Multisliced gausslet basis sets for electronic structure

Steven R. White and E. Miles Stoudenmire

Phys. Rev. B **99**, 081110 — Published 11 February 2019

DOI: 10.1103/PhysRevB.99.081110

Multi-sliced Gausslet Basis Sets for Electronic Structure

Steven R. White[1] and E. Miles Stoudenmire[2]
[1]*Department of Physics and Astronomy, University of California, Irvine, CA 92697-4575 USA*
[2]*Center for Computational Quantum Physics, Flatiron Institute, New York, NY 10010 USA*
(Dated: January 30, 2019)

We introduce highly local basis sets for electronic structure which are very efficient for correlation calculations near the complete basis set limit. Our approach is based on gausslets, recently introduced wavelet-like smooth orthogonal functions. We adapt the gausslets to particular systems using one dimensional coordinate transformations, putting more basis functions near nuclei, while maintaining orthogonality. Three dimensional basis functions are composed out of products of the 1D functions in an efficient way called multislicing. We demonstrate the new bases with both Hartree Fock and density matrix renormalization group (DMRG) calculations on hydrogen chain systems. With both methods, we can go to higher accuracy in the complete basis set limit than is practical for conventional Gaussian basis sets, with errors near 0.1 mH per atom.

Electronic structure calculations on molecules, solids, and biological systems are performed by thousands of groups worldwide and account for a substantial fraction of the world's scientific computing. Strongly correlated systems, for which density functional approaches are inadequate, make up a small but important fraction of these calculations. An almost universal problem with methods for strong correlation is poor computational scaling in both system size and accuracy. For example, coupled cluster or configuration interaction approaches typically scale as N_e^6 or higher for N_e electrons. Since correlation methods must deal with the two-electron interaction directly, scaling of at least N^4 when using N basis functions can appear hard to avoid, since the two-electron interaction terms are described by a tensor V_{ijkl}.

Local representations can reduce the size of the V_{ijkl} tensor. In the extreme case of a grid descretization, the interaction is reduced by a factor of N^2 to a matrix V_{ij}, with $V_{ij} = 1/|\vec{r}_i - \vec{r}_j|$ for $i \neq j$. For the less extreme case of basis sets where the functions have substantial spatial compactness, we say that two basis functions b_i and b_j "overlap" if there is some point $\vec{r}$ where $b_i(\vec{r})b_j(\vec{r})$ is nonnegligible. Terms in V_{ijkl} are negligible unless basis functions i and j overlap and also k and l overlap. However, in three dimensions, even for substantially localized functions, many basis functions overlap, particularly if the functions have been orthogonalized, limiting any increase in the sparseness of V_{ijkl}. This is unfortunate, since basis methods have several advantages over grids, such as the ability to add extra atom-centered core functions to better resolve the nuclear cusps.

Recently one of us introduced a novel basis function approach that has the same favorable scaling of the interaction as a pure grid [1]. This involved two key ingredients: first, the introduction of a wavelet-related set of highly localized, smooth orthogonal basis functions, called *gausslets*, where each function is defined as a sum over an underlying grid of simple Gaussians. Second, it was shown that one can construct an accurate purely diagonal interaction V_{ij} for a gausslet basis. This diagonal interaction for a special type of basis is not new in itself: a basis of sinc functions also allows this construction [2]. However, the extreme delocalization of sinc functions is a severe disadvantage as a basis; the gausslet development in Ref. 1 shows that one can get the diagonal property with much more localized functions, where it is based on the ability of gausslets to integrate like a delta function. But the usefulness of gausslets was previously only demonstrated for 1D toy systems.

Here we generalize the gausslet approach to three dimensions and practical electronic structure calculations. Given a 1D basis, one can always generate a 3D basis as coordinate products, i.e. $G_{ijk}(x, y, z) = f_i(x)g_j(y)h_k(z)$. This simple approach produces overly large basis sets. Instead, we introduce coordinate transformations which put more functions near nuclei, and a procedure called multislicing which allows the use of 1D coordinate transformations rather than more complicated 3D transformations. Our multisliced gausslet (MSG) approach is a generalization of our earlier sliced basis approach [3]. We demonstrate the resulting method on hydrogen chain systems [4], using both Hartree Fock and the density matrix renormalization group (DMRG) [5–7]. In both cases the diagonal property make it possible to use very complete bases. The combination of MSG and DMRG (MSG-DMRG) allows for simultaneously exact correlation and the complete basis limit, going well beyond chemical accuracy in a controlled way, and exhibits linear scaling of the computation time in the length of the chain.

In standard orthogonal wavelet theory, basis sets are made of two types of functions, scaling functions, which carry low momentum, and wavelets, which carry a range of higher momenta. Gausslets are like scaling functions, but they are constructed out of sums of Gaussians for numerical convenience. A set of gausslets are shown in Fig. 1(a), highlighting a single gausslet in the center of the figure. To make a 1D basis, one puts a gausslet at each point on an evenly spaced grid, scaling them to match the grid spacing. The oscillatory tails make the gausslets precisely orthonormal, and they can exactly represent polynomials up to some predetermined order (e.g. 10th order).

Modifying the gausslets with coordinate transformations allows spatially varying resolution. Let $x(u)$ and

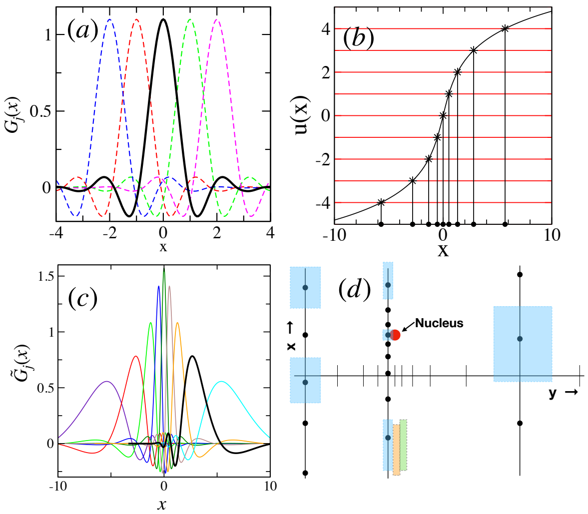

FIG. 1. (a) Array of gausslets, with the gausslet centered at the origin emphasized to show detail. (b) Coordinate transformation function $u(x)$ for a single atom, with $a = s = 0.7$ in Eq. (2), to give gausslets variable resolution. (c) Distorted gausslet basis based on the transformation of (b), which is orthonormal and allows a diagonal approximation. One of the functions is emphasized. (d) Schematic representation of multislicing in 2D. The vertical lines represent slices, with three shown in detail. Each dot is the center of a basis function, and the shaded rectangles illustrate the principle support region of some of the functions, although they have smooth tails well beyond the rectangles. The multicolored shaded rectangles represent long, thin basis functions which one would want to contract at a later stage.

its inverse $u(x)$ define a 1D smooth one-to-one coordinate mapping, which will be used to make the grid narrow and closely spaced near nuclei, and wide and sparse far away. First consider a 1D arrangement, with just one atom at $x = 0$. Define the gausslets on a uniform grid in the u space and then map to x-space, inserting a Jacobian factor to preserve orthonormality. If $G_j(u)$ is a gausslet centered at integer j, define

$$\tilde{G}_j(x) = G_j(u(x))\sqrt{u'(x)}\ . \tag{1}$$

The $\tilde{G}_j$ are orthonormal if the G_j are.

The coordinate mapping we choose for a single atom is given by

$$u(x) = \sinh^{-1}(x/a)/s \tag{2}$$

where the parameter s, the *scale*, sets or adjusts the overall gausslet spacing, and a, the *core cutoff* sets the range in x over which we stop decreasing the gausslet spacing. The smallest gausslet spacing at the nucleus is about $a \cdot s$. This transformation is shown in Fig. 1(b), with the resulting 1D functions shown in Fig. 1(c). In the Supplemental Material, we discuss the motivation for the above form of the transformation $u(x)$, as well as a modified form of the transformation better suited for multiple atoms.

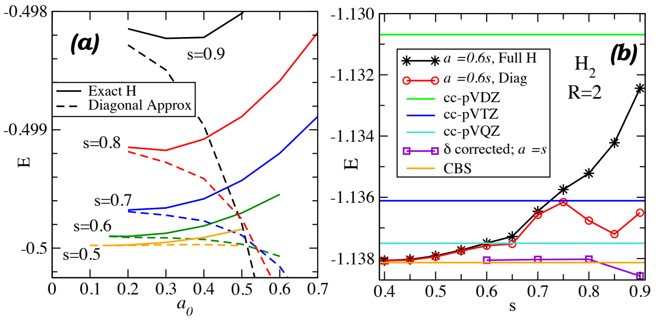

FIG. 2. (a)Energies of a hydrogen atom in an MSG basis as a function of a and s, in Hartrees, using the full Hamiltonian ("exact H") and using the integral diagonal approximation. The exact energy is $-1/2$. (b) Energy of a hydrogen molecule with separation $R = 2$ in standard Gaussian and MSG bases.

For 3D basis functions, a coordinate-product form $\mathcal{G}(x, y, z) = f(x)g(y)h(z)$ greatly simplifies evaluation of integrals defining the Hamiltonian. To keep this form, we apply coordinate transformations to each coordinate separately, in a method we call "multislicing". The coordinate directions are sliced up sequentially, z (which runs along the chain), then y, then x. A first coordinate transformation $u^z(z)$ determines a set of z-values z_k $(k = 1, 2, \ldots)$, with $u^z(z_k) = k$, at which are centered distorted 1D gausslets $\tilde{G}_k(z)$. The plane $z = z_k$ and function $\tilde{G}_k(z)$ together define a z-slice. Next we slice up each z-slice in the y direction, with a coordinate transformation unique to k, $u^y_k(y)$, which defines a set of y-values y_{kj}. A y "slice" (or "subslice" of a "parent" z-slice) is the line $z = z_k$, $y = y_{kj}$, with associated 2D function $\tilde{G}_k(z)\tilde{G}_{kj}(y)$. Finally, for each y-slice, define a unique coordinate transformation $u^x_{kj}(x)$, determining a set of x values x_{kji}, and 3D basis functions $\tilde{G}_k(z)\tilde{G}_{kj}(y)\tilde{G}_{kji}(x)$.

The key point in using this successive procedure is to use of the knowledge of where a slice is, relative to the nuclei, to make subsequent transformations with the lowest density of functions. This is illustrated schematically in 2D in Fig. 1(d). Preserving the product form via multislicing means that some basis functions are long and thin; however, at a later stage on can devise methods to contract such functions with their neighbors, reducing unnecessary degrees of freedom. The details of the coordinate transformations in the multisliced case are discussed in the Supplementary Material.

Each basis function has a well defined center (x_{kji}, y_{kj}, z_k), and we can make a simple rule for which functions to keep: if the basis function is within a distance b of an atom, we keep it. Here $b = 9$ a.u. proved very accurate (< 0.1 mH errors compared to larger b) except for $R = 1$ for H_{10}, where we used $b = 13$.

Figure 2(a) shows energies for a single hydrogen atom for various a and s, using both the standard Hamiltonian and one where a diagonal approximation is made for the single particle potential [1]. Since there are only N^2 single particle terms, using this diagonal approximation barely improves computational efficiency, but one would expect this approximation to mimic the performance of the more important two-particle diagonal approximation. The diagonal approximation is sensitive to the singularity in the potential at the nucleus, but increasing the basis function density near the nucleus by decreasing a, for fixed s, nearly eliminates the diagonal approximation error. A simple procedure to systematically converge to the ground state for this system would be to fix a/s to be a constant, say 0.5-0.6, and then decrease s.

Figure 2(b) shows the energy for a hydrogen molecule, compared to standard basis sets cc-pVxZ, where x=D, T, and Q, and also compared to the exact energy from a treatment in special coordinates [9]. A diagonal approximation for the two particle interaction is used here and in all subsequent MSG bases, since calculations would not be practical with the standard V_{ijkl} form. All results shown are exact (full CI) given the approximate Hamiltonian. The MSG bases systematically converge to the exact results, and the diagonal approximation for the single particle potential closely approximates the full Hamiltonian, particularly for smaller s.

Also shown in Fig. 2(b) is a basis with a special delta-function correction for the nuclear cusp. Increasing the resolution near nuclei by using a small a is inefficient, leading to many basis functions. For example, for the hydrogen atom of Fig. 2(a), taking $a = 0.3$, $s = 0.6$ produced 1179 functions, which resulted in an error of 0.13 mH. Our correction consists of adding a single-particle potential at each atom α of the form $v_\alpha \delta(\vec{r} - \vec{r}_\alpha)$. The parameter v_α is set by "turning off" all nuclear electron potentials for atoms other than α (yet keeping the same set of functions to be used for the entire system), and adjusting v_α so that the one-electron ground state energy is the exact hydrogen atom energy $-1/2$. The errors associated with choosing a too large are localized near the nuclei; the delta function potential alters the terms in the Hamiltonian only for the basis functions overlapping with a nucleus. Most importantly, $v_\alpha \to 0$ as $a \to 0$ or $s \to 0$, so this correction does not change what the basis converges to, only how fast it converges, accelerating the convergence. In Fig. 2 and for the rest of the results, we set $a = s$ and use the delta correction.

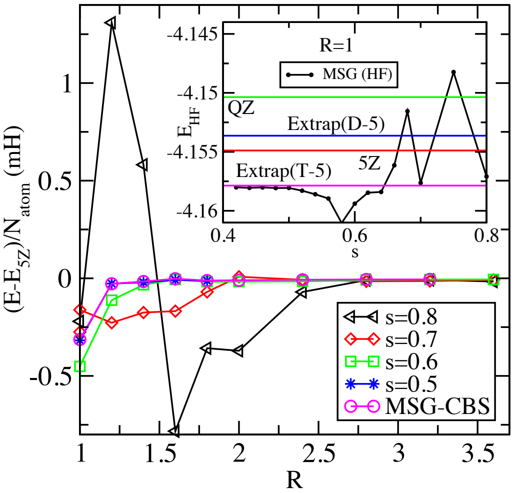

FIG. 3. Hartree Fock energies per atom of H_{10} versus R, relative to the Gaussian basis set cc-pV5Z (5Z)[4]. The connected symbols are MSG-HF results at constant $s = a$, as labeled. For small R, they converge to a small but noticeably different result from cc-pV5Z. The inset shows MSG-HF data versus $s = a$, at $R = 1$. Horizontal lines show Gaussian results, along with two exponential extrapolations, based on (TZ,QZ,5Z) (labeled T-5), and on (DZ,TZ,QZ,5Z) (labeled D-5). The MSG-HF agrees for small s with the T-5 extrapolation, although the D-5 extrapolation was used in Ref. [4].

We now turn to a more challenging system, a linear chain of hydrogen atoms spaced R apart. Hydrogen chain systems were the subject of a recent benchmark study which compared more than a dozen methods in their ability to reach the combined limit of exact correlation, complete basis set, and infinite number of atoms [4]. We first consider unrestricted Hartree Fock (UHF) on H_{10}, shown in Fig. 3. The plot shows HF energy differences relative to those of a large Gaussian basis, cc-pV5Z. The convergence of the MSG basis is irregular because the centers of the gausslets are not aligned with the nuclei; but it is easy to get very accurate results and judge the accuracy. At small R, the Gaussian basis sets have trouble due to linear dependence [4], leading to a small but noticeable discrepancy between the 5Z and MSG results. As shown in the inset, at $R = 1$ the Gaussians converge slowly, and different extrapolations give different results. As a rough comparison of the calculational effort for these very high accuracy calculations: for $R = 1$, $a = s = 0.5$, the MSG basis has just over 13,000 basis functions; the number of two-electron integrals is the square of this, or 1.7×10^8. The 5Z basis has 550 functions, but the number of integrals (N^4, ignoring symmetry) is 9.2×10^{10}. The calculation time of our UHF algorithm, which takes advantage of the diagonal nature of the Hamiltonian, scales as $N^2 N_e$, where N_e is the number of electrons, with the dominant part coming from a Davidson diagonalization, for N_e eigenvectors, of the Fock matrix.

For correlated calculations, to decrease the number of basis functions, one can use the HF occupied orbitals to contract the MSG basis to smaller size. This can be done in a way that maintains the diagonal form of the interactions. One can also extrapolate in a cutoff that controls this contraction, to obtain results for the uncontracted basis. The largest systems needed for a extrapolation are still about a factor of 2 or 3 smaller than the uncontracted basis, and the results below follow this procedure, which is described in the Supplementary Material.

We now turn to MSG-DMRG calculations for H_{10}. Our DMRG implementation uses the matrix product operator compression of our earlier sliced basis DMRG (SB-DMRG) approach [3]. This compression makes the calculation time for fixed accuracy per atom scale linearly in the number of atoms in a hydrogen chain both in SB-DMRG and MSG-DMRG. We are currently limited to

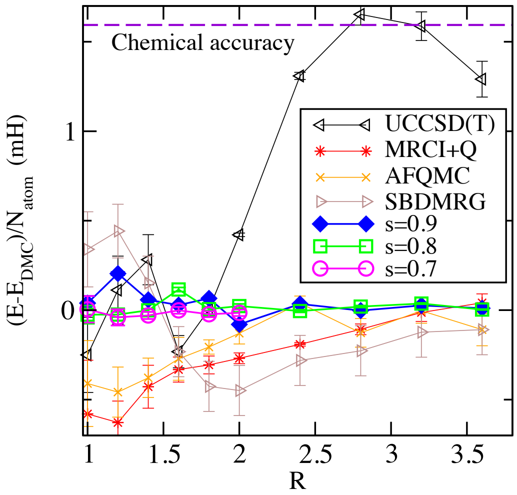

FIG. 4. Complete basis set energies per atom of H_{10} versus R, relative to a diffusion Monte Carlo method, for MSG-DMRG (labeled by s) versus various approaches from Ref. 4.

about 3000-4000 basis functions. (In contrast, standard DMRG in a Gaussian basis—with no diagonal approximation and no compression—is limited to about 100-200 active basis functions.) We find that the DMRG performs very well. For the very high accuracy results shown below, we generally only needed to keep about 200 states for larger R, and up to 400-500 for $R = 1$ (due to its more metallic character). This excellent performance is due to the high locality of the basis, which DMRG and other tensor network methods [10–12] strongly prefer.

We find that the correlation energy converges faster with s than the HF energy. This is not surprising: the representation of the nuclear cusp is poor with a coarse gausslet basis, which is is primarily a single particle effect. Therefore, to get total energies we use the HF energy with very small s, and add to it the correlation energy obtained with a larger s, where the correlation energy is defined by subtracting the unrestricted HF energy from the total energy for the same basis set.

In Fig. 4, we show a comparison of total energies for several methods [4] and our MSG-DMRG for various $s = a$. All methods attempt to reach the CBS limit; for all but the MSG-DMRG and DMC methods, this involved an extrapolation in the basis set. The energy differences here are generally well below chemical accuracy. Often such high accuracy is unnecessary, but studying the high accuracy limit is an excellent way to demonstrate the usefulness of MSG-DMRG. The energies are measured relative to one of the diffusion Monte Carlo methods, LR-DMC-AGP (or DMC). In Ref. [4], at this level of accuracy, none of the best available methods agreed, so it was not known which was best, and reference plots were made relative to MRCI+Q for smaller systems and AFQMC for larger ones. DMRG based on standard Gaussian basis sets could not be done beyond the TZ level, so no CBS results were available. Here, we find systematic convergence of MSG-DMRG to energies agreeing with the LR-DMC-AGP method. Agreement was poorer at small R with a DMC method based on an LDA trial function. There are systematic errors in DMC stemming from the fixed node approximation, which are unusually small in this 1D system, but hard to quantify. Since the nature of errors in DMC and MSG-DMRG are completely different, and since the MSG-DMRG energies converge systematically with a control parameter, we can be rather sure that MSG-DMRG and DMC are both getting the most accurate energies.

The MSG-DMRG errors for fixed $s = a$ are biggest at small R. This is expected; at small R, it would be more natural to scale s with R, keeping the number of basis function more nearly constant. The smallest grid spacings are about $a \cdot s$, or about 0.5 for $s = 0.7$. Small R is challenging to the Gaussian basis set methods because the basis functions become linearly dependent.

In summary, even in this first implementation of the MSG-DMRG method, for the strongly-correlated H_{10} system we surpass the best Gaussian basis approaches in the high-accuracy CBS regime. We believe larger Z systems, not just in the linear geometry of H_{10}, could be treated straightforwardly using pseudopotentials. However, we believe our approach can also be improved so that resorting to pseudopotentials is not necessary for moderate Z. For example, one could add some Gaussians from a standard basis to a gausslet basis, orthogonalizing the Gaussians to the gausslets, to better represent core orbitals. This is very simple to do in principle, but we would also like to find diagonal approximations involving the Gaussians, or develop convenient partially-diagonal approximations, where the number of non-diagonal terms is not too big. The delta correction would likely be eliminated in any of these approaches. Another way to improve fitting core orbitals would involve adapting the gausslets during the slicing to fit 1D Gaussians taken from a standard basis. Regarding how one uses MSG bases, in the hydrogen chains studied here, the linear geometry makes DMRG especially powerful. For less linear molecules or solids, one might couple multisliced gausslets with tensor network states [10–12] or quantum Monte Carlo.

We acknowledge useful conversations with Ryan Babbush, Jarrod McClean, Shiwei Zhang, Mario Motta, Garnet Chan, and Sandro Sorella. We acknowledge support from the Simons Foundation through the Many-Electron Collaboration, and from the U.S. Department of Energy, Office of Science, Basic Energy Sciences under award #DE-SC008696. The Flatiron Institute is a division of the Simons Foundation.

---

[1] Steven R. White, "Hybrid grid/basis set discretizations of the schrodinger equation," The Journal of Chemical Physics **147**, 244102 (2017), https://doi.org/10.1063/1.5007066.

[2] Jeremiah R. Jones, Francois-Henry Rouet, Keith V. Lawler, Eugene Vecharynski, Khaled Z. Ibrahim, Samuel Williams, Brant Abeln, Chao Yang, William McCurdy, Daniel J. Haxton, Xiaoye S. Li, and Thomas N. Rescigno, "An efficient basis set representation for calculating electrons in molecules," Molecular Physics **114**, 2014–2028 (2016), https://doi.org/10.1080/00268976.2016.1176262.

[3] E. Miles Stoudenmire and Steven R. White, "Sliced basis density matrix renormalization group for electronic structure," Phys. Rev. Lett. **119**, 046401 (2017).

[4] Mario Motta, David M. Ceperley, Garnet Kin-Lic Chan, John A. Gomez, Emanuel Gull, Sheng Guo, Carlos A. Jiménez-Hoyos, Tran Nguyen Lan, Jia Li, Fengjie Ma, Andrew J. Millis, Nikolay V. Prokof'ev, Ushnish Ray, Gustavo E. Scuseria, Sandro Sorella, Edwin M. Stoudenmire, Qiming Sun, Igor S. Tupitsyn, Steven R. White, Dominika Zgid, and Shiwei Zhang (Simons Collaboration on the Many-Electron Problem), "Towards the solution of the many-electron problem in real materials: Equation of state of the hydrogen chain with state-of-the-art many-body methods," Phys. Rev. X **7**, 031059 (2017).

[5] Steven R. White, "Density matrix formulation for quantum renormalization groups," Phys. Rev. Lett. **69**, 2863–2866 (1992).

[6] Steven R. White and Richard L. Martin, "Ab initio quantum chemistry using the density matrix renormalization group," The Journal of Chemical Physics **110**, 4127–4130 (1999), https://doi.org/10.1063/1.478295.

[7] Garnet Kin-Lic Chan and Sandeep Sharma, "The density matrix renormalization group in quantum chemistry," Annual Review of Physical Chemistry **62**, 465–481 (2011), pMID: 21219144, https://doi.org/10.1146/annurev-physchem-032210-103338.

[8] See Supplemental Material for more information about computing integrals for a multi-sliced basis and contracting functions together using a technique based on Hartree-Fock. The Supplemental Material also includes the additional references Refs. [13–15].

[9] Wlodzimierz Kolos, Krzysztof Szalewicz, and Hendrik J. Monkhorst, "New bornoppenheimer potential energy curve and vibrational energies for the electronic ground state of the hydrogen molecule," The Journal of Chemical Physics **84**, 3278–3283 (1986), https://doi.org/10.1063/1.450258.

[10] F. Verstraete and J. I. Cirac, "Renormalization algorithms for Quantum-Many Body Systems in two and higher dimensions," eprint arXiv:cond-mat/0407066 (2004), cond-mat/0407066.

[11] Philippe Corboz, "Variational optimization with infinite projected entangled-pair states," Phys. Rev. B **94**, 035133 (2016).

[12] M. J. O'Rourke, Z. Li, and G. Kin-Lic Chan, "Efficient representation of long-range interactions in tensor network algorithms," ArXiv e-prints (2018), arXiv:1807.08378 [cond-mat.str-el].

[13] Lin Lin, Jianfeng Lu, Lexing Ying, and Weinan E, "Adaptive local basis set for kohnsham density functional theory in a discontinuous galerkin framework i: Total energy calculation," Journal of Computational Physics **231**, 2140 – 2154 (2012).

[14] Ryan Babbush, Craig Michael Gidney, Dominic W. Berry, Nathan Wiebe, Jarrod McClean, Alexandru Paler, Austin Fowler, and Hartmut Neven, "Encoding electronic spectra in quantum circuits with linear t complexity," arXiv:1805.03662 (2018).

[15] J. Bezanson, A. Edelman, S. Karpinski, and V. Shah, "Julia: A fresh approach to numerical computing," SIAM Review **59**, 65–98 (2017), https://doi.org/10.1137/141000671.