Hartree-Fock Program Project

C. David Sherrill
School of Chemistry and Biochemistry
Georgia Institute of Technology

Updated 2 April 2020
Based on notes by Dr. Yukio Yamaguchi (University of Georgia)

1 General Information and Outline of the Hartree–Fock Procedure

Goals: Write a computer program to perform a closed-shell restricted Hartree-Fock computation, given nuclear repulsion energy and one- and two-electron integrals. These notes describe three strategies for accomplishing this goal, of which you should choose one: (1) reading the required data from a plain text file, for the specific case of STO-3G H_2O at a specific geometry; or (2) obtaining the required data directly from the Psi4 program package using Psi4's Python front-end and the Psi4's available Python function library.

An advantage of strategy (1), the plain text file for STO-3G H_2O, is that you are free to use any computer programming language that you wish (or Matlab, Mathematica, etc.). The disadvantages are that your program will be limited to working only for this one test case (unless you can independently generate the required integrals), and that you will have to create/find all required subroutines (e.g., matrix diagonalization, matrix multiplication, etc.)

An advantage of strategy (2) is that it will work on general molecules, and you can leverage existing subroutines provided by Psi4. Documentation of all functions required to accomplish the project using Python is provided below. It is also possible to use the Python API, and/or the NumPy interface of Psi4 to write this code, but these notes assume you are using the binary ("Psithon") interface to Psi4. These notes should work for any recent version of Psi4, e.g., version 1.1 or newer (including 1.4).

Additional Information: See chapter 3 (sections 3.4.4 - 3.4.6) of Modern Quantum Chemistry, 1st Ed., Revised, A. Szabo and N. S. Ostlund (McGraw-Hill, New York, 1989).

Procedure: Here we will briefly outline the primary computational steps. In subsequent sections, we will discuss strategies for implementing these steps.

1. Get the nuclear repulsion energy (Enuc).

2. Use PSI to compute one-electron integrals

   (a) Compute the overlap integrals (S).
   (b) Compute the kinetic energy integrals (T).
   (c) Compute the potential energy integrals (V).
   (d) Form the core Hamiltonian (H), via $H_{\mu\nu} = T_{\mu\nu} + V_{\mu\nu}$.

3. Construct the orthogonalizing matrix $S^{-1/2}$

   (a) Diagonalize the S matrix,
   $$U^\dagger S U = \Lambda. \tag{1}$$

   (b) Form the $S^{-1/2}$ matrix,
   $$S^{-1/2} = U\Lambda^{-1/2}U^\dagger. \tag{2}$$

4. Construct an initial (guess) density matrix

   (a) Form the "core" Fock matrix in the orthogonalized basis via
   $$F_0' = (S^{-1/2})^\dagger H S^{-1/2}. \tag{3}$$

   (b) Diagonalize the initial Fock matrix using a standard eigenvalue routine such as the DSYEV routine in the LAPACK library.
   $$C_0'^\dagger F_0' C_0' = \epsilon. \tag{4}$$

   Note: Two steps below this, you will be forming the density matrix using the columns of C that correspond to occupied orbitals. You need to know which columns of C correspond to occupied orbitals and which ones correspond to unoccupied orbitals. Many standard diagonalizers will sort the eigenvalues and eigenvectors for you, and if they sort them in increasing order of the eigenvalues, then you can just take the first $N/2$ columns of C when forming the D matrix, where N is the number of electrons (and hence $N/2$ is the number of doubly-occupied orbitals). But if your matrix diagonalizer does not sort the eigenvalues/eigenvectors, you will have to sort them yourself or else somehow make sure you are using the correct columns of C when forming D below.

   (c) Form the initial SCF eigenvector matrix in the original basis
   $$C_0 = S^{-1/2}C_0'. \tag{5}$$

   (d) Form the initial density matrix, D
   $$D_{\mu\nu} = \sum_i^{N/2} C_{\mu i}C_{\nu i}, \tag{6}$$

   where N is the number of electrons (and hence $N/2$ is the number of doubly-occupied orbitals).

2

5. Perform the SCF iterations

(a) Form the new Fock matrix, F, from the density matrix and the two-electron integrals

$$F_{\mu\nu} = H_{\mu\nu} + \sum_{\rho\sigma}^{AO} D_{\rho\sigma} \left\{ 2[\mu\nu|\rho\sigma] - [\mu\rho|\nu\sigma] \right\}. \tag{7}$$

(b) Calculate the electronic energy

$$E = \sum_{\mu\nu}^{AO} D_{\mu\nu} \left(H_{\mu\nu} + F_{\mu\nu} \right) + E_{nuc}. \tag{8}$$

(c) Transform the Fock matrix to the orthonormal basis

$$F' = (S^{-1/2})^\dagger F S^{-1/2}. \tag{9}$$

(d) Diagonalize the Fock matrix
$$C'^\dagger F' C' = \epsilon. \tag{10}$$

As discussed above, you probably want to ensure the eigenvectors (columns of C) are sorted so that the corresponding eigenvalues are in increasing order.

(e) Construct the new SCF eigenvector matrix

$$C = S^{-1/2} C' \tag{11}$$

(f) Form the new density matrix

$$D_{\mu\nu} = \sum_{i}^{N/2} C_{\mu i} C_{\nu i}, \tag{12}$$

where $N/2$ is the number of doubly-occupied orbitals.

(g) Test for convergence of the energy.

$$\Delta E = E^n - E^{n-1} < \delta_E \tag{13}$$

(h) Optionally, also test convergence of the density and/or Fock matrix. One simple way to do this is just to compute the RMS change in the density matrix:

$$D_{rms} = \left[\sum_{\mu\nu}^{AO} \left(D_{\mu\nu}^n - D_{\mu\nu}^{n-1} \right)^2 \right]^{1/2} < \delta_D \tag{14}$$

This criterion is fine for small systems. For very large systems, it will become harder for this criterion to be met for a fixed RMS cutoff value. An alternative criterion that

is sometimes used is the commutator of the density matrix and the Fock matrix, which will go to zero at convergence.

$$[D, F]_{\mu\nu} = \sum_{\rho\sigma}^{AO} S_{\mu\rho} D_{\rho\sigma} F_{\sigma\nu} - F_{\mu\rho} D_{\rho\sigma} S_{\sigma\nu} \tag{15}$$

$$\|[D, F]\|_F < \delta_D \tag{16}$$

(i) If not converged, do another iteration.

2 Project writeup

Before discussing three different technical strategies for performing the project, if you are doing this project as a class project, below are some suggestions for what to include in your project write-up (if you are just doing this as part of a research experience there is probably no reason to do a formal writeup).

1. Include a *short* introduction giving a brief re-cap of what Hartree-Fock is and how it works, in your own words

2. Give an introduction to how you coded the program: what language (or program like Mathematica) did you use, why did you pick that language, what features of the language were helpful, what things did the language not provide automatically that you had to code up yourself, did you have to use any special tricks, etc.

3. Provide the actual source code of your Hartree-Fock program and any helper functions you had to write

4. Provide at least one sample output from your program (including the corresponding input file, if you are not using the H_2O integrals file from Coding Strategy #1 below).

5. If you did not obtain the correct Hartree-Fock energy, explain what you think might have gone wrong with your program.

3 Coding strategy #1: reading data from a text file for a specific test case

The required data for STO-3G H_2O at its equilibrium geometry is available on the Sherrill group website at http://vergil.chemistry.gatech.edu/h2oints.txt. The file should be self-explanatory.

There are 7 orbitals for this system in the STO-3G basis. Note that the numbering of orbitals in the two-electron integrals starts at 0. Feel free to reformat the file to suit your purposes.

Important note: The printout of the two-electron integrals only lists the *permutationally unique* integrals. Recall that the two-electron integrals have 8-fold permutational symmetry when using real orbitals:

$$[pq|rs] = [qp|rs] = [pq|sr] = [qp|sr] = [rs|pq] = [sr|pq] = [rs|qp] = [sr|qp]. \tag{17}$$

The equations given above for the Fock matrix assume all possible permutations of integrals are available. That means that in forming the Fock matrix you may find yourself needing integral $[21|34]$, and you need to realize that it doesn't appear that way in the list, it appears as $[43|21]$. The integrals listed follow the canonical order that $p \geq q$, $r \geq s$, and with the indices on the left being larger than those on the right (this last requirement can be stated more formally by saying Psi4 requires that the "super-indices" pq and rs satisfy the requirement that $pq \geq rs$, where $pq = (p(p+1)/2) + q$, and analogously for rs, and with orbital numbering beginning with 0).

4 Coding strategy #2: Writing a general program in Python using Psi4

Another option for getting the required integrals is to interface to a program like Psi4. This is relatively easy to do using Psi4's very user/programmer-friendly Python front-end, which allows user input files to include not only molecule/computation information, but also any valid Python code. This means that it is possible to write a Hartree–Fock program in Psi4 using only the Psi4 binary and a single user input file. Python is a high-level language (i.e., it is easy to accomplish a lot in relatively few lines of code), it is widely used, and it is fairly easy to learn.

This approach will require you to install (or have access to) a binary of the Psi4 program. If you do not already have login access to a machine with Psi4 on it, you can either download and install a Psi4 binary (easiest), or download the Psi4 source and compile and install it. To download and install a binary, see the

"Psi4 Downloads".

If you create your program this way, once the Psi4 binary is available, you just need to create your program in an input file according to the template provided in Figure 1 and execute it from Psi4 like this: `psi4 my-program.in`.

Psi4 provides routines that make one- and two-electron integrals available in Python, and

Figure 1: Overall skeleton of a Psi4 program in an input file.

```
# Your program goes first, at the top of the file
def simple_scf(molecule):

  # make sure the molecule object gets correctly populated with
  # the current geometry
  molecule.update_geometry()

  # your Python program goes here, takes a Psi4 molecule as input

# after your program, then specify a molecule using normal
# Psi4 input and run the SCF using your python SCF code and then also
# with the usual Psi4 SCF code, so that you can compare final energies.
# You should be able to pick any small molecule you like
# (although we will assume RHF so make it closed-shell).

molecule mol {
0 1
O
H 1 1.0
H 1 1.0 2 104.5
symmetry c1
}

set {
basis sto-3g
}

simple_scf(mol)
set scf_type direct
energy('scf')
```

a Matrix class that allows various operations like matrix multiplication, matrix diagonalization, etc. It also has a Molecule class that allows one to grab useful information from a user-specified molecule, like its overall charge, etc. All the functions necessary to write a Hartree–Fock program should be given below. However, you can find a complete list of all Psi4 functions available from Python in the Linking C++ and Python section of the Psi4 Manual.

Figure 2 illustrates how to use Psi4's integral generation routines to form all the required one- and two-electron integrals and make them accessible through Psi Matrix objects. Note: these matrices are all stored in core RAM; because the two-electron integrals are an $\mathcal{O}(N^4)$ quantity, you should stick to small test cases with your program. The integrals will be generated for the current default Molecule object; the skeleton program above will correctly specify a `mol` as the default (and only) molecule object, so there is no ambiguity.

Figure 2: Using Psi4's MintsHelper to generate integrals

```
# Integral Generation
wfn = Wavefunction.build(mol, get_global_option("basis"))
mints = MintsHelper(wfn.basisset())
S = mints.ao_overlap()
T = mints.ao_potential()
V = mints.ao_kinetic()
I = mints.ao_eri()
```

`S`, `T`, and `V` are all square matrices of size `nbf` x `nbf`, where `nbf` is the number of AO basis functions, which can be obtained by querrying the matrices about their number of rows (or columns), like this: `nbf = S.rows(0)`. The two-electron integrals are formally represented by a 4-dimensional tensor, but the Psi4 `MintsHelper` object packs them into a (2-dimensional) Matrix for convenience. It accomplishes this by computing a "composite" row index $pq = p * \texttt{nbf} + q$, where p and q are individual orbital indices.

The Psi4 Matrix objects have various built-in capabilities. Some of them are as follows:

1. Create new Matrix X with *nrows* rows and *ncols* cols: `X = Matrix(nrows, ncols)`

2. Get the p, q element of X: `val = X.get(p,q)`

3. Set the p, q element of X to *val*: `X.set(p,q,val)`

4. Add J to F: `F.add(J)`

5. Subtract K from F: `F.subtract(K)`

6. Multiply X by 2: `X.scale(2.0)`

7. Copy X to Y: `Y = X.clone()`

8. Matrix multiplication $C = \beta C + \alpha AB$ (set $\beta = 0$ if the result of αAB is to overwrite C instead of add to it): `C.gemm(transa, transb, alpha, A, B, beta)`, where `transa = True` if we

7

need to transpose matrix A (otherwise `transa = False`), and `transb = True` if we need to transpose matrix B (otherwise `transb = False`). *Note: The input matrices A and B cannot also be used to store the result; that is, you can't do something like `A.gemm(transa, transb, 1.0, A, B, 0)` because you'd need to be writing into `A` at the same time you're reading from `A`.*

9. Diagonalization of a symmetric Matrix:
   `X.diagonalize(evecs, evals, DiagonalizeOrder.Ascending)` will diagonalize symmetric matrix X, placing the eigenvectors in matrix `evecs` (one eigenvector per column) and the corresponding eigenvalues in `evals`, which is a Vector. Both `evecs` and `evals` should have already been created by the user (e.g., `evecs = Matrix(n,n)` and `evals = Vector(n)`). *Note: this routine does not actually change the original matrix X to diagonal form; the original matrix X is left untouched.*

10. Raise a matrix to the p-th power: `X.power(p, 0.0)`

11. Compute an element-wise dot product between two matrices (i.e., $\sum_{pq} A_{pq} \cdot B_{pq}$): `r = A.vector_dot(B)`

For this project you can specify molecular information like the charge, multiplicity, and number of electrons directly in your code, but it is more elegant to parse these out of the Molecule object itself. This is easily done as in Figure 3. Since your program assumes RHF, if you parse the information this way, your program should check to make sure `mult = 1` before it continues.

Figure 3: Using Psi4's Molecule object to get information

```
charge = molecule.molecular_charge()
mult   = molecule.multiplicity()
Z = 0
for A in range(molecule.natom()):
    Z += molecule.Z(A)
ndocc = int(Z / 2) - (charge / 2) # number of doubly-occupied orbitals
Enuc = molecule.nuclear_repulsion_energy()
```