

# Toward quantum-chemical method development for arbitrary basis functions

Michael F. Herbst,<sup>1,a)</sup> Andreas Dreuw,<sup>1,b)</sup> and James Emil Avery<sup>2,c)</sup>

<sup>1</sup>*Interdisciplinary Center for Scientific Computing, Heidelberg University, Im Neuenheimer Feld 205, 69120 Heidelberg, Germany*

<sup>2</sup>*Niels Bohr Institute, University of Copenhagen, Blegdamsvej 17, 2100 København, Denmark*

(Received 15 June 2018; accepted 2 August 2018; published online 27 August 2018)

We present the design of a flexible quantum-chemical method development framework, which supports employing any type of basis function. This design has been implemented in the light-weight program package *molsturm*, yielding a basis-function-independent self-consistent field scheme. Versatile interfaces, making use of open standards like *python*, mediate the integration of *molsturm* with existing third-party packages. In this way, both rapid extension of the present set of methods for electronic structure calculations as well as adding new basis function types can be readily achieved. This makes *molsturm* well-suitable for testing novel approaches for discretising the electronic wave function and allows comparing them to existing methods using the same software stack. This is illustrated by two examples, an implementation of coupled-cluster doubles as well as a gradient-free geometry optimisation, where in both cases, arbitrary basis functions could be used. *molsturm* is open-sourced and can be obtained from <http://molsturm.org>. Published by AIP Publishing. <https://doi.org/10.1063/1.5044765>

## I. INTRODUCTION

The central goal of electronic-structure theory is to find approximate solutions to the electronic wave equation numerically. This requires a discretisation of the electronic wave function. Typically, it is approximated as a linear combination of Slater determinants: anti-symmetrised products of single-particle functions. The latter are in turn constructed by expanding them in a basis set of *a priori* determined single-electron functions. Usually, such basis sets are not complete and introduce basis set errors. Proper choice of the basis function type and size of basis set is thus decisive for an accurate description of the system under investigation. It is also clear from the onset that different basis function types can be more or less suited for a specific problem, suggesting us to conduct investigations across existing basis function types.

Gaussian-based methods are overwhelmingly predominant in computational electronic structure theory, which stems from pragmatic reasons dating back to the founding years.<sup>1,2</sup> It was well-known that bound state electronic wave functions decay exponentially both at short and large distances from the nuclei,<sup>3</sup> but multi-centre electron-repulsion integrals (ERI) of products of exponential-type orbitals (ETO) were impractically difficult to calculate. For Gaussian-type orbitals (GTO), on the other hand, ERI could be calculated efficiently due to the Gaussian product theorem. However, the computational challenges facing quantum chemists have changed since the 1970's, and it may now be worth trading extra

computation per integral for having fewer more accurate basis functions.

In many practical applications, the shortcomings of GTOs are not important, or one is able to compensate by employing specialised contracted basis sets.<sup>4,5</sup> However, even contracted GTOs (cGTO) fail to describe both the nuclear cusp and the exponential decay of the electron density.<sup>6</sup> In addition, no matter the number of GTO basis functions used, the derivatives are always wrong at the nucleus, which causes singularities and computational failure, for example, in quantum Monte Carlo calculations.<sup>7–9</sup> Furthermore the description of some properties such as the nuclear-magnetic resonance (NMR) shielding tensors or a description of Rydberg-like or auto-ionising states<sup>10–13</sup> directly involves the nuclear cusp or the asymptotic tail, making physically accurate basis functions desirable.<sup>14,15</sup>

The name of our implementation, “*molsturm*,” is a portmanteau of *molecular Sturmians*: the project was born as a means to solve the problem of using state-of-the-art quantum chemistry methods together with generalised- and molecular Sturmian basis functions. The many promising results for generalised Sturmians were stranded due to the fact that only electronic structure problems that were small enough to be solved by direct configuration interaction methods could be treated, preventing wider use. The existing mature quantum chemistry software has been developed over hundreds of man-years, and the methods are not easily reimplemented from scratch.

In theory, it should be a simple matter to include new basis function types in existing quantum chemistry software by swapping the integral calculator. In practice, it turned out to be exceedingly difficult due to the very large and complicated code bases of all the investigated quantum chemistry software. Assumptions about the basis function type

<sup>a)</sup>Electronic mail: michael.herbst@iwr.uni-heidelberg.de

<sup>b)</sup>Electronic mail: dreuw@uni-heidelberg.de

<sup>c)</sup>Electronic mail: avery@nbi.ku.dk

scattered around the source code only make this even task more difficult.

Our solution, which is presented in this paper, is to implement a light-weight layer that makes it easy to experiment with many different basis function types and quantum-chemical methods. It is designed for researchers to both build simple stand-alone programs for prototyping and teaching purposes, and make plug-in modules to be hosted in standard quantum chemistry software. To the best of our knowledge, such a framework with the ability to explore quantum-chemical methods across multiple basis function types has been missing up to today.

## A. Alternative basis function types

Many research groups have worked on alternative basis function types. This section will provide a brief overview with particular focus on exponential-type orbitals (ETO). For further details regarding the basis function mentioned, the reader is referred to the cited works.

Efforts on making various types of ETOs computationally feasible were pioneered by Harris *et al.*<sup>16–19</sup> A particular form of complete ETO basis is Coulomb Sturmians<sup>20–24</sup> (CS). Their functional form is identical to the familiar hydrogen-like orbitals, just with all occurrences of the factors  $Z/r$  replaced by the Sturmian exponent  $k$ —a parameter, which is the same for all functions of the basis,

$$\varphi_{nlm}^{\text{CS}}(\mathbf{r}) = k^{3/2} N_{nl} (2kr)^l e^{-kr} L_{n-l-1}^{2l+1}(2kr) Y_l^m(\hat{\mathbf{r}}). \quad (1)$$

In this,  $Y_l^m$  is a spherical harmonic,  $L_{n-l-1}^{2l+1}$  an associated Laguerre polynomial, and

$$N_{nl} = \frac{2}{(2l+1)!} \sqrt{\frac{(l+n)!}{n(n-l-1)!}} \quad (2)$$

a normalisation constant. CS functions proved to be especially easy to work with since their momentum-space representation by hyperspherical harmonics allows efficient calculation of multi-centre integrals, opening the way for efficient molecular calculations.<sup>25–28</sup> The Coulomb Sturmian construction generalises well and generalised Sturmian basis sets preserving many useful Sturmian properties can be constructed. These allow us, for example, to build  $N$ -particle basis functions that include important geometric properties of the physical system under consideration at the level of the basis.<sup>29–37</sup> Similarly,  $d$ -dimensional hyperspherical harmonic basis sets can model collective motions of particles, for example, for treating strongly interacting few-body systems or reactive scattering.<sup>38–42</sup> A particular type of one-particle Sturmians combines a bound-state region and plane-wave asymptotics to model photoionisation in scattering.<sup>43–46</sup> Other directions of research toward alternative discretisation methods include quantum chemistry on numerical real-space grids,<sup>47,48</sup> finite element methods,<sup>49–54</sup> and wavelets.<sup>55–60</sup>

The `molsturm` package is to support such research directions by providing a common platform for development, testing, and analysis of quantum-chemical methods irrespective of the basis function type employed for the discretisation. The goal is for the implementation work of introducing a new basis function type to be reduced to adding an extra integral

back end in `molsturm`, which will then both provide simple stand-alone calculations and a common interface to hook into existing quantum chemistry packages.

## B. Toward basis-type agnostic quantum chemistry

In order to reach a basis-type agnostic design, there are three fundamental components to consider: (i) an integral interface accommodating a wide range of very different basis set types and discretisation, but providing a uniform way of accessing them, (ii) simple discretisation-agnostic implementations of the self-consistent field (SCF) algorithms, and (iii) a flexible interface to employ the resulting SCF orbital basis further in existing third-party code. Once the SCF orbitals have been obtained, the remainder of a calculation, e.g., a Post-Hartree-Fock (Post-HF) method, can usually be formulated entirely in the SCF orbital basis, without reference to the underlying basis functions. Thus, a basis-function independent SCF scheme automatically leads to basis-function independent Post-HF methods as well.

This structure has another advantageous side effect in the context of developing new basis function types, as it allows us to perform comparisons between old and new methods using exactly the same software stack. In other words, one can thus be sure that, apart from the discretisation, all aspects of the calculation, e.g., SCF algorithms or guess methods, are optimised at the same level leading to a fair apples-to-apples comparison between old and new methods.

## C. Paper outline

The remainder of the paper is structured as follows: Sec. II reviews existing projects with similar goals to `molsturm`. Section III provides a theoretical background for the program design choices, which are described in Sec. IV. Section V provides example problems calculated using `molsturm`'s `python` interface, illustrating how to implement new methods on top of `molsturm` in a few lines of `python`. Section VI outlines the current state of `molsturm` and what we hope to achieve in the future.

## II. RELATED QUANTUM-CHEMICAL SOFTWARE PACKAGES

This section reviews existing quantum chemistry software that shares some of the goals of `molsturm`.

The quantum Monte Carlo packages `CASINO`<sup>61</sup> and `QMCPACK`<sup>62</sup> are among the few systems that support many different basis function types. Both programs support discretizations in terms of GTOs, Slater-type orbitals (STOs), plane-waves, and numerical orbitals like splines. Similarly, the packages `CP2K`,<sup>63</sup> `ASE`,<sup>64</sup> and `GPAW`<sup>65,66</sup> can be employed to perform and post-process computations using more than one type of basis function. `GPAW` and `CP2K` further support calculations with hybrid basis sets that mix Gaussian-type orbitals with plane waves. However, to the best of our knowledge, the design of these packages is very specific to the particular combinations of the basis function type and method.

Combining a FORTRAN or C/C++ implementation of the time-critical core with `python` as a high-level interface language works exceedingly well, and this solution has become increasingly popular. Sun *et al.*<sup>67</sup> describe the reasons as follows in their paper about `pyscf`:

- There is no need to learn a particular domain-specific input format.
- All language elements from `python` are immediately available to, e.g., automatise repetitive calculations with loops.
- The code is easily extensible beyond what is available inside `pyscf`, for example, to facilitate plotting or other kinds of analysis.
- Computations can be done interactively, which is helpful for testing or debugging.

We add here that `python` as a high-productivity language often achieves even complicated tasks with few lines of code while remaining easy to read and understand, demonstrated, for example, by the coupled cluster implementation shown in Sec. V. In the context of quantum chemistry, this has the pleasant side effect that a `python` script used for performing calculations and subsequent analysis is typically brief, but still documents the exact procedure followed. All this comes at essentially no downside if `python` is combined with a carefully optimised low-level C or FORTRAN code in the numerical hot spots. Sun *et al.*,<sup>67</sup> for example, claim that `pyscf` is as fast as any other existing quantum chemistry packages written solely in C or FORTRAN.

Even meta-projects like ASE<sup>64</sup> or `cclib`<sup>68</sup> which aim at extending existing packages by a common `python` front end have emerged. Other packages like HORTON,<sup>69</sup> `pyscf`,<sup>67</sup> `pyQuante`,<sup>70</sup> and GPAW<sup>65,66</sup> are written almost exclusively in `python` and only employ low-level C or C++ code for the computationally demanding routines to various extents.

Starting from the opposite direction, Psi4<sup>71</sup> has gradually introduced a more and more powerful `python` interface on top of their existing C++ core over the years. Recently their efforts have led to the Psi4NumPy project,<sup>72</sup> which combines the `python` interface of Psi4 with the tensor operation syntax of `numpy` arrays.<sup>73</sup> The aims of Psi4NumPy are very much in line with `molsturm`, namely, to provide a framework, which yields flexible and easy-to-read codes. It is thus highly suitable for reference implementations, rapid prototyping, or teaching.<sup>72</sup> Unlike `molsturm`, however, Psi4NumPy does not exhibit a basis-type agnostic design and only supports discretisations based on cGTO basis sets.

Another common feature of `pyscf` and Psi4 is their modular design. They use well-established open standards like HDF5<sup>74</sup> or `numpy` arrays<sup>73</sup> for data exchange, such that linking their codes to external projects becomes easy. Psi4, for example, managed to integrate more than 15 external packages into their framework. This includes three completely different back ends for the computation of the required integrals. In the case of `pyscf`, it only took us about a day to link our `molsturm` to the full configuration interaction (FCI) algorithms of `pyscf` via an interface based on `numpy`. Nevertheless the numerical requirements of Gaussian-type orbitals are currently hard-coded inside the optimised C or C++ parts

of both these projects, such that extending them by other types of basis functions could still be difficult.

### III. THEORY

This section briefly discusses the theoretical background and properties of self-consistent field problems in the context of the basis-type independent design aspired for `molsturm`. A more detailed analysis is provided in Ref. 75.

#### A. Self-consistent field schemes

Both Hartree-Fock (HF) as well as Kohn-Sham density-functional theory (DFT) can be viewed as a minimisation procedure of an energy functional with respect to the occupied HF or DFT orbitals.<sup>76–78</sup> After employing a particular basis set for discretisation, this minimisation problem becomes parameterized in the orbital coefficients **C** and the associated Euler-Lagrange equations may be written as

$$\begin{aligned}\mathbf{F}(\mathbf{C})\mathbf{C} &= \mathbf{SCE}, \\ \mathbf{C}^\dagger\mathbf{S}\mathbf{C} &= \mathbf{I},\end{aligned}\tag{3}$$

where **C** is the matrix of occupied orbital coefficients, **S** is the overlap matrix, **I** is the identity matrix, and **E** is the diagonal matrix of orbital energies. The Kohn-Sham or Fock matrix **F**(**C**) itself depends on the solution coefficients **C**, making (3) a non-linear eigenproblem. In the following, our focus will be on the HF problem since `molsturm` currently does not implement any DFT exchange-correlation functional. Due to the structural similarity of both HF and DFT, our approach nevertheless applies to DFT as well.

Because (3) is non-linear, the HF problem must be solved iteratively. Starting from an initial guess **C**<sup>(0)</sup>, the SCF procedure aims to construct a sequence of trial matrices **C**<sup>(1)</sup>, **C**<sup>(2)</sup>, ..., **C**<sup>(n)</sup> converging toward the minimiser of the HF energy functional. Broadly speaking this can be achieved in two ways, either by directly minimising the HF energy functional<sup>79–82</sup> or alternatively by satisfying (3),<sup>83–86</sup> thus obtaining a stationary point on the SCF manifold. On top of that one may alternatively formulate the HF problem, such that instead of the coefficient matrix, the density matrix

$$\mathbf{D}^{(n)} = \mathbf{C}^{(n)}(\mathbf{C}^{(n)})^\dagger\tag{4}$$

is iterated. To distinguish SCF algorithms according to these parameterizations, we will refer to the latter kind of SCF algorithms as *density-matrix-based SCF* schemes, whereas we will use the term *coefficient-based SCF* for the former set of algorithms.<sup>75</sup>

To illustrate, Fig. 1 shows simplified schemes for three SCF algorithms: Roothaan's repeated diagonalisation,<sup>83</sup> the optimal damping algorithm (ODA),<sup>77</sup> and the geometric direct minimisation (GDM) scheme.<sup>82</sup> While Roothaan's algorithm and the GDM are coefficient-based, the ODA is density-matrix-based. Roothaan's algorithm is the simplest representative for solving (3) by repetitively treating the standard eigenproblem that arises from fixing **F**(**C**<sup>(n)</sup>). In contrast to this, the GDM directly minimises the energy functional geometrically. The ODA is a middle ground: It combines a line-search minimisation of the energy with respect to the density matrix with repeated diagonalisation of arising Fock matrices.

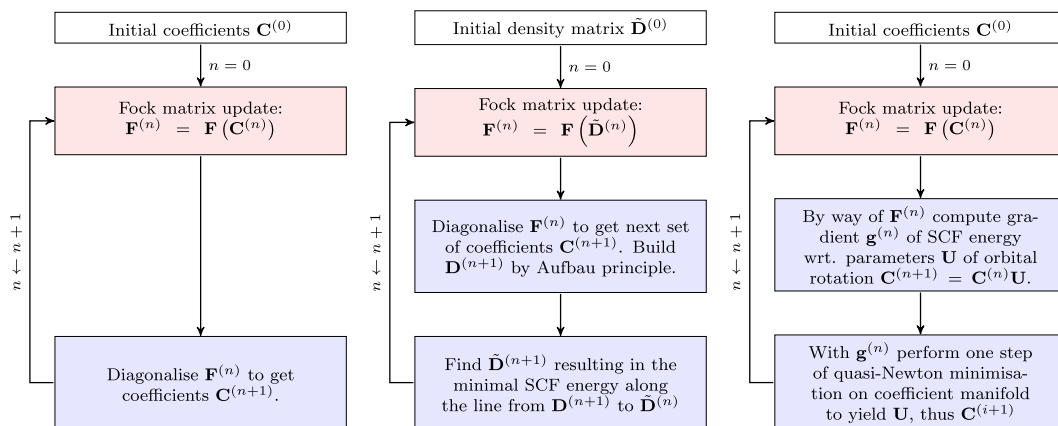


FIG. 1. Schematic overview of a few exemplary self-consistent field algorithms. In each case, the Fock-update step is highlighted in pale red and the steps updating the coefficients or the density matrix are shaded in pale blue. For further details regarding the algorithms, see the indicated references. From left to right: (a) Roothaan repeated diagonalisation.<sup>83</sup> (b) Optimal damping algorithm.<sup>77</sup> (c) Geometric direct minimisation.<sup>82</sup>

The aforementioned algorithms can—on an abstract level—be written as a two-step process, where a *Fock-update* step and a *coefficient-update* or *density-matrix-update* step are iterated. In the former step, a new Fock matrix  $\mathbf{F}^{(n)}$  is constructed from the current set of SCF coefficients  $\mathbf{C}^{(n)}$  or the current density matrix  $\mathbf{D}^{(n)}$  (red boxes in Fig. 1). In the second step, a new set of coefficients  $\mathbf{C}^{(n+1)}$  or a new density matrix  $\mathbf{D}^{(n+1)}$  is found by means of the Fock matrix  $\mathbf{F}^{(n)}$  (blue boxes in Fig. 1). Typically other results obtained in previous iterations are taken into account in this step as well to accelerate convergence. Consider, for example, Pulay’s commutator direct inversion of the iterative subspace (DIIS)<sup>86</sup> scheme forming a linear combination of previous Fock matrices.

Note, that apart from the initial discretisation of the HF or DFT problem, no reference to the basis function type was required in our discussion about SCF procedures. In other words, provided that (i) an SCF algorithm can be brought into two-step form and that (ii) within these steps the details of the basis function can be hidden, the algorithm can be implemented without making explicit reference to the underlying basis. We are unaware of an SCF algorithm which cannot be written in two-step form and will henceforth concentrate primarily on the second point

in our discussion towards a basis-type-independent SCF scheme.

## B. Matrix structure and contraction-based methods

Because different basis types can have very different selection rules and other discretisation properties, the structure of the Fock matrix—as well as the numerical approaches required to efficiently solve the SCF problem—may vary. This in turn affects the requirements we need for the interface to the update steps, which is a challenge for hiding the basis function details from the SCF. This subsection briefly discusses contraction-based methods as a solution to this issue.

Figure 2 shows from left to right the Fock matrices arising if (a) finite elements, piece-wise polynomials on a real-space grid,<sup>88,89</sup> (b) contracted Gaussians, or (c) Coulomb Sturmians have been employed as the basis. While the Coulomb Sturmian and Gaussian matrices are both small, dense, and diagonal dominant, the finite-element-discretised matrix is sparse, but also much larger. In fact, for a description of the beryllium atom density at a relevant accuracy even larger basis sets with  $10^5$ – $10^6$  finite elements are needed.<sup>53</sup>

As a result, for both cGTO as well as CS-based discretisations, direct diagonalisation algorithms, i.e., where the




FIG. 2. Structure of the Fock matrix at the beginning of a Hartree-Fock SCF calculation of beryllium, discretised using finite elements, contracted Gaussians, and Coulomb Sturmians. The elements are colored by the  $\log_{10}$ -scale shown on the right. The (3, 2, 2) CS basis set of calculation (c) contains all CS functions (1) whose quantum numbers  $n, l, m$  satisfy  $n \leq 3, l \leq 2$ , and  $m \leq 2$  and with the exponent parameter chosen as  $k = 2.0$ . More details about construction schemes for CS basis sets can be found in Ref. 75. From left to right: (a)  $Q_2$  finite elements, adaptively refined 3D grid. (b) Contracted Gaussians, pc-2 basis set.<sup>87</sup> (c) Coulomb Sturmians, (3,2,2) basis with  $k = 2.0$ .



full Fock matrix is diagonalised completely, are applicable. For finite-element-based approaches, on the other hand, iterative subspace-based algorithms like Arnoldi's method<sup>90</sup> or Davidson's method<sup>91</sup> are more common, due to the size of the matrix. Recently these have been combined with the so-called matrix-free methods,<sup>92</sup> which avoid building the finite-element problem matrix in memory at all. Instead only an *expression* for the computation of the matrix-vector product is passed to an iterative solver. Since such expressions may in general involve arbitrary tensor contractions, like a contraction over the ERI tensor to compute the Coulomb or the exchange part of the Fock matrix, we will refer to these approaches by the term *contraction-based methods*.<sup>75</sup>

The main driving force for such approaches is usually to reduce the amount of storage required and instead employ well-optimised high-throughput matrix-vector contraction expressions. For cases where this avoids slow storage such as hard drives, runtime may be reduced significantly even though matrix data will effectively be computed over and over. Examples for contraction-based methods in electronic structure theory are efficient implementations of Post-HF methods, like the algebraic diagrammatic construction (ADC) scheme<sup>93–95</sup> as well as modern coupled-cluster algorithms.<sup>96</sup> One should mention that in this context, the contraction expressions are usually called *working equations*.

Note, however, that contraction-based methods may even be favourable for cases where the size of the system matrix allows us to get around using the hard drive and place it in main memory instead.<sup>75</sup> This can be understood by considering modern hardware trends. State-of-the-art central processing units can perform on the order of 1000 floating point operations during the time needed to load data from main memory,<sup>97</sup> a number which is likely going to increase in the future.<sup>98</sup> This implies that more and more algorithms may become bound by memory latency and bandwidth rather than computation. Especially for cases where matrix elements are fast to compute from smaller stored intermediates or even from tractable analytic expressions, contraction-based methods are highly suitable. Additionally a contraction-based approach often allows us to reorder terms in the matrix-vector product or make use of discretisation-specific properties providing additional sources of reducing runtime cost. Such advantageous side-effects apply not only to finite-element-based HF, but to Coulomb-Sturmian-based Hartree-Fock<sup>27,75,92</sup> as well, making contraction-based methods worth considering in a context where multiple basis functions ought to be employed.

Additionally, a contraction-based approach can be readily combined with the two-step SCF described in Subsection III A. Focusing on a coefficient-based SCF for a second, a contraction-based ansatz would implement the Fock matrix  $\mathbf{F}(\mathbf{C})$  as a matrix expression with the current coefficients  $\mathbf{C}$  as some mutable state.<sup>75</sup> The Fock-update step then amounts to transparently replacing the current  $\mathbf{C}$  in the Fock matrix expression, which is a trivial process. Furthermore, both ways to think about the HF problem, namely, to think of it as a non-linear eigenproblem as well as an optimisation problem, can be tackled by iteratively solving appropriate linear problems or eigenproblems. This can be achieved using a wide

range of iterative subspace-based algorithms such as the generalised minimal residual method (GMRES), conjugate-gradient (CG), Arnoldi's method, or Davidson's method.<sup>90,91,99–101</sup> For a density-matrix-based SCF scheme, a contraction-based formulation is possible as well. Since the density matrix and the Fock matrix have similar memory requirements and the density matrix inevitably needs to be stored in a density-matrix-based SCF, the main prospect of contraction-based methods, to avoid the memory bottleneck of storing the Fock matrix, is directly subverted. Our discussion will not consider density-matrix-based SCF schemes further for this reason.

In principle, the Fock update step may be implemented by a conventional re-computation of the full Fock matrix like in a cGTO setting. Similarly, the matrix-vector-product expression may be realised by multiplying the resulting stored matrix with an appropriate vector. Thus the contraction-based SCF scheme is a generalisation of the traditional method that provides additional flexibility to deal with Fock matrices of various structures. To conclude, a single contraction-based interface between the SCF algorithm and Fock matrix is sufficient to provide a contraction-based SCF irrespective of the basis function type and resulting matrix structures.

One should notice, however, that iterative eigensolvers are not appropriate for all systems. For some cGTO discretisations with small and dense Fock matrices stored in memory, direct solver methods perform better than iterative ones and are thus preferable. Optimal performance requires an abstraction layer that—depending on the basis function type and matrix structure—transparently switches (i) between contraction expressions and dense matrices for representing the Fock matrix and (ii) between iterative or direct solver algorithms. As will be discussed in Sec. IV B, the *lazymat* lazy matrix library is used for this purpose to achieve a basis-type-independent SCF code. The details of the solver algorithm switching and the basis-dependent routines for computation are hidden in the abstraction layer of the linear algebra and the contraction expression.

## IV. PROGRAM DESIGN

### A. Design goals

As mentioned above, *molsturm* aims to remove the difficulty in implementing new types of basis functions and discretisations and to simplify experimenting with new computational methods in quantum chemistry. Assessment of new methods and comparison between old and new should be possible within the same framework to ensure treatment on an equal footing. A high-level interface aiding automation of repetitive calculations is desirable, too. Once the trial phase is completed, it should be easy to incorporate the new methods into existing quantum chemistry software and thus make them widely available. This motivates the overall design goals of *molsturm*.

#### 1. Enable rapid development

In the early stages of developing a new quantum-chemical method, it is often not clear how it will perform in practice or which approaches are required to yield efficient and stable algorithms. To simplify implementation, the code should

be high level and close to the physical formulae, and at the same time be flexible enough to enable experimentation with different numerical methods. Sec. IV D discusses further details.

## 2. Plug-and-play implementation of new discretisations

It is a significant barrier, to incorporate new basis function types and discretisation schemes to quantum chemistry, that assumptions about basis function types are scattered around in the—often very large—programs. We have designed *molsturm* to isolate this to the actual electronic integral back ends, and otherwise only where absolutely necessary. The SCF stage and post-HF methods should only know about integrals on an abstract level. Since symmetry, sparsity, selection rules, and recursion rules are basis function dependent, the integral library should be in charge of performing the operations where this information is used. This is primarily in the tensor contractions, for example, the ERI-contractions with molecular coefficients. Sec. IV C discusses how this is done, expanding on the strategies introduced in Sec. III.

## 3. Easy interfacing with existing code

A challenge for new quantum-chemical methods is that they are hard to compare to well-established ones: One is either restricted to toy problems, or faced with the enormous task of implementing advanced methods, refined over hundreds of man-years in the state-of-the-art quantum chemistry packages—clearly a rather daunting task. For this reason, it is explicitly *not* our goal to create yet another general purpose quantum chemistry package and the large ecosystem of functionality needed in such a project, but on the contrary to supply small flexible modules that can both be used on their own for experimentation and teaching, and can be easily incorporated into existing quantum chemistry software by simple interfaces. For details, see Sec. IV D as well as the examples in Sec. V.

## 4. Modular structure with low code complexity

The aspired flexibility requires that individual modules are as independent from each other as possible. We therefore choose a design in *molsturm*, where the five main modules are arranged in layers; see Fig. 3. Dependencies between the modules are only downwards, never sideways or upwards. This aids both reuse of *molsturm*'s modules in external projects as well as restructuring or replacement of the code if this was

required in the future. This is further aided by *molsturm*'s test suite, which employs a range of testing strategies including unit tests, functionality tests, and property-based tests.

Most *molsturm* modules are written in C++, but the top layer of the program is a *python* module defining the user interface of the framework. Below this, *gint*, the *general integral library*, provides a single link to the multiplex between the supported integral calculation back ends and *gscf* implements the contraction-based SCF schemes, following the general two-step update structure mentioned in Sec. III A. Both modules use the library *lazyten*, which defines a generalised linear algebra interface allowing to transparently use dense, sparse, and contraction-based Fock matrices in *gint* and *gscf*; see Sec. IV B. Finally *krims* is *molsturm*'s common utility library, named after the German word “*Krimskrams*” for “odds and ends.” The individual components are discussed further in Secs. IV B–IV D.

## B. Library for contraction-based algorithms

We saw in Sec. III B that contraction-based methods provide a versatile ansatz for self-consistent field algorithms, leading to a basis-function independent formulation of the problem. We noted that when the Fock matrix is small and should be stored in memory, an abstraction layer to switch between dense and iterative solver schemes is needed for maximal efficiency. On top of that an issue with contraction-based methods is that the expressions for computing the matrix-vector products can become complicated, such that these are less intuitive to handle compared to a plain matrix or tensor operations.

Inside *molsturm*, the library *lazyten* solves these challenges by representing matrices by a data structure called a *lazy matrix*.<sup>75</sup> These employ *lazy evaluation*, a rigorous method from programming language theory that allows postponing computation until the moment it is needed.<sup>102</sup>

In contrast to a *stored matrix*, which we define as a dense table, which has all its elements residing in a continuous chunk of memory, this restriction does no longer hold for a *lazy matrix*. It may, for example, follow a particular sparse storage scheme like a compressed-row format,<sup>103</sup> but it may not even be associated with any kind of storage at all. In the most general sense, it can be thought of as an arbitrary contraction expression for computing the matrix elements, which is

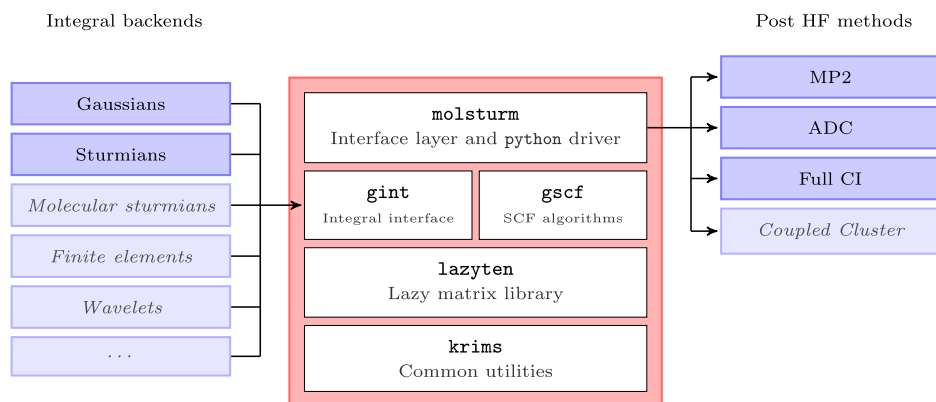


FIG. 3. Structure of the *molsturm* framework: Shown are the five modules of the package, along with third-party integral back ends and post-HF methods to indicate the mediator role of *molsturm*. The grayed-out parts are not yet implemented, but could be supported by the design.

dressed to look like an ordinary matrix from the outside. One may still obtain individual matrix elements and add, subtract, or multiply lazy matrix objects, but not all operations are as efficient as for stored matrices. Most importantly, accessing individual elements of lazy matrices can be costly since the elements may be computed, e.g., from a particular tensor contraction. However, contraction operations of such objects must be fast.

Lazy matrix operations are subject to *lazy evaluation*, explaining the name of these data structures. Lazy evaluation is a prominent concept of functional programming languages, excellently introduced in Ref. 104. In our context, this means that operations between lazy matrices are not directly performed, but delayed until a contraction of the resulting expression with a vector or a stored matrix unavoidably requires evaluation. To illustrate this, consider the instructions

$$\begin{aligned} \mathbf{D} &= \mathbf{A} + \mathbf{B}, \\ \mathbf{E} &= \mathbf{D}\mathbf{C}, \\ \underline{\mathbf{y}} &= \mathbf{E}\underline{\mathbf{x}}, \end{aligned} \quad (5)$$

where  $\mathbf{A}$ ,  $\mathbf{B}$ , and  $\mathbf{C}$  are lazy matrices and  $\underline{\mathbf{x}}$  is a vector stored in memory. The first two lines do not give rise to any computation. They only build an expression tree in the returned lazy matrix  $\mathbf{E}$ , as illustrated in Fig. 4. The final instruction is a matrix-vector product with the stored vector  $\underline{\mathbf{x}}$ , and the actual result should be returned in the vector  $\underline{\mathbf{y}}$ . This triggers the complete expression tree to be worked upon in appropriate order, such that the expression

$$\underline{\mathbf{y}} = (\mathbf{A} + \mathbf{B})\mathbf{C}\underline{\mathbf{x}} \quad (6)$$

is evaluated at once at this very instance. At this point, the full expression could be simplified, and the most efficient evaluation order could be chosen. This is, however, not implemented yet.

Due to lazy evaluation, we can thus build complicated expressions from familiar matrix operations. This allows us to view the lazy matrix framework as a domain-specific language for contraction-based algorithms, which makes working with contraction expressions feel like working with actual matrices. Note, however, that the lazy matrices generalise real matrices as they also allow non-linear transformations of a vector to be represented.

Inside `lazyten`, various kinds of lazy matrices, as well as lazy and stored objects, can be combined transparently.<sup>75</sup> Similarly, `lazyten` provides high-level interfaces for solving linear or eigenproblems<sup>75</sup> where the involved matrices may be stored or lazy. The call passes through a branching layer, which inspects the matrix structure and accordingly selects one

of the available third-party linear algebra back ends for solving the problem. As a result, algorithms programmed using `lazyten` may be called with both lazy and stored matrices, and the solvers will be automatically chosen to match the change in the matrix structure. The user, however, remains in full control: By providing appropriate parameters, all choices made automatically by `lazyten` can be overwritten, as can the parameters passed to the underlying solvers.

Overall, `lazyten` provides intuitive high-level syntax for contraction-based methods in the form of lazy matrices. The library allows the algorithm code to be written only once, but to stay flexible. For example, one may adapt to modern hardware trends or to the deviating numerical requirements imposed by a different basis function type simply by changing the implementation of the lazy matrices passed to the algorithm code. A more in-depth discussion of `lazyten` can be found in Ref. 75.

### C. Self-consistent field methods and integral interface

The lazy matrices of `lazyten` are constructed to be used as a high-level language for implementing basis-type-independent contraction-based SCF algorithms.<sup>75</sup> For example, the linear algebra interfaces of `lazyten` can be employed in the coefficient-update step making implicit use of the automatic switching between dense and iterative diagonalisation methods. The Fock update may be implemented building on top of a similar function from `lazyten`, namely, by altering the coefficient matrix the Fock expression currently refers to.

All the SCF algorithms in `gscf` are implemented in this way: as solvers for a non-linear eigenproblem defined by an input lazy matrix, which represents the SCF problem under study. Since the algorithms only see the final contraction expression and its update function, `gscf` is self-contained and may be applied to *any* non-linear eigenproblem with a structure similar to the HF minimisation problem. This is desirable because quite a few electronic structure theory methods can be thought of as modifications of the HF problem. Examples include the Kohn-Sham matrix arising in the usual density-functional theory (DFT) treatments or additional terms in the problem matrix, arising from modeling an external field, or correction terms due to embedding.

The lazy Fock matrix object describing the problem to be solved is prepared by the upper `molsturm` layer based on the electronic structure method chosen by the user. For example, HF would be the sum of four lazy matrices, which represent the kinetic, nuclear attraction, Coulomb, and exchange matrix.<sup>96</sup> Similarly one would add an exchange-correlation term for DFT or other terms such as an embedding operator. The latter methods are not yet available in `molsturm`, however.

The individual terms of the Fock matrix are obtained from `gint`, which acts as a broker, presenting a common interface for all basis function types and third-party integral back end libraries to the rest of the `molsturm` ecosystem. On calculation start, `molsturm` will take the discretisation parameters supplied by the user and hand them over to `gint`, which—based on these parameters—sets up the selected integral back end library and returns a collection of lazy matrix integral objects. For each basis type and back end, the interface of the returned objects will thus look alike since they are all lazy




FIG. 4. Examples for lazy matrix expression trees. The upper panel represents the instruction  $\mathbf{D} = \mathbf{A} + \mathbf{B}$  and the lower one the multiplication of the result  $\mathbf{D}$  with  $\mathbf{C}$ .

matrices. On call to their respective contraction expressions, however, the required computations will be invoked in the previously selected integral back end. `gint` itself does not implement any routine for computing integral values at all, it just transparently redirects the requests. Notice that the precise kind of parameters needed by `gint` to setup the back end may well vary from discretisation to discretisation. For example, a Coulomb Sturmian basis requires the Sturmian exponent  $k$  and the selection of  $(n, l, m)$ -triples of the basis functions,<sup>75</sup> whereas a contracted Gaussian basis requires the list of angular momentum, exponents, and contraction coefficients.<sup>2,4,5</sup>

At the moment cGTO and CS integrals are the only ones supported in `gint`. For each of these, at least two different implementations are available, however. Adding more back end libraries or basis function types is rather easy since one only needs to implement a collection of lazy matrices, where the contraction expressions initiates the appropriate integral computations in the back end. This collection then needs to be registered as a valid basis type to `gint` to make it available to the upper layers. For example, preliminary support for the contracted Gaussian library `libcint`<sup>105</sup> was added with just two days of work. Notice, that `gint` is designed to allow all of this to be achieved without changing a single line of code inside `gint` itself since the call to the registration function can happen dynamically at runtime. So one can implement a new integral back end in a separate shared library and add it in a plug-in fashion without recompiling `molsturm`.

To summarise, by means of the lazy matrices of `lazyten`, the responsibility for the HF problem has been split between three different well-abstracted modules: `gint`, which provides the interface to the integrals and selects the discretisation, `molsturm`, which builds the lazy matrix expression of the problem to be solved, and `gscf` which uses this expression to solve the SCF problem in a basis-type-independent manner.

#### D. Python interface module

The topmost layer of `molsturm` is the “`molsturm`” `python` module, providing the user interface of the package. This layer assists with setting up a calculation, drives the SCF procedure in `gscf`, and returns the converged results to the user. We chose the scripting language `python` to implement the majority of this interface layer and especially the interface itself.

Our reasoning is related to the arguments of the `pyscf` authors<sup>67</sup> discussed in Sec. II, namely, we wanted to avoid inventing yet another “input format” and “output format” for quantum-chemical calculations. Instead, calculations in `molsturm` can be initiated cleanly and flexibly directly from a host `python` script, which can additionally be used for subsequent analysis. This not only implies that all of `python` and its libraries are available for the calculation setup and analysis, but also that no explicit parsing of program output is required for analyzing the results. This lowers the barriers for people who are new to the field since they do not need to learn both how to write input files that define calculations, as well as the syntax of a scripting language for parsing results. More subtly, the output formats of quantum-chemistry programs

change from time to time, breaking parser scripts or—even worse—producing wrong results without any notice. This is a common problem in the current practice of computational chemistry.

In contrast to this, the SCF results in `molsturm` are returned to the host `python` environment through an interface built on `numpy` arrays. These have become the *de facto* standard for storing and manipulating matrices or tensors in `python`. All `python` packages that are commonly used for plotting or data analysis, such as `matplotlib`,<sup>106</sup> `scipy`,<sup>73,107</sup> or `pandas`,<sup>108</sup> use `numpy` arrays in their interfaces. Consequently, a complete computational procedure may be orchestrated from a single `python` script, which contains all parameters influencing the setup, calculation, analysis, and all decisions taken for presenting the data in plots or tables. Such a script serves as automatic documentation for the full procedure and allows others to reproduce the presented plots or tables without effort: All it takes is to re-run the script.

All parameters for `gint`, the SCF algorithms of `gscf`, and the linear solvers from `lazyten` are made available through the `python` interface. By changing these, the user may directly influence, e.g., the algorithms chosen by `lazyten` for diagonalisation, or how `gscf` switches between SCF solvers. This is particularly handy during method development, where one may run `molsturm` from an interactive IPython<sup>109</sup> shell and use these parameters to control the progress of a calculation. In that way, one can check assertions about intermediate results or visualise such graphically with `matplotlib`.<sup>106</sup> This greatly reduces the feedback loop for small calculations, e.g., during debugging.

Interactive analysis of larger calculations is facilitated by archiving functionality in `molsturm`. The SCF results may be stored either in YAML<sup>110</sup> or HDF5<sup>74</sup> format. In this way, large calculations can be performed in advance over night or on a high-performance computing (HPC) system, then archived and transferred to the workstation. Here, the archive may be loaded in an interactive shell, restoring the full state of the calculation as if it had been performed locally. Next to the SCF results, `molsturm`’s archived state contains the precise set of input parameters which were used to obtain the stored results. These are *not* the parameters provided by the user to start the calculation, but the post-processed parameters which were actually used by the lower layers, including, e.g., default values. This helps make the archive self-documenting, and simplifies setting up a refined calculation building on top of the already obtained results.

Our `numpy`-based interface has already proven to be helpful for linking `molsturm` to other third-party quantum chemistry codes: It allowed us to link `molsturm` to the `python` interfaces of `pyscf`,<sup>67</sup> as well as `adcm`<sup>94</sup> in only a few days. As a result, FCI as well as calculations for computing excited states by the algebraic diagrammatic construction (ADC) scheme<sup>111,112</sup> may be started on top of `molsturm`’s SCF using the respective aforementioned packages. By way of interface generators like Simplified Wrapper and Interface Generator (SWIG),<sup>113</sup> `numpy` arrays can be automatically converted to plain C arrays, such that third-party packages consisting only of low-level C++, C, or FORTRAN code can be linked to `molsturm` in the future.



We see that the interface of `molsturm` not only facilitates rapid development of new algorithms, but also relieves one from the need to re-invent the wheel, i.e., to implement standard quantum-chemical methods over and over for each new basis function type. Instead, existing functionality in external packages can be quickly leveraged for one's own purposes. The aspects described in this section will be demonstrated with practical examples in Sec. V.

## V. EXAMPLES

In this section, we present two examples that demonstrate how the `python` interface of `molsturm` can be combined with existing features of `python` in order to analyze the results or to extend the capabilities of `molsturm`. We concentrate our discussion on molecular computations with contracted Gaussian basis sets. It should be stressed again, however, that due to the basis-function independent nature of `molsturm`, the procedures outlined in the scripts could be easily performed with other types of basis functions as well.

In fact, the design of `molsturm` assures that the discretisation details can be selected at a high level, without affecting the code that performs the actual computation and analysis. Section V B gives an example for which the choice of the basis type is made in the `main` function of the script. This ensures that a script performing a particular modeling task can be easily used as a template for a systematic study of the effect of changing basis function type or integral implementation: All it takes is to iterate over the appropriate list of discretisation parameters and call the calculation for each instance. This greatly simplifies testing a novel basis function type, which

has just become available in `gint`, as well as comparing it to existing methods subject to the test case provided by a script.

### A. Coupled-cluster doubles (CCD)

This example shows how one can extend `molsturm` with novel methods using its high-level `python` interface together with standard functionality from `python/numpy`.<sup>73,107</sup>

Even though `molsturm` right now neither offers any coupled-cluster method nor an interface to any third-party coupled-cluster code, we managed to implement a simple working coupled-cluster doubles (CCD)<sup>114,115</sup> algorithm in only about 100 lines of code and about two days of work, including the time needed for research about the method and the computational procedures. The most relevant part of the implementation, namely, computing the CCD residual for the current guess of the  $T_2$  amplitudes  $t_{ij}^{ab}$ , is shown toward the right of Fig. 5, side-by-side with the expression of the CCD residual.<sup>114</sup> The full CCD code is available as an example in the file `examples/state_interface/coupled_cluster_doubles.py` of the `molsturm` repository.<sup>116</sup> We follow the standard procedure of employing a quasi-Newton minimisation of the CCD residual with respect to the  $T_2$  amplitudes using the orbital energy differences as an approximate Jacobian.<sup>96,114</sup> The guess for the  $T_2$  amplitudes is taken from second order Møller-Plesset perturbation theory (MP2).

The expression of the CCD residual  $r_{ij}^{ab}$  requires the evaluation of a sequence of tensor contractions involving the Fock matrix in the molecular orbital basis,  $f_{pq}$ , the antisymmetrised electron-repulsion integrals,  $\langle pq|lr\rangle$ , as well as the current guess for the  $T_2$  amplitudes,  $t_{ij}^{ab}$ . As usual, we employ the

$$\begin{aligned}
 r_{ij}^{ab} = & \langle ab||ij\rangle \\
 & + \sum_e f_{ae} t_{ij}^{eb} - \sum_e f_{be} t_{ij}^{ea} - \sum_m f_{mi} t_{mj}^{ab} + \sum_m f_{mj} t_{mi}^{ab} \\
 & + \frac{1}{2} \sum_{mn} \langle mn||ij\rangle t_{mn}^{ab} + \frac{1}{2} \sum_{ef} \langle ab||ef\rangle t_{ij}^{ef} \\
 & + \sum_{me} \langle mb||ej\rangle t_{im}^{ae} - \sum_{me} \langle mb||ei\rangle t_{jm}^{ae} \\
 & - \sum_{me} \langle ma||ej\rangle t_{im}^{be} + \sum_{me} \langle ma||ei\rangle t_{jm}^{be} \\
 & - \frac{1}{2} \sum_{mnef} \langle mn||ef\rangle t_{mn}^{af} t_{ij}^{eb} + \frac{1}{2} \sum_{mnef} \langle mn||ef\rangle t_{mn}^{bf} t_{ij}^{ea} \\
 & - \frac{1}{2} \sum_{mnef} \langle mn||ef\rangle t_{in}^{ef} t_{mj}^{ab} + \frac{1}{2} \sum_{mnef} \langle mn||ef\rangle t_{jn}^{ef} t_{mi}^{ab} \\
 & + \frac{1}{4} \sum_{mnef} \langle mn||ef\rangle t_{mn}^{ab} t_{ij}^{ef} + \frac{1}{2} \sum_{mnef} \langle mn||ef\rangle t_{im}^{ae} t_{jn}^{bf} \\
 & - \frac{1}{2} \sum_{mnef} \langle mn||ef\rangle t_{jm}^{ae} t_{in}^{bf} - \frac{1}{2} \sum_{mnef} \langle mn||ef\rangle t_{im}^{be} t_{jn}^{af} \\
 & + \frac{1}{2} \sum_{mnef} \langle mn||ef\rangle t_{jm}^{bf} t_{in}^{af}
 \end{aligned}$$

```

eri_phys = state.eri.transpose((0, 2, 1, 3))
eri = eri_phys - eri_phys.transpose((1, 0, 2, 3))
...
oooo = eri.block("oooo"); vvvv = eri.block("vvvv")
oovv = eri.block("oovv"); ovvo = eri.block("ovvo")
res = (
+ np.einsum("abij->iajb", eri.block("vvoo"))
+ np.einsum("ae,iejb->iajb", state.fock.block("vv"), t2)
- np.einsum("be,ieja->iajb", state.fock.block("vv"), t2)
- np.einsum("mi,majb->iajb", state.fock.block("oo"), t2)
+ np.einsum("mj,maib->iajb", state.fock.block("oo"), t2)

+ 0.5 * np.einsum("mnij,manb->iajb", oooo, t2)
+ 0.5 * np.einsum("abef,iejf->iajb", vvvv, t2)
+ np.einsum("mbej,iame->iajb", ovvo, t2)
- np.einsum("mbei,jame->iajb", ovvo, t2)
- np.einsum("maej,ibme->iajb", ovvo, t2)
+ np.einsum("maei,jbme->iajb", ovvo, t2)

- 0.5 * np.einsum("mnef,manf,iejb->iajb", oovv, t2, t2)
+ 0.5 * np.einsum("mnef,mbnf,ieja->iajb", oovv, t2, t2)
- 0.5 * np.einsum("mnef,ienf,majb->iajb", oovv, t2, t2)
+ 0.5 * np.einsum("mnef,jenf,maib->iajb", oovv, t2, t2)
+ 0.25 * np.einsum("mnef,manb,iejf->iajb", oovv, t2, t2)
+ 0.5 * np.einsum("mnef,iame,jbnf->iajb", oovv, t2, t2)
- 0.5 * np.einsum("mnef,jame,ibnf->iajb", oovv, t2, t2)
- 0.5 * np.einsum("mnef,ibme,janf->iajb", oovv, t2, t2)
+ 0.5 * np.einsum("mnef,jbme,ianf->iajb", oovv, t2, t2)
)

```

FIG. 5. Equation for the coupled-cluster doubles (CCD) residual<sup>114</sup> on the left and excerpt of a CCD implementation using `molsturm` and `numpy` on the right. Equivalent quantities are highlighted in the same colour. The first two lines of code show the computation of the antisymmetrised electron repulsion integrals from the `state.eri` object obtainable from `molsturm`, which is carried out once at the beginning of the algorithm. The remaining lines compute the residual for a particular  $T_2$  amplitude stored in the tensor object `t2`.

convention that indices  $i, j, k, l, \dots$  refer to occupied orbitals, indices  $a, b, c, d, \dots$  to virtual (i.e., unoccupied) orbitals, and indices  $p, q, r, s, \dots$  to either kind of orbitals.

The python implementation (right-hand side of Fig. 5) computes those contractions. For this, it employs the data structures `molsturm` provides in the `state` object, which is returned by the SCF procedure. Our code uses chemists' indexing convention in the electron-repulsion integrals object `state.eri`. The CCD equations, however, are written using the antisymmetrised electron-repulsion integrals  $\langle pq||rs \rangle$ . Therefore the first two lines of the code of Fig. 5 are executed once to perform the antisymmetrisation. The subsequent lines are executed once per CCD iteration and compute the residual tensor `res` by contracting the relevant blocks of the Fock matrix `state.fock`, the `eri` object, and the  $T_2$  amplitudes contained in `t2`. This is implemented using the `einsum` method from `numpy`, which performs tensor contractions expressed in Einstein summation convention. Note how the interplay of `numpy` with the data structures `molsturm` results in a strikingly close resemblance of implementation and actual equation.

The `state` object provides access to more quantities from the SCF procedure than just the Fock matrix and the




FIG. 6. Density plot of the final optimised H<sub>2</sub>O Hartree-Fock geometry with a O–H bond length of 0.95046 Å and a H–O–H bond angle of 106.35°. A geometry optimisation in ORCA<sup>118</sup> employing the same basis set agrees with this result within the convergence tolerance of  $10^{-5}$ .

repulsion integrals. Individual terms of the Fock matrix or quantities like the overlap matrix in terms of the underlying discretisation basis functions may be obtained as well. We provide this data as `numpy` arrays extended with extra

```

6 def optimize_h2o(rHO_guess, angHO_guess, conv_tol,
7                 **params):
8     # Function which computes the cartesian geometry
9     # from r and theta.
10    def geometry(r, theta):
11        rad = theta / 180 * np.pi
12        pos = (r * np.cos(rad), r * np.sin(rad), 0)
13        return molsturm.System(["O", "H", "H"],
14                               [(0, 0, 0), (r, 0, 0), pos])
15
16    def objective_function(args):
17        system = geometry(*args)
18        ret = molsturm.hartree_fock(
19            system, conv_tol=conv_tol/100, **params,
20        )
21        return ret["energy_ground_state"]
22
23    guess = (rHO_guess, angHO_guess)
24    res = scipy.optimize.minimize(
25        objective_function, guess, tol=conv_tol,
26        method="Powell"
27    )
28    return res.x[0], res.x[1]
29
30 def main():
31     r = 2.0          # O-H    radius guess (in au)
32     theta = 120      # H-O-H  angle guess
33
34     # First a crude optimisation with sto-3g
35     r, theta = optimize_h2o(r, theta, conv_tol=5e-4,
36                             basis_type="gaussian",
37                             basis_set_name="sto-3g")
38
39     # Then a more fine optimisation with def2-sv(p)
40     r, theta = optimize_h2o(r, theta, conv_tol=1e-5,
41                             basis_type="gaussian",
42                             basis_set_name="def2-sv(p)")
43     print("optimal H-O bond length (au): ", r)
44     print("optimal H-O-H bond angle:     ", theta)

```

FIG. 7. Example for performing a gradient-free optimisation using Powell's method<sup>119,120</sup> and `molsturm`. python import statements at the top of the script and the explicit call to `main` are skipped.

functionality to simplify implementation of Post-HF quantum-chemical methods, such that the user can employ the SCF results freely and flexibly within the `python` ecosystem. Coupled with the basis-function independence of `molsturm`'s SCF, this allows for rapid development and systematic investigation of Post-HF methods based on arbitrary basis functions.

At the moment, we make no efforts to exploit symmetry or parallelise the computation of the tensor contractions shown in the script of Fig. 5. For this reason, such implementations are not suitable for real-world applications. Nevertheless, the script presented in Fig. 5 may be used for CCD calculations of small molecules with small basis sets. For example, an  $\text{O}_2$  6-31G<sup>117</sup> calculation on a recent laptop took about an hour to converge up to a residual  $l_\infty$ -norm of  $10^{-4}$ . For investigating new methods on top of the `molsturm` framework, or to provide a flexible playground for teaching Post-HF methods to students, such scripts are therefore still well-suited.

## B. Gradient-free geometry optimisation

In order to make a novel basis function type properly accessible to the full range of quantum-chemical methods, a daunting amount of integral routines and computational procedures need to be implemented. For assessing the usefulness of a new discretisation method, it is, however, important to be able to quickly investigate its performance with respect to as many problems as possible. Undoubtedly, a very important application of computational chemistry is structure prediction, i.e., geometry optimisation. Performing such calculations requires the appropriate integral derivatives for the chosen basis function type. Since implementing these in the integral library can be as difficult as implementing the integrals required for the SCF scheme itself, one would much rather skip this step at first, and concentrate only on what is required for the SCF.

This example extends `molsturm` with a gradient-free geometry-optimisation procedure, implemented with building blocks readily available from `python`. This sidesteps the need for nuclear derivatives on the side of the integral library and facilitates simple structure optimisations, even without nuclear gradients—neither analytical nor numerical.

Figure 6 shows the script, which performs a geometry optimisation of water based on Powell's gradient-free optimisation algorithm<sup>119,120</sup> as implemented in the `scipy` library.<sup>73,107</sup> The optimal structure is found in a two-step procedure. First, a cheap STO-3G<sup>2</sup> basis set is used to obtain a reasonable guess. Then, the final geometry is found by minimising to a lower convergence threshold in the more costly def2-SV(P)<sup>121</sup> basis.

The time required to code the script was only about 30 min, showing the great power of a flexible design. Nevertheless, convergence to the equilibrium geometry shown in Fig. 7 was achieved in a couple of minutes. In line with what was discussed above, a novel basis function type, for which one just implemented the SCF integrals in `gint`, can be directly used for geometry optimisations. Only the discretisation parameters need to be changed, in lines 36 and 41 of the outermost `main` function.

## VI. CURRENT STATE AND FUTURE OF MOLSTURM

After about two years of development, `molsturm` allows us to solve the Hartree-Fock (HF) equations basis-function independently, following a contraction-based self-consistent field (SCF) ansatz. All aspects of the calculation, including the diagonalisation algorithm and the basis function type of the discretisation, may be fully controlled via a `python` module. This module integrates well into the existing `python` ecosystem, simplifying repetitive calculations as well as analysis of obtained results.

At present, `molsturm`'s integral library `gint` supports calculations employing either Coulomb Sturmians or contracted Gaussian basis functions, both in multiple implementations. For contracted Gaussians, the third-party `libint`<sup>122,123</sup> or `libcint`<sup>105</sup> libraries can be used, and Coulomb Sturmians are available via our own `sturmint`<sup>124</sup> library. In the future, we plan to add support for further basis function types in `gint` and `molsturm`, in particular, molecular and generalised Sturmians.<sup>25–28</sup>

Via `gscf`, multiple SCF schemes are available, namely, Roothaan's repeated diagonalisation,<sup>83</sup> Pulay's commutator direct inversion of the iterative subspace (DIIS),<sup>86</sup> and the truncated optimal damping algorithm (tODA),<sup>75</sup> an approximation of the optimal damping algorithm,<sup>77</sup> which is more suitable for the contraction-based interface of `gscf`. During the SCF procedure, `molsturm` automatically switches between the available schemes, trying to balance the convergence rate and expense of the individual algorithms.

Once an SCF computation has finished, the results can be archived in either in YAML<sup>110</sup> plain text or in HDF5<sup>74</sup> binary files. Such an archive not only contains the full final state of the calculation but also the precise parameters which were used in the SCF procedure, making the archive file self-documenting.

For treating electron correlation, `molsturm` only implements second order Møller-Plesset perturbation theory (MP2). Further methods, however, can be easily called via interfaces to third-party libraries. Full configuration interaction (FCI) is available via `pyscf`, and a range of excited states-methods based on the algebraic diagrammatic construction (ADC) scheme via `adcmn`,<sup>94</sup> namely, ADC(1), ADC(2), ADC(2)-x,<sup>111</sup> and ADC(3).<sup>112</sup>

The extension of `molsturm` to other methods or packages is easily accomplished by `molsturm`'s `python` interface; see Sec. V. Along these lines, closer integration with projects such as `pyscf`,<sup>67</sup> `Psi4`,<sup>71</sup> or `Psi4NumPy`<sup>72</sup> could be promising since these already provide high-level `python` interfaces to many state-of-the-art Post-HF methods. In this way, configuration-interaction, coupled-cluster, multi-configurational self-consistent field, or density matrix renormalisation group approaches could be used directly from `molsturm`'s SCF. With manageable development time, these methods would thus become available for all basis function types implemented in `gint`. Similarly, the extension of `molsturm`'s SCF toward Kohn-Sham density-functional theory is possible employing third-party libraries such as `libxc`<sup>125</sup> or `xcfun`<sup>126</sup> for computing the required exchange-correlation integrals.

As discussed in Sec. IV B, a contraction expression inside `molsturm`'s SCF is evaluated whenever the Fock matrix is applied to a trial vector. This proceeds by working on the expression tree, which represents the Fock matrix. In `lazyten`, the corresponding computations are right now neither parallelised, nor are symmetries or repetitive terms in the expression tree exploited. This currently limits the applicability of `molsturm`'s contraction-based SCF to small basis set sizes. Both automatic parallelisation of linear algebra expressions and finding optimal evaluation schemes for expression trees is ongoing research,<sup>127–134</sup> however. By integrating such efforts into `lazyten`, a direct improvement of `molsturm`'s SCF could be achieved without changing any other code.

## VII. DISCUSSION AND CONCLUSION

Implementation of quantum-chemical methods using novel types of basis functions often requires unusual numerical techniques as well. Implementing these into existing quantum chemistry packages can be a large task, as these are highly optimised for the methods they already accommodate and are typically not flexible enough to meet other requirements.

The `molsturm` framework presented here tries to fill this gap by providing a light-weight package designed with a range of different basis functions in mind. The key ingredient to reach the necessary flexibility is a contraction-based self-consistent field (SCF) scheme, which we employ for solving the Hartree-Fock problem. In a contraction-based ansatz, the numerical algorithms are formulated without requiring any explicit reference to the Fock matrix memory. Instead, the SCF iterations are driven by contractions of the Fock matrix with other vectors. The details how this matrix-vector product is computed can be varied flexibly, matching the numerical properties of the discretisation at hand. In this way, we have reached a design where the code for the SCF algorithm is separated from the code performing the linear algebra computation. Thus, changes to the SCF scheme can be made without affecting other parts of `molsturm` and the SCF code itself becomes basis-function independent.

On top of that, the interfaces of our SCF are easy-to-use and readily extensible. This allows quick incorporation of functionality of third-party packages and extensions of `molsturm` in ways we as the authors would have never thought of. Right now, `molsturm` may be used to perform calculations based on contracted Gaussians<sup>2</sup>—using the integral libraries `libint`<sup>122,123</sup> or `libcint`<sup>105</sup>—and based on Coulomb Sturmians<sup>20,21</sup>—using `sturmint`.<sup>124</sup> Selected Post-HF methods from `pyscf`<sup>67</sup> as well as excited states methods from `adcm`<sup>94</sup> are available on top. Extending the set of basis function types available inside `molsturm` can be achieved in a plug-and-play fashion, namely, by implementing a single well-defined interface class in our integral library `gint`. Thereafter such basis functions are available for the full `molsturm` ecosystem including the Post-HF methods provided by the third-party libraries mentioned above.

The abilities of `molsturm` have been demonstrated by two practical examples with particular emphasis on the way our `python` interface integrates with existing `python` packages. We showed how to aid repetitive calculations, implement

novel quantum-chemical methods, or rapidly amend functionality in a preliminary way, where a proper implementation would be much more involved. We hinted how systematic comparisons with established basis functions as well as subsequent graphical analysis is convenient to perform by the means of our readily scriptable interface. We hope that in this manner, `molsturm` will be a useful package to rapidly try novel basis function types and get a feeling for their range of applicability.

## ACKNOWLEDGMENTS

The authors express their thanks to Adrian L. Dempwolff and Maximilian Scheurer for many fruitful discussions during the preparation of the work. Michael F. Herbst gratefully acknowledges funding by the Heidelberg Graduate School of Mathematical and Computational Methods for the Sciences (GSC220). Last but not least the authors wish to commemorate their former collaborator, supervisor, and friend Dr. Michael Wormit, whose ideas have survived in this project as well as in countless others.

- <sup>1</sup>S. F. Boys, *Proc. R. Soc. London, Ser. A* **200**, 542 (1950).
- <sup>2</sup>W. J. Hehre, R. F. Stewart, and J. A. Pople, *J. Chem. Phys.* **51**, 2657 (1969).
- <sup>3</sup>J. D. Morgan, *J. Phys. A: Math. Gen.* **10**, L91 (1977).
- <sup>4</sup>F. Jensen, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **3**, 273 (2013).
- <sup>5</sup>J. G. Hill, *Int. J. Quantum Chem.* **113**, 21 (2013).
- <sup>6</sup>T. Kato, *Commun. Pure Appl. Math.* **10**, 151 (1957).
- <sup>7</sup>W. M. C. Foulkes, L. Mitás, R. J. Needs, and G. Rajagopal, *Rev. Mod. Phys.* **73**, 33 (2001).
- <sup>8</sup>A. Ma, M. D. Towler, N. D. Drummond, and R. J. Needs, *J. Chem. Phys.* **122**, 224322 (2005).
- <sup>9</sup>P.-F. Loos, A. Schemama, and M. Caffarel, “Electron-nucleus cusp corrections for molecular orbitals,” <http://www.irsamc.upstlse.fr/loos/poster/o24.pdf>, 2017, cMMSE 2017.
- <sup>10</sup>H. Feshbach, *Ann. Phys.* **5**, 357 (1958).
- <sup>11</sup>H. Feshbach, *Ann. Phys.* **19**, 287 (1962).
- <sup>12</sup>U. V. Riss and H.-D. Meyer, *J. Phys. B: At., Mol. Opt. Phys.* **26**, 4503 (1993).
- <sup>13</sup>R. Santra and L. S. Cederbaum, *Phys. Rep.* **368**, 1 (2002).
- <sup>14</sup>M. Güell, J. M. Luis, M. Solà, and M. Swart, *J. Phys. Chem. A* **112**, 6384 (2008).
- <sup>15</sup>P. E. Hoggan, in *Self-Organization of Molecular Systems: From Molecules and Clusters to Nanotubes and Proteins*, edited by N. Russo, V. Y. Antonchenko, and E. S. Kryachko (Springer Netherlands, Dordrecht, 2009), pp. 199–219.
- <sup>16</sup>F. Harris and H. Michels, *Adv. Chem. Phys.* **13**, 205 (1967).
- <sup>17</sup>E. Steinborn, in *Methods in Computational Molecular Physics*, edited by G. Dierksen and S. Wilson (Reidel, 1983).
- <sup>18</sup>E. Weniger and E. Steinborn, *J. Chem. Phys.* **78**, 6121 (1983).
- <sup>19</sup>*International Conference on ETO Multicenter Integrals*, edited by C. Weatherford and H. Jones (Reidel, Dordrecht, 1982).
- <sup>20</sup>H. Shull and P.-O. Löwdin, *J. Chem. Phys.* **30**, 617 (1959).
- <sup>21</sup>M. Rotenberg, in *Advances in Atomic and Molecular Physics*, edited by D. Bates and I. Esterman (Academic Press, 1970), Vol. 6, pp. 233–268.
- <sup>22</sup>A. Vincenzo, C. Andrea, and C. Simonetta, *Int. J. Quantum Chem.* **92**, 99 (2003).
- <sup>23</sup>C. Coletti, D. Calderini, and V. Aquilanti, *Adv. Quantum Chem.* **67**, 73 (2013).
- <sup>24</sup>D. Calderini, S. Cavalli, C. Coletti, G. Grossi, and V. Aquilanti, *J. Chem. Sci.* **124**, 187 (2012).
- <sup>25</sup>J. S. Avery and J. E. Avery, *Mol. Phys.* **110**, 1593 (2012).
- <sup>26</sup>J. E. Avery, *Adv. Quantum Chem.* **67**, 129 (2013).
- <sup>27</sup>J. E. Avery and J. S. Avery, *Adv. Quantum Chem.* **70**, 265 (2015).
- <sup>28</sup>J. E. Avery and J. S. Avery, *Adv. Quantum Chem.* **76**, 133 (2017).
- <sup>29</sup>V. Aquilanti, S. Cavalli, C. Coletti, and G. Grossi, *Chem. Phys.* **209**, 405 (1996).
- <sup>30</sup>V. Aquilanti, S. Cavalli, and C. Coletti, *Chem. Phys.* **214**, 1 (1997).
- <sup>31</sup>V. Aquilanti, S. Cavalli, and C. Coletti, *Phys. Rev. Lett.* **80**, 3209 (1998).



- <sup>32</sup>J. S. Avery and J. E. Avery, *Adv. Quantum Chem.* **43**, 185 (2003).
- <sup>33</sup>J. S. Avery, J. E. Avery, V. Aquilanti, and A. Caligiana, *Adv. Quantum Chem.* **47**, 157 (2004).
- <sup>34</sup>J. E. Avery and J. S. Avery, *Generalized Sturmians and Atomic Spectra* (World Scientific, 2006).
- <sup>35</sup>J. E. Avery and J. S. Avery, *J. Math. Chem.* **46**, 164 (2009).
- <sup>36</sup>D. Calderini, C. Coletti, G. Grossi, and V. Aquilanti, in *Computational Science and Its Applications—ICCSA 2013*, edited by B. Murgante, S. Misra, M. Carlini, C. M. Torre, H.-Q. Nguyen, D. Taniar, B. O. Apduhan, and O. Gervasi (Springer, Berlin, Heidelberg, 2013), pp. 32–45.
- <sup>37</sup>A. Abdouraman, A. Frapiccini, A. Hamido, F. Mota-Furtado, P. O'Mahony, D. Mitnik, G. Gasaneo, and B. Piroux, *J. Phys. B: At., Mol. Opt. Phys.* **49**, 235005 (2016).
- <sup>38</sup>J. S. Avery, *Hyperspherical Harmonics: Applications in Quantum Theory* (Springer, 1989).
- <sup>39</sup>V. Aquilanti and S. Cavalli, in *Few-Body Problems in Physics*, edited by C. Ciofi degli Atti, E. Pace, G. Salmè, and S. Simula (Springer Vienna, Vienna, 1992), pp. 573–580.
- <sup>40</sup>V. Aquilanti, A. Lombardi, and R. G. Littlejohn, *Theor. Chem. Acc.* **111**, 400 (2004).
- <sup>41</sup>J. E. Avery and J. S. Avery, *Hyperspherical Harmonics and Their Physical Applications* (World Scientific, 2018).
- <sup>42</sup>T. K. Das, *Hyperspherical Harmonics Expansion Techniques* (Springer, 2016).
- <sup>43</sup>J. M. Randazzo, L. U. Ancarani, G. Gasaneo, A. L. Frapiccini, and F. D. Colavecchia, *Phys. Rev. A* **81**, 042520 (2010).
- <sup>44</sup>D. M. Mitnik, F. D. Colavecchia, G. Gasaneo, and J. M. Randazzo, *Comput. Phys. Commun.* **182**, 1145 (2011).
- <sup>45</sup>J. M. Randazzo, D. Mitnik, G. Gasaneo, L. U. Ancarani, and F. D. Colavecchia, *Eur. Phys. J. D* **69**, 189 (2015).
- <sup>46</sup>C. M. Granados-Castro, L. U. Ancarani, G. Gasaneo, and D. M. Mitnik, *Adv. Quantum Chem.* **73**, 3 (2016).
- <sup>47</sup>J. M. Soler, E. Artacho, J. D. Gale, A. García, J. Junquera, P. Ordejón, and D. Sánchez-Portal, *J. Phys.: Condens. Matter* **14**, 2745 (2002).
- <sup>48</sup>L. Frediani and D. Sundholm, *Phys. Chem. Chem. Phys.* **17**, 31357 (2015).
- <sup>49</sup>E. Tsuchida and M. Tsukada, *Phys. Rev. B* **52**, 5573 (1995).
- <sup>50</sup>L. Lehtovaara, V. Havu, and M. Puska, *J. Chem. Phys.* **131**, 054103 (2009).
- <sup>51</sup>R. Alizadegan, K. J. Hsia, and T. J. Martinez, *J. Chem. Phys.* **132**, 034101 (2010).
- <sup>52</sup>J. E. Avery, “New computational methods in the quantum theory of nanostructures,” Ph.D. thesis, University of Copenhagen, 2011.
- <sup>53</sup>D. Davydov, T. D. Young, and P. Steinmann, *Int. J. Numer. Methods Eng.* **106**, 863 (2015).
- <sup>54</sup>N. M. Boffi, M. Jain, and A. Natan, *J. Chem. Theory Comput.* **12**, 3614 (2016).
- <sup>55</sup>F. A. Bischoff and E. F. Valeev, *J. Chem. Phys.* **134**, 104104 (2011).
- <sup>56</sup>F. A. Bischoff, R. J. Harrison, and E. F. Valeev, *J. Chem. Phys.* **137**, 104103 (2012).
- <sup>57</sup>F. A. Bischoff and E. F. Valeev, *J. Chem. Phys.* **139**, 114106 (2013).
- <sup>58</sup>F. A. Bischoff, *J. Chem. Phys.* **141**, 184106 (2014).
- <sup>59</sup>F. A. Bischoff, *J. Chem. Phys.* **141**, 184105 (2014).
- <sup>60</sup>F. A. Bischoff, *J. Chem. Phys.* **146**, 124126 (2017).
- <sup>61</sup>R. J. Needs, M. D. Towler, N. D. Drummond, and P. L. Ríos, *J. Phys.: Condens. Matter* **22**, 023201 (2010).
- <sup>62</sup>J. Kim, K. P. Esler, J. McMinis, M. A. Morales, B. K. Clark, L. Shulenburg, and D. M. Ceperley, *J. Phys.: Conf. Ser.* **402**, 012008 (2012).
- <sup>63</sup>J. Hutter, M. Iannuzzi, F. Schiffmann, and J. VandeVondele, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **4**, 15 (2014).
- <sup>64</sup>A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dulak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermode, J. R. Kitchin, E. L. Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. B. Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, and K. W. Jacobsen, *J. Phys.: Condens. Matter* **29**, 273002 (2017).
- <sup>65</sup>J. J. Mortensen, L. B. Hansen, and K. W. Jacobsen, *Phys. Rev. B* **71**, 035109 (2005).
- <sup>66</sup>J. Enkovaara, C. Rostgaard, J. J. Mortensen, J. Chen, M. Dulak, L. Ferrighi, J. Gavnholt, C. Glinsvad, V. Haikola, H. A. Hansen, H. H. Kristoffersen, M. Kuisma, A. H. Larsen, L. Lehtovaara, M. Ljungberg, O. Lopez-Acevedo, P. G. Moses, J. Ojanen, T. Olsen, V. Petzold, N. A. Romero, J. Stausholm-Møller, M. Strange, G. A. Tritsarlis, M. Vanin, M. Walter, B. Hammer, H. Häkkinen, G. K. H. Madsen, R. M. Nieminen, J. K. Nørskov, M. Puska, T. T. Rantala, J. Schiøtz, K. S. Thygesen, and K. W. Jacobsen, *J. Phys.: Condens. Matter* **22**, 253202 (2010).
- <sup>67</sup>Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li, J. Liu, J. McClain, E. R. Sayfutyarova, S. Sharma, S. Wouters, and G. K.-L. Chan, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **8**, e1340 (2017).
- <sup>68</sup>N. M. O'boyle, A. L. Tenderholt, and K. M. Langner, *J. Comput. Chem.* **29**, 839 (2008).
- <sup>69</sup>T. Verstraelen, P. Tecmer, F. Heidar-Zadeh, C. E. González-Espinoza, M. Chan, T. D. Kim, K. Boguslawski, S. Fias, S. Vandenbrande, D. Berrocal, and P. W. Ayers, *Horton 2.1.0*, 2017.
- <sup>70</sup>R. Muller, “Pyquante: Python quantum chemistry,” <http://pyquante.sourceforge.net/>; accessed 26 November 2017.
- <sup>71</sup>R. M. Parrish, L. A. Burns, D. G. A. Smith, A. C. Simmonett, A. E. DePrince, E. G. Hohenstein, U. Bozkaya, A. Y. Sokolov, R. Di Remigio, R. M. Richard, J. F. Gonthier, A. M. James, H. R. McAlexander, A. Kumar, M. Saitow, X. Wang, B. P. Pritchard, P. Verma, H. F. Schaefer, K. Patkowski, R. A. King, E. F. Valeev, F. A. Evangelista, J. M. Turney, T. D. Crawford, and C. D. Sherrill, *J. Chem. Theory Comput.* **13**, 3185 (2017).
- <sup>72</sup>D. G. A. Smith, L. A. Burns, D. A. Sirianni, D. R. Nascimento, A. Kumar, A. M. James, J. B. Schriber, T. Zhang, B. Zhang, A. S. Abbott, E. J. Berquist, M. H. Lechner, L. A. Cunha, A. G. Heide, J. M. Waldrop, T. Y. Takeshita, A. Alenaizan, D. Neuhauser, R. A. King, A. C. Simmonett, J. M. Turney, H. F. Schaefer, F. A. Evangelista, A. E. DePrince, T. D. Crawford, K. Patkowski, and C. D. Sherrill, *J. Chem. Theory Comput.* **14**, 3504 (2018).
- <sup>73</sup>S. van der Walt, S. C. Colbert, and G. Varoquaux, *Comput. Sci. Eng.* **13**, 22 (2011).
- <sup>74</sup>The HDF Group, *HDF5 Reference Manual* (The HDF Group, 2011), release 1.8.8.
- <sup>75</sup>M. F. Herbst, “Development of a modular quantum-chemistry framework for the investigation of novel basis functions,” Ph.D. thesis, Ruprecht-Karls-Universität Heidelberg, 2018, [https://michael-herbst.com/publications/2018.05\\_phd\\_corrected.pdf](https://michael-herbst.com/publications/2018.05_phd_corrected.pdf).
- <sup>76</sup>B. Sutcliffe, E. Cancès, M. Caffarel, R. Assaraf, G. Turinici, I. Catto, P.-L. Lions, C. L. Bris, O. Bokanowski, B. Grébert, N. J. Mauser, X. Blanc, M. Defranceschi, V. Louis-Achille, B. Mennucci, J. Dolbeault, M. J. Esteban, E. Séré, T. Saue, and H. J. A. Jensen, in *Mathematical Models and Methods for Ab Initio Quantum Chemistry*, Lecture Notes in Chemistry, edited by M. Defranceschi and C. L. Bris (Springer-Verlag, 2000), Vol. 74.
- <sup>77</sup>E. Cancès and C. Le Bris, *Int. J. Quantum Chem.* **79**, 82 (2000).
- <sup>78</sup>E. Cancès and C. Le Bris, *ESAIM: Math. Modell. Numer. Anal.* **34**, 749 (2000).
- <sup>79</sup>R. McWeeny, *Proc. R. Soc. A* **235**, 496 (1956).
- <sup>80</sup>A. Igawa and H. Fukutome, *Prog. Theor. Phys.* **54**, 1266 (1975).
- <sup>81</sup>R. Seeger and J. A. Pople, *J. Chem. Phys.* **65**, 265 (1976).
- <sup>82</sup>T. van Voorhis and M. Head-Gordon, *Mol. Phys.* **100**, 1713 (2002).
- <sup>83</sup>C. C. J. Roothaan, *Rev. Mod. Phys.* **23**, 69 (1951).
- <sup>84</sup>V. R. Saunders and I. H. Hillier, *Int. J. Quantum Chem.* **7**, 699 (1973).
- <sup>85</sup>P. Pulay, *Chem. Phys. Lett.* **73**, 393 (1980).
- <sup>86</sup>P. Pulay, *J. Comput. Chem.* **3**, 556 (1982).
- <sup>87</sup>F. Jensen, *J. Phys. Chem. A* **111**, 11198 (2007).
- <sup>88</sup>C. Großmann and H.-G. Roos, *Numerik Partieller Differentialgleichungen*, Teubner Studienbücher Mathematik, 2nd ed. (Vieweg+Teubner Verlag, 1992).
- <sup>89</sup>S. C. Brenner and L. R. Scott, *The Mathematical Theory of Finite Element Methods*, 3rd ed. (Springer-Verlag, 2008).
- <sup>90</sup>W. E. Arnoldi, *Q. Appl. Math.* **9**, 17 (1951).
- <sup>91</sup>E. R. Davidson, *J. Comput. Phys.* **17**, 87 (1975).
- <sup>92</sup>M. Kronbichler and K. Kormann, *Comput. Fluids* **63**, 135 (2012).
- <sup>93</sup>M. Wormit, “Development and application of reliable methods for the calculation of excited states: From light-harvesting complexes to medium-sized molecules,” Ph.D. thesis, Universität Frankfurt, 2009.
- <sup>94</sup>M. Wormit, D. R. Rehn, P. H. Harbach, J. Wenzel, C. M. Krauter, E. Epifanovsky, and A. Dreuw, *Mol. Phys.* **112**, 774 (2014).
- <sup>95</sup>A. Dreuw and M. Wormit, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **5**, 82 (2014).
- <sup>96</sup>T. Helgaker, J. Olsen, and P. Jorgensen, *Molecular Electronic-Structure Theory*, 1st ed. (Wiley, 2013).
- <sup>97</sup>Latency numbers every programmer should know, <https://gist.github.com/hellerbarde/2843375>; accessed 08 February 2018.
- <sup>98</sup>D. Cheney, Five things that make Go fast, Presentation at Gocon2014, Tokyo, Japan, 2014.

- <sup>99</sup>Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed., edited by Y. Saad (SIAM Publishing, 2003).
- <sup>100</sup>P. Arbenz, Lecture notes on solving large scale eigenvalue problems, Lecture Notes, ETH Zürich, 2010.
- <sup>101</sup>Y. Saad, *Numerical Methods for Large Eigenvalue Problems*, 2nd ed. (SIAM Publishing, 2011).
- <sup>102</sup>P. Hudak, *ACM Comput. Surv.* **21**, 359 (1989).
- <sup>103</sup>A. Buluç, J. T. Fineman, M. Frigo, J. R. Gilbert, and C. E. Leiserson, in *Symposium on Parallelism in Algorithms and Architectures (SPAA)* (ACM, 2009), pp. 233–244.
- <sup>104</sup>J. Hughes, in *Research Topics in Functional Programming* (Addison-Wesley, 1990), Chap. 2, pp. 17–42.
- <sup>105</sup>Q. Sun, *J. Comput. Chem.* **36**, 1664 (2015).
- <sup>106</sup>J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- <sup>107</sup>E. Jones, T. Oliphant, P. Peterson *et al.*, SciPy: Open source scientific tools for Python (2001–); accessed 09 March 2017.
- <sup>108</sup>W. McKinney, J. Reback *et al.*, PANDAS: Python data analysis library (2008–); accessed 09 March 2017.
- <sup>109</sup>F. Pérez and B. E. Granger, *Comput. Sci. Eng.* **9**, 21 (2007).
- <sup>110</sup>O. Ben-Kiki, C. Evans, and I. döt Net, YAML Ain't markup language (YAML) version 1.2, <http://www.yaml.org/spec/1.2/spec.html> (2009), accessed 03 December 2017.
- <sup>111</sup>J. Schirmer, *Phys. Rev. A* **26**, 2395 (1982).
- <sup>112</sup>A. B. Trofimov, G. Stelter, and J. Schirmer, *J. Chem. Phys.* **111**, 9982 (1999).
- <sup>113</sup>D. Beazley, L. Ballabio, W. Fulton, M. Gossage, M. Köppe, J. Lenz, M. Matus, J. Stewart, A. Yerkes, S. Yoshiki, S. Singhi, X. Delacour, O. Betts, and D. Z. Gang, SWIG: Simplified wrapper and interface generator, accessed 09 March 2018.
- <sup>114</sup>R. J. Bartlett and G. D. Purvis, *Int. J. Quantum Chem.* **14**, 561 (1978).
- <sup>115</sup>A. C. Hurley, *Electron Correlation in Small Molecules* (Academic Press, London, New York, 1976).
- <sup>116</sup>M. F. Herbst and J. E. Avery, <https://molsturm.org>.
- <sup>117</sup>W. J. Hehre, R. Ditchfield, and J. A. Pople, *J. Chem. Phys.* **56**, 2257 (1972).
- <sup>118</sup>F. Neese, *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2**, 73 (2012).
- <sup>119</sup>M. J. D. Powell, *Comput. J.* **7**, 155 (1964).
- <sup>120</sup>W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes* (Cambridge University Press, 1992).
- <sup>121</sup>F. Weigend and R. Ahlrichs, *Phys. Chem. Chem. Phys.* **7**, 3297 (2005).
- <sup>122</sup>E. F. Valeev, Libint: A library for the evaluation of molecular integrals of many-body operators over Gaussian functions, <http://libint.valeev.net> (2017), version 2.3.1.
- <sup>123</sup>E. Valeev, J. Calvin, D. Lewis, J. Dullea, C. Peng, K. Nishimra, J. T. Fermann, J. D. Whitfield, O. Čertík, M. F. Herbst, S. Y. Willow, and D. Williams-Young, *evaleev/libint*: 2.3.1, 2017.
- <sup>124</sup>J. E. Avery and M. F. Herbst, <https://molsturm.org/sturmint>.
- <sup>125</sup>S. Lehtola, C. Steigemann, M. J. Oliveira, and M. A. Marques, *SoftwareX* **7**, 1 (2018).
- <sup>126</sup>U. Ekström, R. Bast, S. S. Reine, C. Jacob, J. Juselius, E. Rebolini, R. D. Remigio, A. S. P. Gomes, S. Reimann, A. Borgoo, M. Ilias, and yurivict, Xcfun DFT library, <http://dftlibs.org/xcfun>; accessed 11 June 2018.
- <sup>127</sup>G. Baumgartner, A. Auer, D. Bernholdt, A. Bibireata, V. Choppella, D. Cociorva, X. R. H. Gao, S. Hirata, S. Krishnamoorthy, S. Krishnan, C. Lam, Q. Lu, M. Nooijen, R. Pitzer, J. Ramanujam, P. Sadayappan, and A. Sibiryakov, *Proc. IEEE* **93**, 276–292 (2005).
- <sup>128</sup>E. Solomonik, D. Matthews, J. Hammond, J. Stanton, and J. Demmel, *J. Parallel Distrib. Comput.* **74**, 3176 (2014).
- <sup>129</sup>E. Peise, D. Fabregat-Traver, and P. Bientinesi, in *High Performance Computing Systems. Performance Modeling, Benchmarking, and Simulation*, edited by S. A. Jarvis, S. A. Wright, and S. D. Hammond (Springer-Verlag, Cham, 2015), pp. 193–212.
- <sup>130</sup>J. A. Calvin, C. A. Lewis, and E. F. Valeev, in *Proceedings of the 5th Workshop on Irregular Applications: Architectures and Algorithms, IA3 '15* (ACM, New York, NY, USA, 2015), pp. 4–1–4–8.
- <sup>131</sup>B. Huber and S. Wolf, Xerus—A general purpose tensor library, <https://libxerus.org> (2014–2017); accessed 10 March 2018.
- <sup>132</sup>M. R. B. Kristensen, S. A. F. Lund, T. Blum, and J. E. Avery, in *Proceedings of the 2016 International Conference on Parallel Architectures and Compilation, PACT 16* (ACM, 2016), pp. 71–85.
- <sup>133</sup>M. R. B. Kristensen, J. E. Avery, T. Blum, S. A. F. Lund, and B. Vinter, in *International Conference on High Performance Computing* (Springer, 2016), Vol. 9945.
- <sup>134</sup>E. Epifanovsky, M. Wormit, T. Kuś, A. Landau, D. Zuev, K. Khistyayev, P. Manohar, I. Kaliman, A. Dreuw, and A. I. Krylov, *J. Comput. Chem.* **34**, 2293 (2013).