FEAST fundamental framework for electronic structure calculations

Alan Levin, Deyin Zhang, and Eric Polizzi*

Department of Electrical and Computer Engineering,

University of Massachusetts, Amherst, Massachusetts 01003, USA

(Dated: June 21, 2011)

Abstract

In a recent article [1], FEAST has been presented as a general purpose eigenvalue algorithm which is ideally suited for addressing the numerical challenges in electronic structure calculations. Here, FEAST is presented furthermore as a fundamental modeling framework which can naturally address the original numerical complexity of the electronic structure problem as formulated by Slater in 1937 [2]. In doing so, the FEAST framework is capable of bypassing the motivations and needs for the approximation techniques used in first-principle calculations nowadays.

PACS numbers: 02.60.x,02.70.c,31.15.p,71.15.Dx

*Electronic address: polizzi@ecs.umass.edu

Since the 1930's, progress in electronic structure calculations has always been tied together with advances in numerical strategies for addressing the eigenvalue problem. In particular, several attempts have been undertaken to reduce the complexity of the eigenvalue problem in self-consistent calculations by dissociating, removing or screening the effect of the core electrons. Most used techniques include [3]: muffin-tin approximations along with augmented plane wave (APW) [2] and linearized APW, muffin-tin orbitals (MTO) and linearized MTO, KKR methods and pseudopotential approaches [4, 5]. The conceptual approach of the former consists in partitioning the real space into spheres around each atom, allowing different discretization and solving strategies to take place in separate regions in space. Therefore, the atom-centered regions can benefit from specific discretization schemes (i.e. basis sets) that are both suitable to capture the highly localized core states around the nuclei and considerably reduce the effective size of the resulting eigenvalue problem in the interstitial region. This approach can be cast as a domain decomposition method in modern days which is most suitable for parallel computing since calculations on all these sub-domains can also be performed independently. Once the eigenvalue problem is reformulated using domain decomposition strategies, however, the resulting (and still exact) problem takes now the form of a non-linear one in the interstitial region (i.e. $H(E)\psi = E\psi$) since the boundary conditions at the interface with the atom-centered regions are energy dependent. The major difficulty of solving this non-linear eigenvalue problem has been largely avoided by the mainstream approaches to electronic structure calculations to rely most entirely on approximations ranging from direct linearization techniques (e.g. LAPW, LMTO, etc.) to pseudopotential techniques. Linear eigenvalue problems can indeed be obtained alternatively from pseudopotential techniques using smooth but non-local potential in atom-centered regions that eliminate the core states.

This paper presents a fundamental strategy for performing electronic structure calculations which bypasses the need for the traditional approximation techniques above by relying entirely on the capabilities of the new FEAST algorithm framework for solving the eigenvalue problem. Using FEAST, solving the original eigenvalue problem within a given search interval is mainly reformulated into solving a set of well-defined independent linear systems along a complex energy contour [1]. As a result, we show here that FEAST can naturally address two main fundamental issues for electronic structure calculations: (i) muffin-tin domain decomposition type approaches used to partition the real-space can act directly on the

linear systems; therefore, the eigenvalue problem does not need to be explicitly formulated into a non-linear one. Stated otherwise, the choice of "pivot energies", needed to evaluate the Green's function in the interstitial region using linearization techniques, becomes now a well-defined choice which follows a mathematically sounded approach for solving the electronic structure problem in the entire system; (ii) several FEAST algorithms can operate in parallel to obtain core and valence electrons independently spanning different energy ranges.

I. FROM LINEAR TO NON-LINEAR EIGENVALUE PROBLEM

In electronic structure calculations, one considers solving the Schrodinger-type equation in an entire domain Ω which can be finite, periodic or Bloch periodic:

$$H\Psi(\mathbf{x}) = E\Psi(\mathbf{x}), \quad \mathbf{x} \in \Omega \tag{1}$$

where $\{E_i, \Psi_i\}$ are the resulting eigenpairs (also parametrized by k in the case of bandstructure calculation using a Bloch periodic system). Thereafter, any discretization schemes in Ω would give rise to the generalized and linear eigenvalue problem:

$$\mathbf{H\Psi} = E\mathbf{S\Psi}, \tag{2}$$

where $\mathbf{S}$ is a positive definite matrix (mass matrix) obtained using non-orthogonal basis functions ($\mathbf{S} = \mathbf{I}$ otherwise), and $\mathbf{\Psi}$ contains the unknown components of the wave function (e.g. basis set coefficients, nodal values, etc.). Without any loss of generality, Figure 1 illustrates the essence of the muffin-tin domain decomposition strategy.

Formally, the solutions $\{E_i, \Psi_i\}$ that satisify the continuum model (1), can also be obtained from a Schrödinger equation in the interstitial region Ω_0 alone provided that appropriate boundary conditions are imposed at the interfaces Γ_j with the atom-centered region Ω_j i.e.

$$H_0\Psi(\mathbf{x}) = E\Psi(\mathbf{x}), \quad \mathbf{x} \in \Omega_0 \tag{3}$$

where H_0 is the Hamiltonian in Ω_0. A general mathematical form for these boundary conditions on Γ_j supplies a relation between the normal derivative of the solution and their boundary values ($\forall j$):

$$\frac{\hbar^2}{2m}\frac{\partial\Psi(\mathbf{x})}{\partial\eta_j} = \int_{\Gamma_j} d\mathbf{x}^{'}\ \Sigma_j(E, \mathbf{x}, \mathbf{x}^{'})\ \Psi(\mathbf{x}^{'}), \quad \mathbf{x} \in \Gamma_j, \tag{4}$$

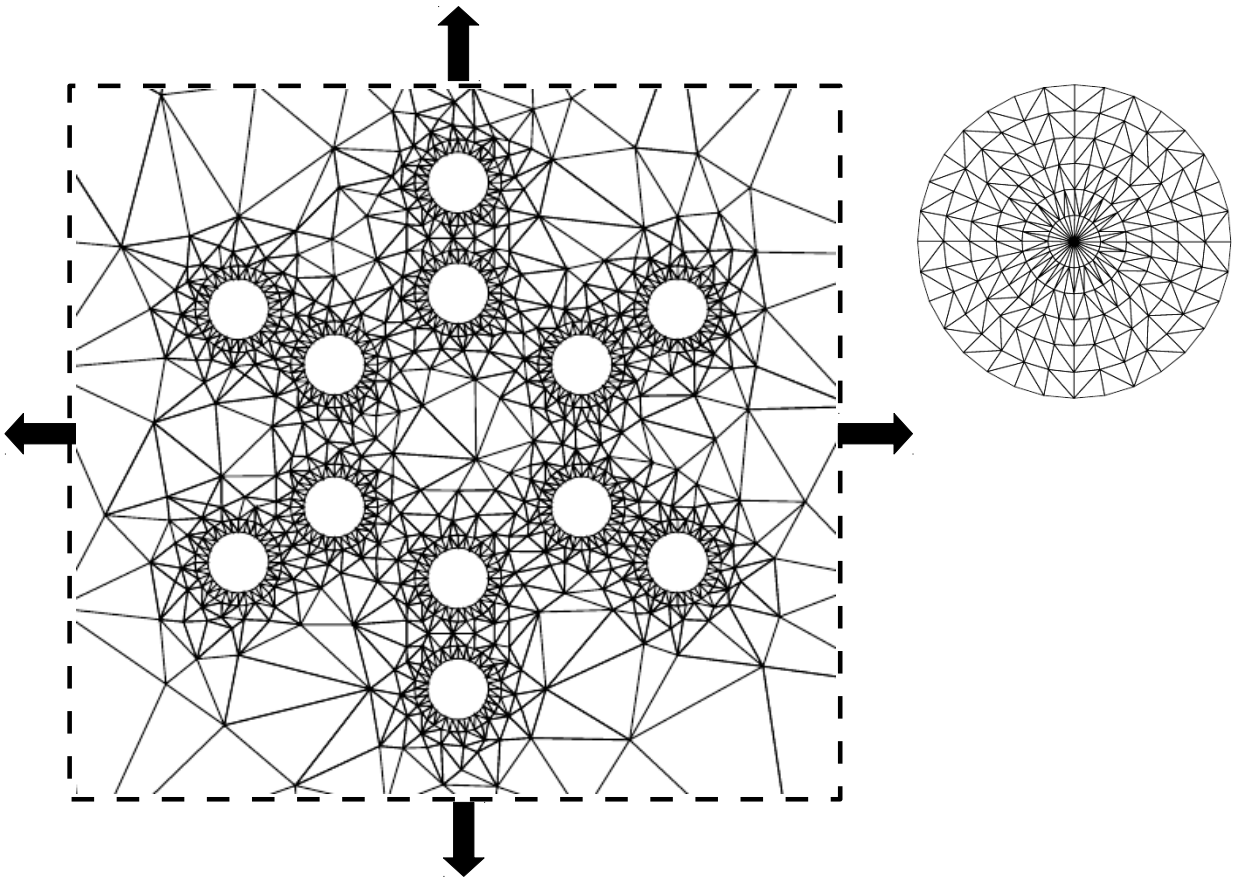

FIG. 1: *Using a muffin-tin domain-decomposition method, the whole simulation domain* Ω *is separated into multiple atom-centered regions* Ω_j *(*$j = 1, \dots$*) and one large interstitial region* Ω_0*. Different basis-sets can be used independently to describe the different regions. The figures represent a 2D cross-section of local finite element discretization using a coarse mesh for* Ω_0 *connecting all of the atoms of a Benzene molecule, and a much finer mesh for the* Ω_j *regions suitable to capture the highly localized core states around the nuclei.*

where η_j represents here the external normal at Γ_j, and Σ_j is a non-local and energy dependent operator (i.e. self-energy) which can be derived from the the atom-centered Green's function G_j in Ω_j. This later is given by $(\forall j)$:

$$\Big(E - H_j\Big) G_j(E, \mathbf{x}, \mathbf{x}^{'}) = \delta(\mathbf{x} - \mathbf{x}^{'}), \quad \mathbf{x}, \mathbf{x}^{'} \in \Omega_j, \tag{5}$$

where H_j is the Hamiltonian in Ω_j, and G_j can be constructed with arbitrary boundary conditions at Γ_j. For instance, by choosing the Green's function G_j to have zero derivative on Γ_j (i.e. homogeneous Neumann boundary conditions), one can obtain from the Green's

identity, a simple expression for Σ_j (inverse of the surface Green's function):

$$\Sigma_j(E, \mathbf{x}, \mathbf{x}') = G_j^{-1}(E, \mathbf{x}, \mathbf{x}'), \quad \mathbf{x}, \mathbf{x}' \in \Gamma_j. \tag{6}$$

This derivation was originally introduced in [6] as an embedding potential technique for the Schrödinger equation. Alternatively, another simple expression for Σ_j has been derived in [7] using homogeneous Dirichlet boundary conditions for G_j on Γ_j. After discretization of (3) using the condition (4) (and usually performed on the variational form of the problem), the resulting non-linear eigenvalue problem in Ω_0 takes the general form

$$\big(\mathbf{H_0} - \sum_j \mathbf{\Sigma_j}(E)\big)\mathbf{\Psi_0} = E\mathbf{S_0}\mathbf{\Psi_0}, \tag{7}$$

where $\mathbf{S_0}$ is the mass matrix in Ω_0, and $\mathbf{\Psi_0}$ contains the unknown components of the solution in Ω_0 as well.

In the specific case of real-space mesh discretization (see Figure 1), $\mathbf{\Sigma_j}(E)$ is non-zero only for a small number of matrix elements coupling all the n_j unknowns on Γ_j. Using (5) and (6), and assuming a particular ordering of the matrix elements (for clarity), it results $(\forall j)$:

$$\hat{\mathbf{\Sigma}}_\mathbf{j}(E) = \left([\mathbf{I}_{n_j}\ \mathbf{0}\dots\mathbf{0}]\ (E\mathbf{S_j} - \mathbf{H_j})^{-1}\ [\mathbf{I}_{n_j}\ \mathbf{0}\dots\mathbf{0}]^T\right)^{-1}, \tag{8}$$

where $\mathbf{S_j}$ is the mass matrix in Ω_j, and the matrix $\hat{\mathbf{\Sigma}}_\mathbf{j}$ of size n_j contains all the non-zero elements of $\mathbf{\Sigma_j}$.

Alternatively to a continuum treatment of the problem (1), one could directly replace the unknown components of $\mathbf{\Psi_j}$ belonging to the interior subdomains Ω_j from the system matrix (2) by the following self-energy [8]

$$\mathbf{\Sigma_j}(E) = \tau_\mathbf{j}\ \mathbf{G_j}\ \tau_\mathbf{j}^\dagger, \tag{9}$$

where $\tau_\mathbf{j}$ describes the interaction between Ω_0 and the atom-centered region Ω_j. In linear algebra, this non-overlapping domain decomposition procedure gives rise to a reduced coupled system identical to (7) which is known as the Schur complement.

II. FEAST FRAMEWORK

FEAST is both a new numerical algorithm [1] and a new general purpose high-performance numerical library [9] for solving the standard or generalized eigenvalue problem

of type (2), and obtaining all the eigenvalues and eigenvectors within a given search interval $[E_{min}, E_{max}]$. The algorithm offers many important and unique capabilities for achieving accuracy, robustness, high-performance and scalability on parallel computing architectures. The FEAST basic computing steps for solving the system (2) of size N, are briefly summarized in the following.

Starting from a set of M_0 linearly independent random vectors $\mathbf{Y}_{N\times M_0} = \{\mathbf{y_1}, \mathbf{y_2}, ..\mathbf{y_{M_0}}\}$, where M_0 is chosen greater than the number of the eigenvalues M in the search interval (i.e. M_0 represents then an over-estimation of M which is not known *a priori*), a new set of vectors $\mathbf{Q}_{N\times M_0} = \{\mathbf{q_1}, \mathbf{q_2}, ..\mathbf{q_{M_0}}\}$ is obtained as follows:

$$\mathbf{Q}_{N\times M_0} = -\frac{1}{2\pi\imath}\int_{\mathcal{C}} dZ\ \mathbf{G}(Z)\mathbf{Y}_{N\times M_0}, \tag{10}$$

where $\mathcal{C}$ represents a complex contour from E_{min} to E_{max}. In practice, the vectors $\mathbf{Q}$ in (10) can be computed using a high-order numerical integration where only very few linear systems $\mathbf{G}(Z)\mathbf{Y}$ need to be solved along the complex contour $\mathcal{C}$ i.e.

$$(Z\mathbf{S} - \mathbf{H})\mathbf{Q}^{(\mathbf{Z})} = \mathbf{Y}, \tag{11}$$

where $\mathbf{Q}^{(\mathbf{Z})}$ denotes the set of responses at a given pivot energy Z for a given set of excitations $\mathbf{Y}$ in Ω.

Thereafter, by computing

$$\mathbf{H}_{\mathbf{Q}_{M_0\times M_0}} = \mathbf{Q}^\dagger\mathbf{H}\mathbf{Q} \quad \text{and} \quad \mathbf{S}_{\mathbf{Q}_{M_0\times M_0}} = \mathbf{Q}^\dagger\mathbf{S}\mathbf{Q}, \tag{12}$$

a projected reduced dense eigenvalue problem of size M_0 can be formed:

$$\mathbf{H}_\mathbf{Q}\mathbf{\Phi} = \epsilon\mathbf{S}_\mathbf{Q}\mathbf{\Phi}, \tag{13}$$

This reduced problem can be solved using standard eigenvalue routines for dense systems to obtain all the eigenpairs $(\epsilon_m, \mathbf{\Phi}_m)$. By setting $E_m = \epsilon_m$ and computing $\mathbf{\Psi}_{N\times M_0} = \mathbf{Q}_{N\times M_0}\mathbf{\Phi}_{M_0\times M_0}$, it follows that if E_m lies inside the contour, it is an eigenvalue solution and its eigenvector is $\mathbf{\Psi}_m$ (the m^{th} column of $\mathbf{\Psi}$). The eigenvectors $\mathbf{\Psi}$ are also naturally $\mathbf{S}$-orthonormal, if the eigenvectors of the reduced problem are $\mathbf{S_Q}$-orthonormal. In order to improve the accuracy, a new set of initial guess vectors $\mathbf{Y} = \mathbf{S}\mathbf{\Psi}$ can also be used iteratively up until convergence.

Finally, efficient parallel implementations for FEAST can be addressed at three different levels: (i) many search intervals can be run independently (no overlap), (ii) each linear system (11) can be solved independently along the complex contour $\mathcal{C}$, and (iii) the linear system can also be solved in parallel (the multiple right sides can be parallelized as well). Consequently, one can show that if enough parallel computing power is available at hand, the main computational cost of FEAST for solving the eigenvalue problem and even for capturing millions of eigenpairs, can be ultimately reduced to solving only one linear system (11).

III. IMPLICIT TREATMENT OF THE NON-LINEAR PROBLEM

Using the muffin-tin domain decomposition framework, solving explicitly the non-linear eigenvalue problem (7) could be possible but very challenging (difficulties would include in particular: absence of orthogonality for $\mathbf{\Psi_0}$ in Ω_0, and a non-linear reduced system (13) using FEAST).

All these issues, however, can be addressed implicitly within the FEAST framework. At first, starting from a set of excitations $Y(\mathbf{x})$ in the continuum domain, the set of responses $Q^{(Z)}$ can also be obtained by solving the Schrödinger equation (3) in Ω_0 alone:

$$(z - H_0)Q^{(Z)}(\mathbf{x}) = Y(\mathbf{x}), \quad \mathbf{x} \in \Omega_0, \tag{14}$$

where the boundary condition for $Q^{(Z)}$ on Γ_j should formally satisfy (4) but augmented by a source term $F_j^{(Z)}(\mathbf{x})$ (to add to the right hand side) which accounts for the effects of the excitations $Y(\mathbf{x})$ within the atom-centered regions Ω_j. For instance, using Neumann boundary condition for G_j, the self-energy Σ_j are defined in (6) and one can show that ($\forall j$):

$$\begin{aligned} F_j^{(Z)}(\mathbf{x}) = & \int_{\Gamma_j} d\mathbf{x}' \; G_j^{-1}(z, \mathbf{x}, \mathbf{x}') * \\ & \Big[\int_{\Omega_j} d\mathbf{x}'' \; G_j(z, \mathbf{x}', \mathbf{x}'') \; Y(\mathbf{x}'') \Big], \quad \mathbf{x} \in \Gamma_j. \end{aligned} \tag{15}$$

Once $Q^{(Z)}$ is known in Ω_0 and hence on all the Γ_j interfaces, the solution in the Ω_j can be independently retrieved for $\forall j$ by solving the linear systems

$$(z - H_j)Q^{(Z)}(\mathbf{x}) = Y(\mathbf{x}), \quad \mathbf{x} \in \Omega_j, \tag{16}$$

with Dirichlet boundary conditions.

After discretization of (14), (4) and (15), solving (11) in the entire domain Ω can then be replaced by solving the following linear system for the unknown components of the solutions $\mathbf{Q_0^{(Z)}}$ in Ω_0

$$\Big(Z\mathbf{S_0} - \mathbf{H_0} + \sum_j \mathbf{\Sigma_j}(Z)\Big)\mathbf{Q_0^{(Z)}} = \mathbf{Y_0} + \sum_{\mathbf{j}} \mathbf{F_j^{(Z)}}, \tag{17}$$

and a series of independent sub-problems (16) to obtain the unknown components of the solutions $\mathbf{Q_j^{(Z)}}$ in the atom-centered regions Ω_j. Using Neumann boundary conditions for G_j, in particular, Σ_j is given in (6) and from the discretization of (15), it comes:

$$\mathbf{F_j^{(Z)}} = \mathbf{\Sigma_j}(Z)\mathbf{G_j}(Z)\mathbf{Y_j}. \tag{18}$$

Thereafter, the subspace $\mathbf{Q}$ (10) is obtained by integration of the set of solutions $\mathbf{Q_0^{(Z)}}$ and all the $\mathbf{Q_j^{(Z)}}$ over the complex contour $\mathcal{C}$. In practice, it is possible to construct $\mathbf{H_q}$ and $\mathbf{S_q}$ in (12) directly from the projection of $\mathbf{H_0}$ and $\mathbf{S_0}$ for Ω_0 and $\mathbf{H_j}$ and $\mathbf{S_j}$ for all Ω_j.

As a result of (17) which is solved only for specific complex pivot energies Z, the non-linearity of the Schrödinger equation (7) in Ω_0 is then explicitly removed. It is also important to mention that the additional computational costs by pivot energy Z for obtaining $\mathbf{F_j^{(Z)}}$ and retrieving the solution $\mathbf{Q_j^{(Z)}}$ in Ω_j, can be made minimal. Indeed, most of the efforts that have been devoted for obtaining $\mathbf{\Sigma_j}(Z)$ do not need to be repeated (e.g. factorization of the matrix $(Z\mathbf{S_j} - \mathbf{H_j})$, computations of some key elements of $\mathbf{G_j}$). For instance, in the specific case of real-space mesh discretization (see Figure 1), only n_j columns of $\mathbf{G_j}$ associated to the nodes at the boundary Γ_j are needed to compute both $\mathbf{\Sigma_j}(Z)$ in (8) and $\mathbf{F_j^{(Z)}}$ in (18) and this independently on the number of nodes inside Ω_j. Alternatively, $\mathbf{F_j^{(Z)}}$ can also be obtained using the Schur complement technique used to construct $\mathbf{\Sigma_j}$ in (9).

IV. DISCUSSIONS

In 1937, Slater originally derived a non-linear electronic structure problem by introducing the APW method using a muffin-tin domain decomposition, he then stated [2] *"Of course, we cannot solve this exactly, and we must look for methods of approximations"*. Indeed, the non-linear type problem (7) cannot be handled by traditional linear eigenvalue algorithms and its explicit treatment would appear very challenging. These limitations have historically motivated the development of a wide spectrum of approximation techniques ranging from direct linearization to pseudopotential methods.

Within the framework of the FEAST algorithm, however, this problem benefits now from a numerical exact treatment which consists of removing the non-linearity of the Schrödinger equation (7) in the interstitial region Ω_0, by considering only certain complex pivot energies Z (17). In contrast to linear approximations (including LAPW, LMTO, linearized embedding method [6], etc.), these pivot energies are explicitly provided by FEAST to guaranteed convergence of the solutions of the Schrödinger equation in the whole simulation domain Ω. The approach that has been presented here, is free from any particular form for the potential in the atom-centered regions (such as spherically symmetric potential, etc.). Since the size of linear system (17) in the interstitial region Ω_0 is independent of the discretization schemes for the atom-centered regions Ω_j, the approach can ideally deal with full potential within self-consistent calculations (i.e. all-electron calculations).

The development of techniques such as pseudopotential have also been originally motivated to ease several numerical difficulties that one can encounter with all-electron calculations in the atom-centered regions [4, 5]. Let us then outline how some of these main issues are naturally addressed within the FEAST all-electron framework:

(i) Since FEAST can act independently on different energy ranges, the number of states in a search interval can be narrowed as desired, and the frozen-core approximation does not need to be considered within self-consistent iterations.

(ii) In contrast to pseudopotential, a much finer level of discretization for the FEAST all-electron framework is indeed needed to capture the (true) wave functions in Ω_j. The linear eigenvalue system obtained using pseudopotential can either be seen as a much smaller size system as compared to (2), or a linearized version of (7) where Σ_j represents then the pseudopotential which is now fully non-local over Ω_j. This pseudopotential system can also ideally be solved using the FEAST algorithm; however, the resulting system matrix (11) would end up (paradoxically) being larger and much less scalable than (17). As mentioned above, the extra-arithmetic costs for obtaining the vectors $\mathbf{F_j}$ and retrieving the solution $\mathbf{Q_j}$ on each subdomain are minimal. Moreover, these computations along with the ones for obtaining the self-energy $\mathbf{\Sigma_j}$, can be fully parallelized $\forall j$. Alternatively and similarly to the development of the atomic pseudopotential databases, one can envision precomputed spectral decomposition for the Green's function G_j built using all-electron self-consistent calculations for each atomic configuration. As a result, one could obtain very good and inexpensive preconditioner for computing $\mathbf{\Sigma_j}$, $\mathbf{F_j}$ and $\mathbf{Q_j}$ using iterative linear system solvers.

Finally, the FEAST fundamental framework for first-principle electronic structure calculations can be used independently of the choice for the physical model (e.g. Density-Functional Theory or Hartree-Fock), the nature of the atomistic system (e.g. isolated or Bloch periodic), or the choice for the basis set (e.g. PW, atomic orbitals, real-space mesh).

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No ECCS 0846457.

---

[1] E. Polizzi, Phys. Rev. B **79**, 115112 (2009).

[2] J. C. Slater, Phys. Rev. **51**, 846 (1937).

[3] R. M. Martin, *Electronic Structure: Basic Theory and Practical Methods* (Cambridge University Press, 2004).

[4] H. Hellmann, J. Chem. Phys. **3**, 61 (1935).

[5] J. C. Phillips and L. Kleinman, Phys. Rev. **116**, 287 (1959).

[6] J. E. Inglesfield, J. Phys. C **14**, 3795 (1981).

[7] A. J. Fisher, Journal of Physics: Condensed Matter **2**, 6079 (1990).

[8] J.M. Thijssen and J. E. Inglesfield, Europhysics Let. **27**, 65 (1994).

[9] E. Polizzi, *The feast eigenvalue solver*, URL `http://www.ecs.umass.edu/~polizzi/feast/`.