

# A Primer to Electronic Structure Computation

Nick Schafer

Fall 2006

## Abstract

A brief overview of some literature the author read as a part of an independent study on Electronic Structure Computation is given. The Hartree Fock Approach and Density Functional Theory are emphasized. Only some important aspects of each are presented and results are stated without proof. Some of the difficulties involved are discussed along with recently proposed solutions. The treatment is qualitative. The intent is to introduce the problem to someone who has the physics and mathematics background of an advanced undergraduate but no prior experience in Electronic Structure Computation. It should be read as an overview before and a guide during a more in depth study. The interested reader should see the references for further reading. Some familiarity with basic Quantum Mechanics and Linear Algebra is assumed, which can be obtained by following the references.

## Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
1.1	Purpose . . . . .	2
1.2	Definition . . . . .	2
<b>2</b>	<b>Hartree Fock Approach</b>	<b>3</b>
2.1	Basic Procedure . . . . .	3
2.2	Schrodinger Equation . . . . .	4
2.3	Variational Principle . . . . .	4
2.4	Born Oppenheimer Approximation . . . . .	5
2.5	Antisymmetry Principle . . . . .	5
2.6	Hartree Fock Approximation . . . . .	6
2.6.1	Spatial and Spin Orbitals . . . . .	6
2.6.2	Hartree Products . . . . .	6
2.6.3	Slater Determinants . . . . .	7
2.6.4	Fock Operator . . . . .	7
2.6.5	Hartree Fock Equations . . . . .	8
2.7	Computational Procedure . . . . .	9
2.7.1	Basis sets . . . . .	9

2.7.2	Slater Type . . . . .	9
2.7.3	Gaussian Type . . . . .	10
2.7.4	Procedure . . . . .	10
<b>3</b>	<b>Density Functional Theory</b>	<b>11</b>
3.1	Electron Correlation . . . . .	11
3.2	Electron Density as the Fundamental Quantity . . . . .	12
3.2.1	Electron Density . . . . .	12
3.2.2	Pair Density . . . . .	12
3.2.3	Fermi and Coulomb Holes . . . . .	13
3.3	Hohenberg-Kohn Theorems . . . . .	13
3.3.1	Proof of Existence . . . . .	13
3.3.2	Variational Principle . . . . .	14
3.4	Kohn-Sham Approach . . . . .	14
3.4.1	Approximate Functionals . . . . .	15
<b>4</b>	<b>Recent Work on Reducing Computational Intensity</b>	<b>15</b>
4.1	Pseudopotentials and Plane Waves . . . . .	16
4.2	Real Space Mesh . . . . .	16
4.3	Reconstruction . . . . .	17
<b>5</b>	<b>Acknowledgements</b>	<b>17</b>
<b>6</b>	<b>References</b>	<b>17</b>
<b>7</b>	<b>Contact</b>	<b>18</b>

# 1 Introduction

## 1.1 Purpose

The properties of materials are of intense interest to many branches of science and industry. Unfortunately, direct measurements of these properties is not always feasible or reliable. However, these properties are known to be heavily dependent upon the electronic structure of the material. If the electronic structure can somehow be obtained then, in principle, all properties of the material can be calculated (8, p.11). In addition to static properties, dynamic simulations can be carried out using an iterative process to update the electronic structure and nuclear configuration successively (8, p.317). Dynamic simulation allows us to probe the movements of molecules on time and length scales inaccessible by any other means.

## 1.2 Definition

Electronic Structure Computation is the calculation of information about electrons near nuclei from which properties of a material can be derived. Quantum

mechanics tells us that the description of an electron as a spherical charge with a definite trajectory is not always appropriate. In particular, electrons near nuclei become delocalized and are best described by a probability distribution, which gives the probability of finding an electron at a certain place and time (9, p.18). It is exactly this distribution that we refer to as the electronic structure and from which we derive the properties of a material. In Hartree Fock Theory the information computed is the many electron wave function.

Density Functional Theory has taken on several forms since its inception. Very early attempts used relatively simple functionals of the electron density to compute the total energy. Another group of methods computes wave functions of a fictitious noninteracting system of electrons which allows it to resolve most of the energy contribution quickly and accurately. The remaining energy terms are computed using the approximate functionals. Unfortunately, very often it is this remaining energy term which contains most of the interesting information. These methods are known as Kohn-Sham methods. Just recently there have been attempts to revive the direct use of functionals. These methods go under the name Orbital-Free.

The primary and most basic property that can be derived from this information is the ground state energy, as it can be used to obtain the equilibrium configuration of a system.

## 2 Hartree Fock Approach

The Hartree Fock Approach is the starting point for many more sophisticated approaches to Electronic Structure Computation. It has been extended and corrected in many ways, but most of the core ideas and assumptions remain the same and basic knowledge of these is necessary for understanding any of the more modern approaches to Electronic Structure Computation.

### 2.1 Basic Procedure

Anyone who has done any work with numerical methods will recognize the basic procedure for computing electronic structure using the Hartree Fock method. In the broadest sense, there are four steps:

1. Fix the necessary inputs
2. Make an initial guess at the electron density
3. Perform calculations to update the electron density
4. Compare the “new” electron density to the “old” electron density

If the two electron densities differ by more than a specified tolerance, then the “new” electron density becomes the “old” electron density and steps three and four are repeated. In most cases, an average of the “new” and “old” electron densities is used to improve the convergence of the solution.

This may sound simple and in a sense it is. However, there are many subtleties that go into each step and understanding them all takes many years of study. In fact, the details of how each step should be carried out so that the

final electron density is obtained accurately and efficiently is still very much an open question. This paper covers only the basics of some well established techniques, as well as some more recent attempts at making the process more efficient.

## 2.2 Schrodinger Equation

Considering the procedure above, it is helpful to imagine that we are dealing with a classical system. If that was the case, step three would almost certainly involve the use of Newton's second law, which governs the dynamics of classical bodies. However, our system is composed of delocalized electrons interacting with each other and with nuclei. This is not a classical situation and we must look for the appropriate dynamic equation.

The Schrodinger Equation governs the evolution of non-relativistic quantum mechanical systems. In particular, it governs the evolution of wave functions of the electrons. It is a linear second order partial differential equation (9, p.31). When treating materials, the sheer number of variables on which it depends precludes the use of analytical solution techniques. Finding approximate solutions to the Schrodinger Equation is the primary concern of basic Quantum Chemistry. Approximate solutions are necessary because complete analytic solutions are only available for a few simple systems, such as the hydrogen atom (one electron orbiting a nucleus) where the interaction is sufficiently simple (9, p.215).

It seems convenient to note here that all calculations mentioned herein are done in the non-relativistic limit: where the speed of the electrons is much less than the speed of light. This approximation turns out to be appropriate for atomic species lighter than silver, but begins to fail if applied to heavier species.

## 2.3 Variational Principle

The Variational Principle states that any normalized wave function corresponds to an energy, obtained by taking the expectation value of the Hamiltonian, that is greater than or equal to the ground state energy of the system. Also, if this energy is equal to the ground state energy, then that wave function is the ground state wave function of the system. Equivalently, it is not possible to find a normalized wave function corresponding to an energy lower than the ground state energy. This result is physically reasonable in the sense that the ground state is completely "relaxed" in the energetic sense (1, p.31).

It may not be immediately obvious how this would be useful. Consider a quantum mechanical system for which you only know the Hamiltonian but not the ground state wave function or the ground state energy. Now, you make two guesses at the ground state wave function of the system. The guess which has a lower expectation value for the Hamiltonian is a better guess at the ground state wave function according to the Variational Principle. A common trick for finding a good estimate of the ground state energy of a system is to try a form of the wave function that depends on one or more parameters, calculate the

expectation value of the Hamiltonian, then minimize the resulting expression for the energy with respect to the parameters. Inserting these parameters into the original expression gives the best possible guess at the ground state wave function of that particular form (1, p.37).

This method works well for systems where you already know approximately what the wave function should look like. This is not always the case, however, and a more general method is necessary. Guessing an arbitrary form for a wave function is often not sufficient, as it may not be “flexible” enough to mimic the actual ground state wave function. The situation is particularly bad if the properties which you are trying to derive depend sensitively on the value of the wave function at its tails. However, the wave function, like any other function, can be expanded in terms of basis functions in a certain space. The parameters are then the weighting coefficients of the basis functions. If the basis is complete, then the exact ground state wave function can be found. In practice, however, complete basis sets are often infinite and infinite basis sets are impractical. To reduce computational complexity a good basis set should have functions that already resemble the form of the function which you are trying to expand and contain as few functions as possible (1, p.181).

## 2.4 Born Oppenheimer Approximation

The first important simplifying assumption employed by the Hartree Fock Approach is the Born Oppenheimer Approximation. Ordinary materials are made of atoms that are in turn made up of nuclei and electrons. Nuclei are made of protons and neutrons which are many times more massive than electrons. The nuclei move much more slowly than the electrons because they are so much heavier. When computing the electronic structure of a material, the simplifying assumption that the electrons are moving in a field of fixed (unmoving) nuclei is often made. This is known as the Born Oppenheimer Approximation (1, p.43). This simple assumption has important computational consequences. The kinetic energy part of Hamiltonian of the nuclei is completely neglected because the nuclei are assumed to be fixed. The nuclei-nuclei interaction energy is simply a constant determined by the relative positions and charges of the nuclei and need only be computed once.

## 2.5 Antisymmetry Principle

The fact that in a quantum mechanical system no two electrons can ever simultaneously be described by the same quantum numbers is an experimental observation that goes by the name of the Pauli Exclusion Principle. It is equivalent to saying that no two electrons can simulatenously occupy the same state. The Pauli Exclusion Principle is a specific case of the more general Antisymmetry Principle (1, p.45). Its consequence in Electronic Structure Computation is that the many electron wave function must be antisymmetric, i.e. it must change sign upon the interchange of any two electrons’ four (three spatial and one spin) coordinates. This consequence may not be obvious, but the proof in the

case of three dimensions is simple and available in most introductory quantum mechanics and quantum chemistry textbooks (9, p. 314).

## 2.6 Hartree Fock Approximation

The Hartree Fock Approximation underlies many modern Electronic Structure Computation techniques. In a typical system there are many (more than two) electrons interacting. Under the Hartree Fock Approximation, when computing the wave function of an electron in a system of  $N$  electrons, the effect of the remaining  $N-1$  electrons is replaced by an average field which depends only on the coordinates of the electron whose wave function is being computed (1, p.53).

This assumption turns out to be equivalent to the molecular orbital approximation which assumes that electrons occupy orbitals. This idea should be familiar to anyone who has taken a general chemistry course where the molecular orbital approximation is often taught not as an approximation but rather as the basic way which electrons behave. To better understand this approximation's origins, usefulness and limitations, as well as how it is used in Electronic Structure Computation, familiarity with a few basic concepts is necessary.

### 2.6.1 Spatial and Spin Orbitals

It has already been stated that the Schrodinger equation describes the evolution of a quantum mechanical system. In this case we are considering a system of electrons interacting with fixed nuclear point charges. A solution of the Schrodinger equation is known as a wave function. Multiplying a wave function by its complex conjugate gives the differential probability distribution which can be integrated to obtain the electron density. Let an orbital be defined as the wave function of a single electron. A spatial orbital contains information only about the values of the wave function at different points in space. This alone, however, does not completely describe an electron. We know from our knowledge of the Pauli Exclusion Principle that a complete description of an electron must contain information about its spin. A spatial orbital together with a spin orientation is known as a spin orbital (1, p.46).

### 2.6.2 Hartree Products

Spin orbitals represent the wave function for a single electron. Multiplying a spin orbital by its complex conjugate and integrating out the spin would give you the probability of finding an electron occupying that orbital at any place in space. The probability of finding a particular electron, however, is not of real interest. We must look for a way to express the combined wave function of all the electrons in the system, called the many electron wave function.

The simplest way to do this is known as a Hartree Product. It is obtained by multiplying all of the spin orbitals together. The Hartree Product is what is known as an independent wave function because the probability distribution obtained from it is equal to the product of all the independent, single electron

probabilities. This is also known as an uncorrelated wave function because it acts as if the motion of all the electrons are not correlated. This representation deviates from reality in two important respects. Firstly, the movement of electrons are certainly correlated because they experience a mutual coulomb repulsion and therefore the presence of an electron decreases the probability of finding another electron nearby. Secondly, real electrons are indistinguishable and obey the antisymmetry principle (1, p.47).

Because of these deficiencies Hartree Products cannot be used as many electron wave functions. However, they are used as a starting point for a more satisfactory type of many electron wave function known as a Slater Determinant.

### 2.6.3 Slater Determinants

Hartree Products are appealing as many electron wave functions because they are simple, but they do not satisfy the antisymmetry principle. Therefore it is natural to seek a way to systematically “antisymmetrize” a Hartree Product. This is accomplished in Quantum Chemistry literature by the Slater Determinant.

Consider a matrix of spin orbitals. Each column is filled with a different spin orbital and each row of spin orbitals is occupied by a different electron. We always consider the same amount of spin orbitals as there are electrons in the system, so the matrix is necessarily square. Taking the determinant of this matrix and multiplying by the appropriate normalization factor gives an antisymmetric many electron wave function. The fact that it satisfies the antisymmetry principle is easy to see by recalling a fact from linear algebra. The determinants of two matrices which differ only by the exchange of two columns are equal in magnitude and opposite in sign. This is precisely what is required: exchanging the coordinates of two electrons should cause the wave function to change sign (1, p.49).

This type of antisymmetrization of the wave function has some important consequences. Most noticeably, the motion of the electrons are no longer completely uncorrelated. Electrons with parallel spins exhibit exchange correlation effects and a Fermi hole is said to exist around each electron. It is a “hole” in the sense that the probability of finding two electrons with parallel spins at the same point in space is always zero. One must note, however, that, unlike real electrons, two electrons with opposite spins described in this framework have completely uncorrelated motion. Some more advanced wave function based methods improve on Hartree Fock by addressing this very discrepancy.

### 2.6.4 Fock Operator

The Fock Operator is central to the Hartree-Fock method. It is analogous to a one electron Hamiltonian. Its eigenfunctions are spin orbitals. It contains three terms: a kinetic energy term, a potential energy interaction between the electron and all nuclei, and an average field term which represents the electron’s interaction with all of the other electrons and is known as the Hartree Fock

potential (1, p.54). The average field is calculated using the spin orbitals of all the electrons. Because the Fock Operator depends on its eigenfunctions, solving for spin orbitals is a nonlinear eigenfunction problem and must be resolved iteratively. In this context, this procedure is known as a Self Consistent Field (SCF) procedure. Once the average fields of two successive iterations agree to a specified tolerance, the procedure is complete.

The Hartree Fock potential is often written as a sum over the Coulomb and Exchange operators which are integrals over all space that represent the Coulomb repulsion and nonclassical exchange terms arising from the antisymmetric properties of the Slater determinant, respectively. The Coulomb is said to be local (or diagonal) because the value of the operator on the  $i^{th}$  orbital is proportional to the  $i^{th}$  orbital. In contrast, the action of the Exchange operator on the  $i^{th}$  orbital is proportional to a different orbital, so it is said to be non-local (or non-diagonal) (1, p.114).

### 2.6.5 Hartree Fock Equations

The Hartree Fock equations are a set of integro-differential equations that are solved to give the best set of spin orbitals as determined by the variational principle (1, p.111). Recall that the many electron wave function is represented by a Slater determinant. The energy of the wave function is determined by taking the expectation value of the Hamiltonian between the Slater determinant. The Hamiltonian is given as a sum over all single electron hamiltonians, known as the Fock Operators. The ground state energy and wave function are then obtained by minimizing the energy with respect to variation in the spin orbitals.

There are several ways in which the Hartree Fock Equations can be formulated. These variations correspond to different degrees of complexity and accuracy. Practically they arise from the choice of basis set used to represent the solution of the Hartree Fock equations.

The Antisymmetry principle says that two electrons with parallel spin cannot occupy the same point in space simultaneously. However, with electrons of opposite spin this is not the case. A 'restricted' calculation assumes that each pair of electrons (one spin up and the other spin down) occupy the same spatial orbital. In reality, if the system can relax by allowing the spatial orbitals of the spin up and spin down electrons to differ then they will. When a calculation allows for the spin up and spin down spatial wave functions to differ, it is known as an 'unrestricted' calculation.

One important special case that is often considered is when there are an even number of electrons and these electrons all form spin up and down pairs. This is known as a 'closed-shell' calculation. Allowing for the existence of an unpaired electron is, not suprisingly, called an 'open-shell' calculation.

These two variations can be used in any combination (i.e. closed-shell restricted, open-shell unrestricted, etc.). The different combinations are useful in different situations and so the choice often depends on the application. Closed-shell restricted calculations are often the simplest choice, but do not always accurately capture certain phenomena, such as the breaking of chemical bonds.



Open-shell unrestricted calculations are without a doubt the most general, but can be unnecessarily complicated depending on the situation (1, p.131 and p.205).

## 2.7 Computational Procedure

Up until this point I have made very little mention of how Electronic Structure Computation is carried out in practice. To completely understand how the electron density is computed using the Hartree Fock method, one would have to know all of what has been mentioned here and more. Computer code that is capable of performing these calculations has taken many people many years to develop. Nevertheless, I think it is useful to outline the procedure using the basic understanding we have built up to here so that during the course of learning the Hartree Fock method in more rigor, one does not get lost in the matrices and lose sight of the big picture.

### 2.7.1 Basis sets

A common way of representing an unknown function is as a sum of coefficients times known functions. A basis set in this context is exactly that set of known functions. Solutions to the Hartree Fock equations are in practice obtained by introducing a set of basis functions. Expressing the spatial part of the spin orbitals in terms of basis functions changes the integro-differential equations into matrix equations, which can be solved using any of the many available methods. Accurate and efficient numerical methods for solving matrix equations have been developed because they apply to a large class of problems spanning many disciplines beyond quantum mechanics.

In theory any complete set of functions could be used as a basis set. However, it is advantageous to choose functions which resemble the solutions to the equations which you are trying to solve. If the basis functions are chosen correctly then it will only be necessary to sum a few of them to obtain an accurate representation of the solution.

There are two types of basis functions that are commonly used when solving the Hartree Fock equations: Slater Type and Gaussian Type.

### 2.7.2 Slater Type

The first type of basis function is known as the Slater Type. It is an exponential function that falls off with the first power of the radius. It has its origin in the analytic solution for the wave function of an electron in the 1s orbital of Hydrogen. As such, it is a very good first guess for the form of an arbitrary electron wave function near an atom. The main disadvantage of using a Slater Type basis is computational in nature. To understand this, however, we must compare it with the other common type of basis function, the Gaussian Type (1, p.153).

### 2.7.3 Gaussian Type

Gaussian Type basis functions are used because, in some ways, they resemble Slater Type functions and are much easier to perform computations on. The Gaussian Type function is an exponential function that falls off with the second power of the radius. It differs most drastically from the Slater Type function at the origin, where it has a zero slope rather than a finite slope, and at large radii where it falls off much more quickly. In order to maintain the ease of computation and mimic the functional form of Slater Type functions, sums over several Gaussian functions with carefully chosen parameters are used as basis functions. These are known as Contracted Gaussian functions.

Despite these attempts to mimic the functional form of the Slater Type functions, larger basis sets are still necessary when Gaussian Type functions are used. Nevertheless, the amount of analytical reduction in the calculation of two Gaussian functions, which is necessary when calculating electron-electron interactions, makes them computationally favorable.

### 2.7.4 Procedure

Using the language we have built up so far I would like to elaborate on Section 1.2 which outlines in the roughest way possible the basic Self Consistent Field procedure for computing electronic structure using the Hartree Fock Approach.

The first step in computing the electronic structure is specifying the positions and atomic numbers of the nuclei. Remember that under the Born Oppenheimer approximation the nuclei remain fixed throughout the computation. The number of electrons and the set of basis functions to be used must also be specified. As discussed earlier the choice of basis sets will likely depend on the system. It is necessary to introduce the basis set in order to convert the integro-differential Hartree Fock equations into matrix equations.

Next, an initial guess at the electron density must be made. Often this is done by ignoring the effect of the electron-electron interactions; however, using this method to make an initial guess does not guarantee that the procedure will converge onto the correct solution. There are other more complicated ways of choosing an initial density which improve the chance of convergence. In practice the density matrix is obtained from a matrix which contains the values of the expansion coefficients that multiply the basis functions.

Once a guess at the electron density is in place it is possible to calculate the Hartree Fock potential which, as noted before, approximates electron-electron interactions by an average field that depends only on one electron's coordinates. Calculating this average field is the most computationally intense part of the Hartree Fock procedure because an integral must be done for each unique pair of electrons in the system. Remember, however, that we are not explicitly computing the interactions between all the electron pairs. The heart of the Hartree Fock method is to first calculate the average potential and then solve the eigenvalue problem for the electron orbitals with the electron-electron interactions implicit in the average field. The interactions between electrons and nuclei are

also computed, but tend to be much less expensive because the nuclei are point charges with no spatial extent.

All of the interactions that have been computed in the previous step are added together to obtain the matrix representation of the Fock operator. The Fock operator is like a single electron Hamiltonian; i.e. its eigenfunctions are the electron wave functions. Solving this system in practice involves diagonalizing the Fock matrix. The only constraint on this minimization process is that the orbitals remain orthonormal. Once the spin orbitals have been computed, a new density matrix can be computed and compared to the old density matrix. If the two densities disagree by more than a specified tolerance then the whole procedure is repeated until the solution has converged.

Once the solution is converged, we can say a few things about our results. First of all, we have not computed the exact many electron ground state wave function. Instead, we have followed the variational principle to obtain the best possible guess to the wave function that can be expressed in terms of our chosen basis functions. This distinction is important when considering the fact that changing the basis functions, although seemingly physically superficial, can have a very real effect on the ground state energy that is calculated (1, p.146).

### 3 Density Functional Theory

Density Functional Theory started as something of a hunch, based more on simple physical intuition than any formal platform. The first attempts to use the electron density as a fundamental quantity for calculating properties of a system were made around 1927 by Thomas and Fermi. The basic idea was this: why deal with a complicated many electron wave function when the quantity that you are really after, the electron density, is much simpler? It has taken some clever thinking to formally justify this approach and the search for accurate functionals continues to occupy the minds of many researchers. Nevertheless, Density Functional Theory has already found many applications and is widely considered the method of choice for computing electronic structure in many contexts.

#### 3.1 Electron Correlation

As discussed earlier, the most difficult part of computing the electronic structure is calculating the interactions between electrons (as opposed to nuclei-electron or nuclei-nuclei interactions). When the presence or movement of an electron can influence the movement of another electron, those electrons are said to be correlated. How the movements of electrons influence other electrons is called the electron correlation.

In Hartree Fock Theory there is an error associated with using a Slater Determinant and an average potential for representing electron-electron interactions. Because the interactions between electrons are not computed explicitly, there is no way of accounting for the instantaneous repulsion between the electrons.

The difference between the exact energy and the energy calculated by Hartree Fock theory is known as the electron correlation energy. Many methods, such as Configuration Interaction, have been devised to account for this difference in the framework of Hartree Fock theory. Density Functional Theory, however, is not one of them. To see how Density Functional Theory deals with the problem of electron correlation we must first define some relevant terms and clarify a few points.

## 3.2 Electron Density as the Fundamental Quantity

### 3.2.1 Electron Density

The electron density is a probability distribution that gives the chance of finding an electron in a particular region of space. It is often described as a smearing of electrons inside a volume which evolves in time more like a fluid or gas than a particle in a classical mechanical system, which is confined to a definite trajectory.

In some sense the electron density is the primary quantity of interest, even in Hartree Fock theory. The many electron wave function is used as a means of obtaining the electron density because the dynamics of electron waves have been studied since the beginning of Quantum Mechanics and, for the most part, are well understood. However, from a computational perspective it seems like a shame to compute a complex valued function of  $3N$  variables (where  $N$  is the number of electrons in the system) and then immediately convert it to a real valued function of only 3 variables. It is from this simpler function that one derives all the properties of the material being simulated.

Electron density does not, in itself, tell you anything about the way in which electrons interact, i.e., about electron correlation. There is, however, a variation on the idea of electron density that does relate to electron correlation, which I introduce presently.

### 3.2.2 Pair Density

Imagine a system of  $N$  electrons all interacting with one another. The electron density gives the probability of finding a single electron in a certain region of space. The pair density gives the probability of simultaneously finding two electrons in two specified regions in space. You might intuit that this function would contain information about electron correlation, and indeed it does (2, p.20). A form of the pair density which is normalized by the electron density, known as the conditional probability, gives the probability of finding an electron at a position when another electron is known to be at a different position. The conditional probability minus the uncorrelated probability gives what is called the exchange correlation hole. Not surprisingly the conditional probability is lower than the uncorrelated probability in most cases, which justifies the name exchange correlation "hole".

### 3.2.3 Fermi and Coulomb Holes

The above discussion about pair density is a somewhat roundabout way of stating what you might already expect to be true: the presence of an electron reduces the probability of finding another electron in the vicinity. Roughly speaking there are two reasons why this is true.

The first reason applies only to electrons with parallel spin. I have already mentioned Pauli Exclusion in the context of Hartree Fock theory and those ideas apply equally well here. The antisymmetric nature of the wave function leads to a reduced probability of finding two electrons of parallel spin close to each other. It might not be so obvious, but this effect actually has a larger influence on the conditional probability than the Coulomb repulsion between electrons (2, p.25). This fact helps to explain why Hartree Fock theory can still obtain accurate ground state energies by using an antisymmetric wave function (in the form of a Slater Determinant) but not accounting for the instantaneous repulsion between electrons.

Any two electrons carry the same charge and two like charges repel. This causes a smaller but still significant reduction in the probability of finding an electron in the vicinity of another electron.

Both of these ideas could, at least potentially, be addressed in the framework of Hartree Fock theory, so why wait to talk about them now? There is a fundamental difference in the way which Hartree Fock and Density Functional Theory attempts to deal with these effects. The Hartree Fock Approach computes an average field based on spin orbitals. Density Functional Theory seeks a functional of the electron density itself. How do we know that such a functional could even exist? Theoretical justification for Density Functional Theorists came in the form of the Hohenberg-Kohn Theorems.

## 3.3 Hohenberg-Kohn Theorems

I said before that Density Functional Theory started as more of a hunch than rigorously justified approach. This might seem like an exaggeration, but the truth is that formal justification in the form of the Hohenberg-Kohn theorems came some forty years after the earliest attempts to use Density Functional Theory to compute electronic structure. After the theorems were published, Density Functional Theory came to the forefront of electronic structure computation. The statement of the theorems is remarkably simple.

### 3.3.1 Proof of Existence

The first Hohenberg-Kohn theorem says that the energy of a system is a unique functional of the total charge density. That is to say that a given charge density cannot produce two different ground state energies (2, p.33). It is important to note that this theorem only says that a functional which maps charge densities onto energies exists. It does not tell you how to obtain the functional, but does tell you that the search for such a functional might not be futile.

### 3.3.2 Variational Principle

The second Hohenberg-Kohn theorem is about which charge densities map to the lowest (ground state) energy. It says that functional which maps charge densities to energies yields the lowest energy if and only if the ground state density is used (2, p.36).

With these two theorems in place, we know that a functional exists which maps the ground state charge density to the ground state energy of a system. This must have been good news to anyone looking to use Density Functional Theory to compute electronic structure. Only a few years after these theorems were published, a new approach to Density Functional Theory emerged that would dominate the scene for years to come.

### 3.4 Kohn-Sham Approach

In the next few paragraphs I will explain roughly the procedure for computing the electronic structure using the Kohn-Sham Approach. Firstly, however, I will give you the punch line. Kohn-Sham essentially tries to compute exactly as much of the energy associated with a given charge density as possible and estimate the remainder with approximate functionals. It differs from the Hartree Fock Approach in that the approximation is made at the beginning rather than the end. With Hartree Fock the many electron wave function is approximated as a Slater Determinant, which we have already found to be lacking because it does not take into account the instantaneous repulsion between the electrons.

Kohn and Sham realized that early attempts to use functionals to map charge densities onto energies failed mostly because they were incapable of accurately accounting for the kinetic energy contribution. In contrast, the Hartree Fock Approach was capable of calculating the kinetic energy of the noninteracting electrons exactly. Their idea was to use a hybrid method where orbitals are used to compute as much of the energy contribution as accurately as possible and use functionals to deal with the unknown exchange and correlation effects.

In order to do this, first they constructed a noninteracting reference system that is much like the orbitals of Hartree Fock but differs in one important respect. The average potential in this case is chosen such that the electron density exactly equals the ground state density of the fully interacting system (2, p.41). It is not immediately obvious how you would obtain such a potential, however. What is done in practice is to minimize the Kohn-Sham functional with respect to the electron density under the constraint that the total charge integrate to the correct number of electrons. This process leads to a self consistent procedure which is very much analogous to the one in the Hartree Fock method. From this noninteracting system the kinetic energy is calculated in the usual way. Similarly the classical contributions of the potential energy due to nuclei-electron and electron-electron coulomb repulsion are calculated using the Kohn-Sham orbitals. If it were known, the remainder of the potential energy of the system due to exchange and correlation would be calculated using the appropriate functional. In practice only approximate functionals are available.

### 3.4.1 Approximate Functionals

The part of the contribution to the energy which cannot be computed exactly using the Kohn-Sham approach includes many different parts including exchange, self-interaction and part of the kinetic energy. Finding a local functional that captures all these accurately might seem futile. There is no way to systematically improve the results of a Density Functional Theory calculation (2, p.65) like there is in Hartree Fock theory where you can simply expand your basis set to obtain a more accurate ground state wavefunction (at the cost of additional computation time).

To date the most popular way to approximate the exchange-correlation energy is by calculating it for a very simple system, the uniform electron gas. The uniform electron gas is just what it sounds like: a uniform smear of charge density. The exchange-correlation contribution to the energy can be calculated very accurately for this simple system. Of course real systems do not have uniform charge density everywhere. However, if you estimate the exchange-correlation energy at each point in a real system to be what the contribution would be from a uniform electron gas of density equal to the density at that point in the real system, the results are remarkably good for any system that has slowly (spatially) varying charge densities. This is known as the Local Density Approximation (2, p.70).

The Local Density Approximation works well for systems with slowly varying charge densities. This is the case for most metals which exist in a regular lattice form. For molecules, however, this is not a good approximation. Molecules are not spatially regular and can have very rapidly changing charge densities as a result. More complicated approximations, such as the Generalized Gradient Approximation (2, p.75), have been formulated in order to deal with these situations. These approximations contain information about the charge density and its spatial gradients. The existence of these approximations allowed Density Functional Theory to be extended from material science applications to more general chemical problems where the electron density varies rapidly. Because there is no systematic way of improving functionals, they are often tested against experimental data. Another way of testing functionals is to compare them to more expensive but reliably accurate wave function based methods.

## 4 Recent Work on Reducing Computational Intensity

Although there still remains certain classes of problems that are resistant to resolution by either method, the Hartree Fock Approach and Density Functional Theory have both been successful in producing accurate electronic structures and predicting material properties. Each method has some advantages and some disadvantages, but they both share one difficulty. Computing the electronic structure of systems of more than just a few atoms can be very computationally intensive. Much recent work has focused on decreasing the computational work-

load necessary to perform accurate calculations. Here I briefly describe just a few.

## 4.1 Pseudopotentials and Plane Waves

Pseudopotentials and Plane Waves are two separate techniques that tend to go hand in hand (3 and 4, p.1052). The basic idea is that rapid fluctuations in the solution reflect rapid fluctuations in the underlying potential. This is not an issue when the problem is analytically soluble. However, as we have seen, numerical solutions rely on a superposition of basis functions to express the wave function. Rapidly varying solutions require many basis functions to approximate, effectively increasing the amount of calculation necessary for that portion of the space. There would be no way around this if capturing these rapid oscillations were essential to predicting whatever material properties you are trying to calculate. Fortunately they are not. This allows the use of pseudopotentials to be used that are smoother than the actual potentials. Physically this corresponds to letting the nucleus and a chosen set of core electrons act together as a pseudocore, with the core electrons effectively “shielding” the valence electrons from some of the nuclear charge. As a result the wave functions of the valence electrons near the ions are smoother and require less computation. In order to maintain consistency, the wave function near the pseudocore is required to match the wave function of the actual system at a given distance from the nucleus.

The smoothness of the solution given by pseudopotentials lends itself to a plane wave basis. Plane waves have the advantage of being independent of the positions of the nuclei and being very easy to perform computations on. Also, a complete set of plane waves is complete and any solution can be refined by adding more plane waves into the basis.

## 4.2 Real Space Mesh

Real Space Mesh techniques are essentially a Finite Element or Finite Difference formulation of the electronic structure problem (5). The solution of differential equations is carried out on a mesh of points around the area of interest which can be made more or less dense depending on the resolution that is required. For instance, it is obvious that the charge density will be higher and more rapidly varying near a nucleus than away from a nucleus. An appropriate mesh would be more dense around the nucleus and more sparse in the empty space where the values of the electron density are both low and slowly varying.

These methods were originally developed in the context of engineering problems that were too difficult to solve analytically. Since then they have been extended to solve many problems that rely on the numerical solution of partial differential equations.



### 4.3 Reconstruction

Reconstruction is a very recent attempt at reducing the computational workload involved in computing the electronic structure by reducing the total number of degrees of freedom in a system (6). Consider a slab of atoms with nuclei fixed under the Born Oppenheimer Approximation and the electron density being evaluated on a mesh like the one described above. If the arrangement of the nuclei is periodic it is easy to predict that the electronic structure will also have some periodicity. The Reconstruction approach attempts to leverage this periodicity by only explicitly computing the electronic structure around a few strategically chosen atoms in the slab and using those results to interpolate the electronic structure everywhere else.

This idea can be extended to include both phases of a dynamic simulation. Once the electronic structure has been computed, the forces on the nuclei can be computed. The nuclei then move according to Newton's second law. Here again, instead of explicitly computing the movement of every nucleus, only compute the displacement of a few nuclei and interpolate where the other nuclei on the slab will move. In structural engineering, this is known as the "quasicontinuum" approach (7). This approach has shown promise in reducing the computation time while maintaining accurate results for periodic and semiperiodic systems.

## 5 Acknowledgements

I would like to thank Dr. Dan Negrut for formulating the curriculum for the independent study on which this paper is based and for talking with me regularly about the material throughout the semester. I would also like to thank Dr. L. W. Bruch for carefully reading the paper over and providing many helpful suggestions. Of course, any mistakes that remain are entirely my own responsibility.

## 6 References

1. Szabo, A., & Ostlund, N. S. (1989). *Modern quantum chemistry: introduction to advanced electronic structure theory*. New York: McGraw-Hill.
2. Koch, W., & Holthausen, M. C. (2001). *A chemist's guide to density functional theory*. Weinheim: Wiley-VCH.
3. Segal, M. D., et al. (2002). First-principles simulation: ideas, illustrations and the CASTEP code. *J. Phys.: Condens. Matter* 14, 2717-2744.
4. Payne, M. C., et al. (1992). Iterative minimization techniques for ab initio total-energy calculations: molecular dynamics and conjugate gradients. *Rev. Modern Phys.* 64(4), 1045-1097.
5. Beck, T. L. (2000). Real-space mesh techniques in density-functional theory. *Rev. Modern Phys.* 72(4), 1041-1080.

6. Negrut, D., et al. (2007). Quasi-continuum like reduction of DFT calculations of nanostructures. *Journal of Nanoscience and Nanotechnology*, accepted.
7. Tadmor, E. B., et al. (1996). Quasicontinuum analysis of defects in solids. *Philosoph. mag., A, Phys. condens. matter, Struct. defects mech. prop.* 73(6), 1529-1563.
8. Martin, R. M. (2004). *Electronic structure: basic theory and practical methods*. Cambridge: Cambridge University Press.
9. Basdevant, J. & Dalibard, J. (2005). *Quantum mechanics*. Heidelberg: Springer.

## 7 Contact

If you have any comments, corrections, or questions please forward them to [npschafer@wisc.edu](mailto:npschafer@wisc.edu) and I will try my best to reply as soon as possible.