Journal Pre-proof

Spectral quadrature for the first principles study of crystal defects: Application to magnesium

Swarnava Ghosh and Kaushik Bhattacharya

PII: S0021-9991(22)00097-3
DOI: <https://doi.org/10.1016/j.jcp.2022.111035>
Reference: YJCPH 111035

To appear in: *Journal of Computational Physics*

Received date: 26 November 2020
Revised date: 16 November 2021
Accepted date: 29 January 2022

Please cite this article as: S. Ghosh and K. Bhattacharya, Spectral quadrature for the first principles study of crystal defects: Application to magnesium, *Journal of Computational Physics*, 111035, doi: <https://doi.org/10.1016/j.jcp.2022.111035>.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2022 Published by Elsevier.

Highlights

- An efficient computational method for first principles density functional theory study of crystal defects in crystalline solids.
- Spectral quadrature formulation of density functional theory, and two separate spectral quadrature methods – Gauss quadrature for the electronic fields and Clenshaw-Curtis quadrature for the atomic forces – for efficient computation.
- New insights into defect interaction in magnesium, a promising lightweight structural material.

Spectral quadrature for the first principles study of crystal defects: Application to magnesium

Swarnava Ghosh^a and Kaushik Bhattacharya^b

^a National Center for Computational Sciences, Oak Ridge National Laboratory, Oak Ridge, TN 37830

^b Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125

February 3, 2022

Abstract

We present an accurate and efficient finite-difference formulation and parallel implementation of Kohn-Sham Density (Operator) Functional Theory (DFT) for non periodic systems embedded in a bulk environment. Specifically, employing non-local pseudopotentials, local reformulation of electrostatics, and truncation of the spatial Kohn-Sham Hamiltonian, and the Linear Scaling Spectral Quadrature method to solve for the pointwise electronic fields in real-space and the non-local component of the atomic force, we develop a parallel finite difference framework suitable for distributed memory computing architectures to simulate non-periodic systems embedded in a bulk environment. Choosing examples from magnesium-aluminum alloys, we first demonstrate the convergence of energies and forces with respect to spectral quadrature polynomial order, and the width of the spatially truncated Hamiltonian. Next, we demonstrate the parallel scaling of our framework, and show that the computation time and memory scale linearly with respect to the number of atoms. Next, we use the developed framework to simulate isolated point defects and their interactions in magnesium-aluminum alloys. Our findings conclude that the binding energies of divacancies, Al solute-vacancy and two Al solute atoms are anisotropic and are dependent on cell size. Furthermore, the binding is favorable in all three cases.

Keywords: Spectral quadrature, Linear-Scaling, Density Functional Theory, Defects, Magnesium

1 Introduction

Crystal defects, even when present in small concentrations are crucial in determining macroscopic properties of materials. Vacancies, present in parts per million are fundamental to creep, spall and radiation aging. Dislocations, whose density is as small as 10^{-8} per atomic row, are the primary carriers of plasticity in metals, solutes present in parts per hundred are responsible for strengthening by interacting with the motion of dislocations, further solutes can aggregate to nucleate a precipitate.

First principles calculations based on Kohn-Sham density functional theory (DFT)[23, 29] have become central to computational materials research, thereby providing fundamental insights into materials properties and behavior. The success of DFT can be attributed to its excellent predictive power with low accuracy to cost ratio compared to other wavefunction based electronic structure methods [9, 6, 59]. In spite of its success and widespread use, the efficient solution of Kohn-Sham equations is still computationally daunting, thereby restricting the range of physical systems that can be investigated. In particular, crystal defects have been particularly challenging because they lead to long-range interactions, the reason why they influence mechanical behavior at small concentrations.

These challenges have led to the development of a number of linear-scaling DFT methods. However, many of them assume exponential decay of the off-diagonal components of the density matrix γ , and truncate them to a finite width. While this is reasonable for insulators (since it requires the existence of a band-gap) and metals at finite electronic temperature, additional subtleties such as method of truncation, and numerical stability are of concern. An alternative linear-scaling approach – the linear scaling spectral Gauss

quadratures (LSSGQ) – was introduced by Suryanarayana, Bhattacharya and Ortiz [56]. The key idea is to write the density matrix as a (Reimann-Stieltjes) integral over the spectrum (energy states) of the linearized Hamiltonian and then to approximate it using quadratures. In particular, Gauss quadratures allows one to use the Lanczos algorithm to evaluate the diagonal components of the density matrix at $O(1)$ effort at each spatial point leading to a linear scaling algorithm when one has local pseudopotentials. Further, the local aspect allows one to introduce variable resolution where fine resolution is maintained where necessary and sub-grid sampling is used in regions of uniform deformation [56, 42, 43]. However, it is computationally difficult to compute off-diagonal components of the density matrix using Gauss quadratures, and this makes LSSGQ difficult to extend to nonlocal pseudopotentials.

Suryanarayana [54] subsequently showed that LSSGQ may be considered a generalization of the Fermi operator expansion (FOE) method using Lagrange polynomials. This work also proved error estimates for FOE using a polynomial basis, and showed that Gauss quadratures were the most efficient. It also showed that purification methods using Chebyshev polynomials are equivalent to the use of Clenshaw-Curtis quadratures. This understanding led to a series of efficient algorithms using spectral quadratures [45, 58] and applications to first-principles molecular dynamics [58, 67, 50].

The first goal of this work is to retain the efficiency of LSSGQ, but also extend it to nonlocal pseudopotentials. This requires a methodological advance that combines ideas of Ponga *et al.* [42] and Suryanarayana *et al.* [58]. As noted above, LSSGQ efficiently computes the energy and electron density using the Lanczos iteration since these only depend on the diagonal component of the density matrix. This is true for both local and non-local pseudopotentials, and therefore we continue to use LSSGQ to compute the energy and electron density for given atomic position. However, the Lanczos algorithm is unable to compute the off-diagonal components that are necessary to compute the non-local forces on the atoms. Further, Fermi operator expansion using Legendre polynomials associated with Gauss quadratures to evaluate off diagonal terms is inefficient. Therefore, we use Clenshaw-Curtis quadratures to compute the non-local contributions to the forces on the atom. We develop a structure that enables us to efficiently combine these methods.

We also provide some advances in the implementation of these methods. The parallelization approach of Ponga *et al.* [42] required example-specific assignment; instead we use domain decomposition enabling easier implementation on HPC platforms. The implementation of Suryanarayana *et al.* [58] was limited to periodic boundary conditions and cuboidal domains. These are not appropriate for the study of defects. Therefore, our implementation extends to general domains and Dirichlet boundary conditions.

The second goal of this work is to use this algorithm to study defects in magnesium. Magnesium is abundantly available on the earth's crust, the lightest among all commonly used structural metals, and has among the highest strength to weight ratio [25, 37, 41]. Aluminum is a commonly used alloying element, and the relative strength of AZ class of magnesium alloys can be attributed to the hexagonal closed-packed (HCP) structure of the magnesium matrix and the β phase $\text{Mg}_{17}\text{Al}_{12}$ precipitates with body-centred cubic structure (space group $I\bar{4}3m$) [37, 25, 41, 35].

We study isolated vacancy and isolated substitutional aluminum solute in a magnesium lattice along with defect pairs – vacancy pairs, solute-vacancy and solute pairs. These pairs play an important role in determining the mechanical behavior and processing of magnesium and its alloys. Vacancy clusters can give rise to prismatic dislocation loops [43, 15] or serve as nuclei for voids which in turn are important for spall [34, 20, 14]. Such clusters can only form if vacancies in fact can bind. Similarly, aluminum has limited solubility in magnesium and the resulting $\text{Mg}_{17}\text{Al}_{12}$ precipitates play a critical role in strengthening magnesium alloys [41, 25, 35, 36, 16]. This in turn requires both the diffusion of aluminum in a magnesium lattice and an accumulation of aluminum. The diffusion is greatly aided by the formation of solute-vacancy pairs while the accumulation of aluminum is aided by the binding of aluminum solutes. Finally, vacancy diffusion is important for dislocation climb, a critical mechanism in creep, and the formation of solute-vacancy pairs are again important. Previous DFT studies have shown a solute vacancy binding energy that is significantly smaller than that experimentally observed, and non-binding of Al solute pairs raising questions

about the mechanism of formation of Al-rich precipitates. We show that the study of these defects require large computational cells and non-periodic boundary conditions of the type afforded by our algorithm, and suggests that previous contradictory results may have been artifacts of small computational cells.

We introduce our method in Section 2 and describe the numerical implementation in Section 3. We study the convergence and performance of our implementation in Section 4. We study defects in Section 5, and close with brief comments in Section 6.

2 Methodology

2.1 Density Operator Formulation of Kohn Sham Density Functional Theory

We consider a cuboidal domain Ω with N atoms and N_e valence electrons. Let $\mathbf{R} = \{\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_N\}$ be the positions of the nuclei with valence charges $\{Z_1, Z_2, \dots, Z_N\}$, respectively. The free energy of this system in Kohn-Sham Density Functional Theory (DFT) [23, 29] and expressed in terms of density operator [38, 64] is

$$\mathcal{F}(\gamma, \mathbf{R}) = 2 \text{Tr} \left(-\frac{1}{2} \nabla^2 \gamma \right) + E_{xc}(\rho) + 2 \text{Tr} (\mathcal{V} \gamma) + E_{el}(\rho, \mathbf{R}) - \theta S(\gamma), \quad (1)$$

where γ is the density operator, $\text{Tr}(\cdot)$ denotes the trace of an operator, $\rho(\mathbf{x}) := 2\gamma(\mathbf{x}, \mathbf{x})$ is the electron density. The first term in (1) is the kinetic energy of non-interacting electrons and E_{xc} is the exchange correlation energy in the local density approximation (LDA). Specifically we use the Perdew-Wang parameterization [40] of the correlation energy calculated by Ceperley-Alder [10]. The third term, $2 \text{Tr} (\mathcal{V} \gamma)$, is the contribution of the non-local pseudopotential to the free energy. $\mathcal{V}$ is the non-local pseudopotential operator given by

$$\mathcal{V}(\mathbf{x}, \mathbf{x}') = \sum_J \mathcal{V}_J(\mathbf{x}, \mathbf{x}') = \sum_J \sum_{lm} c_{Jl} \chi_{Jlm}(\mathbf{x}, \mathbf{R}_J) \chi_{Jlm}(\mathbf{x}', \mathbf{R}_J)$$

where $\mathcal{V}_J$ is the contribution to the non-local pseudopotential operator from the J^{th} atom, and the summation over J is over all nuclei whose supports of the non-local projectors χ_{Jlm} overlap with domain Ω . The fourth term is the electrostatic energy which is expressed within the local reformulation framework [57, 39, 17] as

$$E_{el}(\rho, \mathbf{R}) = \sup_{\phi} \left\{ -\frac{1}{8\pi} \int_{\Omega} |\nabla \phi(\mathbf{x}, \mathbf{R})|^2 d\mathbf{x} + \int_{\Omega} (\rho(\mathbf{x}) + b(\mathbf{x}, \mathbf{R})) \phi(\mathbf{x}, \mathbf{R}) d\mathbf{x} \right\} + E_{self}(\mathbf{R}), \quad (2)$$

where ϕ denotes the electrostatic potential, b represents the total pseudocharge density of the nuclei and $E_{self}(\mathbf{R})$ is the self energy of the nuclei

$$E_{self}(\mathbf{R}) = -\frac{1}{2} \sum_I \int_{\Omega} b_I(\mathbf{x}, \mathbf{R}) V_I(\mathbf{x}, \mathbf{R}) d\mathbf{x}, \quad (3)$$

where b_I is the pseudocharge density of the I^{th} nucleus generating the potential V_I (the summation over I is over all nuclei in $\mathbb{R}^3$ whose pseudocharge densities overlap with Ω). The final term in (1) is the electronic entropy arising from the partial occupancies of the electronic states at a finite electronic temperature θ with

$$S(\gamma) = -2k_B \text{Tr} [(\gamma \log \gamma + (\mathcal{I} - \gamma) \log (\mathcal{I} - \gamma))] \quad (4)$$

and $\mathcal{I}$ the identity operator.

The *ground state* in DFT is obtained by minimizing the functional $\mathcal{F}(\gamma, \mathbf{R})$ over all atomic positions $\mathbf{R}$ and all density operators γ associated with N_e electrons. It is convenient to nest this minimization problem by first calculating the *electronic ground state*:

$$\hat{\mathcal{F}}(\mathbf{R}) = \inf_{\{\gamma \text{ s.t. } 2 \text{Tr}(\gamma) = N_e\}} \mathcal{F}(\gamma, \mathbf{R}) \quad (5)$$

and then relaxing over all atomic configurations

$$\mathcal{F}_0 = \inf_{\mathbf{R}} \hat{\mathcal{F}}(\mathbf{R}). \quad (6)$$

The Euler-Lagrange equation to the variational problem (5) is a nonlinear fixed-point problem:

$$\gamma = g(\mathcal{H}, \lambda_f; \theta) = \left(1 + \exp \left(\frac{\mathcal{H} - \lambda_f \mathcal{I}}{k_B \theta} \right) \right)^{-1} \quad (7)$$

where $k_B \theta$ is the electronic smearing, the Fermi energy λ_f is the Lagrange multiplier employed to enforce the constraint on the number of electrons, and the Hamiltonian $\mathcal{H}$ is

$$\mathcal{H} = -\frac{1}{2} \nabla^2 + V_{xc} + \phi + \mathcal{V}, \quad (8)$$

with $V_{xc} = \delta E_{xc} / \delta \rho$ the exchange-correlation potential and ϕ the solution of the Poisson equation

$$-\frac{1}{4\pi} \nabla^2 \phi(\mathbf{x}, \mathbf{R}) = \rho(\mathbf{x}) + b(\mathbf{x}, \mathbf{R}) \quad (9)$$

subject to appropriate boundary conditions. Note that $V_{eff} = V_{xc} + \phi$ is local (diagonal operator) and hence its action on a function f is given by $(V_{eff} f)(\mathbf{x}) = V_{eff}(\mathbf{x}) f(\mathbf{x})$. The action of the non-local pseudopotential operator $\mathcal{V}$ is given by

$$(\mathcal{V} f)(\mathbf{x}) = \sum_J \mathcal{V}_J f = \sum_J \sum_{lm} c_{Jl} \chi_{Jlm}(\mathbf{x}) \int_{\Omega} \chi_{Jlm}(\mathbf{x}', \mathbf{R}_J) f(\mathbf{x}') d\mathbf{x}'. \quad (10)$$

After evaluating the electronic ground state, the free energy is calculated using the functional:

$$\begin{aligned} \hat{\mathcal{F}}(\mathbf{R}) = & U + E_{xc}(\rho) - \int_{\Omega} V_{xc}(\rho) \rho(\mathbf{x}) d\mathbf{x} + \frac{1}{2} \int_{\Omega} (b(\mathbf{x}, \mathbf{R}) - \rho(\mathbf{x})) \phi(\mathbf{x}, \mathbf{R}) d\mathbf{x} \\ & - \theta S + E_{self}(\mathbf{R}), \end{aligned} \quad (11)$$

where $U = 2 \text{Tr}(\gamma \mathcal{H})$ is the *band structure energy*. The atomic force on the J^{th} atom is calculated using the expression

$$\mathbf{f}_J = \frac{\partial \hat{\mathcal{F}}(\mathbf{R})}{\partial \mathbf{R}_J} = \int_{\Omega} \nabla b_J(\mathbf{x}, \mathbf{R}_J) \phi(\mathbf{x}, \mathbf{R}) d\mathbf{x} - 4 \text{Tr}(\mathcal{V}_J \nabla \gamma) \quad (12)$$

where the first term is local (recall that $\phi(\mathbf{x}, \mathbf{R})$ depends only on the local or diagonal components $\rho(\mathbf{x}) = \gamma(\mathbf{x}, \mathbf{x})$ of the density operator) while the second term is non-local and requires the off-diagonal terms of the density operator.

2.2 Linear Scaling Spectral Quadrature

We follow the Spectral Quadrature (SQ) method [42, 43, 56, 54, 45, 58] for solving the DFT problem. The key idea is to write the ground state quantities as (Reimann-Stieltjes) integrals over the spectrum of the Hamiltonian. Given the Fermi level λ_F and the Hamiltonian $\mathcal{H}$, we can use (7) to write

$$(\eta, \gamma \eta) = \int_{\sigma} g(\mathcal{H}, \lambda_f; \theta) d\mu_{\eta, \eta}(\lambda) = \int_{\sigma} g(\lambda, \lambda_f; \theta) d\mu_{\eta, \eta}(\lambda) \quad (13)$$

for any function η where σ is the spectrum of $\mathcal{H}$, and $\mu_{\eta, \eta}$ is the spectral measure of $\mathcal{H}$ contracted with η . We use the spectral theorem [48] to obtain the second equality. We can now use quadratures to approximate the integral. In this work we use Gauss quadratures to find the electronic ground state and Clenshaw-Curtis quadratures to find the force on an atom. We note that the coarse-grained implementation of Ponga et. al. [42, 43] solely employed Gauss Quadrature for the electronic ground state and forces, while the implementation of Pratapa et.al. solely employed the Clenshaw-Curtis quadrature to calculate the electronic ground state and the atomic forces. As shown by Suryanarayana [54], the degree of the polynomial required by Clenshaw-Curtis quadrature is almost twice the degree of polynomial required for Gauss quadrature.

Spectral Gauss quadrature We follow the linear scaling spectral Gauss quadrature (LSSGQ) method of Suryanarayana, Bhattacharya and Ortiz [56] that exploits the structure of Gauss quadratures to evaluate the electronic ground state. In Gauss quadratures, we approximate any function $f(\lambda)$ in terms of Lagrange polynomials $P_k^\eta(\lambda)$ as

$$f(\lambda) \approx \sum_{k=1}^K f(\lambda_k^\eta) P_k^\eta(\lambda) \quad (14)$$

where K is the degree of the expansion and λ_k^η are the spectral nodes. We can use this expansion to approximate the integral of the function f over the spectrum of $\mathcal{H}$,

$$I[f] = \int_{\sigma} f(\lambda) d\mu_{\eta,\eta}(\lambda) \approx \sum_{k=1}^K f(\lambda_k^\eta) \left(\int_{\sigma} P_k^\eta(\lambda) d\mu_{\eta,\eta}(\lambda) \right) = \sum_{k=1}^K w_k^\eta f(\lambda_k^\eta) \quad (15)$$

where the spectral weight w_k^η denotes the integral $\int_{\sigma} P_k^\eta(\lambda) d\mu_{\eta,\eta}(\lambda)$.

Now, consider a discretization of the computational domain using a regular finite difference grid with N_d points, and let $\{\eta_q\}_{q=1}^{N_d}$ be a set of orthonormal basis functions associated with this discretization such that η_q is compactly supported near the q^{th} grid point. We can then approximate the integrals that make up the electronic ground state (cf. (11)) as

$$U = \sum_{q=1}^{N_d} u^q, \quad S = \sum_{q=1}^{N_d} s^q, \quad N_e = \sum_{q=1}^{N_d} \rho^q, \quad (16)$$

where

$$u^q = 2 \int_{\sigma} \lambda g(\lambda, \lambda_f; \theta) d\mu_{\eta_q, \eta_q}(\lambda) \approx \sum_{k=1}^K w_k^{\eta_q} \lambda_k^{\eta_q} g(\lambda_k^{\eta_q}, \lambda_f; \theta), \quad (17)$$

$$s^q = -2k_B \int_{\sigma} s(\lambda, \lambda_f; \theta) d\mu_{\eta_q, \eta_q}(\lambda) \approx -2k_B \sum_{k=1}^K w_k^{\eta_q} s(\lambda_k^{\eta_q}, \lambda_f; \theta), \quad (18)$$

$$\rho^q = 2 \int_{\sigma} g(\lambda, \lambda_f; \theta) d\mu_{\eta_q, \eta_q}(\lambda) \approx 2 \sum_{k=1}^K w_k^{\eta_q} g(\lambda_k^{\eta_q}, \lambda_f; \theta), \quad (19)$$

$g(\lambda, \lambda_f; \theta)$ is the Fermi-Dirac function

$$g(\lambda, \lambda_f; \theta) = \frac{1}{1 + \exp(\frac{\lambda - \lambda_f}{k_B \theta})}, \quad (20)$$

and $s(\lambda, \lambda_f; \theta)$ is the pointwise contribution to the electronic entropy

$$s(\lambda, \lambda_f; \theta) = g(\lambda, \lambda_f; \theta) \log g(\lambda, \lambda_f; \theta) + (1 - g(\lambda, \lambda_f; \theta)) \log(1 - g(\lambda, \lambda_f; \theta)). \quad (21)$$

In Gauss quadrature, the weights $\{w_k^\eta\}_{k=1}^K$ are fixed apriori, and the spectral nodes $\{\lambda_k^\eta\}_{k=1}^K$ are treated as unknowns. An efficient way of evaluating the spectral weights and nodes at any grid point q is by employing the Lanczos iteration

$$\begin{aligned} b_{k+1}^q v_{k+1}^q &= (\mathcal{H} - a_{k+1}^q) v_k^q - b_k^q v_{k-1}^q, \quad k = 0, 1, \dots, K-1 \\ v_{-1}^q &= 0, \quad v_0^q = \eta_q, \quad b_0^q = 1, \end{aligned} \quad (22)$$

where $a_{k+1}^q = (\mathcal{H}v_k^q, v_k^q)$, $k = 0, 1, \dots, K-1$, and b_k^q is computed such that $\|v_k^q\| = 1$, $k = 0, 1, \dots, K-1$. We denote the Jacobi matrix $\hat{J}_K^q$ as

$$\hat{J}_K^q = \begin{pmatrix} a_1^q & b_1^q & & & \\ b_1^q & a_2^q & b_2^q & & \\ & \ddots & \ddots & \ddots & \\ & & b_{K-2}^q & a_{K-1}^q & b_{K-1}^q \\ & & & b_{K-1}^q & a_K^q \end{pmatrix} \quad (23)$$

The nodes $\{\lambda_k^{\eta q}\}_{k=1}^K$ are calculated as the eigenvalues of $\hat{J}_K^q$, and the weights $\{w_k^{\eta q}\}_{k=1}^K$ are the squares of the first elements of the normalized eigenvectors of $\hat{J}_K^q$. The aforesaid form of Gauss quadrature with unknown nodes integrates polynomials upto degree $2K-1$. When the nodes are fixed apriori, the Gauss quadrature integrates polynomials upto degree $K-1$ [54].

The key observations of LSSGQ are that (i) the number of quadrature points K is independent of system size and (ii) the vectors v_k^q remains zero except for a small region around the grid point q during Lanczos iteration (22). Therefore, the evaluation of the spectral nodes at each grid point is $O(1)$ as is the evaluation of all the electronic quantities of interest (cf. (11) and (16)). This enables us to evaluate the electronic ground state at linear cost.

We further observe that this limited support of the vectors v_k^q enables the restriction of the Hamiltonian $\mathcal{H}$ to an appropriate subspace of the real-space computational domain Ω . This enables us to use domain decomposition in our numerical implementation discussed in Section. 3.

Spectral Clenshaw-Curtis quadrature While the Gauss quadrature provides a very efficient approach to evaluating the electronic ground state since it depends only on the diagonal terms of the density matrix, it is unable to evaluate quantities like the contribution of the non-local pseudopotential to the atomic force that depends on the off-diagonal components. Therefore, we use the spectral Clenshaw-Curtis quadrature [54, 45, 58].

The Hamiltonian is first scaled and shifted such that its spectrum lies in the interval $[-1, 1]$

$$\hat{\mathcal{H}} = \frac{\mathcal{H} - \omega\mathcal{I}}{\tau} \quad (24)$$

where $\omega = (\lambda^{max} + \lambda^{min})/2$, $\tau = (\lambda^{max} - \lambda^{min})/2$, and λ^{max} , λ^{min} are the maximum and minimum eigenvalues of $\mathcal{H}$, respectively. Given any function η , we can rewrite (13) using the scaled and shifted Hamiltonian $\hat{\mathcal{H}}$:

$$(\eta, \gamma\eta) = \int_{\sigma} g(\mathcal{H}, \lambda_f; \theta) d\mu_{\eta, \eta}(\lambda) = \int_{-1}^1 g(\hat{\mathcal{H}}, \hat{\lambda}_f; \hat{\theta}) d\mu_{\eta, \eta}(\hat{\lambda}) = \int_{-1}^1 g(\hat{\lambda}, \hat{\lambda}_f; \hat{\theta}) d\mu_{\eta, \eta}(\hat{\lambda}) \quad (25)$$

where $\hat{\lambda}_f = (\lambda_f - \omega)/\tau$ and $\hat{\theta} = \theta/\tau$ are the scaled Fermi energy and the scaled electronic temperature respectively. In the Clenshaw-Curtis variant of the spectral quadrature, any function $f(\hat{\lambda})$ is expanded in terms of Chebyshev polynomials $T_k^{\eta}(\hat{\lambda})$ as:

$$f(\hat{\lambda}) \approx \sum_{k=0}^K f(\hat{\lambda}_k^{\eta}) T_k^{\eta}(\hat{\lambda}) \quad (26)$$

where K is the degree of the expansion, $\hat{\lambda}_k^{\eta}$ are the spectral nodes which are fixed in Clenshaw-Curtis

quadrature. The non-local component of the atomic force [45, 58] as

$$\begin{aligned}
 4 \operatorname{Tr}(\mathcal{V}_J \nabla \gamma) &\approx 4 \sum_{p=1}^{N_d} (\eta_q, (\mathcal{V}_J \nabla g(\mathcal{H}, \lambda_f; \theta)) \eta_q) = 4 \sum_{q=1}^{N_d} (\eta_q, (\mathcal{V}_J \nabla g(\hat{\mathcal{H}}, \hat{\lambda}_f; \hat{\theta})) \eta_q) \\
 &\approx 4 \sum_{q=1}^{N_d} \sum_{k=0}^K c_k^q \eta_q^* \mathcal{V}_J \nabla v_k^q
 \end{aligned} \tag{27}$$

where c_k^q are constants

$$c_k^q = \frac{2}{\pi} \int_{-1}^1 \frac{g(\lambda, \hat{\lambda}_f; \hat{\theta}) T_k(\lambda)}{\sqrt{1 - \lambda^2}} d\lambda, \tag{28}$$

and $v_{k+1}^q = T_k(\hat{\mathcal{H}}) \eta_q$ are functions which are determined by the three term recurrence relation for Chebyshev polynomials:

$$\begin{aligned}
 v_{k+1}^q &= 2\hat{\mathcal{H}}v_k^q - v_{k-1}^q \\
 v_1^q &= \hat{\mathcal{H}}\eta_q, v_0^q = \eta_q.
 \end{aligned} \tag{29}$$

Once again, the number of quadratures is independent of the system size and therefore this evaluation scales linearly.

We highlight some key aspects about the two quadratures. First, the asymptotic convergence rate of Gauss spectral quadrature is almost twice the convergence rate of Clenshaw Curtis quadrature[54]. Second, the dominant cost of both Gauss and Clenshaw-Curtis quadratures are the K nodal Hamiltonian-vector multiplications, which scale at $\mathcal{O}(K)$ at each nodal point. However, in the Clenshaw-Curtis quadrature, the minimum and maximum eigenvalues of the nodal Hamiltonian needs to be computed for scaling and shifting the Hamiltonian. This is typically calculated using Lanczos iteration, thereby requiring additional matrix-vector products. The diagonalization of the tridiagonal matrix in Gauss quadrature or the evaluation of the Chebyshev coefficients in Clenshaw-Curtis quadrature, take only a small fraction of the total time. Third, the Gauss quadrature yields only the diagonal components of the local density matrix, whereas in Clenshaw-curtis quadrature, the off diagonal components can be expanded in terms of v_k^q in Eq. (29). Therefore, the typical cost associated with Clenshaw-Curtis quadrature is more than the cost associated with Gauss quadrature. In view of this, we use the Gauss quadrature to solve for the electronic ground state and calculate the local forces. We only use the Clenshaw-Curtis quadrature to calculate the nonlocal forces where it is necessary.

Though the spectral quadrature method is a generalization of the Fermi Operator Expansion (FOE) method (see [7] for a review of such methods), such as expanding the Fermi-Dirac function with a polynomial basis, it is worthwhile to note some key differences [55]. First, In FOE, the complete density matrix is computed, whereas in spectral quadrature, only the diagonal components are computed. Second, in FOE, the full Hamiltonian is employed and global sparse matrix-matrix products are computed, whereas in spectral quadrature, nodal Hamiltonian is employed and local sparse matrix-vector products are computed. Therefore, in spectral quadrature, the interprocessor communication is performed only once every SCF iteration, whereas in FOE, it is required during every global matrix-matrix product. Third, the Fermi energy in the spectral quadrature method is directly determined from Eqns. (16), (19) and (20), whereas in FOE, the Fermi energy is typically determined by another outer loop and requires storing the Chebyshev matrices. Fourth, the spectral quadrature method typically require lower storage than FOE.

3 Numerical Implementation

Fig. 1 shows the flowchart of the scheme employed to solve the DFT problem. The non-linear fixed point problem (Eqn. 7) is solved using the self-Consistent field (SCF) iteration. Briefly, given a charge density

Figure 1: Flowchart of the Self Consistent Field iteration for solving the DFT problem.

and electrostatic potential, we construct a linearized Hamiltonian which is used to compute the spectral weights nodes using Lanczos iteration from which the updated electronic states including an updated charge density and electrostatic potential are determined; the process iterates till convergence. Once the electronic ground state is established, the atomic relaxation for the overall ground state (Eqn. 6) is achieved using the Polak-Ribiere variant of the Non-Linear Conjugate Gradient (NLCG) method [51].

We discuss further details of the key components of the implementation below. In this paper, we are interested in isolated defects. Therefore, we consider a computational domain that is embedded in an infinite perfect crystal. The method can be extended to other situations including periodic boundary conditions and isolated clusters surrounded by vacuum.

Discretization Let $\Omega = L_1 \times L_2 \times L_3$ be a cuboidal computational domain aligned with the coordinate axis, and let it be discretized using a regular $n_1 \times n_2 \times n_3$ grid so that the grid spacing is h_i in the i^{th} coordinate where $L_i = n_i h_i$. We index the grid points with q and let K_Ω be the set of all grid points. We discretize the gradient and Laplace operators using finite differences on this grid [19, 18].

Hamiltonian at each grid point In each iteration of SCF, we need to determine the spectral nodes and weights at each grid point. As evident from (22) and (29), this requires the calculation of the action of Hamiltonian on vectors v_k^q . These vectors are compactly supported around a ball centered at the q^{th} grid point. Therefore, it is sufficient to work with the Hamiltonian $\mathcal{H}^q$ projected on to a smaller subspace around the q^{th} grid point. Specifically, we work with a cuboidal region of side $2R_{cut}$ and centered at the q^{th} grid point which we call *region of influence*. This is shown schematically in Fig. 2a. The controllable parameter R_{cut} depends on the order of the quadrature which in turn depends on material properties and electronic smearing temperature θ [45, 58].

For grid points close to the edge of Ω , the nodal Hamiltonian can extend beyond the computational domain for grid points which are near the boundary of Ω . However, since we consider problems where

Figure 2: (a) Region of influence associated with the q^{th} grid point and extended domain. (b) Neighbor communicator in domain decomposition. The orange region is the union of partitioned domains that influence the process p .

our computational domain is embedded in an infinite crystal, we have to compute the Hamiltonian on an extended region Ω' of size $(L_1 + 2R_{cut}) \times (L_2 + 2R_{cut}) \times (L_3 + 2R_{cut})$, and centered at the origin. The values of the Hamiltonian associated with the annular $\Omega' \setminus \Omega$ is obtained using precomputed electronic fields (ρ and ϕ) for the perfect crystal.

Domain decomposition We use domain decomposition for parallel implementation. The computational domain is partitioned into disjoint domains, $\Omega = \bigcup_{p=1}^{N_p} \Omega_p$, where Ω_p denotes the domain local to the p^{th} MPI rank, and N_p is the total number of MPI ranks. The collection of grid points belonging to the p^{th} MPI rank is denoted by K_Ω^p , where $K_\Omega = \bigcup_{p=1}^{N_p} K_\Omega^p$, and $K_\Omega^p \cap K_\Omega^q = \emptyset$ (null set) for $p \neq q$. In our implementation, we use the DMDA class available through the Portable, Extensible Toolkit for Scientific Computation (PETSc) [3, 4] for mesh management. The communication between processes is handled by Message Passing Interface (MPI) libraries [21, 22].

The region of influence of an grid point q owned by a process p (i.e. $q \in K_\Omega^p$) may extend to the spatial regions owned by neighboring processes. In such a situation, the values of the effective potential V_{eff} from neighboring processes are communicated to the process p using an MPI communicator. In Fig. 2b we schematically illustrate the parallel communication pattern involved for communicating the effective potential V_{eff} from neighboring processes. This reduces the number of MPI related calls otherwise required for matrix vector products.

Spectral weights, nodes and electronic fields In each SCF iteration, the spectral weights $\{w_k^q\}_{k=1}^K$ and nodes $\{\lambda_k^q\}_{k=1}^K$ are computed at every grid point $q \in K_\Omega$ from the projected Hamiltonian $\mathcal{H}^q$. Overall, this

is computationally the most expensive part of the method. However, the computation is local and with no MPI related calls. Further, we do not explicitly store the matrices, and their multiplication with a vector is performed in a matrix-free manner. These lead to excellent parallel efficiency.

Once the spectral weights $\{w_k^q\}_{k=1}^K$ and spectral nodes $\{\lambda_k^q\}_{k=1}^K$ are computed at all the grid points $q \in K_\Omega^p$ for all the processes p , we first solve the following equation for the Fermi energy λ_f :

$$N_e = 2 \sum_{p=1}^{N_p} \sum_{q \in K_\Omega^p} \sum_{k=1}^K w_k^q g(\lambda_k^q, \lambda_f; \theta) . \quad (30)$$

We utilize Brent's algorithm [8] for this purpose. The summation across the polynomial degree k and the grid point q is performed locally in each processes, and the summation across the processes p is performed with one global MPI communication call.

Once the Fermi energy is calculated, we calculate the point-wise band structure energy u^q , the point-wise entropy s^q and the point-wise electron density ρ^q , following (17), (18), (19). These are all local.

Electrostatic and effective potential Once the electron density ρ is calculated at the grid points, we calculate the total electrostatic potential ϕ at the grid points by solving the Poisson equation (9) on Ω subject to Dirichlet boundary conditions obtained from the perfect crystal outside using the generalized minimal residual algorithm (GMRES) [49]. Once the electrostatic potential is calculated at every grid point $q \in \Omega$, we calculate the effective potential V_{eff} at every grid point

$$V_{eff}^q = V_{xc}(\rho^q) + \phi^q , \quad (31)$$

where $V_{xc}(\rho^q)$ is the exchange correlation potential.

The convergence of the SCF iteration is accelerated by employing mixing schemes. In every SCF iteration, we mix the effective potential V_{eff} , where we have the option of employing Anderson mixing [2], Pulay mixing scheme [46] and its periodic [5] and restarted [44] variants.

We check the convergence of SCF iteration by calculating the normalized error in the effective potential.

Free energy The free energy is computed once the SCF iteration has converged using a discrete version of Eqn. 11

$$\hat{\mathcal{F}}(\mathbf{R}) \approx \hat{\mathcal{F}}^h(\mathbf{R}) = \sum_{p=1}^{N_p} \sum_{q \in K_\Omega^p} \left[u^q + h_1 h_2 h_3 \left\{ (\varepsilon_{xc}^q - V_{xc}^q) \rho^q + \frac{1}{2} (b^q - \rho^q) \phi^q \right\} - 2\theta s^q \right] + E_{self}^h , \quad (32)$$

where ε_{xc}^q is the contribution of the exchange correlation energy at the q^{th} grid point, and E_{self}^h is the discrete representation of the self energy (3):

$$E_{self}^h = -\frac{1}{2} (h_1 h_2 h_3) \sum_{p=1}^{N_p} \sum_{q \in K_\Omega^p} \sum_I b_I^q V_I^q \quad (33)$$

In evaluating them, the sum over the grid points are carried out locally on each MPI process followed by a global sum across all the MPI processes p .

Figure 3: Convergence of (a) energy and (b) force vs. mesh size

Atomic forces The final step is the computation of the atomic forces (Eqn. 12). This has two parts, the contributions of the local pseudopotential and the non-local pseudopotentials. The contribution of the local pseudopotential to the atomic force is calculated as

$$\mathbf{f}_{J,l} = \int_{\Omega} \nabla b_J(\mathbf{x}, \mathbf{R}_J) \phi(\mathbf{x}, \mathbf{R}) d\mathbf{x} \approx h_1 h_2 h_3 \sum_{p=1}^{N_p} \sum_{q \in K_{\Omega}^p} (\nabla_h b_J)^q \phi^q \quad (34)$$

where ∇_h is the gradient operator in the discrete setting. The summation over the the grid points q is local to every process, followed by one summation over the MPI processes.

The non-local contribution to the atomic force is calculated by employing the Clenshaw-Curtis quadrature described in Section 2.2. At each grid point $q \in K_{\Omega}$, we calculate the discrete Chebyshev vectors $\mathbf{v}_q^k$ using the iterative scheme (29), and the Clenshaw-Curtis quadrature weights c_q^k using (28) with the discrete nodal Hamiltonian $\mathbf{H}_q$. The non-local force is given by

$$\mathbf{f}_{J,nl} = -4 \text{Tr}(\mathcal{V}_J \nabla \gamma) \approx -4 \sum_{p=1}^{N_p} \sum_{q \in K_{\Omega}^p} \mathbf{e}_q^T \mathcal{V}_J^q \nabla_h \left(\sum_{k=0}^K c_q^k \mathbf{v}_k^q \right). \quad (35)$$

The calculation of the atomic forces scales linearly with the number of atoms.

4 Convergence and Performance

We set the electronic temperature to be $k_B \theta = 0.03333$ Ha, and use Troullier-Martins non-local pseudopotentials [60]. These yield lattice parameters of $a = 6.043$ Bohr, $c = 9.848$ Bohr and c/a ratio of 1.629 for hexagonal closed packed (HCP) magnesium, and $a=19.58$ Bohr for body centered cubic (BCC) $\text{Mg}_{17}\text{Al}_{12}$. These agree with the previously reported values of the lattice constants: $a=5.972$ Bohr, $c/a=1.61$ (DFT) [11], $a=6.066$ Bohr, $c/a=1.623$ (experiment) [28, 30] for Mg, and the equilibrium lattice constant of the $\text{Mg}_{17}\text{Al}_{12}$ or MgAl phase is $a = 19.96$ Bohr (DFT) [13], and $a = 19.653$ Bohr (experiment) [66] for $\text{Mg}_{17}\text{Al}_{12}$.

Convergence with mesh size We first verify the convergence with respect to mesh size. For this study, we choose HCP Mg, and randomly perturb the atoms from the ground state. Fig. 3a shows the convergence of

Figure 4: Convergence of (a) energy and (b) force as a function of degree of spectral quadrature.

the energy and Fig. 3b shows the convergence of atomic forces with h . The error is computed with respect to the reference solutions with mesh size $h \approx 0.2$ Bohr. From Figs. 3a and 3b, we see that a mesh size $h \approx 0.6$ Bohr is sufficient to obtain errors of 1×10^{-5} Ha/atom in energy and 1×10^{-5} Ha/Bohr in forces. The difference in energy and forces between a $h \approx 0.6$ Bohr calculation and a converged ABINIT calculation is 5×10^{-5} Ha/atom and 6×10^{-5} Ha/Bohr, respectively.

Convergence of spectral quadrature Next, we verify the convergence with respect to spectral quadrature. We take the same degree K for both the Gauss and the Clenshaw-Curtis quadrature. We study HCP Mg and BCC MgAl. For each of these systems, we randomly perturb the atoms from the ground state to obtain the test configuration, and use a mesh of $h \approx 0.6$ Bohr for pure Mg and $h \approx 0.5$ Bohr for MgAl.

Fig. 4a shows the convergence of the energy and Fig. 4b shows the convergence of atomic forces with K . The error is computed with respect to the reference that is taken to be the solutions for $K=240$. The decay is not monotone because neither the free energy functional (Eqn. 1) nor the atomic forces (Eqn. 12) is variational with respect to K . However, it broadly follows an exponential decay ($\text{Error} \approx Ae^{-\beta K}$). The best fit β for the various cases is shown in Table 1. The pure Mg system has a higher rate of convergence than the MgAl system. Furthermore, in both the cases, the rate of convergence in total energy and the local component of the atomic force is similar, whereas the rate of convergence in the total force, which require calculating the non-local force component is smaller by almost a factor 2. This is so because the local component of the atomic force and energy use only Gauss quadrature which depends on Lagrange polynomials whereas the non-local component of atomic forces depends on Clenshaw Curtis quadrature with Chebyshev polynomials. The former is known to be more accurate than the latter [54].

Since we require an accuracy of about 1×10^{-3} and 1×10^{-4} Ha/atom in our energy, we need K between 60 and 80 for Mg, and between 60 and 90 for MgAl. Similarly, since we require an accuracy between 1×10^{-3} and 1×10^{-4} Ha/Bohr in the total force, we need between 50 and 100 for Mg and MgAl. These guide our further calculations.

Convergence with radius of truncation A loose upper bound of the size R of Lanczos vectors in (22) is given by $2R \leq hn_oK$, where h is the average mesh spacing, and n_o is the finite difference order. This suggests a choice of $2R_{\text{cut}} = hn_oK$. Using the choice of parameters used for this study, $R_{\text{cut}} \approx 288$ Bohr. This would make the calculations extremely expensive. However, we now show that it is not necessary by

Table 1: Rates of convergence β of the total energy, local force and the total force with the spectral quadrature polynomial order K .

System	Energy	Force (local)	Force (total)
MgAl	0.070	0.075	0.036
Mg	0.109	0.088	0.042

Figure 5: Convergence of (a) energy and (b) total atomic force as a function of truncation radius R_{cut} of the nodal Hamiltonian.

studying the error as a function of R_{cut} .

For the Mg and MgAl systems, we use a Lanczos polynomial degree $K_L = 80$ and a Chebyshev polynomial degree $K_C = 100$, and vary R_{cut} from 4.0 to 14.0 Bohr. Fig. 5a and 5b shows the error in energy and atomic force as a function of R_{cut} . The reference values of energy and atomic force is calculated using an $R_{cut} = 16$ Bohr. From these figures, we observe that as R_{cut} increases, the error in energy and forces decrease. We once again observe generally exponential decay, though it is not monotone (also see [45]). This allows us to treat R_{cut} as a controllable approximation parameter.

Scaling and performance Next, we turn to the scaling and efficiency of the formulation and parallel implementation developed in this work. We choose magnesium crystal with one aluminum solute atom, placed at the center of the computational domain. The simulation parameters used are $h \approx 0.6$ Bohr, $R_{cut} = 12.0$ Bohr, $K_L = 80$ and $K_C = 100$ with a desired accuracy of 2×10^{-4} Ha/atom in energy and 2×10^{-4} Ha/Bohr in atomic force. We perform our simulations on a computer cluster with the following configuration: 16-core Intel Skylake processors (2.1 GHz), 32 cores per node, 192 GB memory/node and EDR Infiniband.

We first perform a strong scaling study with a 600 atom system, varying the number of cores from 1 to 420, and use 1 MPI rank per core. The wall times for each SCF iteration is presented in Fig 6a. The parallel efficiency on 420 cores relative to 1 core is 92 percent. This data is plotted in terms of speed up in Fig. 6b. The minimum wall time achieved was on 420 cores. When the number of cores is increased beyond 420, the wall time deviates from the ideal strong scaling line. Next, we perform a weak scaling study by selecting systems with sizes ranging from 64 atoms to 2560 atoms, while increasing the number of cores from 30 to 900. We choose these such that the number of atoms per MPI process is between two and three. Fig 6c shows that the variation of CPU time for one SCF iteration versus the number of atoms is linear ($\approx \mathcal{O}(N)$). Fig. 6d

Figure 6: Scaling and performance of the framework. The dash dot line is the ideal linear scaling behavior.

shows that the memory required also scales linearly with respect to the number of atoms. All of this shows excellent scaling of our algorithm. This is due in part to restricting much of the parallel communication to the MPI processes that are neighbors, and keeping the global communication to a minimum.

We note that the spectral quadrature step accounts for greater than 98 percent of the total time in each SCF iteration. The prefactor of spectral quadrature can be significantly reduced by incorporating reduced basis methods such as Discrete Discontinuous Basis Projection ([65]). Further, the number of SCF iterations to achieve a fixed target SCF error increases with system size in metallic systems due to charge sloshing [26]. The introduction of real space preconditioning schemes [52, 31] is likely to reduce this for large metallic systems.

5 Defects in magnesium

We now study isolated point defects and defect pairs in magnesium. Of particular interest are the formation energy of isolated defects and binding energy of defect pairs. Let $\mathcal{E}(M, n, m)$ be the energy of a crystal with M solvent atoms, n solute atoms and m vacancies. The formation energy of a defect cluster with n

solute atoms and m vacancies is the excess energy of the crystal with defect cluster over the those of perfect crystals of the host and solute:

$$E_{n,m}^f = \mathcal{E}(M - n - m, n, m) - \frac{M - n - m}{M} \mathcal{E}(M, 0, 0) - n\bar{\mathcal{E}}, \quad (36)$$

where $\bar{\mathcal{E}}$ is the energy per atom of the solute in its perfect crystalline state. Note that when we have an isolated solute and no vacancies, the formation energy $E_{1,0}^f$ is referred to as dilute impurity energy. Further, the binding energy of this cluster is

$$E_{n,m}^b = nE_{1,0}^f + mE_{0,1}^f - E_{n,m}^f. \quad (37)$$

Note that the defect is stable when the formation energy is negative, and the defect cluster has favorable binding when the binding energy is positive.

Isolated point defect We calculate the formation energy of a monovacancy for various computational domain size from 63 atoms to 1151 atoms, and the results are shown in Fig. 7a. We observe that the formation energy strongly depends on cell size, and converges at cell sizes of approximately 1000 atoms to 0.8456 eV. This is broadly in agreement with values reported in the literature: calculated values 0.768 - 0.779 eV using local pseudo-potential and a coarse grained approach with 1024 to 1 billion atoms [42] and measured values of 0.58 - 0.90 eV [33, 24, 61, 63]. Fig. 7b shows the electron density on the basal plane in the vicinity of the vacancy. Unsurprisingly, it is depleted at the vacancy. An interesting feature is that the electron density does not display the reflection symmetry of the basal plane (e.g. about the red dashed line). This is because the three dimensional crystal breaks this symmetry at the plane at a height $c/2$ above and below this plane. This is emphasized by indicating the atoms on this upper and lower planes with a ‘*’ in the figure. This observation plays a role in binding.

Next we compute the dilute impurity energy of an aluminum solute atom, and this is shown in Fig. 7c. We again observe that this energy depends on the cell size and converges at a few hundred atoms. The electron density on the basal plane is shown Fig. 7d. The electron density is elevated near the Al solute and the distribution is more symmetric.

Defect pairs We now study the interaction between a pair of point defects: vacancy-vacancy, vacancy-solute and solute-solute. Fig. 8a, shows the six possible nearest neighbor positions on a HCP lattice where one defect is placed at the site marked 0 while the other defects are placed in the sites marked 1 through 6 denote the first six nearest neighbor configurations.

Table 2 shows the computed binding energy (cf. (37)) for various defect pairs. Recall that a positive binding energy indicates an energetic propensity for the defects to bind. We observe that all three types of defect pairs – vacancy pair, vacancy-solute and solute-pair – have an energetic preference for binding¹. This is especially strong for vacancy-solute pairs and solute pairs. Further the binding energy increases with decreasing inter-defect distance in all three cases, indicating that the binding is stronger with decreasing distance. These results are consistent with the observation that voids form readily in Mg, Al has a high diffusivity in Mg and that Al-rich precipitates form readily in Mg. Our results agree quantitatively with the only case (first nearest-neighbor vacancy-solute-pair) where we are aware of experimental observation [47] – Table 2.

The table also shows that while our results agree with those of Ponga, Bhattacharya and Ortiz [42] in the case of vacancy pairs, they do not agree with other prior DFT calculations. Specifically, our results predict much stronger binding compared to other results ([62] for vacancy pairs, [1, 53] for vacancy-solute, and

¹We have found no noticeable change in the calculated binding energies when the radius of truncation R_{cut} is increased from 12 to 14 and the polynomial orders K_L and K_C are increased from 80 to 100 and 100 to 120, respectively.

Figure 7: Vacancy (a,b) and an aluminum solute (c,d) in magnesium. The calculated vacancy formation energy (a) and dilute impurity energy (c) for various computational cell size and computed electron density along the basal plane for a vacancy (b) and solute (d). The “*” marks in (b,d) indicate the projected positions of the atoms in the basal plane at a height $c/2$ above and below this plane.

[27, 32] for solute pairs). In fact, previous results [27, 32] for solute pairs also predict negative binding energies in some instances.

The previous calculations [62, 1, 53, 27, 32] all used relatively small computational cells (64-192 atoms) compared to the current results (1151 atoms). Fig. 8 shows the binding energy of various defect pairs computed for various sizes of the computational domain. We see a significant dependence of the binding energy on the size of the computational cell with convergence only at around a 1000 atoms. Further, there is no common trend. For a solute pair, small cells predict a lack of binding while large cells predict binding. The vacancy-vacancy binding energy is large for small computational domains, but decreases with domain size. For vacancy-solute, we see that the binding energy is small for small computational domains, but increases with domain size. This tells us that the disturbance in the electronic and atomistic fields decay slowly. Further, the previous studies use periodic boundary conditions which means that the defects interact if the unit cell is too small. We have separately verified that a periodic calculation with a 96 atom unit cell (using SPARC[18] which uses a similar high order finite difference formulation as the current work)

Figure 8: (a) Nearest neighbor positions for two point defects. The first defect is at the point 0, and the second defect at one of the sites marked 1 through 6 in increasing order distance. Computed (b) Vacancy-vacancy, (c) Vacancy-solute and (d) Solute-solute binding energies as a function of computational cell size.

Table 2: Defect pair binding energies.

NN	Other works (eV)	This work (eV)
Vacancy-vacancy		
1	0.06 ^a , 0.106 ^b	0.127
2	0.07 ^a , 0.096 ^b	0.095
3	-0.01 ^a , 0.057 ^b	0.064
4	0.01 ^a , 0.048 ^b	0.033
5	0.01 ^a , 0.038 ^b	0.033
6	0.01 ^a , 0.034 ^b	0.033
Vacancy-solute		
1	0.06 ^c , 0.29 ^d , 0.316 ^h	0.254
2	0.03 ^e	0.195
3		0.206
4		0.125
5		0.125
6		0.125
Solute-solute		
1	-0.021 ^f , 0.199 ^h	0.238
2	-0.21 ^g , -0.018 ^f	0.223
3	-0.24 ^g , 0.024 ^f	0.219
4	0.014 ^f	0.191
5		0.169
6		0.138

a. DFT, 96 atoms using Troullier-Martins pseudopotential [62]
b. DFT, 3456 atoms local pseudopotential [42]
c. DFT, 192 atoms using Projector Augmented Wave (PAW) pseudopotential [1]
d. Experiment [47]
e. DFT, 64 atoms using ultrasoft pseudopotential [53]
f. DFT, 96 atoms using Projector Augmented Wave (PAW) pseudopotential [27]
g. DFT, 64 atoms using Projector Augmented Wave (PAW) pseudopotential [32]
h. DFT, 1152 atoms using local pseudopotential

reproduces the result of Uesugi *et al.* [62] who used the same pseudopotential. Therefore, the discrepancy between the binding energies from previous works and our calculations with small cell sizes is due to the different choice of boundary conditions.

We turn to the issue of pseudopotentials. We see good agreement between the results of Ponga, Bhattacharya and Ortiz [42] (results marked with superscript *b* in Table 2) who use a local pseudopotential and the current work with non-local pseudopotentials in the case of vacancy-vacancy interaction. In other words, local pseudopotentials seem to suffice in cases like magnesium whose electron distribution is close to that of an electron gas. Table 2 also shows, this is not true when we introduce a solute. We repeated our calculation using a local pseudopotential for the nearest neighbor vacancy-solute and solute-solute interactions (marked with superscript *h* in Table 2); the results differ significantly from those conducted with a non-local pseudopotential. In short, while local pseudopotentials seem to suffice in cases like magnesium whose electron distribution is close to that of an electron gas, they do not provide accurate results in general.

To understand the nature of the binding and explore further the decay, we examine the difference in

(a) vacancy pair, basal plane

(b) vacancy pair, prismatic plane

(c) vacancy-solute, basal plane

(d) vacancy-solute, prismatic plane

(e) solute pair, basal plane

(f) solute pair, prismatic plane

Figure 9: Electron density difference around a defect pair on the basal and prismatic plane. The blue circles mark the position of the vacancies and the red circles that of the solutes. The square box in (a,c,e) marks a 64 atom basal cell.

electron density between that of a defect pair with that of a pair of defects:

$$\Delta\rho(\mathbf{x}) = \rho^0(\mathbf{x}) + \rho^n(\mathbf{x}) - \rho_p(\mathbf{x}) - \rho^{0n}(\mathbf{x}) \quad (38)$$

where ρ^0 is the electron density in a system with a defect placed at site 0, ρ^n is that with a defect at site n , ρ_p is that of the perfect crystal, and ρ^{0n} is that of a defect pair (we need ρ_p to make the system charge neutral). These are plotted for second nearest-neighbor pair (i.e, a basal pair) in Fig. 9. We see that the core of the vacancy pair is severely depleted of electron density while the center of the solute pair has excess electron density. The vacancy-solute complex increases electronic density and depletes it at the distal region of the solute. Further, we see that oscillations spread over several atomic spacings before decaying. Furthermore, we see that the electron density in the case of vacancy and solute pairs are not symmetric with respect to the axis through the defects. This is because of the influence of atoms on the basal plane just above and below the plane of observations. Together, these mean that we need large computational cells for accurately computing the binding energies to accurately account for this decay.

The spectral quadrature formulation has several advantages in simulating the systems discussed in this work. The Chebyshev Filtering method in real-space [19, 18, 68, 12] can achieve smaller walltimes than the spectral quadrature implementation. However, they require considerably more storage, which scale quadratically with the number of atoms. The simulations of metallic systems are restricted to a few thousand atoms with Chebyshev filtering [18, 68, 12]², whereas the spectral quadrature formulation allows simulations of even larger systems. Next the spectral quadrature method allows the application of specified Dirichlet boundary conditions, which is not clearly known in an orbital based solver such as Chebyshev Filtering. In view of this, the wall time for an SCF iteration of 25,000 Mg atoms with a prismatic dislocation loop calculated using spectral quadrature on 2400 MPI ranks is 4 hours, which can significantly decrease with the use of accelerators such as Graphics Processing Unit for performing the matrix-vector products. The 25,000 Mg showed significant charge-sloshing, which can be mitigated by real-space preconditioning schemes [31, 52].

6 Concluding Remarks

In this work, we have developed a computational framework that combines the computational efficiency of LSSGQ and enables the use of non-local pseudopotentials. Further our implementation shows excellent strong and weak scaling. This enables the application of DFT to problems that require a large computational domain such as defects in metals. We have demonstrated the accuracy, efficiency and applicability of the framework by studying defects in magnesium. In particular, we have shown that vacancy pairs, solute-vacancy pairs and solute pairs all bind strongly. Importantly, the electron density decays slowly around these defects thereby necessitating large computational cells.

Our ongoing efforts include the development of an optimized implementation which is amenable to heterogeneous computer architectures. This is motivated from the observation that the calculation of spectral quadrature weights and nodes – the computationally dominant part of the method – is local to each process, and therefore can be accelerated on Graphics Processing Units (GPUs). The implementation of real-space preconditioning schemes [31, 52] can further reduce the number of Self Consistent Field iterations for DFT simulations of large metallic systems. Finally, the incorporation of coarse-graining strategies [42, 43] will further reduce the scaling to sublinear with the number of atoms. Furthermore an accelerated calculation of the electronic predictor in the coarse grained framework will reduce the prefactor associated with scaling. Overall, these advances will render the first principles simulations of the interactions of extended defects such as dislocations tractable.

²The semi finalist of the ACM Gordon Bell prize in 2019 performed a 10,508 Mg atom simulation with a screw dislocation [12].

Acknowledgements

This research was performed when S.G. held a position at the California Institute of Technology. We thank the anonymous reviewers for their valuable comments and suggestions. The work was sponsored by the Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-12-2-0022. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein. The computations presented here were conducted on the Caltech High Performance Cluster partially supported by a grant from the Gordon and Betty Moore Foundation, and the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. Part of this research used resources of the Oak Ridge Leadership Computing Facility, a DOE Office of Science User Facility operated by the Oak Ridge National Laboratory under contract DE-AC05-00OR22725.

References

- [1] R. Agarwal and D. R. Trinkle. Ab initio magnesium-solute transport database using exact diffusion theory. *Acta Materialia*, 150:339 – 350, 2018.
- [2] D. G. Anderson. Iterative procedures for nonlinear integral equations. *Journal of the ACM (JACM)*, 12(4):547–560, 1965.
- [3] S. Balay, J. Brown, , K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, and H. Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.4, Argonne National Laboratory, 2013.
- [4] S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press, 1997.
- [5] A. S. Banerjee, P. Suryanarayana, and J. E. Pask. Periodic pulay method for robust and efficient convergence acceleration of self-consistent field iterations. *Chemical Physics Letters*, 647:31–35, 2016.
- [6] A. D. Becke. Perspective: Fifty years of density-functional theory in chemical physics. *The Journal of chemical physics*, 140(18):18A301, 2014.
- [7] D. Bowler and T. Miyazaki. $O(N)$ methods in electronic structure calculations. *Reports on Progress in Physics*, 75(3):036503, 2012.
- [8] R. P. Brent. An algorithm with guaranteed convergence for finding a zero of a function. *The Computer Journal*, 14(4):422–425, 01 1971.
- [9] K. Burke. Perspective on density functional theory. *The Journal of chemical physics*, 136(15):150901, 2012.
- [10] D. M. Ceperley and B. J. Alder. Ground state of the electron gas by a stochastic method. *Physical review letters*, 45(7):566, 1980.
- [11] M. Chou and M. L. Cohen. Ab initio study of the structural properties of magnesium. *Solid State Communications*, 57(10):785 – 788, 1986.

- [12] S. Das, P. Motamarri, V. Gavini, B. Turcksin, Y. W. Li, and B. Leback. Fast, scalable and accurate finite-element based ab initio calculations using mixed precision computing: 46 pflops simulation of a metallic dislocation system. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–11, 2019.
- [13] Y.-H. Duan, Y. Sun, M.-J. Peng, and Z.-Z. Guo. First principles investigation of the binary intermetallics in pb–mg–al alloy: stability, elastic properties and electronic structure. *Solid State Sciences*, 13(2):455–459, 2011.
- [14] V. Gavini. Role of the defect core in energetics of vacancies. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 465(2110):3239–3266, 2009.
- [15] V. Gavini, K. Bhattacharya, and M. Ortiz. Vacancy clustering and prismatic dislocation loop formation in aluminum. *Phys. Rev. B*, 76:180101, Nov 2007.
- [16] S. Ghosh and K. Bhattacharya. Influence of thermomechanical loads on the energetics of precipitation in magnesium aluminum alloys. *Acta Materialia*, 193:28 – 39, 2020.
- [17] S. Ghosh and P. Suryanarayana. Higher-order finite-difference formulation of periodic orbital-free density functional theory. *Journal of Computational Physics*, 307:634–652, 2016.
- [18] S. Ghosh and P. Suryanarayana. Sparc: Accurate and efficient finite-difference formulation and parallel implementation of density functional theory: Extended systems. *Computer Physics Communications*, 216:109–125, 2017.
- [19] S. Ghosh and P. Suryanarayana. Sparc: Accurate and efficient finite-difference formulation and parallel implementation of density functional theory: Isolated clusters. *Computer Physics Communications*, 212:189–204, 2017.
- [20] S. Ghosh and P. Suryanarayana. Electronic structure study regarding the influence of macroscopic deformations on the vacancy formation energy in aluminum. *Mechanics Research Communications*, 99:58–63, 2019.
- [21] W. Gropp, W. D. Gropp, E. Lusk, A. Skjellum, and A. D. F. E. E. Lusk. *Using MPI: portable parallel programming with the message-passing interface*, volume 1. MIT press, 1999.
- [22] W. Gropp, T. Hoefler, R. Thakur, and E. Lusk. *Using advanced MPI: Modern features of the message-passing interface*. MIT Press, 2014.
- [23] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Physical Review*, 136(3B):B864–B871, 1964.
- [24] C. Janot, D. Malléjac, and B. George. Vacancy-formation energy and entropy in magnesium single crystals. *Physical Review B*, 2(8):3088, 1970.
- [25] K. U. Kainer and B. L. Mordike. *Magnesium alloys and their applications*. Wiley-VCH Weinheim, Germany:, 2000.
- [26] G. Kerker. Efficient iteration scheme for self-consistent pseudopotential calculations. *Physical Review B*, 23(6):3082, 1981.
- [27] H. Kimizuka and S. Ogata. Predicting atomic arrangement of solute clusters in dilute mg alloys. *Materials Research Letters*, 1(4):213–219, 2013.

- [28] C. Kittel et al. *Introduction to solid state physics*, volume 8. Wiley New York, 1976.
- [29] W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Physical Review*, 140(4A):A1133–A1138, 1965.
- [30] W. Koster and H. Franz. Poisson’s ratio for metals and alloys. *Metallurgical Reviews*, 6(1):1–56, 1961.
- [31] S. Kumar, Q. Xu, and P. Suryanarayana. On preconditioning the self-consistent field iteration in real-space density functional theory. *Chemical Physics Letters*, page 136983, 2019.
- [32] G. Liu, J. Zhang, and Y. Dou. First-principles study of solute–solute binding in magnesium alloys. *Computational Materials Science*, 103:97–104, 2015.
- [33] C. Mairy, J. Hillairet, and D. Schumacher. Energie de formation et concentration d’équilibre des lacunes dans le magnésium. *Acta metallurgica*, 15(7):1258–1261, 1967.
- [34] M. A. Meyers and C. T. Aimone. Dynamic fracture (spalling) of metals. *Progress in Materials Science*, 28(1):1–96, 1983.
- [35] E. Nembach. *Particle strengthening of metals and alloys*. Wiley, New York, 1997.
- [36] J. F. Nie. Effects of precipitate shape and orientation on dispersion strengthening in magnesium alloys. *Scripta Materialia*, 48(8):1009–1015, 2003.
- [37] J.-F. Nie. Precipitation and hardening in magnesium alloys. *Metallurgical and Materials Transactions A*, 43(11):3891–3939, Nov 2012.
- [38] R. G. Parr and W. Yang. Density functional theory of atoms and molecules. *Oxford University Press*, 1:989, 1989.
- [39] J. Pask and P. Sterne. Real-space formulation of the electrostatic potential and total energy of solids. *Physical Review B*, 71(11):113101, 2005.
- [40] J. P. Perdew and Y. Wang. Accurate and simple analytic representation of the electron-gas correlation energy. *Phys. Rev. B*, 45:13244–13249, Jun 1992.
- [41] I. J. Polmear. Magnesium alloys and applications. *Materials Science and Technology*, 10(1):1–16, 1994.
- [42] M. Ponga, K. Bhattacharya, and M. Ortiz. A sublinear-scaling approach to density-functional-theory analysis of crystal defects. *Journal of the Mechanics and Physics of Solids*, 95:530 – 556, 2016.
- [43] M. Ponga, K. Bhattacharya, and M. Ortiz. Large scale ab-initio simulations of dislocations. *Journal of Computational Physics*, page 109249, 2020.
- [44] P. P. Pratapa and P. Suryanarayana. Restarted pulay mixing for efficient and robust acceleration of fixed-point iterations. *Chemical Physics Letters*, 635:69–74, 2015.
- [45] P. P. Pratapa, P. Suryanarayana, and J. E. Pask. Spectral quadrature method for accurate o (n) electronic structure calculations of metals and insulators. *Computer Physics Communications*, 2016.
- [46] P. Pulay. Convergence acceleration of iterative sequences. the case of scf iteration. *Chemical Physics Letters*, 73(2):393–398, 1980.

- [47] A. R. Rao and C. Suryanarayana. Solute-vacancy binding energies in magnesium alloys. *physica status solidi (a)*, 45(2):K131–K133, 1978.
- [48] W. Rudin. *Functional analysis*. McGraw-Hill, New York, 1973.
- [49] Y. Saad and M. H. Schultz. Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on scientific and statistical computing*, 7(3):856–869, 1986.
- [50] A. Sharma, S. Hamel, M. Bethkenhagen, J. E. Pask, and P. Suryanarayana. Real-space formulation of the stress tensor for o (n) density functional theory: Application to high temperature calculations. *The Journal of Chemical Physics*, 153(3):034112, 2020.
- [51] J. R. Shewchuk. An introduction to the conjugate gradient method without the agonizing pain. 1994.
- [52] Y. Shiihara, O. Kuwazuru, and N. Yoshikawa. Real-space kerker method for self-consistent calculation using non-orthogonal basis functions. *Modelling and Simulation in Materials Science and Engineering*, 16(3):035004, 2008.
- [53] D. Shin and C. Wolverton. First-principles study of solute–vacancy binding in magnesium. *Acta Materialia*, 58(2):531–540, 2010.
- [54] P. Suryanarayana. On spectral quadrature for linear-scaling density functional theory. *Chemical Physics Letters*, 584:182–187, 2013.
- [55] P. Suryanarayana. On nearsightedness in metallic systems for o (n) density functional theory calculations: A case study on aluminum. *Chemical Physics Letters*, 679:146–151, 2017.
- [56] P. Suryanarayana, K. Bhattacharya, and M. Ortiz. Coarse-graining kohn-sham density functional theory. *Journal of the Mechanics and Physics of Solids*, 61(1):38 – 60, 2013.
- [57] P. Suryanarayana, V. Gavini, T. Blesgen, K. Bhattacharya, and M. Ortiz. Non-periodic finite-element formulation of kohn-sham density functional theory. *Journal of the Mechanics and Physics of Solids*, 58(2):256 – 280, 2010.
- [58] P. Suryanarayana, P. P. Pratapa, A. Sharma, and J. E. Pask. Sqdft: Spectral quadrature method for large-scale parallel o (n) kohn–sham calculations at high temperature. *Computer Physics Communications*, 224:288–298, 2018.
- [59] A. Szabo and N. S. Ostlund. *Modern quantum chemistry: introduction to advanced electronic structure theory*. Courier Corporation, 2012.
- [60] N. Troullier and J. L. Martins. Efficient pseudopotentials for plane-wave calculations. *Physical review B*, 43(3):1993, 1991.
- [61] P. Tzanetakis, J. Hillairet, and G. Revel. The formation energy of vacancies in aluminium and magnesium. *physica status solidi (b)*, 75(2):433–439, 1976.
- [62] T. Uesugi, M. Kohyama, and K. Higashi. Ab initio study on divacancy binding energies in aluminum and magnesium. *Physical Review B*, 68(18):184103, 2003.
- [63] A. Vehanen and K. Rytölä. *Proceedings of the International School of Physics*. North-Holland, Utrecht, 1981.

- [64] X.-C. Wang, T. Blesgen, K. Bhattacharya, and M. Ortiz. A Variational Framework for Spectral Approximations of Kohn–Sham Density Functional Theory. *Arch. Rat. Mech. Anal.*, 221:1035–1075, 2016.
- [65] Q. Xu, P. Suryanarayana, and J. E. Pask. Discrete discontinuous basis projection method for large-scale electronic structure calculations. *The Journal of chemical physics*, 149(9):094104, 2018.
- [66] M.-X. Zhang and P. M. Kelly. Edge-to-edge matching and its applications: part ii. application to mg–al, mg–y and mg–mn alloys. *Acta materialia*, 53(4):1085–1096, 2005.
- [67] S. Zhang, A. Lazicki, B. Militzer, L. H. Yang, K. Caspersen, J. A. Gaffney, M. W. Däne, J. E. Pask, W. R. Johnson, A. Sharma, P. Suryanarayana, D. D. Johnson, A. V. Smirnov, P. A. Sterne, D. Erskine, R. A. London, F. Coppari, D. Swift, J. Nilsen, A. J. Nelson, and H. D. Whitley. Equation of state of boron nitride combining computation, modeling, and experiment. *Phys. Rev. B*, 99:165103, Apr 2019.
- [68] Y. Zhou, Y. Saad, M. L. Tiago, and J. R. Chelikowsky. Parallel self-consistent-field calculations via chebyshev-filtered subspace acceleration. *Physical Review E*, 74(6):066704, 2006.

Spectral quadrature for the first principles study of crystal defects: Application to magnesium
Ghosh and Bhattacharya

Declaration of interests

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

Ghosh: Conceptualization, Methodology, Software, Writing. Bhattacharya: Conceptualization, Methodology, Software, Writing.