AFRL-AFOSR-VA-TR-2021-0101

Tensor-structured techniques for large-scale electronic-structure calculations

Gavini, Vikram
REGENTS OF THE UNIVERSITY OF MICHIGAN
503 THOMPSON ST
ANN ARBOR, MI, 48109
USA

08/17/2021
Final Technical Report

Air Force Research Laboratory
Air Force Office of Scientific Research
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
17-08-2021	Final	01 Feb 2020 - 31 Jul 2020

4. TITLE AND SUBTITLE
Tensor-structured techniques for large-scale electronic-structure calculations

5a. CONTRACT NUMBER

5b. GRANT NUMBER
FA9550-17-1-0172

5c. PROGRAM ELEMENT NUMBER
61102F

6. AUTHOR(S)
Vikram Gavini

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
REGENTS OF THE UNIVERSITY OF MICHIGAN
503 THOMPSON ST
ANN ARBOR, MI 48109
USA

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AF Office of Scientific Research
875 N. Randolph St. Room 3112
Arlington, VA 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/AFOSR RTA2

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
AFRL-AFOSR-VA-TR-2021-0101

12. DISTRIBUTION/AVAILABILITY STATEMENT
A Distribution Unlimited: PB Public Release

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This report summarizes the research objectives achieved in this project during the period 02-01-2017 to 07-31-2020. Computational techniques based on low-rank tensor decomposition and tensor representation have been developed that enable large-scale, reduced-order scaling electronic structure calculations based on Kohn-Sham density functional theory. In particular, the various components of the developed techniques include (i) develop an additive separable approximation to the Kohn-Sham Hamiltonian and use this to generate a reduced-order, and yet a systematically convergent tensor-structured basis adapted to the Kohn-Sham Hamiltonian to represent the electronic structure; (ii) develop localization techniques to transform the basis into a localized basis while preserving the tensor structure and accuracy of representation of the electronic structure; (iii) develop ecient algorithms to project the Kohn-Sham problem into the localized tensor-structured basis and solve the resulting discrete Kohn-sham eigenvalue problem;

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			FARIBA FAHROO
U	U	U	UU	9	**19b. TELEPHONE NUMBER** *(Include area code)* 426-8429

Final Report

Tensor-Structured Techniques for Large-Scale Electronic-Structure Calculations

Principal Investigator:
Vikram Gavini
Department of Mechanical Engineering
University of Michigan, Ann Arbor

Program Manager:
Dr. Fariba Fahroo
Computational Mathematics
Air Force Office of Scientific Research

Abstract

This report summarizes the research objectives achieved in this project during the period 02-01-2017 to 07-31-2020. Computational techniques based on low-rank tensor decomposition and tensor representation have been developed that enable large-scale, reduced-order scaling electronic structure calculations based on Kohn-Sham density functional theory. In particular, the various components of the developed techniques include (i) develop an additive separable approximation to the Kohn-Sham Hamiltonian and use this to generate a reduced-order, and yet a systematically convergent tensor-structured basis adapted to the Kohn-Sham Hamiltonian to represent the electronic structure; (ii) develop localization techniques to transform the basis into a localized basis while preserving the tensor structure and accuracy of representation of the electronic structure; (iii) develop efficient algorithms to project the Kohn-Sham problem into the localized tensor-structured basis and solve the resulting discrete Kohn-sham eigenvalue problem; (iv) develop a scalable parallel implementation of DFT code based on the aforementioned ideas, and developing a GPU implementation of the code. The benchmark studies suggest that the proposed approach is systematically convergent providing exponential convergence in the ground-state energy with the rank of the tensor-structured basis. Further, due to the slow growth of the rank with system-size, the scaling with system size was subquadratic for both systems with and without a gap. This has enabled systematically convergent large-scale DFT calculations that substantially outperform plane-wave DFT codes beyond system-sizes containing 5,000 electrons. In addition, the computational framework developed in his work was instrumental in tackling a long-standing numerical challenge of accurately solving the inverse DFT problem for polyatomic systems.

Summary of Research Objectives Achieved

In this project we have conducted a program of research to develop computational techniques based on low-rank tensor decomposition and tensor representation that have enabled systematically convergent, large-scale, reduced-order scaling density functional theory calculations. In addition, parts of the computational framework built in this work was instrumental in providing an accurate solution to the inverse DFT problem. A comprehensive summary of the important research objectives achieved in this project is presented below, and has appeared in the following articles (Motamarri

et al., 2016; Kanungo et al., 2019) and under review in Lin et al. (2020a). Another manuscript reporting the GPU porting of the numerical implementation is under preparation (Lin et al., 2020b).

Tucker tensor approach to accurate, efficient and reduced-order scaling large-scale DFT calculations

The ground state energy in Kohn-Sham DFT (spin independent formulation) of a materials system with N_a atoms and N_e electrons is computed by solving a non-interacting single-particle Schrödinger equation in a mean field determined by the effective potential $V_{\text{eff}}(\rho; \mathbf{R}_{nu})$ (Hohenberg & Kohn, 1964; Kohn & Sham, 1965):

$$\left(-\frac{1}{2}\nabla^2 + V_{\text{eff}}(\rho; \mathbf{R}_{nu})\right)\Psi_i = \epsilon_i \Psi_i, \quad i \in \{1, ..., N\}. \tag{1}$$

Equation [1] represents a non-linear eigenvalue problem with $\mathcal{H} := -\frac{1}{2}\nabla^2 + V_{\text{eff}}$ being the Kohn-Sham Hamiltonian, and ϵ_i denoting the ith Kohn-Sham eigenvalue and Ψ_i denoting the corresponding Kohn-Sham orbital (eigenvector). The electron density $\rho(\mathbf{x})$ is computed in terms of the Kohn-Sham eigenstates as $\rho(\mathbf{x}) = 2\sum_{i=1}^{N} f(\epsilon_i; \mu)\left|\Psi_i(\mathbf{x})\right|^2$, where $f(\epsilon_i; \mu)$ denotes the orbital occupancy factor given by the Fermi-Dirac distribution. The effective potential $V_{\text{eff}}(\rho; \mathbf{R}_{nu})$, parametrized by the coordinates of nuclei $\mathbf{R}_{nu} = \{\mathbf{R}_1, \mathbf{R}_2, \ldots, \mathbf{R}_{N_a}\}$, is given by

$$V_{\text{eff}}(\rho; \mathbf{R}_{nu}) = \frac{\delta E_{\text{H}}}{\delta \rho} + \frac{\delta E_{\text{XC}}}{\delta \rho} + V_{\text{ext}}(\mathbf{x}; \mathbf{R}_{nu}). \tag{2}$$

E_{H} is the Hartree energy, which represents the Coulomb electrostatic interaction defined as (for a non-periodic setting)

$$E_{\text{H}} = \frac{1}{2}\int_{\mathbb{R}^3}\int_{\mathbb{R}^3} \frac{\rho(\mathbf{x})\rho(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} d\mathbf{x} d\mathbf{x}' = \frac{1}{2}\int_{\mathbb{R}^3} \rho(\mathbf{x}) V_{\text{H}}(\rho) d\mathbf{x}, \tag{3}$$

where $V_{\text{H}}(\rho) := \frac{\delta E_{\text{H}}}{\delta \rho} = \int_{\mathbb{R}^3} \frac{\rho(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} d\mathbf{x}'$ is the Hartree potential. E_{XC} is the exchange-correlation energy, with the exchange-correlation potential $V_{\text{XC}}(\rho) = \frac{\delta E_{\text{XC}}}{\delta \rho}$. $V_{\text{ext}}(\mathbf{x}; \mathbf{R}_{nu})$ is the effective external potential of the ions, and in the case of pseudopotential calculations accounts for the combined effect of nuclei and the core electrons.

Equation [1] represents a non-linear eigenvalue problem, as the Kohn-Sham Hamiltonian depends on ρ, which in turn depends on the eigenstates. Thus, Equation [1] is solved self-consistently via a self-consistent field (SCF) iteration. Upon achieving self-consistency, the ground state energy of the materials system can be obtained as

$$\begin{aligned} E_{\text{tot}} = 2\sum_{i=1}^{N} f(\epsilon_i; \mu)\epsilon_i + E_{\text{XC}} - \int_{\mathbb{R}^3} \rho V_{\text{XC}}(\rho) d\mathbf{x} \\ - \frac{1}{2}\int_{\mathbb{R}^3} \rho V_{\text{H}}(\rho) d\mathbf{x} + \sum_{I=1}^{N_a}\sum_{J>I}^{N_a} \frac{Z_I Z_J}{|\mathbf{R}_I - \mathbf{R}_J|}, \end{aligned} \tag{4}$$

where Z_I denotes the effective nuclear charge of the I-th nucleus.

Construction of reduced order tensor-structured basis

For a domain Ω spanned by tensor product of 1-D domains $\omega_{k=1,2,3}$, we seek to construct an additive separable approximation to the Kohn-Sham Hamiltonian ($\mathcal{H}_1(x_1) + \mathcal{H}_2(x_2) + \mathcal{H}_3(x_3) \approx \mathcal{H}(\mathbf{x})$) that retains some features of the Hamiltonian, and thus presents a useful operator to generate reduced-order basis functions. The eigenfunctions of the additive separable approximation to the Hamiltonian, which constitute a Tucker tensor basis (Kolda & Bader, 2009) formed from the 1-D eigenfunctions of the separable parts of the Hamiltonian ($\mathcal{H}_k$, $k = 1, 2, 3$), presents a systematically improvable and efficient reduced-order basis for solving the Kohn-Sham eigenvalue problem. To this end, in order to construct the additive separable approximation to the Kohn-Sham Hamiltonian, we consider the rank-1 approximation of the eigenfunction corresponding to the lowest eigenvalue as $\Psi'(\mathbf{x}) = \psi_1(x_1)\psi_2(x_2)\psi_3(x_3)$. The problem of computing the smallest eigenvalue of the Kohn-Sham Hamiltonian using the rank-1 approximation is given by the variational problem

$$\min_{\psi_k} L(\Psi') \qquad \text{subject to: } \langle \Psi' | \Psi' \rangle = 1 \tag{5}$$

with the Lagrangian

$$\begin{aligned} L(\Psi') &= \left\langle \Psi' \right| -\frac{1}{2}\nabla^2 + V_{\text{eff}}(\rho; \mathbf{R}_{nu}) \left| \Psi' \right\rangle \\ &= \int_\Omega \Big[\sum_{\ell=1}^{3} \frac{1}{2} \Big| \frac{d\psi_\ell(x_\ell)}{dx_\ell} \Big|^2 \prod_{m\neq\ell}^{3} \psi_m^2(x_m) + V_{\text{eff}}^{loc}(\mathbf{x}) \prod_{\ell=1}^{3} \psi_\ell^2(x_\ell) + \prod_{\ell=1}^{3} \psi_\ell(x_\ell) V_{\text{ext}}^{nl}(\mathbf{x}) \prod_{\ell=1}^{3} \psi_\ell(x_\ell) \Big] d\mathbf{x}\,, \end{aligned} \tag{6}$$

where V_{eff}^{loc} and V_{ext}^{nl} are the local and the non-local contributions to the effective potential, respectively. The variations of the functional with respect to ψ_1, ψ_2 and ψ_3, provide three simultaneous 1-D eigenvalue problems

$$\begin{aligned} &\mathcal{H}_k \psi_k = -(\lambda + a_k)\psi_k \\ &\mathcal{H}_k = -\frac{1}{2}\frac{d^2}{dx_k^2} + v_k^{loc}(x_k; \psi_{l\neq k}) + v_k^{nl}(x_k; \psi_{l\neq k})\,, \end{aligned} \tag{7}$$

where $v_k^{loc}(x_k)$, $v_k^{nl}(x_k)$, and a_k are quantities depending on ψ_1, ψ_2, ψ_3, and λ is the Lagrange multiplier corresponding to the orthonormality constraint in equation [5]. $v_k^{loc}(x_k)$ and $v_k^{nl}(x_k)$ are 1-D potentials corresponding to the local and non-local contribution. These potentials are parametrized by ψ_l ($l \neq k$) in the other two dimensions. Equation [7] represents simultaneous 1-D eigenvalue problems, which are solved self-consistently. Upon achieving self-consistency, the 1-D Hamiltonians ($\mathcal{H}_k$) we obtain represent the additive separable approximation of the Kohn-Sham Hamiltonian that we seek. The eigenfunctions of this additive separable approximation to the Hamiltonian, which can be obtained as the tensor product of the 1-D eigenfunctions of $\mathcal{H}_k$ ($k = 1, 2, 3$), constitute a complete basis, thus provides systematic convergence as will be demonstrated.

Construction of localized Tucker tensor basis

The tensor-structured basis computed using the 1-D eigenfunctions of $\mathcal{H}_k$ represents an efficient basis. However, the global nature of the basis limits the computational efficiency and scaling (with system size) of solution to the Kohn-Sham equations. To this end, we use an L_1 localization approach (Ozlins et al., 2013) to construct a spatially localized tensor-structured basis that is a close approximation to the original tensor-structured basis. The localized basis is obtained by

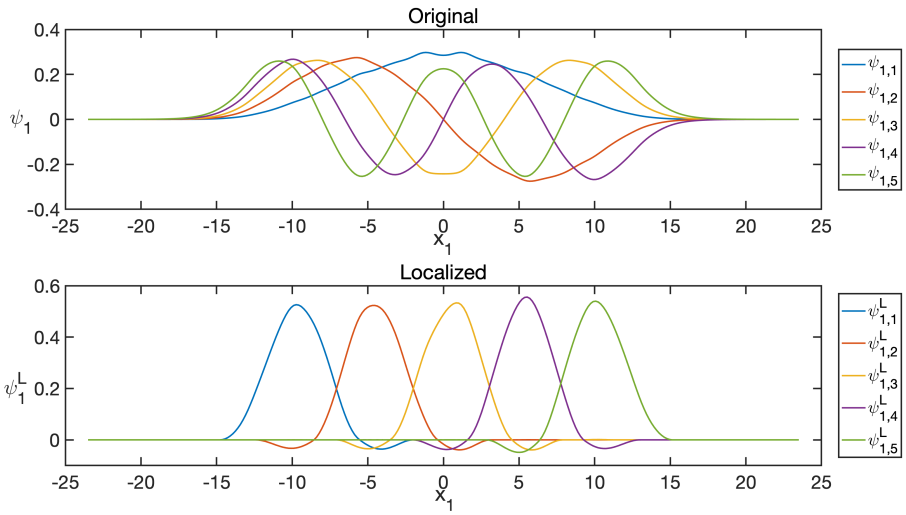

Figure 1: 1-D functions in x_1 direction constructed from the additive separable approximation of the Kohn-Sham Hamiltonian for Al_{147} nano-partice. Top: Lowest five eigenfunctions of $\mathcal{H}_1$. Bottom: The corresponding L_1 localized 1-D functions.

solving the variational problem (for $k = 1, 2, 3$)

$$\min_{\mathbf{\Psi}'_k \in \mathbb{R}^{n \times N_k}} \frac{1}{\mu} \left| \mathbf{\Psi}'_k \right| + \mathrm{Tr}({\mathbf{\Psi}'_k}^T \mathbf{H}_k \mathbf{\Psi}'_k) \quad \text{s.t. } {\mathbf{\Psi}'_k}^T \mathbf{\Psi}'_k = I, \tag{8}$$

where $\mathbf{H}_k$ is matrix representation of $\mathcal{H}_k$ in a suitable orthogonal basis with dimension n, $\mathbf{\Psi}'_k$ denotes the representation of N_k trial localized functions in the chosen basis, and μ is a parameter controlling the trade-off between the representability of the original eigensubspace and the locality of the 1-D functions, with $|\cdot|$ denoting the L_1 norm of the matrix. The minimizer of this variational problem, denoted as $\mathbf{\Psi}^L_k$, provides localized functions whose span closely approximates the eigensubspace of the lowest N_k eigenfunctions of $\mathcal{H}_k$, as will be demonstrated.

In order to demonstrate the locality and the approximation properties of the localized Tucker tensor basis constructed from the 1-D localized functions obtained by solving the variational problem in equation (8), we consider the benchmark problem involving an Aluminum nanoparticle, Al_{147}. We compute the additive separable approximation of the Kohn-Sham Hamiltonian for this nano-particle, and, using this, compute the lowest 70 eigenstates of $\mathcal{H}_k$. We subsequently use the L_1 localization approach to compute the localized functions that are a close approximation to the eigensubspace. Figure 1 shows the lowest 5 eigenfunctions of $\mathcal{H}_1$ (one of the 1-D separable Hamiltonian) (top) and the corresponding 1-D localized functions (bottom). As is evident the localization approach provides the desired locality in the 1-D functions. In order to assess the eigensubspace representability of the resulting localized 1-D functions, we consider the first 70 eigenstates of $\mathcal{H}_1$ and the eigenvalues of the matrix $K_{ij} = \left\langle \psi^L_{1,i} \middle| \mathcal{H}_1 \middle| \psi^L_{1,j} \right\rangle$, $1 \leq i, j, \leq 70$. Figure 2 shows the eigenvalues of $\mathcal{H}_1$ and the eigenvalues of K_{ij}. As is evident, the localized Tucker tensor basis presents a close approximation of the eigensubspace of the additive separable decomposition of the Kohn-Sham Hamiltonian.

Solution of the discrete Kohn-Sham problem in the localized Tucker tensor basis

The 1-D localized functions whose span is a close approximation to the subspace spanned by the 1-D eigenfunctions of $\mathcal{H}_k$ are subsequently used to construct the 3-D Tucker tensor basis. Denoting

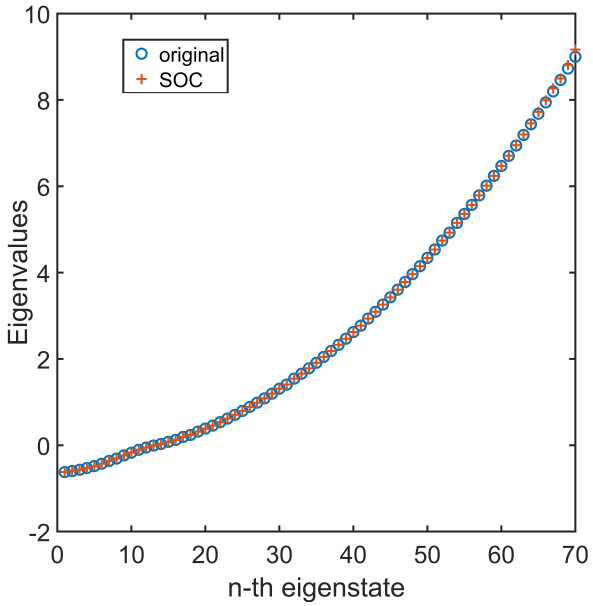

Figure 2: Comparison of the eigenvalues of the 1-D separable Hamiltonian in x_1 direction of Al_{147} nano-particle (marked with blue circle) with the eigenvalues of $K_{ij} = \left\langle \psi^L_{1,i} \right| \mathcal{H}_k \left| \psi^L_{1,j} \right\rangle$ (marked with red cross).

the 1-D localized functions as $\psi^L_{1,i_1}(x_1)$, $\psi^L_{2,i_2}(x_2)$, $\psi^L_{3,i_3}(x_3)$, the 3-D localized tensor-structured basis functions T^L_I are given by

$$T^L_I = \psi^L_{1,i_1}(x_1)\psi^L_{2,i_2}(x_2)\psi^L_{3,i_3}(x_3), \tag{9}$$

where $1 \leq i_d \leq R_d$ and I is the composite index $I = (i_1, i_2, i_3)_{1 \leq i_d \leq R_d}$. The rank of the Tucker tensor basis is given by (R_1, R_2, R_3) which denotes the number of localized 1-D functions in each direction. The Kohn-Sham Hamiltonian projected onto the subspace spanned by the 3-D tensor-structured localized basis functions, is given by

$$\tilde{H}^L_{I,J} = \left\langle T^L_I \right| -\frac{1}{2}\nabla^2 + \tilde{V}_{\text{eff}}(\rho; \mathbf{R}) \left| T^L_J \right\rangle . \tag{10}$$

We employ a low-rank Tucker approximation (Kolda & Bader, 2009) of the effective potential V_{eff} to efficiently evaluate the Hamiltonian matrix elements. We also introduce a truncation tolerance below which the hamiltonian elements are set to zero to obtain a sparse matrix. The truncation tolerance is chosen such that the error is significantly lower than the desired chemical accuracy. The discrete Kohn-Sham eigenvalue problem in the localized Tucker tensor basis is given by

$$\mathbf{H}^L \mathbf{\Psi}_i = \epsilon_i \mathbf{\Psi}_i \,, \quad i \in \{1, \ldots, N\} \tag{11}$$

where $\mathbf{H}^L$ denotes the truncated sparse Kohn-Sham Hamiltonian matrix.

We use the Chebyshev filtering based subspace iteration (ChFSI) to efficiently solve the Kohn-Sham equations. We refer to Motamarri et al. (2013); Motamarri & Gavini (2014) for the specific details of our implementation. In brief, the ChFSI technique represents a modified power iteration in every SCF iteration that exploits the rapid growth of Chebyshev polynomials in $(-\infty, -1)$ to magnify the relevant eigenspectrum (Zhou et al., 2006), thus providing a good subspace for the solution of the Kohn-Sham eigenvalue problem. Upon computing the Chebyshev filtered subspace,

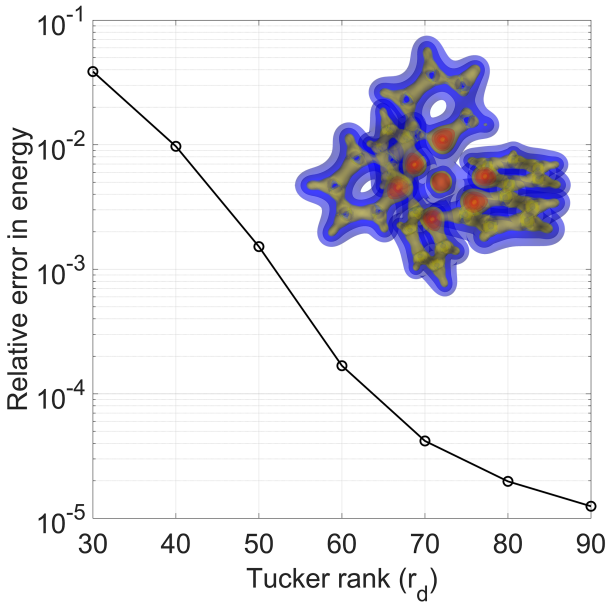

Figure 3: Convergence with respect to the Tucker rank for tris (bipyridine) ruthenium. The electron density isocontours are provided in the inset.

we orthogonalize the Chebyshev filtered vectors and project the discrete Hamiltonian onto the subspace and compute the eigenvalue and eigenvectors of the subspace projected Hamiltonian. The ChFSI method has been demonstrated to be efficient with good parallel scalability for real-space implementations of DFT (Motamarri et al., 2020; Das et al., 2019).

We now present benchmark studies to demonstrate the systematic convergence and accuracy afforded by the proposed tensor-structured approach for solving the Kohn-Sham equations. In order to demonstrate the accuracy and systematic convergence, we consider tris (bipyridine) ruthenium, a transition metal complex, which has no tensor structure symmetry and presents itself as a stringent test of the proposed approach. The ground-state energy for these molecules is computed for various Tucker tensor ranks R ($R_1 = R_2 = R_3 = R$), and the error is measured with respect to a well-converged Quantum Espresso result. The converged Quantum Espresso ground-state energies for tris (bipyridine) ruthenium is taken to be -118.2128 eV/atom ($E_{cut} = 65$ Ha). Figure 3 shows the relative error in the ground-state energy for the various ranks of the Tucker tensor basis. It is evident from these results that the Tucker tensor basis constructed using our approach provides an exponential convergence in the ground-state energy with increasing Tucker rank. Further, the proposed tensor-structured approach provides systematic convergence with high accuracy and is capable of handling generic materials systems, including those involving transition metals.

In order to assess the computational efficiency and scaling with system size afforded by the proposed approach, we considered two classes of benchmark systems: (i) Aluminum nano-particles of various sizes ranging from 13 atoms to 6,525 atoms; (ii) Silicon quantum dots with system sizes ranging from 26 atoms to 7,355 atoms. These benchmark systems constitute materials systems with and without a gap, thus allowing us to assess the system-size scaling for both classes of materials. We also compared the computational efficiency of the proposed approach with widely used plane-wave DFT calculations (Quantum Espresso code). For the sake of estimating the computational efficiency, the energy cutoff for Quantum Espresso and the Tucker rank is chosen such that the ground-state energy is converged to within 10 meV/atom measured with respect to a highly converged reference calculations.

Nano-particle	Rank	Tucker (node-hrs)	QE (node-hrs)
Al_{13}	40	0.00067	0.00022
Al_{147}	70	0.069	0.028
Al_{561}	85	0.69	1.24
Al_{2057}	120	7.96	66.68
Al_{6525}	150	55.08	-

Table 1: Computational time (in node-hrs) per SCF iteration for various system sizes of aluminum nano-particles using the proposed tensor-structured approach (Tucker) and Quantum Espresso (QE).

Si quantum dot	Rank	Tucker (node-hrs)	QE (node-hrs)
$Si_{10}H_{16}$	40	0.0065	0.00014
$Si_{220}H_{144}$	80	0.094	0.096
$Si_{525}H_{276}$	90	0.96	1.12
$Si_{1214}H_{504}$	100	3.85	20.01
$Si_{6047}H_{1308}$	140	67.49	-

Table 2: Computational time (in node-hrs) per SCF iteration for various system sizes of silicon quantum dots using the proposed tensor-structured approach (Tucker) and Quantum Espresso (QE).

The computational times per SCF iteration in node-hrs using the proposed tensor-structured approach and Quantum Espresso are provided in Table 1 for the various aluminum nano-particles with icosahedral symmetry considered in this work. As is evident, the proposed approach starts to substantially outperform for larger system sizes, with the tensor-structured approach being 8-fold more efficient for Al_{2057}. From these computational times, we determine the scaling with system size to be $\mathcal{O}(N_e^{1.78})$. It is notable that we obtain subquadratic scaling with system size for metallic systems involving many thousands of atoms. Table 2 compares the computational performance of the proposed tensor-structured approach with Quantum Espresso for a wide range of silicon quantum dots passivated with hydrogen. As in the case of aluminum nano-particles, the proposed approach starts competing with Quantum Espresso beyond a few hundred atoms, and substantially outperforms for larger systems. The scaling with system size for the tensor-structured algorithm, for a range of system-sizes with the largest containing 7,355 atoms, is estimated to be $\mathcal{O}(N_e^{1.8})$.

Parallel scaling and GPU porting

The proposed algorithms have been implemented in parallel and the scaling has been demonstrated on up to 1000 MPI tasks. The proposed tensor-structured approach is amenable to GPU acceleration and we are currently pursuing the GPU porting of this code, which is almost complete. Preliminary studies suggest a $\sim 5\times$ speed-up compared to CPUs from a node-hrs perspective. A manuscript is currently under preparation, which reports the performance of the proposed algorithm on GPUs.

1 Publications supported by this grant

1 Motamarri, P., Blesgen, T., Gavini, V., Tucker-tensor algorithm for large-scale Kohn-Sham density functional theory calculations, *Phys. Rev. B* **93**, 125104 (2016).

2 Kanungo, B., Zimmerman, P., Gavini, V., Exact exchange-correlation potentials from ground-state electron densities, Nature Communications 10, 4497 (2019).

3 Lin, C.-C., Motamarri, P., Gavini, V., Tensor-structured algorithm for redued-order scaling large-scale Kohn-Sham DFT calculations, submitted to npj Computational Materials (2020).

4 Lin, C.-C., Gavini, V., GPU accelerated Tucker DFT code for fast DFT calculations, under preparation (2020).

References

Das, S., Motamarri, P., Gavini, V., Turcksin, B., Li, Y.-W., Leback, B., Fast, scalable and accurate finite-element based ab initio Calculations using mixed precision computing: 46 PFLOPS simulation of a metallic dislocation system, Proceedings of SC19, The International Conference for High Performance Computing, Networking, Storage, and Analysis (2019).

Hohenberg, P., Kohn, W., Inhomogeneous electron gas. *Phys. Rev.* **136**, B864?B871 (1964).

Kanungo, B., Zimmerman, P., Gavini, V., Exact exchange-correlation potentials from ground-state electron densities, *Nature Communications* **10**, 4497 (2019).

Kolda, T. G., Bader, B. W., Tensor Decompositions and Applications. *SIAM Rev.* **51**, 455-500, (2009).

Lin, C.-C., Motamarri, P., Gavini, V., Tensor-structured algorithm for reduced-order scaling large-scale Kohn-Sham DFT calculations, submitted to npj Computational Materials (2020).

Lin, C.-C., Gavini, V., GPU accelerated Tucker DFT code for fast DFT calculations, under preparation (2020).

Motamarri, P., Nowak, M.R., Leiter, K., Knap, J., Gavini, V., A higher order adaptive finite-element discretization of Kohn-Sham density functional theory, *J. Comput. Phys.* **253**, 308-343 (2013).

Motamarri, P., Gavini, V., A subquadratic-scaling subspace projection method for large-scale Kohn-Sham DFT calculations using spectral finite-element discretization, *Phys. Rev. B* **90**, 115127 (2014).

Motamarri, P., Blesgen, T., Gavini, V., Tucker-tensor algorithm for large-scale Kohn-Sham density functional theory calculations, *Phys. Rev. B* **93**, 125104 (2016).

Motamarri, P., Das, S., Rudraraju, S., Ghosh, K., Davydov, D., Gavini, V., DFT-FE: A massively parallel adaptive finite-element code for large-scale density functional theory calculations, *Comput. Phys. Commun.* **246**, 106853 (2020).

Kohn, W., Sham, L. J., Self-consistent equations including exchange and correlation effects. *Phys. Rev.* **140**, A1133?A1138 (1965).

Ozolins, V., Lai, R., Caflisch, R., Osher, S., Compressed modes for variational problems in mathematics and physics. *Proc. Natl. Acad. Sci.* **110**, 18368?18373 (2013).

Zhou, Y., Saad, Y., Tiago, M.L., Chelikowsky, J.R., Parallel self- consistent-field calculations via Chebyshev-filtered subspace acceleration *Phys. Rev. E*, **74** 066704, (2006).