4.5) to compute the den-  
888 sity matrix with the `BLACS_DENSE` format and the PEXSI solver, the input  
889 Hamiltonian and overlap matrices are not in the correct format for PEXSI.  
890 The `elsi_dm_real` subroutine internally converts the input Hamiltonian ma-  
891 trix to the `PEXSI_CSC` format, and converts the density matrix computed  
892 by PEXSI back to the original format. The overlap matrix is converted as  
893 well, albeit only in the first iteration of an SCF cycle. The performance  
894 of the Hamiltonian matrix conversion from `BLACS_DENSE` to `PEXSI_CSC`  
895 and the density matrix conversion from `PEXSI_CSC` to `BLACS_DENSE` are  
896 shown and compared to the PEXSI computation time in Fig. 8. For matrix  
897 sizes ranging from 25,200 (1,800 atoms) to 100,800 (7,200 atoms), Fig. 8 (a)  
898 shows that the wall clock time for both conversions with 1,920 MPI tasks is  
899 always below 10% of the PEXSI computation time (red lines in Fig. 8). Fig.  
900 8 (b) shows that the data redistribution time is consistently below 10% of  
901 the computation time, when using 1,920 to 9,600 MPI tasks for a problem  
902 of dimension 70,000. The `BLACS_DENSE` to `PEXSI_CSC` conversion stops  
903 scaling at 9,600 MPI tasks. Further optimization of the conversion using  
904 MPI point-to-point communications is planned as a future work direction.

### 905 *5.3. SCF Initialization*

906 Computational steps that are only required in the first one or few SCF  
907 iterations have some impact on the overall performance of an SCF cycle. Here  
908 we discuss three such steps: (1) Cholesky factorization of the overlap matrix  
909 in Eq. 10, which is used to transform the generalized eigenvalue problem  
910 to the standard form. This is a mandatory step for ELPA and an optional  
911 step for libOMM. The Cholesky factorization of a dense matrix in ELSI is  
912 performed using subroutines provided in ELPA. (2) Symbolic factorization  
913 that provides PEXSI necessary information of the sparsity pattern of the  
914 Hamiltonian and overlap matrices before numerical factorization and selected  
915 inversion are carried out. The symbolic factorization of a sparse matrix  
916 is performed using subroutines provided in the SuperLU\_DIST library [79,  
917 80]. (3) Inertia counting that quickly estimates the chemical potential of the  
918 system according to Sylvester’s Inertia Law theorem [71]. This reasonable  
919 initial guess of the chemical potential is essential to the fast convergence of  
920 PEXSI.

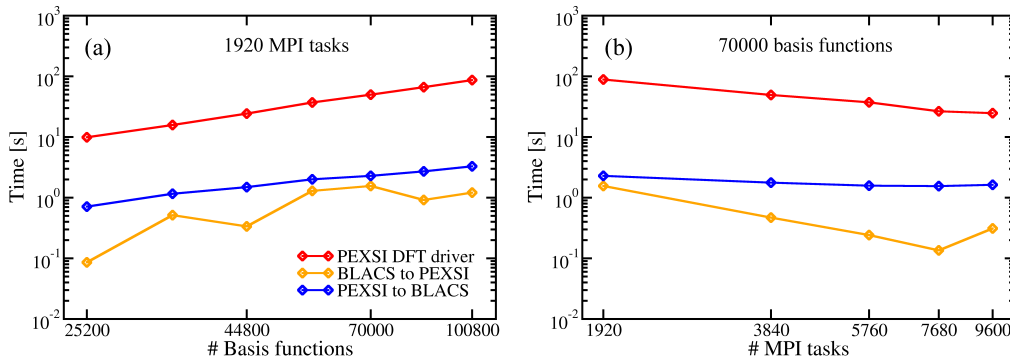


Figure 8: Scaling of matrix redistribution with respect to (a) the number of basis functions and (b) the number of MPI tasks. The number of MPI tasks in (a) is 1,920. The number of basis functions in (b) is 70,000. BLACS to PEXSI: redistribution of the Hamiltonian matrix from 2D block-cyclic dense storage (BLACS\_DENSE) to 1D block CSC sparse storage (PEXSI\_CSC). PEXSI to BLACS: redistribution of the density matrix from 1D block CSC sparse storage (PEXSI\_CSC) to 2D block-cyclic dense storage (BLACS\_DENSE). The overlap matrix is redistributed only once per SCF cycle, hence its absence here.

921 Fig. 9 (a) shows the wall clock time of the three initialization steps as  
 922 a function of the system size. The Cholesky factorization of a dense matrix  
 923 using ELPA subroutines scales cubically with the system size, whereas the  
 924 symbolic factorization and inertia counting scale linearly. The scaling dif-  
 925 ference among these preprocessing steps helps explain why PEXSI is more  
 926 favorable for large systems. In the strong scaling plot shown in Fig. 9 (b),  
 927 the dense Cholesky factorization is shown to scale up to 9,600 MPI tasks. Be-  
 928 cause the symbolic factorization implemented in SuperLU\_DIST is not stable  
 929 when executed on multiple processors, we used a sequential version of the  
 930 symbolic factorization in the experiment, which obviously does not scale. We  
 931 are in the process of developing a more robust and scalable implementation  
 932 of the symbolic factorization procedure as part of the development of a new  
 933 parallel sparse Cholesky (and LDLT) factorization library called symPACK  
 934 [81].

#### 935 5.4. ELPA

936 To analyze the performance of the ELPA eigensolver for the graphene  
 937 problem solved here, the solution of a generalized eigenproblem (red lines  
 938 in Fig. 7) is divided into three steps: the transformation of the generalized  
 939 eigenproblem to the standard form (Eq. 11), the solution of the standard

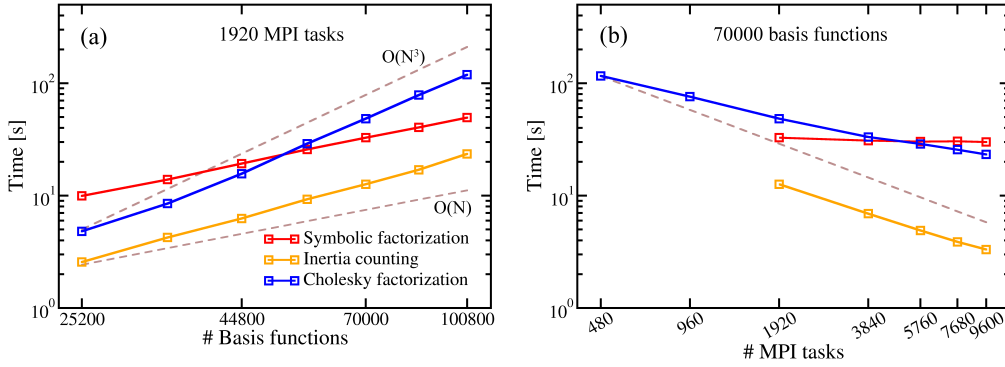


Figure 9: Scaling of symbolic factorization using SuperLU\_DIST, inertia counting using PEXSI, and Cholesky factorization using ELPA, with respect to (a) the number of basis functions and (b) the number of MPI tasks. The number of MPI tasks in (a) is 1,920. The number of basis functions in (b) is 70,000. Symbolic factorization is performed in serial. Ideal scaling is indicated by the dashed lines.

940 eigenproblem (Fig. 1), and the back-transformation of the eigenvectors. Fig.  
 941 10 (a) and (b) show the scaling of the three steps with respect to the number  
 942 of basis functions and the number of MPI tasks, respectively. All these  
 943 steps scale cubically with respect to the system size. Solving the standard  
 944 eigenproblem is more expensive than the transformation steps. While the  
 945 three steps show similar strong scaling up to 9,600 MPI tasks, the solution  
 946 of a standard eigenproblem dominates the total computation time. In Fig.  
 947 10 (c) and (d), the solution time of a standard eigenproblem using ELPA  
 948 2-stage solver is further decomposed into five steps illustrated in Fig. 1.  
 949 These plots show that the current bottlenecks in terms of both computation  
 950 time and parallel efficiency are the first step, i.e. the transformation of a full  
 951 matrix to a banded form, and the fifth step, i.e. the back-transformation of  
 952 the eigenvectors from a banded form to a full form. The fourth step, back-  
 953 transformation of the eigenvectors to the banded form, is not the most time  
 954 consuming step of the computation. In fact, the computational complexity  
 955 of the third, fourth, and fifth steps is roughly proportional to the number of  
 956 eigenvectors to compute, as only these eigenvectors need to be calculated in  
 957 the third step and transformed in the fourth and fifth steps.

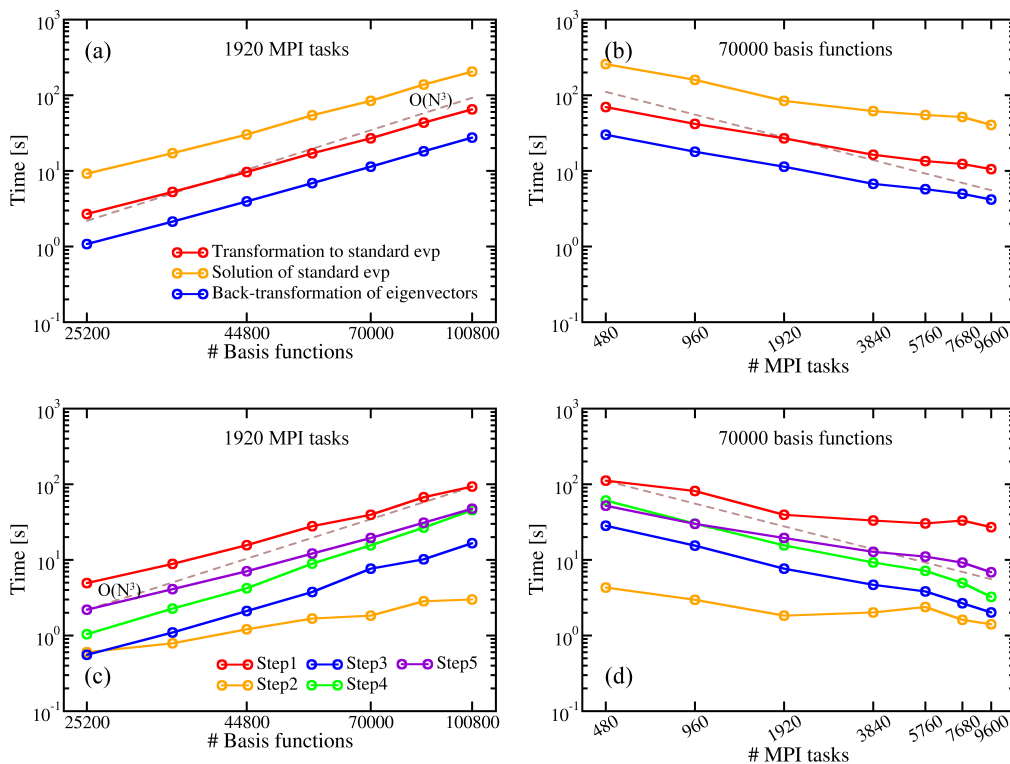


Figure 10: Scaling of the key computational steps of the ELPA eigensolver with respect to (a,c) the number of basis functions and (b,d) the number of MPI tasks. The number of MPI tasks in (a) is 1,920. The number of basis functions in (b) is 70,000. Ideal scaling is indicated by the dashed lines. The upper panel (a,b) focuses on the transformation from a generalized eigenproblem to its standard form, the solution of a standard problem, and the back-transformation of the eigenvectors to the original generalized problem. The lower panel (c,d) further decomposes the solution of a standard eigenproblem using the ELPA 2-stage solver into 5 substeps, as illustrated in Fig. 1.

958 5.5. *libOMM*

959 The performance of the iterative OMM method depends significantly on  
 960 the convergence rate of the CG minimization. The prototype OMM im-  
 961 plementation in libOMM generates random numbers as the initial guess for  
 962 the coefficients of Wannier functions used in the first SCF iteration, con-  
 963 sequently leading to a large and unpredictable number of iterations in the  
 964 CG line search scheme. Then, the convergence of line search is dramatically  
 965 accelerated as the SCF cycle proceeds, as the Wannier functions coefficients  
 966 calculated in the current iteration are reused as the initial guess in the next

967 iteration. Inspired by the connection between the Wannier functions and the  
 968 basis functions in Eq. 14, a better idea is to use the eigenfunctions corre-  
 969 sponding to the occupied space computed by ELPA as the initial guess for  
 970 OMM. In ELSI, this is achieved automatically, controlled by the `n_elpa_steps`  
 971 parameter (see Section 4.6). Table 2 reflects how `n_elpa_steps` affects the CG  
 972 convergence of OMM in the  $(n\_elpa\_steps + 1)^{\text{th}}$  SCF iteration, by showing  
 973 the number of CG line searches in the basic and Cholesky flavors of OMM  
 974 as a function of the number of ELPA steps. In general, more ELPA steps  
 975 lead to faster CG convergence. In this particular test case with 5,000 carbon  
 976 atoms and 70,000 basis functions, 6 ELPA steps are sufficient to reduce the  
 977 number of CG line searches in libOMM to 1 for both tested flavors.

Table 2: Number of conjugate gradient (CG) line search steps required by libOMM to minimize the OMM energy functional. The benchmark system here is the graphene  $50 \times 50 \times 1$  supercell model containing 5,000 atoms and 70,000 basis functions. In the table, “Basic” refers to the method that directly operates on the generalized eigenproblem; “Cholesky” refers to the method that applies Cholesky factorization to transform the generalized problem to the standard form. “x” in the second column means that the minimization cannot converge within the maximum allowed number of CG iterations (5000).

# ELPA steps	0	1	2	3	4	5	6	7
# CG (Basic)	x	185	255	153	72	7	1	1
# CG (Cholesky)	254	27	36	24	13	5	1	1

978 Compared in the second and third rows of Table 2 is another factor that  
 979 has an impact on the number of CG line searches in libOMM, i.e., the method  
 980 used to minimize the OMM functional. The basic algorithm directly follows  
 981 the recipe in Eq. 16, but Eq. 16 can also be minimized by first transform-  
 982 ing the generalized eigenproblem to a standard problem based on Cholesky  
 983 factorization. As shown in Table 2, minimizing the OMM functional in the  
 984 context of a standard eigenproblem (Cholesky, the third row in the table)  
 985 contributes to a decrease in the number of line searches. This acceleration of  
 986 the CG line search, however, comes at the price of the additional complexity  
 987 required by the eigenproblem transformation. Fig. 11 shows the compari-  
 988 son of the computational time of one CG line search in libOMM with the  
 989 basic flavor versus the Cholesky flavor. The two flavors scale similarly, with  
 990 respect to both the number of basis functions (Fig. 11 (a), from 25,200 to  
 991 100,800 basis functions) and the number of MPI tasks (Fig. 11 (b), from  
 992 480 to 9,600 MPI tasks). The Cholesky flavor is consistently slower than the



993 basic flavor by a factor of  $2 \sim 4$ , due to the eigenproblem transformation  
 994 and the corresponding back-transformation of Wannier function coefficients.  
 995 Also reflected in Fig. 11 is the shortest time to compute the density matrix  
 996 using OMM, which is the basic flavor that converges in one CG line search.  
 997 Indicated by Table 2 and Fig. 11, the most promising approach that could be  
 998 used in practical calculations is the combination of a few ELPA steps followed  
 999 by the basic flavor of OMM, whose convergence is guaranteed within one CG  
 1000 iteration. To further improve the performance of this solver, future work will  
 1001 include the inclusion in the ELSI interface of a preconditioned libOMM flavor,  
 1002 which has already proven to efficiently speed up the line search convergence  
 1003 [16, 82]; a spectral slicing method to separately evaluate the eigenstates near  
 1004 the Fermi level and thus to enable the proper handling of fractional occu-  
 1005 pation numbers; the sparse linear algebra via routines implemented in the  
 1006 PSPBLAS (Parallel SParse BLAS) library [83]; and ultimately the extension  
 1007 of OMM to a linear scaling solver as originally proposed [61, 62, 63, 64].

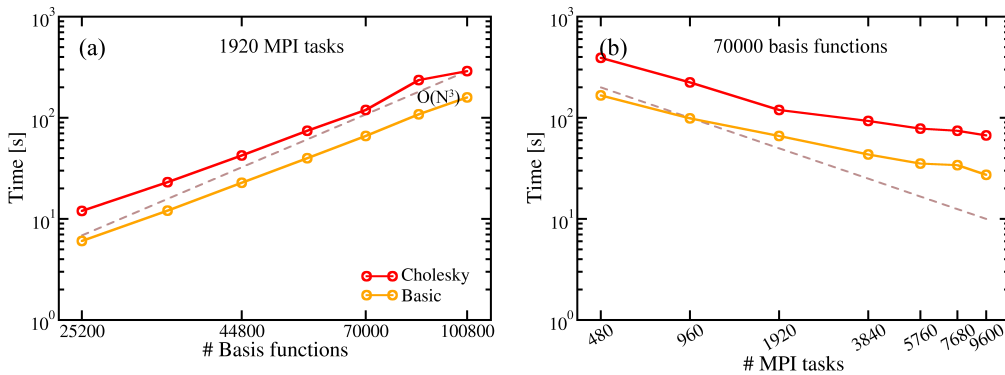


Figure 11: Scaling of the computation of the density matrix using orbital minimization method, with respect to (a) the number of basis functions and (b) the number of MPI tasks. The number of MPI tasks in (a) is 1,920. The number of basis functions in (b) is 70,000. Shown here is the ideal case of OMM, where the CG line search of the OMM energy functional minimum requires only one step to converge. In practical SCF calculations, the number of line searches in OMM can only be reduced to one after several SCF steps. “Basic” refers to the method that directly handles the generalized eigenproblem. “Cholesky” refers to the method that applies Cholesky factorization to transform the generalized problem to the standard form before minimization. Ideal scaling is indicated by the dashed line.

1008 *5.6. PEXSI*

1009 As noted in Section 3.3, PEXSI exploits two levels of parallelization: the  
1010 first level is the parallel evaluation of each pole in the pole expansion (Eq.  
1011 19), and the second level is the parallel numerical factorization and selected  
1012 inversion at each pole. MPI tasks are divided into several groups with one  
1013 pole assigned to each group. Fig. 12 (a) shows that both steps scale as  
1014  $O(N^{1.5})$  for the graphene model, which is in agreement with the theoretical  
1015 prediction for quasi-2D systems. The selected inversion step is slightly more  
1016 expensive than the numerical factorization step. As shown in the strong  
1017 scaling plot in Fig. 12 (b), both the numerical factorization and the selected  
1018 inversion scale almost ideally to at least 9,600 MPI tasks. The number of  
1019 MPI tasks shown in Fig. 12 (b) should be divided by the number of poles,  
1020 80, to reflect the scaling of numerical factorization and selected inversion at  
1021 each pole. Since PEXSI has been shown to scale to several thousands of  
1022 MPI tasks [70], the performance reported in Fig. 12 (b), which measures  
1023 scalability up to 120 tasks per pole, is still far from the scalability limit. To  
1024 further demonstrate the strong scaling of the PEXSI solver, Fig. 13 shows the  
1025 wall clock time used by PEXSI for a graphene model consisting of 6,480 atoms  
1026 (97,200 basis functions) using 2,592 to 31,104 MPI tasks (Fig. 13 (a)) and  
1027 a graphene model consisting of 11,520 atoms (172,800 basis functions) using  
1028 2,304 to 110,592 MPI tasks (Fig. 13 (b)). These tests are performed using  
1029 the ELSI interface as implemented in the DGDFT software package. The  
1030 ELPA eigensolver is also included as a reference. For both models, PEXSI  
1031 exhibits nearly ideal strong scaling and eventually outperforms ELPA as the  
1032 number of MPI tasks becomes sufficiently large. The ELPA solver ceases to  
1033 scale beyond 18,432 MPI tasks for the 172,800-atom system.

1034 **6. Conclusions**

1035 Materials simulations based on Kohn-Sham density-functional theory re-  
1036 quire solving an eigenvalue problem repeatedly in an iterative procedure de-  
1037 signed to obtain the ground state electron density of a poly-atomic system.  
1038 Although this is a well studied subject in numerical linear algebra, it consti-  
1039 tutes the bottleneck in large-scale calculations. A number of new approaches  
1040 have emerged in the last few years. These approaches have different features  
1041 and performance characteristics. Proper use of these approaches requires a  
1042 good understanding of the pros and cons of each approach, and the input  
1043 and output of specific algorithms. ELSI is designed to provide a common

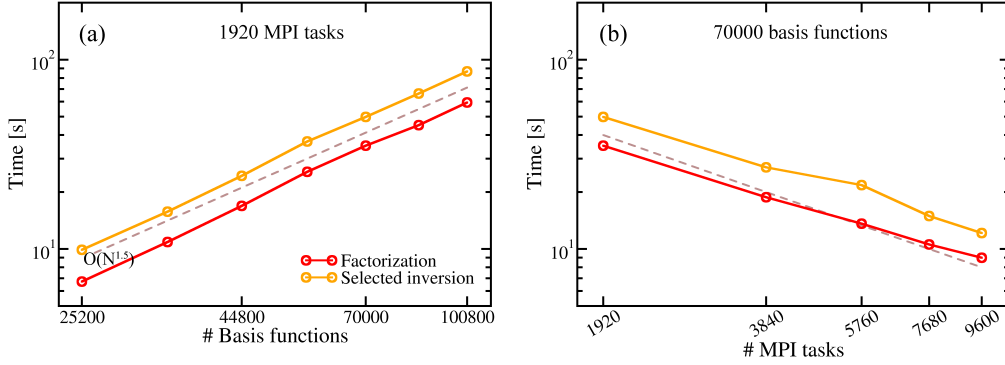


Figure 12: Scaling of the two key computational steps of the PEXSI DFT driver, namely the numerical factorization and the selected inversion, with respect to (a) the number of basis functions and (b) the number of MPI tasks. The number of MPI tasks in (a) is 1,920. The number of basis functions in (b) is 70,000. Ideal scaling is indicated by the dashed lines. The 80 poles employed for the pole expansion in Eq. 19 are independently evaluated in parallel. The numerical factorization and selected inversion of each pole are carried out using  $1920/80 = 24$  MPI tasks in (a), and  $\# \text{ MPI tasks}/80$  in (b).

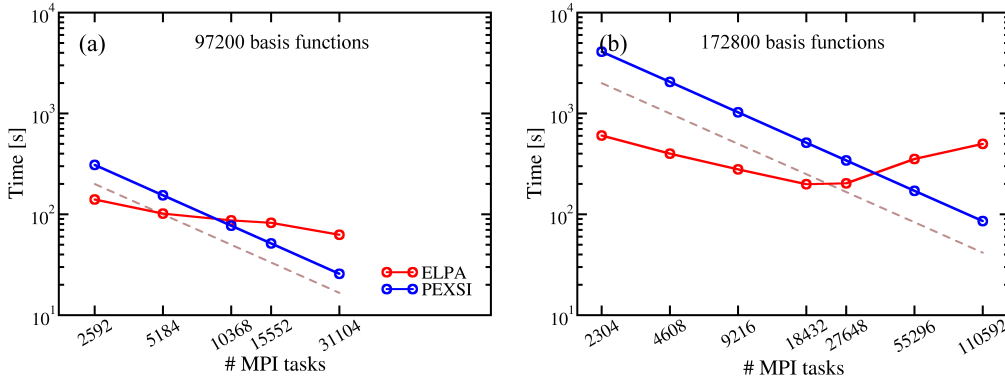


Figure 13: Comparison of the strong scaling of the ELPA and PEXSI solvers. The number of basis functions is 97,200 in (a) and 172,800 in (b). The matrices are from DGDFT code. There are 48 poles employed in the PEXSI pole expansion. Ideal scaling is indicated by the dashed lines.

1044 interface that allows users to easily choose an appropriate solver. Although  
 1045 the choice of the best solver often depends on a number of factors such as  
 1046 the problem size and the available computational resource, the benchmark  
 1047 results presented in this paper provide some general guidance on how to make

1048 these choices. In particular, we have shown different regimes in which one  
1049 approach outperforms others and the crossover points between these regimes.

1050 Finally, we demonstrated how different solvers can be organized in a com-  
1051 mon framework to enable easy integration with a vast number of electronic  
1052 structure software packages. We anticipate that the number of new ap-  
1053 proaches to solving eigenvalue problems related to KS-DFT will continue  
1054 to increase. We hope that ELSI will become a focal point for the commu-  
1055 nity to integrate, comparatively assess and, ultimately, adopt this diverse  
1056 ecosystem in a simple, effective fashion.

## 1057 7. Acknowledgments

1058 This work was supported by the National Science Foundation under grant  
1059 number 1450280. We thank the National Energy Research Scientific Comput-  
1060 ing Center (NERSC) for the computational resources. This work also used  
1061 resources of the Argonne Leadership Computing Facility, which is a DOE Of-  
1062 fice of Science User Facility supported under contract DE-AC02-06CH11357.  
1063 Regarding the establishment and improvement of the ELSI API, we especially  
1064 appreciate the fruitful discussions with and the feedback from many fellow  
1065 researchers in the electronic structure community, including developers of the  
1066 BigDFT, CP2K, DGDFE, FHI-aims, SIESTA code projects and of CECAM’s  
1067 Electronic Structure Library (<http://esl.cecam.org>). AG thanks EU H2020  
1068 grant 676598 (‘MaX: Materials at the eXascale’ CoE), Spain’s MINECO  
1069 (FIS2012-37549-C05-05, FIS2015-64886-C5-4-P and the ‘Severo Ochoa’ pro-  
1070 gram grant SEV-2015-0496), and GenCat (2014 SGR 301). The work of  
1071 JL was partially supported by the National Science Foundation under grant  
1072 number DMS-1127914 to the Statistical and Applied Mathematical Sciences  
1073 Institute. VB particularly acknowledges the experiences shared by many of  
1074 the co-authors of the FHI-aims code over many years, for instance, regarding  
1075 parallelization strategies or the handling of ill-conditioned overlap matrices,  
1076 from which we learned during the development of the ELSI interface - espe-  
1077 cially Dr. Ville Havu (Aalto University) and Dr. Rainer Johanni (Munich;  
1078 now deceased).

## 1079 Appendix A. Technical Settings in FHI-aims Calculations

1080 The benchmark calculations reported in Figs. 7, 8, 9 10, 11, and 12 in Sec-  
1081 tion 5 are KS-DFT calculations performed with the FHI-aims code [33, 77],

1082 PBE exchange-correlation functional [76], “tier1” numeric atom-centered or-  
 1083 bital (NAO) basis set (see Table 1 in Ref. [33], “light” numerical settings,  
 1084 and a  $1 \times 1 \times 1$  k-grid ( $\Gamma$  point). In order to place the timings reported in Fig.  
 1085 7 into perspective with respect to the other parts of a KS-DFT calculation,  
 1086 Fig. A.14 shows timings for all other important computational steps in the  
 1087 corresponding FHI-aims calculations, obtained on the same hardware and in  
 1088 the same runs as the results shown in Fig. 7. The main additional steps  
 1089 are executed on a real-space grid and include the Hartree potential evaluation,  
 1090 the numerical integrations of the Hamiltonian matrix elements, and the  
 1091 update of the electron density and its gradients, all implemented in a near  
 1092  $O(N)$  fashion and efficiently parallelized in FHI-aims. Refs. [33, 77] provide  
 1093 a more detailed account of the algorithms involved.

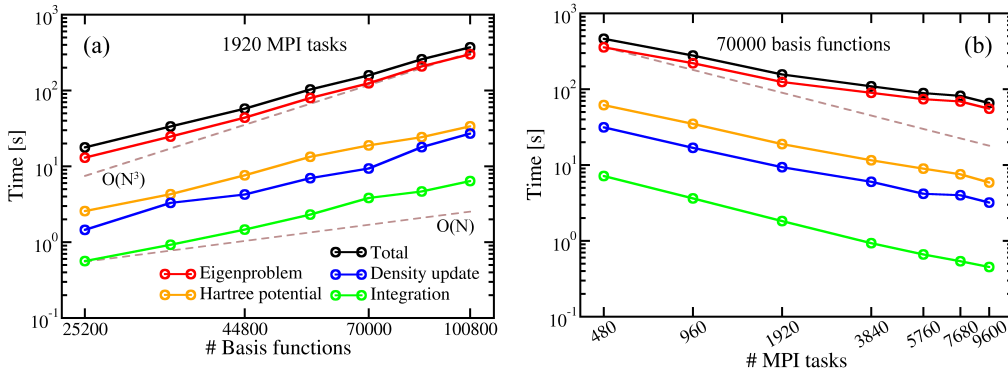


Figure A.14: Scaling of the key computational steps of the DFT-PBE calculations in FHI-aims, with respect to (a) the number of basis functions and (b) the number of MPI tasks. The number of MPI tasks in (a) is 1,920. The number of basis functions in (b) is 70,000. Ideal scaling is indicated by the dashed lines. The key steps are the evaluation of the Hartree potential, the numerical integrations of the Hamiltonian matrix elements, the update of the electron density and its gradient, and solving the Kohn-Sham eigenproblem using the ELPA eigensolver library. Shown here are timings corresponding to one SCF iteration, not accumulated timings in a complete SCF cycle.

## 1094 Appendix B. Technical Settings in DGDFE Calculations

1095 The benchmark calculations reported in Figs. 13 in Section 5 are KS-  
 1096 DFT calculations performed with DGDFE [51, 52] using the PBE exchange-  
 1097 correlation functional [76]. The global system is partitioned into  $36 \times 36$   
 1098 and  $48 \times 48$  elements for the system containing 6,480 and 11,520 atoms,

1099 respectively. The number of adaptive local basis functions (ALB) per atom  
1100 is 15, which is sufficient for the error of the total energy per atom and the  
1101 maximum error of the force to be below  $10^{-3}$  Hartree and  $10^{-3}$  Hartree/Bohr,  
1102 respectively. The DG penalty parameter is chosen to be 5.0, and the kinetic  
1103 energy cutoff to generate the ALBs is set to 40 Hartree. The number of poles  
1104 used by PEXSI is 48.

- 1105 [1] W. Kohn, L. J. Sham, Self-consistent equations including exchange and  
1106 correlation effects, *Physical Review* 140 (4A) (1965) 1133–1138.
- 1107 [2] A. D. Becke, A new mixing of Hartree-Fock and local density-functional  
1108 theories, *The Journal of Chemical Physics* 98 (2) (1993) 1372–1377.
- 1109 [3] A. Seidl, A. Görling, P. Vogl, J. A. Majewski, M. Levy, Generalized  
1110 Kohn-Sham schemes and the band-gap problem, *Physical Review B* 53  
1111 (1996) 3764–3774.
- 1112 [4] L. M. Ghiringhelli, C. Carbogno, S. Levchenko, F. Mohamed, G. Huhs,  
1113 M. Lueders, M. Oliveira, M. Scheffler, Towards a common format  
1114 for computational materials science data, *Psi-k Scientific Highlight*  
1115 July (131).
- 1116 [5] R. Haunschuld, A. Barth, W. Marx, Evolution of DFT studies in view  
1117 of a scientometric perspective, *Journal of Cheminformatics* 8 (1) (2016)  
1118 52.
- 1119 [6] P. Mavropoulos, P. Dederichs, Statistical data about density functional  
1120 calculations, *Psi-k Scientific Highlight* April (135).
- 1121 [7] S. Goedecker, Linear scaling electronic structure methods, *Reviews of*  
1122 *Modern Physics* 71 (1999) 1085–1123.
- 1123 [8] D. R. Bowler, T. Miyazaki,  $O(N)$  methods in electronic structure calcu-  
1124 lations, *Reports on Progress in Physics* 75 (3) (2012) 036503.
- 1125 [9] C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, M. C. Payne, Introduc-  
1126 ing ONETEP: Linear-scaling density functional simulations on parallel  
1127 computers, *The Journal of Chemical Physics* 122 (8) (2005) 084119.
- 1128 [10] D. R. Bowler, T. Miyazaki, Calculations for millions of atoms with den-  
1129 sity functional theory: Linear scaling shows its potential, *Journal of*  
1130 *Physics: Condensed Matter* 22 (7) (2010) 074207.

- 1131 [11] J. VandeVondele, U. Borstnik, J. Hutter, Linear scaling self-consistent  
1132 field calculations with millions of atoms in the condensed phase, *Journal*  
1133 *of Chemical Theory and Computation* 8 (10) (2012) 3565–3573.
- 1134 [12] H. F. Schurkus, C. Ochsenfeld, Communication: An effective linear-  
1135 scaling atomic-orbital reformulation of the random-phase approxima-  
1136 tion using a contracted double-laplace transformation, *The Journal of*  
1137 *Chemical Physics* 144 (3) (2016) 031101.
- 1138 [13] A. Luenser, H. F. Schurkus, C. Ochsenfeld, Vanishing-overhead linear-  
1139 scaling random phase approximation by Cholesky decomposition and an  
1140 attenuated coulomb-metric, *Journal of Chemical Theory and Computa-*  
1141 *tion* 13 (4) (2017) 1647–1655.
- 1142 [14] T. Auckenthaler, V. Blum, H.-J. Bungartz, T. Huckle, R. Johanni,  
1143 L. Kramer, B. Lang, H. Lederer, P.-R. Willems, Parallel solution of  
1144 partial symmetric eigenvalue problems from electronic structure calcu-  
1145 lations, *Parallel Computing* 37 (12) (2011) 783–794.
- 1146 [15] A. Marek, V. Blum, R. Johanni, V. Havu, B. Lang, T. Auckenthaler,  
1147 A. Heinecke, H.-J. Bungartz, H. Lederer, The ELPA library: Scalable  
1148 parallel eigenvalue solutions for electronic structure theory and compu-  
1149 tational science, *Journal of Physics: Condensed Matter* 26 (21) (2014)  
1150 213201.
- 1151 [16] F. Corsetti, The orbital minimization method for electronic structure  
1152 calculations with finite-range atomic basis sets, *Computer Physics Com-*  
1153 *munications* 185 (3) (2014) 873–883.
- 1154 [17] L. Lin, J. Lu, L. Ying, R. Car, W. E, Fast algorithm for extracting  
1155 the diagonal of the inverse matrix with application to the electronic  
1156 structure analysis of metallic systems, *Communications in Mathematical*  
1157 *Sciences* 7 (3) (2009) 755–777.
- 1158 [18] L. Lin, M. Chen, C. Yang, L. He, Accelerating atomic orbital-based  
1159 electronic structure calculation via pole expansion and selected inver-  
1160 sion, *Journal of Physics: Condensed Matter* 25 (29) (2013) 295501.
- 1161 [19] L. Lin, A. García, G. Huhs, C. Yang, SIESTA-PEXSI: Massively parallel  
1162 method for efficient and accurate ab initio materials simulation without

- 1163 matrix diagonalization, *Journal of Physics: Condensed Matter* 26 (2014)  
1164 305503.
- 1165 [20] S. Mohr, D. Caliste, M. Wagner, L. Genovese, Efficient computation of  
1166 sparse matrix functions for large scale electronic structure calculations:  
1167 The CheSS library, arXiv:1704.00512.
- 1168 [21] M. Keceli, H. Zhang, P. Zapol, D. A. Dixon, A. F. Wagner, Shift-and-  
1169 invert parallel spectral transformation eigensolver: Massively parallel  
1170 performance for density-functional based tight-binding, *Journal of Com-  
1171 putational Chemistry* 37 (4) (2016) 448–459.
- 1172 [22] S. R. Jensen, S. Saha, J. A. Flores-Livas, W. Huhn, V. Blum,  
1173 S. Goedecker, L. Frediani, The elephant in the room of density func-  
1174 tional theory calculations, *The Journal of Physical Chemistry Letters*  
1175 8 (7) (2017) 1449–1457.
- 1176 [23] A. Szabo, N. S. Ostlund, *Modern quantum chemistry: Introduction to  
1177 advanced electronic structure theory*, Dover Publications, 1989.
- 1178 [24] J. Hutter, M. Iannuzzi, F. Schiffmann, J. VandeVondele, CP2K: Atom-  
1179 istic simulations of condensed matter systems, *Wiley Interdisciplinary  
1180 Reviews: Computational Molecular Science* 4 (1) (2014) 15–25.
- 1181 [25] M. J. Frisch, G. W. Trucks, H. B. Schlegel, G. E. Scuseria, J. R. C.  
1182 M A Robb, G. Scalmani, V. Barone, G. A. Petersson, H. Nakatsuji,  
1183 X. Li, M. Caricato, A. Marenich, J. Bloino, B. G. Janesko, R. Gomperts,  
1184 B. Mennucci, H. P. Hratchian, J. V. Ortiz, A. F. Izmaylov, J. L. Son-  
1185 nenberg, D. Williams-Young, F. Ding, F. Lipparini, F. Egidi, J. Goings,  
1186 B. Peng, A. Petrone, T. Henderson, D. Ranasinghe, V. G. Zakrzewski,  
1187 J. Gao, N. Rega, G. Zheng, W. Liang, M. Hada, M. Ehara, K. Toyota,  
1188 R. Fukuda, J. Hasegawa, M. Ishida, T. Nakajima, Y. Honda, O. Kitao,  
1189 H. Nakai, T. Vreven, K. Throssell, J. A. M. Jr, J. E. Peralta, F. Ogliaro,  
1190 M. Bearpark, J. J. Heyd, E. Brothers, K. N. Kudin, V. N. Staroverov,  
1191 T. Keith, R. Kobayashi, J. Normand, K. Raghavachari, A. Rendell, J. C.  
1192 Burant, S. S. Iyengar, J. Tomasi, M. Cossi, J. M. Millam, M. Klene,  
1193 C. Adamo, R. Cammi, J. W. Ochterski, R. L. Martin, K. Morokuma,  
1194 O. Farkas, J. B. Foresman, , D. J. Fox, *Gaussian09 Revision A.02*, Gaus-  
1195 sian Inc. Wallingford.



- 1196 [26] M. Valiev, E. J. Bylaska, N. Govind, K. Kowalski, T. P. Straatsma,  
1197 H. J. J. van Dam, D. Wang, J. Nieplocha, E. Apra, T. L. Windus,  
1198 W. A. de Jong, NWChem: A comprehensive and scalable open-source  
1199 solution for large scale molecular simulations, *Computer Physics Com-*  
1200 *munications* 181 (9) (2010) 1477–1489.
- 1201 [27] Y. Shao, Z. Gan, E. Epifanovsky, A. T. B. Gilbert, M. Wormit, J. Kuss-  
1202 mann, A. W. Lange, A. Behn, J. Deng, X. Feng, D. Ghosh, M. Goldey,  
1203 P. R. Horn, L. D. Jacobson, I. Kaliman, R. Z. Khaliullin, T. Kua, A. Lan-  
1204 dau, J. Liu, E. I. Proynov, Y. M. Rhee, R. M. Richard, M. A. Rohrdanz,  
1205 R. P. Steele, E. J. Sundstrom, H. L. W. III, P. M. Zimmerman, D. Zuev,  
1206 B. Albrecht, E. Alguire, B. Austin, G. J. O. Beran, Y. A. Bernard,  
1207 E. Berquist, K. Brandhorst, K. B. Bravaya, S. T. Brown, D. Casanova,  
1208 C.-M. Chang, Y. Chen, S. H. Chien, K. D. Closser, D. L. Crittenden,  
1209 M. Diedenhofen, R. A. D. Jr, H. Do, A. D. Dutoi, R. G. Edgar, S. Fatehi,  
1210 L. Fusti-Molnar, A. Ghysels, A. Golubeva-Zadorozhnaya, J. Gomes,  
1211 M. W. D. Hanson-Heine, P. H. P. Harbach, A. W. Hauser, E. G. Hohen-  
1212 stein, Z. C. Holden, T.-C. Jagau, H. Ji, B. Kaduk, K. Khistyayev, J. Kim,  
1213 J. Kim, R. A. King, P. Klunzinger, D. Kosenkov, T. Kowalczyk, C. M.  
1214 Krauter, K. U. Lao, A. D. Laurent, K. V. Lawler, S. V. Levchenko, C. Y.  
1215 Lin, F. Liu, E. Livshits, R. C. Lochan, A. Luenser, P. Manohar, S. F.  
1216 Manzer, S.-P. Mao, N. Mardirossian, A. V. Marenich, S. A. Maurer,  
1217 N. J. Mayhall, E. Neuscamman, C. M. Oana, R. Olivares-Amaya, D. P.  
1218 O'Neill, J. A. Parkhill, T. M. Perrine, R. Peverati, A. Prociuk, D. R.  
1219 Rehn, E. Rosta, N. J. Russ, S. M. Sharada, S. Sharma, D. W. Small,  
1220 A. Sodt, T. Stein, D. Staeck, Y.-C. Su, A. J. W. Thom, T. Tsuchimochi,  
1221 V. Vanovschi, L. Vogt, O. Vydrov, T. Wang, M. A. Watson, J. Wenzel,  
1222 A. White, C. F. Williams, J. Yang, S. Yeganeh, S. R. Yost, Z.-Q. You,  
1223 I. Y. Zhang, X. Zhang, Y. Zhao, B. R. Brooks, G. K. L. Chan, D. M.  
1224 Chipman, C. J. Cramer, W. A. G. III, M. S. Gordon, W. J. Hehre,  
1225 A. Klamt, H. F. S. III, M. W. Schmidt, C. D. Sherrill, D. G. Truhlar,  
1226 A. Warshel, X. Xu, A. Aspuru-Guzik, R. Baer, A. T. Bell, N. A. Besley,  
1227 J.-D. Chai, A. Dreuw, B. D. Dunietz, T. R. Furlani, S. R. Gwaltney,  
1228 C.-P. Hsu, Y. Jung, J. Kong, D. S. Lambrecht, W. Liang, C. Ochsen-  
1229 feld, V. A. Rassolov, L. V. Slipchenko, J. E. Subotnik, T. V. Voorhis,  
1230 J. M. Herbert, A. I. Krylov, P. M. W. Gill, M. Head-Gordon, *Advances*  
1231 *in molecular quantum chemistry contained in the Q-Chem 4 program*  
1232 *package*, *Molecular Physics* 113 (2) (2015) 184–215.

- 1233 [28] F. Furche, R. Ahlrichs, C. Hattig, W. Klopper, M. Sierka, F. Weigend,  
1234 Turbomole, Wiley Interdisciplinary Reviews: Computational Molecular  
1235 Science 4 (2) (2014) 91–100.
- 1236 [29] J. C. Slater, Atomic shielding constants, Physical Review 36 (1930) 57–  
1237 64.
- 1238 [30] G. te Velde, F. M. Bickelhaupt, E. J. Baerends, C. F. Guerra, S. J. A.  
1239 van Gisbergen, J. G. Snijders, T. Ziegler, Chemistry with ADF, Journal  
1240 of Computational Chemistry 22 (9) (2001) 931–967.
- 1241 [31] <http://www.quantumwise.com> (Accessed: 2017-04-27).
- 1242 [32] B. Delley, An all-electron numerical method for solving the local density  
1243 functional for polyatomic molecules, The Journal of Chemical Physics  
1244 92 (1) (1990) 508.
- 1245 [33] V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter,  
1246 M. Scheffler, Ab initio molecular simulations with numeric atom-  
1247 centered orbitals, Computer Physics Communications 180 (2009) 2175–  
1248 2196.
- 1249 [34] K. Koepernik, H. Eschrig, Full-potential nonorthogonal local-orbital  
1250 minimum-basis band-structure scheme, Physical Review B 59 (3) (1999)  
1251 1743.
- 1252 [35] T. Ozaki, Variationally optimized atomic orbitals for large-scale elec-  
1253 tronic structures, Physical Review B 67 (15) (2003) 155108.
- 1254 [36] J. M. Soler, E. Artacho, J. D. Gale, A. Garcia, J. Junquera, P. Ordejon,  
1255 D. Sanchez-Portal, The SIESTA method for ab initio order-N materials  
1256 simulation, Journal of Physics: Condensed Matter 14 (11) (2002) 2745–  
1257 2779.
- 1258 [37] D. J. Singh, Planewaves, pseudopotentials and the LAPW method,  
1259 Springer US, 1994.
- 1260 [38] <http://elk.sourceforge.net> (Accessed: 2017-04-27).
- 1261 [39] A. Gulans, S. Kontur, C. Meisenbichler, D. Nabok, P. Pavone, S. Rig-  
1262 amonti, S. Sagmeister, U. Werner, C. Drax, Exciting: A full-potential

- 1263 all-electron package implementing density-functional theory and many-  
 1264 body perturbation theory, *Journal of Physics: Condensed Matter* 26 (36)  
 1265 (2014) 363202.
- 1266 [40] S. Blügel, G. Bihlmayer, D. Wortmann, C. Friedrich, M. Heide,  
 1267 M. Lezaic, F. Freimuth, M. Betzinger, The Jülich FLEUR project  
 1268 (1987).
- 1269 [41] P. Blaha, K. Schwarz, G. K. H. Madsen, D. Kvasnicka, J. Luitz,  
 1270 WIEN2k: An augmented plane wave + local orbitals program for cal-  
 1271 culating crystal properties.
- 1272 [42] J. E. Pask, B. M. Klein, P. A. Sterne, C. Y. Fong, Finite-element meth-  
 1273 ods in electronic-structure theory, *Computer Physics Communications*  
 1274 135 (1) (2001) 1–34.
- 1275 [43] R. M. Martin, *Electronic structure: Basic theory and practical methods*,  
 1276 Cambridge University Press, 2004.
- 1277 [44] X. Gonze, B. Amadon, P.-M. Anglade, J.-M. Beuken, F. Bottin,  
 1278 P. Boulanger, F. Bruneval, D. Caliste, R. Caracas, M. Côté, T. Deutsch,  
 1279 L. Genovese, P. Ghosez, M. Giantomassi, S. Goedecker, D. R. Hamann,  
 1280 P. Hermet, F. Jollet, G. Jomard, S. Leroux, M. Mancini, S. Mazevet,  
 1281 M. J. T. Oliveira, G. Onida, Y. Pouillon, T. Rangel, G.-M. Rignanese,  
 1282 D. Sangalli, R. Shaltaf, M. Torrent, M. J. Verstraete, G. Zerah, J. W.  
 1283 Zwanziger, ABINIT: First-principles approach to material and nanosys-  
 1284 tem properties, *Computer Physics Communications* 180 (12) (2009)  
 1285 2582–2615.
- 1286 [45] S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. I. J. Probert,  
 1287 K. Refson, M. C. Payne, First principles methods using CASTEP,  
 1288 *Zeitschrift für Kristallographie - Crystalline Materials* 220 (2005) 567–  
 1289 570.
- 1290 [46] P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavaz-  
 1291 zoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo, A. D. Corso,  
 1292 S. de Gironcoli, S. Fabris, G. F. R. Gebauer, U. Gerstmann, C. Gougous-  
 1293 sis, A. Kokalj, M. Lazzeri, L. Martin-Samos, N. Marzari, F. Mauri,  
 1294 R. Mazzarello, S. Paolini, A. Pasquarello, L. Paulatto, C. Sbraccia,  
 1295 S. Scandolo, G. Sclauzero, A. P. Seitsonen, A. Smogunov, P. Umari,

- 1296 R. M. Wentzcovitch, QUANTUM ESPRESSO: A modular and open-  
1297 source software project for quantum simulations of materials, *Journal of*  
1298 *Physics: Condensed Matter* 21 (39) (2009) 395502.
- 1299 [47] G. Kresse, J. Furthmüller, Efficient iterative schemes for ab initio total-  
1300 energy calculations using a plane-wave basis set, *Physical Review B* 54  
1301 (1996) 11169–11186.
- 1302 [48] S. Mohr, L. E. Ratcliff, L. Genovese, D. Caliste, P. Boulanger,  
1303 S. Goedecker, T. Deutscha, Accurate and efficient linear scaling DFT  
1304 calculations with universal applicability, *Physical Chemistry Chemical*  
1305 *Physics* 17 (47) (2015) 31360–31370.
- 1306 [49] R. J. Harrison, G. Beylkin, F. A. Bischoff, J. A. Calvin, G. I. Fann,  
1307 J. Fosso-Tande, D. Galindo, J. R. Hammond, R. Hartman-Baker, J. C.  
1308 Hill, J. Jia, J. S. Kottmann, M.-J. Y. Ou, J. Pei, L. E. Ratcliff, M. G.  
1309 Reuter, A. C. Richie-Halford, N. A. Romero, H. Sekino, W. A. Shelton,  
1310 B. E. Sundahl, W. S. Thornton, E. F. Valeev, A. Vázquez-Mayagoitia,  
1311 N. Vence, T. Yanai, Y. Yokoi, MADNESS: A multiresolution, adaptive  
1312 numerical environment for scientific simulation, *SIAM Journal on Sci-*  
1313 *entific Computing* 38 (5) (2016) S123–S142.
- 1314 [50] <http://mrchemdoc.readthedocs.org> (Accessed: 2017-05-21).
- 1315 [51] L. Lin, J. Lu, L. Ying, W. E, Adaptive local basis set for Kohn-Sham  
1316 density functional theory in a discontinuous Galerkin framework I: Total  
1317 energy calculation, *Journal of Computational Physics* 231 (4) (2012)  
1318 2140–2154.
- 1319 [52] W. Hu, L. Lin, C. Yang, DGDFT: A massively parallel method for large  
1320 scale density functional theory calculations, *The Journal of Chemical*  
1321 *Physics* 143 (2015) 124110.
- 1322 [53] J. Bernholc, M. Hodak, W. Lu, Recent developments and applications of  
1323 the real-space multigrid method, *Journal of Physics: Condensed Matter*  
1324 20 (29) (2008) 294205.
- 1325 [54] L. Kronik, A. Makmal, M. L. Tiago, M. M. G. Alemany, M. Jain,  
1326 X. Huang, Y. Saad, J. R. Chelikowsky, PARSEC: The pseudopotential  
1327 algorithm for real-space electronic structure calculations: Recent

- 1328 advances and novel applications to nano-structures, *Physica Status So-*  
1329 *lidi* (b) 243 (5) (2006) 1063–1079.
- 1330 [55] G. H. Golub, C. F. V. Loan, *Matrix Computations*, Johns Hopkins Stud-  
1331 *ies in the Mathematical Sciences*, Johns Hopkins University Press, 2013.
- 1332 [56] C. H. Bischof, B. Lang, X. Sun, Algorithm 807: The SBR toolbox - soft-  
1333 *ware for successive band reduction*, *ACM Transactions on Mathematical*  
1334 *Software* 26 (4) (2000) 602–616.
- 1335 [57] J. J. M. Cuppen, A divide and conquer method for the symmetric tridi-  
1336 *agonal eigenproblem*, *Numerische Mathematik* 36 (2) (1980) 177–195.
- 1337 [58] C. Bischof, X. Sun, B. Lang, Parallel tridiagonalization through two-step  
1338 *band reduction*, in: *Proceedings of IEEE Scalable High Performance*  
1339 *Computing Conference*, 1994, pp. 23–27.
- 1340 [59] L. S. Blackford, J. Choi, A. Cleary, E. D’Azevedo, J. Demmel, I. Dhillon,  
1341 *J. Dongarra, S. Hammarling, G. Henry, A. Petitet, et al., ScaLAPACK*  
1342 *users’ guide*, SIAM, 1997.
- 1343 [60] E. Anderson, A. Benzoni, J. Dongarra, S. Moulton, S. Ostrouchov,  
1344 *B. Tourancheau, R. van de Geijn, Basic linear algebra communication*  
1345 *subprograms*, in: *Distributed Memory Computing Conference*, 1991.  
1346 *Proceedings.*, The Sixth, IEEE, 1991, pp. 287–290.
- 1347 [61] F. Mauri, G. Galli, R. Car, Orbital formulation for electronic-structure  
1348 *calculations with linear system-size scaling*, *Physical Review B* 47 (15)  
1349 (1993) 9973.
- 1350 [62] P. Ordejón, D. A. Drabold, M. P. Grumbach, R. M. Martin, Uncon-  
1351 *strained minimization approach for electronic computations that scales*  
1352 *linearly with system size*, *Physical Review B* 48 (1993) 14646–14649.
- 1353 [63] F. Mauri, G. Galli, Electronic-structure calculations and molecular-  
1354 *dynamics simulations with linear system-size scaling*, *Physical Review*  
1355 *B* 50 (1994) 4316–4326.
- 1356 [64] P. Ordejón, D. A. Drabold, R. M. Martin, M. P. Grumbach, Linear  
1357 *system-size scaling methods for electronic-structure calculations*, *Phys-*  
1358 *ical Review B* 51 (3) (1995) 1456.

- 1359 [65] M. P. Teter, M. C. Payne, D. C. Allan, Solution of schrödinger's equation  
1360 for large systems, *Physical Review B* 40 (1989) 12255.
- 1361 [66] M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, J. D. Joannopoulos,  
1362 Iterative minimization techniques for ab initio total-energy calculations: molecular dynamics and conjugate gradients, *Reviews of Modern  
1363 Physics* 64 (1992) 1045–1097.  
1364
- 1365 [67] J. Kim, F. Mauri, G. Galli, Total-energy global optimizations using  
1366 nonorthogonal localized orbitals, *Physical Review B* 52 (1995) 1640–  
1367 1648.
- 1368 [68] N. D. Mermin, E. Canel, Long wavelength oscillations of a quantum  
1369 plasma in a uniform magnetic field, *Annals of Physics* 26 (2) (1964)  
1370 247–273.
- 1371 [69] L. Lin, C. Yang, J. Meza, J. Lu, L. Ying, W. E, SelInv - An algorithm  
1372 for selected inversion of a sparse symmetric matrix, *ACM Transactions  
1373 on Mathematical Software* 37 (2011) 40.
- 1374 [70] M. Jacquelin, L. Lin, C. Yang, PSelInv - A distributed memory parallel  
1375 algorithm for selected inversion: The symmetric case, *ACM Transactions  
1376 on Mathematical Software* 43 (3) (2016) 21.
- 1377 [71] J. J. Sylvester, A demonstration of the theorem that every homogeneous  
1378 quadratic polynomial is reducible by real orthogonal substitutions to the  
1379 form of a sum of positive and negative squares, *Philosophical Magazine*  
1380 4 (1852) 138–142.
- 1381 [72] C.-L. Fu, K.-M. Ho, First-principles calculation of the equilibrium  
1382 ground-state properties of transition metals: Applications to Nb and  
1383 Mo, *Physical Review B* 28 (1983) 5480–5486.
- 1384 [73] P. Pulay, Convergence acceleration of iterative sequences. The case of  
1385 SCF iteration, *Chemical Physics Letters* 73 (2) (1980) 393–398.
- 1386 [74] J. Nieplocha, B. Palmer, V. Tipparaju, M. Krishnan, H. Trease, E. Aprà,  
1387 Advances, applications and performance of the Global Arrays shared  
1388 memory programming toolkit, *The International Journal of High Per-  
1389 formance Computing Applications* 20 (2) (2006) 203–231.

- 1390 [75] The ELSI team, ELSI Users' Guide, <http://elsi-interchange.org>, 2017.
- 1391 [76] J. P. Perdew, K. Burke, M. Ernzerhof, Generalized gradient approxima-  
1392 tion made simple, *Physical Review Letters* 77 (1996) 3865–3868.
- 1393 [77] V. Havu, V. Blum, P. Havu, M. Scheffler, Efficient integration for all-  
1394 electron electronic structure calculation using numeric basis functions,  
1395 *Journal of Computational Physics* 228 (22) (2009) 8367–8379.
- 1396 [78] L. Nemeč, V. Blum, P. Rinke, M. Scheffler, Thermodynamic equilibrium  
1397 conditions of graphene films on SiC, *Physical Review Letters* 111 (6)  
1398 (2013) 065502.
- 1399 [79] X. S. Li, J. W. Demmel, SuperLU\_DIST: A scalable distributed-memory  
1400 sparse direct solver for unsymmetric linear systems, *ACM Transactions*  
1401 *on Mathematical Software* 29 (2) (2003) 110–140.
- 1402 [80] L. Grigori, J. W. Demmel, X. S. Li, Parallel symbolic factorization for  
1403 sparse LU with static pivoting, *SIAM Journal on Scientific Computing*  
1404 29 (3) (2007) 1289–1314.
- 1405 [81] M. Jacquelin, Y. Zheng, E. Ng, K. Yelick, An asynchronous task-based  
1406 fan-both sparse Cholesky solver, arXiv:1608.00044.
- 1407 [82] J. Lu, H. Yang, Preconditioning orbital minimization method for  
1408 planewave discretization, *Multiscale Modeling and Simulation* 15 (1)  
1409 (2017) 254–273.
- 1410 [83] H. Yang, J. Lu, PSPBLAS: A new framework for distributed memory  
1411 sparse BLAS, in preparation (2017).