

Semidefinite programming formulation of linear-scaling electronic structure theories

Srikant Veeraraghavan and David A. Mazziotti*

Department of Chemistry and The James Franck Institute, The University of Chicago, Chicago, Illinois 60637, USA

(Received 27 September 2014; published 24 August 2015)

We present a linear-scaling approach based on semidefinite programs (SDPs) to compute the density matrix for effective one-electron theories. Traditional methods constrain the density matrix to represent a Slater determinant and hence rely on parameterization or purification. We eliminate the need for such a constraint by performing an energy minimization over all the convex combinations of density matrices representing Slater determinants. By not relying on purification, the SDP approach not only eliminates accumulation error present in some methods but also reduces the amount of truncation error. Sparsity in the Hamiltonian can be exploited to make the SDP approach scale linearly with system size. Crossovers in computational time with a cubically scaling algorithm are demonstrated for one-dimensional hydrogen chains ranging from H_{50} to H_{1500} .

DOI: [10.1103/PhysRevA.92.022512](https://doi.org/10.1103/PhysRevA.92.022512)

PACS number(s): 31.10.+z

I. INTRODUCTION

Conventional formulations of electronic structure methods based on effective one-electron theories like Hartree-Fock, density-functional, and tight-binding theories scale cubically with system size. This scaling is due to the diagonalization of the effective one-electron Hamiltonian matrix which is an $O(r^3)$ operation. However, there are methods which achieve linear scaling by computing the one-electron reduced density matrix (henceforth referred to as the density matrix) directly [1–5]. All of these methods constrain the density matrix to represent a Slater determinant by either parametrization or purification [1,6]. We present an approach based on semidefinite programs (SDPs) [7,8] which eliminates the need for such a constraint by performing an energy minimization over all the convex combinations of density matrices representing Slater determinants. This enables us to compute an optimal ensemble density matrix for a molecular system with a band gap with no accumulation error and minimal truncation error.

II. THEORY

Density-matrix-based linear-scaling methods, applicable to effective one-electron theories, aim to solve the following minimization problem over the set of Hermitian matrices (\mathbb{H}^r) for an effective one-electron Hamiltonian H of a given N -electron system in a basis of size r :

$$\text{minimize}_{D \in \mathbb{H}^r} E(D) = \text{Tr}(HD), \quad (1)$$

$$\text{subject to } \text{Tr}(D) = N, \quad (2)$$

$$D^2 = D. \quad (3)$$

The trace constraint in Eq. (2) ensures that the density matrix corresponds to the correct number of electrons N . The idempotency constraint in Eq. (3) ensures that the density matrix D corresponds to a Slater determinant preimage. It would be desirable to be able to obtain the global minimum for this minimization, but one can never be sure that an obtained minimum is global due to the fact that this minimization is

nonconvex and hence has multiple local minima. Since the objective function and the other constraints are linear and hence convex, the idempotency constraint is solely responsible for the nonconvexity. We can, however, relax the minimization from the nonconvex set of idempotent density matrices to the convex set of ensemble N -representable density matrices [9,10] which is the set of density matrices representable by at least one ensemble N -electron density matrix and is also the convex hull of idempotent density matrices. Therefore, we relax the nonconvex nonlinear program in Eqs. (1)–(3) to the following convex SDP whose solution yields the global minimum:

$$\text{minimize}_{D, Q \in \mathbb{H}_+^r} E(D) = \text{Tr}(HD), \quad (4)$$

$$\text{subject to } \text{Tr}(D) = N, \quad (5)$$

$$D + Q = I, \quad (6)$$

where $Q = I - D$ is the one-hole density matrix. The notation $D, Q \in \mathbb{H}_+^r$ means that D and Q belong to the set of positive semidefinite Hermitian matrices.

In order to obtain a linear-scaling algorithm to solve Eqs. (4)–(6), we must exploit the sparse structure of the density matrix which is typically more significant in a local basis. Since local bases are typically nonorthogonal, we will first have to generalize the SDP for a nonorthogonal basis as follows:

$$\text{minimize}_{D, Q \in \mathbb{H}_+^r} E(D) = \text{Tr}(S^{-1}HS^{-1}D), \quad (7)$$

$$\text{subject to } \text{Tr}(S^{-1}D) = N, \quad (8)$$

$$D + Q = S, \quad (9)$$

where D , Q , H , and S represent the density, hole density, Hamiltonian, and overlap matrices in a given nonorthogonal local basis. However, since matrix inversion scales cubically, it would be desirable to eliminate S^{-1} from the minimization. Such an elimination can be performed by defining \bar{D} such that $D = S\bar{D}S$ [11] and reformulating Eqs. (7)–(9) as follows:

$$\text{minimize}_{\bar{D}, \bar{Q} \in \mathbb{H}_+^r} E(\bar{D}) = \text{Tr}(H\bar{D}), \quad (10)$$

$$\text{subject to } \text{Tr}(S\bar{D}) = N, \quad (11)$$

$$S(\bar{D} + \bar{Q}) = I, \quad (12)$$

*damazz@uchicago.edu

where $\bar{D}, \bar{Q} \in \mathbb{H}_+^r$ is equivalent to $D, Q \in \mathbb{H}_+^r$. Therefore, solving the above SDP (\bar{D} -SDP) allows us to obtain the globally optimal solution for a nonorthogonal basis. Solution of \bar{D} -SDP by a first-order algorithm scales as $O(r^3)$ [12,13]. In order to formulate a linear-scaling algorithm we exploit the sparsity of \bar{D} and \bar{Q} in the local basis. Specifically, in the present work we reduce \bar{D} and \bar{Q} to smaller blocks whose size is independent of system size. For example, if \bar{D} and \bar{Q} are expected to be block diagonal, then the positive semidefinite constraint can be expressed by enforcing positive semidefiniteness on each of the diagonal blocks. However, even for more general sparsity patterns where nonzero blocks are scattered throughout the \bar{D} and \bar{Q} , positive semidefiniteness of selected blocks can be shown to be necessary though not sufficient for \bar{D} and \bar{Q} to be positive semidefinite. Exploitation of the sparsity in the enforcement of the semidefinite constraints allows us to obtain an approximate energy by solving a SDP with only the nonzero matrix blocks as the variables, which we shall call linear-scaling SDP (LS-SDP). Solution of LS-SDP would scale as $O(r)$, provided the size of the largest nonzero matrix blocks r_b is substantially less than r ($r_b \ll r$) and the number of such blocks scales as r .

Before presenting an example for a particular sparsity pattern, let us compare the LS-SDP formulation to other linear-scaling methods. Being a density-matrix minimization, it is not affected by the accumulation error of spectral projection methods like the trace-correction [14] and trace-resetting [15] methods. Furthermore, experience shows that successive iterations of the spectral projection methods reduce the sparsity of the density matrix. If the density matrix is pruned at every iteration to retain sparsity, there is an increase in truncation and accumulation errors which prevents the methods from converging in a nonorthogonal basis.

Density-matrix minimization methods [16–18], most of which are based on the Li-Nunes-Vanderbilt [19] method, enforce idempotency by utilizing McWeeny’s purification transformation, $\tilde{D} = 3\bar{D}S\bar{D} - 2\bar{D}S\bar{D}S\bar{D}$ [20,21], in the energy expression as follows:

$$E(\bar{D}) = \text{Tr}[H(3\bar{D}S\bar{D} - 2\bar{D}S\bar{D}S\bar{D})]. \quad (13)$$

Changing the energy expression in this fashion introduces global minima which are spurious solutions and correspond to $-\infty$ energies. This forces one to start with an initial guess whose eigenvalues lie in $[-0.5, 1.5]$ in order to obtain finite-energy solutions, which are local minima for this modified energy expression. In the SDP formulation spurious solutions are not introduced, and by convexity any local minimum is the global minimum that solves the problem. Furthermore, the SDP solver [12] allows us to start with a random matrix as the initial guess. Finally, the energy expression in Eq. (13) has a higher number of matrix multiplications than Eq. (10), which not only increases the computational cost but also either reduces the sparsity or increases the truncation error for the same level of sparsity. By not employing any purification, the LS-SDP formulation retains the linear energy expression which is the optimal choice for reducing the truncation error or retaining as much sparsity as possible. Among density matrix minimization methods, only the curvy-steps [22,23] approach




FIG. 1. Performing an energy minimization over blocks of size $r_b \ll r$ enables a $O(r)$ computation of the density matrix. The SDP constrains $(r - r_b + 1)$ blocks of dimension r_b from \bar{D} and an equivalent number of blocks from \bar{Q} to be positive semidefinite.

enforces idempotency using an exponential parametrization, and hence does not formally introduce any spurious nonidempotent solutions. However, since idempotency is a nonconvex constraint, global optimization is not ensured.

As an example of LS-SDP, we present a formulation which exploits a banded sparsity structure, i.e., $\bar{X}_j^i = 0$ for $|i - j| > r_b$, which is common in large linear systems. Each nonzero block is of size r_b along the diagonal and is overlapping with other blocks as shown in Fig. 1. The number of these blocks is $r - r_b + 1$ per banded matrix. That makes solving the resulting SDP scale as $2(r - r_b + 1)r_b^3$ which is $O(r)$ provided $r_b \ll r$.

Approximately enforcing positive semidefiniteness on \bar{D} and \bar{Q} by enforcing the constraint only on the nonzero blocks is a relaxation which should yield a lower bound on the exact energy. However, this also has the effect of constraining elements which are farther than r_b from the diagonal to be zero, which when applied alone would yield an upper bound on the exact energy, for the same reason that the Li-Nunes-Vanderbilt [19] and related methods yield an upper bound. This is because any $r_b < r$ can be seen as imposing additional constraints on the structure of \bar{D} and \bar{Q} , which will force the minimum to be higher than the minimum obtained without those additional constraints. Therefore, the net effect is that the energy obtained is guaranteed to be neither a lower bound nor an upper bound (unless H has a band $\leq r_b$, in which case the “upper bounding effect” is eliminated and overall a lower bound to the exact energy is guaranteed). In practice, we find that the relaxation of N -representability constraints in Eq. (12) outweighs the constraint of zeroing elements that are farther than r_b from the diagonal, yielding a value for the energy which is lower than the exact energy.

III. APPLICATIONS

To demonstrate that the LS-SDP formulation scales linearly with system size, we perform calculations on linear hydrogen chains with a 1-Å spacing between the hydrogen atoms in a minimal basis of Slater-type orbitals expanded in six Gaussian functions denoted STO-6G. Like the traditional density-matrix minimization methods LS-SDP requires a band gap for linear scaling. We use a modified form of the tight-binding

Hamiltonian proposed by Absar and Coleman [24,25] defined as follows:

$$H_j^i = {}^1K_j^i + \frac{N}{r} \sum_{kl} {}^2V_{jl}^{ik} S^{-1}{}_l^k, \quad (14)$$

$${}^1K_j^i = \langle i | \hat{h} | j \rangle, \quad (15)$$

$${}^2V_{jl}^{ik} = \frac{1}{2} (\langle ij | kl \rangle - \langle ij | lk \rangle). \quad (16)$$

The one-electron Hamiltonian operator \hat{h} contains the kinetic-energy operator and electron-nuclei potential; $\langle ij | kl \rangle$ represents the electron-electron repulsion integrals; and the indices i, j, k , and l denote the orbitals in the one-electron basis set of rank r . More generally, any effective one-electron Hamiltonian from tight-binding, Hartree-Fock, or density-functional theories can be employed.

Figure 2 shows the central processing unit (CPU) time required to solve the \bar{D} -SDP (10) and LS-SDP formulations as a function of basis functions (also equal to the number of hydrogen atoms in this case), fitted to a cubic polynomial and a line, respectively. As is evident, the computational cost required to solve LS-SDP indeed scales linearly as compared to solving the \bar{D} -SDP, which scales cubically as expected. We observe a crossover at about $r = 1300$, where solving the LS-SDP becomes computationally less expensive than solving the \bar{D} -SDP. While the use of r overlapping blocks for an $r \times r$ matrix is an intuitive blocking structure, a more sophisticated blocking structure could be used to decrease significantly the basis-set size at which the crossover occurs. For example, in the present example, a more efficient scheme could be achieved by using r/n blocks where the integer n is determined by selecting an offset between the overlapping blocks. Other more complicated sparsity patterns could also be incorporated within the framework of the outlined approach. Additional examples will be given in future work.

Figure 3 shows the absolute error in energy per basis function plotted as a function of the number of basis functions.




FIG. 2. (Color online) The CPU times required to solve LS-SDP and \bar{D} -SDP, fitted to a line and a cubic polynomial, respectively, are shown as functions of the number of basis functions r . The quality of the fit indicates that the LS-SDP method is $O(r)$. Although solving \bar{D} -SDP is $O(r^3)$, it is guaranteed to obtain the global minimum.




FIG. 3. (Color online) The absolute error in energy (calculated relative to the exact energy obtained by solving \bar{D} -SDP) per basis function is plotted as a function of the number of basis functions r . By $r = 400$ the error saturates to a nearly constant value independent of r , thus showing that solving LS-SDP is indeed $O(r)$ without loss of accuracy.

The error is calculated with respect to the exact energy obtained by solving \bar{D} -SDP. Even for H_{1500} the absolute error per basis function stays below 0.00035 a.u. Although the error increases rapidly initially, it quickly saturates to a nearly constant value, proving that the LS-SDP method is $O(r)$.

IV. CONCLUSIONS

We have presented a linear-scaling SDP approach to electronic structure theories that does not constrain the density matrix to represent a Slater determinant either by purification or parametrization. The present work is important in the context of linear-scaling Hartree-Fock and Kohn-Sham density-functional theories for two key reasons: (1) the robustness of convex optimization and (2) the interplay between idempotency and truncation error.

First, the SDP approach to linear scaling relaxes the usual nonconvex optimization of the density matrix, in which the density matrix is forced to be an extreme point of the nonconvex set of idempotent density matrices, to a convex optimization of the density matrix, in which the density matrix can become any density matrix that obeys the Fermi statistics of N electrons. A convex optimization like SDP has advantages over nonconvex optimization including (i) computable conditions that guarantee global optimality and (ii) a direct path toward the minimum unlike the nonconvex methods such as the curvy-step method which restrict their steps to the idempotent manifold.

Second, the SDP approach permits a relaxation of the idempotency constraint in the presence of truncation error. Previously developed linear-scaling methods require exact idempotency of the density matrix even when a sparsity structure (or assumption) is being imposed that affects the idempotency. The SDP approach relaxes the idempotency criterion. In the absence of truncation error, the solution density matrix will be idempotent, but in the presence of

truncation error, the solution density matrix need not be precisely idempotent. Therefore, practical solutions from the SDP method will differ from the conventional solutions by not being strictly idempotent. Advantages include the following: (1) the nonidempotent solution density matrix can potentially be a more accurate representation of the exact idempotent density matrix (for example, upon matrix completion, that is, upon recovering the “missing information” in the sparse representation of the density matrix, the resulting density matrix may in fact correspond to an idempotent density matrix) and (2) the deviation of the density matrix from idempotency can be used to assess the magnitude of the truncation error.

We have demonstrated linear scaling and a crossover in computational time compared to a $O(r^3)$ algorithm. The

LS-SDP does not suffer from an accumulation error like the spectral projection methods and preserves sparsity because it does not use purification. Since better algorithms to exploit sparsity and further improvements in sparse SDP solvers can reduce the prefactor, the SDP approach has potential to provide a highly efficient alternative to existing linear-scaling electronic structure methods.

ACKNOWLEDGMENTS

D.A.M. gratefully acknowledges the NSF, the US Army Research Office (ARO), and the Keck Foundation for their generous support.

-
- [1] R. McWeeny, *Rev. Mod. Phys.* **32**, 335 (1960).
 - [2] W. Yang and T.-S. Lee, *J. Chem. Phys.* **103**, 5674 (1995).
 - [3] S. Goedecker, *Rev. Mod. Phys.* **71**, 1085 (1999).
 - [4] J. Kussmann, M. Beer, and C. Ochsenfeld, *WIREs: Comp. Mol. Sci* **3**, 614 (2013).
 - [5] D. A. Mazziotti, *J. Chem. Phys.* **115**, 8305 (2001).
 - [6] D. A. Mazziotti, *Phys. Rev. E* **68**, 066701 (2003).
 - [7] L. Vandenberghe and S. Boyd, *SIAM Rev.* **38**, 49 (1996).
 - [8] R. M. Erdahl and B. Y. Jin, in *Many-Electron Densities and Reduced Density Matrices*, edited by J. Cioslowski (Kluwer, Boston, 2000), pp. 57–84.
 - [9] A. J. Coleman, *Rev. Mod. Phys.* **35**, 668 (1963).
 - [10] D. A. Mazziotti, *Phys. Rev. Lett.* **108**, 263002 (2012).
 - [11] R. W. Nunes and D. Vanderbilt, *Phys. Rev. B* **50**, 17611 (1994).
 - [12] D. A. Mazziotti, *Phys. Rev. Lett.* **93**, 213001 (2004).
 - [13] D. A. Mazziotti, *Phys. Rev. Lett.* **106**, 083001 (2011).
 - [14] A. M. N. Niklasson, *Phys. Rev. B* **66**, 155115 (2002).
 - [15] A. M. N. Niklasson, C. J. Tymczak, and M. Challacombe, *J. Chem. Phys.* **118**, 8611 (2003).
 - [16] E. Hernández, M. J. Gillan, and C. M. Goringe, *Phys. Rev. B* **53**, 7147 (1996).
 - [17] J. M. Millam and G. E. Scuseria, *J. Chem. Phys.* **106**, 5569 (1997).
 - [18] M. Challacombe, *J. Chem. Phys.* **110**, 2332 (1999).
 - [19] X.-P. Li, R. W. Nunes, and D. Vanderbilt, *Phys. Rev. B* **47**, 10891 (1993).
 - [20] R. McWeeny, *Proc. R. Soc. A* **235**, 496 (1956).
 - [21] R. McWeeny, *Phys. Rev.* **114**, 1528 (1959).
 - [22] T. Helgaker, H. Larsen, J. Olsen, and P. Jørgensen, *Chem. Phys. Lett.* **327**, 397 (2000).
 - [23] Y. Shao, C. Saravanan, M. Head-Gordon, and C. A. White, *J. Chem. Phys.* **118**, 6144 (2003).
 - [24] I. Absar and A. J. Coleman, *Chem. Phys. Lett.* **39**, 609 (1976).
 - [25] I. Absar and A. J. Coleman, *Int. J. Quantum Chem.* **10**, 319 (1976).