1

What are Theory, Computation, and Modeling?

1.1 Definition of Terms

A clear definition of terms is critical to the success of all communication. Particularly in the area of computational chemistry, there is a need to be careful in the nomenclature used to describe predictive tools, since this often helps clarify what approximations have been made in the course of a modeling 'experiment'. For the purposes of this textbook, we will adopt a specific convention for what distinguishes theory, computation, and modeling.

In general, 'theory' is a word with which most scientists are entirely comfortable. A theory is one or more rules that are postulated to govern the behavior of physical systems. Often, in science at least, such rules are quantitative in nature and expressed in the form of a mathematical equation. Thus, for example, one has the theory of Einstein that the energy of a particle, E, is equal to its relativistic mass, m, times the speed of light in a vacuum, c, squared,

$$E = mc^2 \tag{1.1}$$

The quantitative nature of scientific theories allows them to be tested by experiment. This testing is the means by which the applicable range of a theory is elucidated. Thus, for instance, many theories of classical mechanics prove applicable to macroscopic systems but break down for very small systems, where one must instead resort to quantum mechanics. The observation that a theory has limits in its applicability might, at first glance, seem a sufficient flaw to warrant discarding it. However, if a sufficiently large number of 'interesting' systems falls within the range of the theory, practical reasons tend to motivate its continued use. Of course, such a situation tends to inspire efforts to find a more *general* theory that is not subject to the limitations of the original. Thus, for example, classical mechanics can be viewed as a special case of the more general quantum mechanics in which the presence of macroscopic masses and velocities leads to a simplification of the governing equations (and concepts).

Such simplifications of general theories under special circumstances can be key to getting anything useful done! One would certainly *not* want to design the pendulum for a mechanical

clock using the fairly complicated mathematics of quantal theories, for instance, although the process would ultimately lead to the same result as that obtained from the simpler equations of the more restricted classical theories. Furthermore, at least at the start of the twenty-first century, a generalized 'theory of everything' does not yet exist. For instance, efforts to link theories of quantum electromagnetics and theories of gravity continue to be pursued.
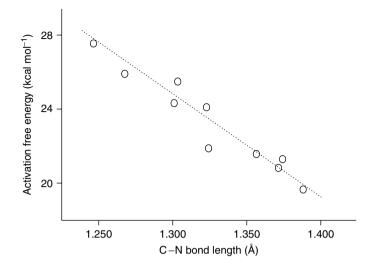
Occasionally, a theory has proven so robust over time, even if only within a limited range of applicability, that it is called a 'law'. For instance, Coulomb's law specifies that the energy of interaction (in arbitrary units) between two point charges is given by

$$E = \frac{q_1 q_2}{\varepsilon r_{12}} \tag{1.2}$$

where q is a charge, ε is the dielectric constant of a homogeneous medium (possibly vacuum) in which the charges are embedded, and r_{12} is the distance between them. However, the term 'law' is best regarded as honorific – indeed, one might regard it as hubris to imply that experimentalists *can* discern the laws of the universe within a finite span of time.

Theory behind us, let us now move on to 'model'. The difference between a theory and a model tends to be rather subtle, and largely a matter of intent. Thus, the goal of a theory tends to be to achieve as great a generality as possible, irrespective of the practical consequences. Quantum theory, for instance, has breathtaking generality, but the practical consequence is that the equations that govern quantum theory are intractable for all but the most ideal of systems. A model, on the other hand, typically involves the deliberate introduction of simplifying approximations into a more general theory so as to extend its practical utility. Indeed, the approximations sometimes go to the extreme of rendering the model deliberately qualitative. Thus, one can regard the valence-shell-electron-pair repulsion (VSEPR; an acronym glossary is provided as Appendix A of this text) model familiar to most students of inorganic chemistry as a drastic simplification of quantum mechanics to permit discrete choices for preferred conformations of inorganic complexes. (While serious theoreticians may shudder at the empiricism that often governs such drastic simplifications, and mutter gloomily about lack of 'rigor', the value of a model is not in its intrinsic beauty, of course, but in its ability to solve practical problems; for a delightful cartoon capturing the hubris of theoretical dogmatism, see Ghosh 2003.)

Another feature sometimes characteristic of a *quantitative* 'model' is that it incorporates certain constants that are derived wholly from experimental data, i.e., they are empirically determined. Again, the degree to which this distinguishes a model from a theory can be subtle. The speed of light and the charge of the electron are fundamental constants of the universe that appear either explicitly or implicitly in Eqs. (1.1) and (1.2), and we know these values only through experimental measurement. So, again, the issue tends to be intent. A model is often designed to apply specifically to a restricted volume of what we might call chemical space. For instance, we might imagine developing a model that would predict the free energy of activation for the hydrolysis of substituted β-lactams in water. Our motivation, obviously, would be the therapeutic utility of these species as antibiotics. Because we are limiting ourselves to consideration of only very specific kinds of bond-making and bond-breaking, we may be able to construct a model that takes advantage of a few experimentally known free energies of activation and correlates them with some other measured or predicted

Figure 1.1   Correlation between activation free energy for aqueous hydrolysis of β-lactams and lactam C−N bond lengths as determined from X-ray crystallography (data entirely fictitious)

quantity. For example, we might find from comparison with X-ray crystallography that there is a linear correlation between the aqueous free energy of activation, $\Delta G^{\ddagger}$, and the length of the lactam C−N bond in the crystal, r_{CN} (Figure 1.1). Our 'model' would then be

$$\Delta G^{\ddagger} = ar_{CN} + b \tag{1.3}$$

where a would be the slope (in units of energy per length) and b the intercept (in units of energy) for the empirically determined correlation.

Equation (1.3) represents a very simple model, and that simplicity derives, presumably, from the small volume of chemical space over which it appears to hold. As it is hard to imagine deriving Eq. (1.3) from the fundamental equations of quantum mechanics, it might be more descriptive to refer to it as a 'relationship' rather than a 'model'. That is, we make some attempt to distinguish between correlation and causality. For the moment, we will not parse the terms too closely.

An interesting question that arises with respect to Eq. (1.3) is whether it may be more broadly applicable. For instance, might the model be useful for predicting the free energies of activation for the hydrolysis of γ-lactams? What about amides in general? What about imides? In a statistical sense, these chemical questions are analogous to asking about the degree to which a correlation may be trusted for extrapolation vs. interpolation. One might say that we have derived a correlation involving two axes of multi-dimensional chemical space, activation free energy for β-lactam hydrolysis and β-lactam C−N bond length. Like any correlation, our model is expected to be most robust when used in an interpolative sense, i.e., when applied to newly measured β-lactam C−N bonds with lengths that fall within the range of the data used to derive the correlation. Increasingly less certain will be application of Eq. (1.3) to β-lactam bond lengths that are *outside* the range used to derive the correlation,

or assumption that other chemical axes, albeit qualitatively similar (like γ-lactam C$-$N bond lengths), will be coincident with the abscissa.

Thus, a key question in one's mind when evaluating any application of a theoretical model should be, 'How similar is the system being studied to systems that were employed in the development of the model?' The generality of a given model can only be established by comparison to experiment for a wider and wider variety of systems. This point will be emphasized repeatedly throughout this text.

Finally, there is the definition of 'computation'. While theories and models like those represented by Eqs. (1.1), (1.2), and (1.3), are not particularly taxing in terms of their mathematics, many others can only be efficiently put to use with the assistance of a digital computer. Indeed, there is a certain synergy between the development of chemical theories and the development of computational hardware, software, etc. If a theory cannot be tested, say because solution of the relevant equations lies outside the scope of practical possibility, then its utility cannot be determined. Similarly, advances in computational technology can permit existing theories to be applied to increasingly complex systems to better gauge the degree to which they are robust. These points are expanded upon in Section 1.4. Here we simply close with the concise statement that 'computation' is the use of digital technology to solve the mathematical equations defining a particular theory or model.

With all these definitions in hand, we may return to a point raised in the preface, namely, what is the difference between 'Theory', 'Molecular Modeling', and 'Computational Chemistry'? To the extent members of the community make distinctions, 'theorists' tend to have as their greatest goal the development of new theories and/or models that have improved performance or generality over existing ones. Researchers involved in 'molecular modeling' tend to focus on target systems having particular chemical relevance (e.g., for economic reasons) and to be willing to sacrifice a certain amount of theoretical rigor in favor of getting the right answer in an efficient manner. Finally, 'computational chemists' may devote themselves not to chemical aspects of the problem, *per se*, but to computer-related aspects, e.g., writing improved algorithms for solving particularly difficult equations, or developing new ways to encode or visualize data, either as input to or output from a model. As with any classification scheme, there are no distinct boundaries recognized either by observers or by individual researchers, and certainly a given research endeavor may involve significant efforts undertaken within all three of the areas noted above. In the spirit of inclusiveness, we will treat the terms as essentially interchangeable.

1.2   Quantum Mechanics

The postulates and theorems of quantum mechanics form the rigorous foundation for the prediction of observable chemical properties from first principles. Expressed somewhat loosely, the fundamental postulates of quantum mechanics assert that microscopic systems are described by 'wave functions' that completely characterize all of the physical properties of the system. In particular, there are quantum mechanical 'operators' corresponding to each physical observable that, when applied to the wave function, allow one to predict the probability of finding the system to exhibit a particular value or range of values (scalar, vector,

etc.) for that observable. This text assumes prior exposure to quantum mechanics and some familiarity with operator and matrix formalisms and notation.

However, many successful chemical models exist that do not necessarily have obvious connections with quantum mechanics. Typically, these models were developed based on intuitive concepts, i.e., their forms were determined inductively. In principle, any successful model *must* ultimately find its basis in quantum mechanics, and indeed *a posteriori* derivations have illustrated this point in select instances, but often the form of a good model is more readily grasped when rationalized on the basis of intuitive chemical concepts rather than on the basis of quantum mechanics (the latter being desperately non-intuitive at first blush).

Thus, we shall leave quantum mechanics largely unreviewed in the next two chapters of this text, focusing instead on the intuitive basis for classical models falling under the heading of 'molecular mechanics'. Later in the text, we shall see how some of the fundamental approximations used in molecular mechanics can be justified in terms of well-defined approximations to more complete quantum mechanical theories.

1.3   Computable Quantities

What predictions can be made by the computational chemist? In principle, if one can measure it, one can predict it. In practice, some properties are more amenable to accurate computation than others. There is thus some utility in categorizing the various properties most typically studied by computational chemists.

1.3.1   Structure

Let us begin by focusing on isolated molecules, as they are the fundamental unit from which pure substances are constructed. The minimum information required to specify a molecule is its molecular formula, i.e., the atoms of which it is composed, and the manner in which those atoms are connected. Actually, the latter point should be put more generally. What is required is simply to know the relative positions of all of the atoms in space. Connectivity, or 'bonding', is itself a property that is open to determination. Indeed, the determination of the 'best' structure from a chemically reasonable (or unreasonable) guess is a very common undertaking of computational chemistry. In this case 'best' is defined as having the lowest possible energy given an overall connectivity roughly dictated by the starting positions of the atoms as chosen by the theoretician (the process of structure optimization is described in more detail in subsequent chapters).

This sounds relatively simple because we are talking about the modeling of an isolated, single molecule. In the laboratory, however, we are much more typically dealing with an equilibrium mixture of a very large number of molecules at some non-zero temperature. In that case, *measured* properties reflect thermal averaging, possibly over multiple discrete stereoisomers, tautomers, etc., that are structurally quite different from the idealized model system, and great care must be taken in making comparisons between theory and experiment in such instances.

1.3.2   Potential Energy Surfaces

The first step to making the theory more closely mimic the experiment is to consider not just one structure for a given chemical formula, but all possible structures. That is, we fully characterize the potential energy surface (PES) for a given chemical formula (this requires invocation of the Born−Oppenheimer approximation, as discussed in more detail in Chapters 4 and 15). The PES is a hypersurface defined by the potential energy of a collection of atoms over all possible atomic arrangements; the PES has $3N − 6$ coordinate dimensions, where N is the number of atoms ≥ 3. This dimensionality derives from the three-dimensional nature of Cartesian space. Thus each structure, which is a point on the PES, can be defined by a vector $\mathbf{X}$ where

$$\mathbf{X} \equiv (x_1, y_1, z_1, x_2, y_2, z_2, \ldots, x_N, y_N, z_N) \tag{1.4}$$

and x_i, y_i, and z_i are the Cartesian coordinates of atom i. However, this expression of $\mathbf{X}$ does not *uniquely* define the structure because it involves an arbitrary origin. We can reduce the dimensionality without affecting the structure by removing the three dimensions associated with translation of the structure in the x, y, and z directions (e.g., by insisting that the molecular center of mass be at the origin) and removing the three dimensions associated with rotation about the x, y, and z axes (e.g., by requiring that the principal moments of inertia align along those axes in increasing order).

A different way to appreciate this reduced dimensionality is to imagine constructing a structure vector atom by atom (Figure 1.2), in which case it is most convenient to imagine the dimensions of the PES being internal coordinates (i.e., bond lengths, valence angles, etc.). Thus, choice of the first atom involves no degrees of geometric freedom – the atom defines the origin. The position of the second atom is specified by its distance from the first. So, a two-atom system has a single degree of freedom, the bond length; this corresponds to $3N − 5$ degrees of freedom, as should be the case for a linear molecule. The third atom must be specified either by its distances to each of the preceding atoms, or by a distance to one and an angle between the two bonds thus far defined to a common atom. The three-atom system, if collinearity is not enforced, has 3 total degrees of freedom, as it should. Each additional atom requires three coordinates to describe its position. There are several ways to envision describing those coordinates. As in Figure 1.2, they can either be a bond length, a valence angle, and a dihedral angle, or they can be a bond length and two valence angles. Or, one can imagine that the first three atoms have been used to create a fixed Cartesian reference frame, with atom 1 defining the origin, atom 2 defining the direction of the positive x axis, and atom 3 defining the upper half of the xy plane. The choice in a given calculation is a matter of computational convenience. Note, however, that the *shapes* of particular surfaces necessarily depend on the choice of their coordinate systems, although they will map to one another in a one-to-one fashion.

Particularly interesting points on PESs include local minima, which correspond to optimal molecular structures, and saddle points (i.e., points characterized by having no slope in any direction, downward curvature for a single coordinate, and upward curvature for all of the other coordinates). Simple calculus dictates that saddle points are lowest energy barriers
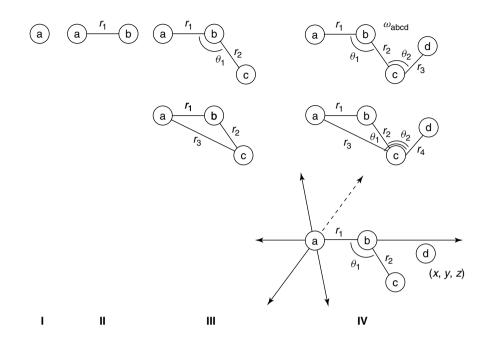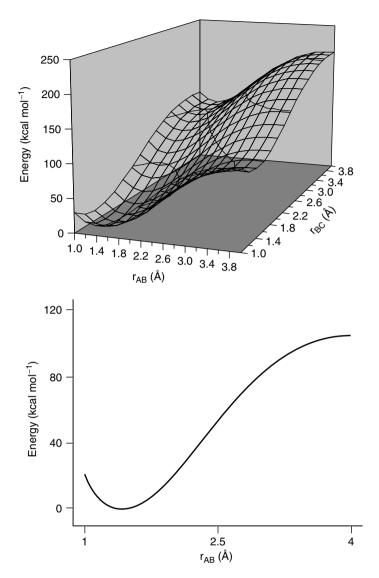
Figure 1.2   Different means for specifying molecular geometries. In frame **I**, there are no degrees of freedom as only the nature of atom 'a' has been specified. In frame **II**, there is a single degree of freedom, namely the bond length. In frame **III**, location of atom 'c' requires two additional degrees of freedom, either two bond lengths or a bond length and a valence angle. Frame **IV** illustrates various ways to specify the location of atom 'd'; note that in every case, three new degrees of freedom must be specified, either in internal or Cartesian coordinates

on paths connecting minima, and thus they can be related to the chemical concept of a transition state. So, a complete PES provides, for a given collection of atoms, complete information about all possible chemical structures and all isomerization pathways interconnecting them.

Unfortunately, complete PESs for polyatomic molecules are very hard to visualize, since they involve a large number of dimensions. Typically, we take slices through potential energy surfaces that involve only a single coordinate (e.g., a bond length) or perhaps two coordinates, and show the relevant reduced-dimensionality energy curves or surfaces (Figure 1.3). Note that some care must be taken to describe the nature of the slice with respect to the *other* coordinates. For instance, was the slice a hyperplane, implying that all of the non-visualized coordinates have fixed values, or was it a more general hypersurface? A typical example of the latter choice is one where the non-visualized coordinates take on values that minimize the potential energy given the value of the visualized coordinate(s). Thus, in the case of a single visualized dimension, the curve attempts to illustrate the minimum energy path associated with varying the visualized coordinate. [We must say 'attempts' here, because an actual continuous path connecting any two structures on a PES may involve any number of structures all of which have the same value for a single internal coordinate. When that

Figure 1.3 The full PES for the hypothetical molecule ABC requires four dimensions to display ($3N - 6 = 3$ coordinate degrees of freedom plus one dimension for energy). The three-dimensional plot (top) represents a hyperslice through the full PES showing the energy as a function of two coordinate dimensions, the AB and BC bond lengths, while taking a fixed value for the angle ABC (a typical choice might be the value characterizing the global minimum on the full PES). A further slice of this surface (bottom) now gives the energy as a function of a single dimension, the AB bond length, where the BC bond length is now also treated as frozen (again at the equilibrium value for the global minimum)

path is projected onto the dimension defined by that single coordinate (or any reduced number of dimensions including it) the resulting curve is a non-single-valued function of the dimension. When we arbitrarily choose to use the lowest energy point for each value of the varied coordinate, we may introduce discontinuities in the actual structures, even though the curve may appear to be smooth (Figure 1.4). Thus, the generation and interpretation of such 'partially relaxed' potential energy curves should involve a check of the individual structures to ensure that such a situation has not arisen.]
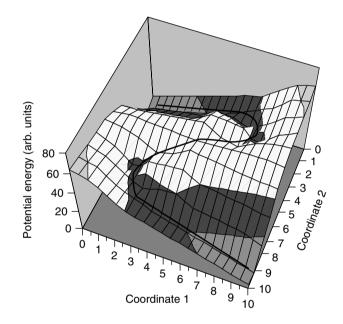

Figure 1.4   The bold line in (a) traces out a lowest-energy path connecting two minima of energy 0, located at coordinates (0,1) and (10,9), on a hypothetical three-dimensional PES – shaded regions correspond to contour levels spanning 20 energy units. Following the path starting from point (0,1) in the upper left, coordinate 1 initially smoothly increases to a value of about 7.5 while coordinate 2 undergoes little change. Then, however, because of the coupling between the two coordinates, coordinate 1 begins *decreasing* while coordinate 2 changes. The 'transition state structure' (saddle point) is reached at coordinates (5,5) and has energy 50. On this PES, the path downward is the symmetric reverse of the path up. If the full path is projected so as to remove coordinate 2, the two-dimensional potential energy diagram (b) is generated. The solid curve is what would result if we only considered lowest energy structures having a given value of coordinate 1. Of course, the solid curve is discontinuous in coordinate 2, since approaches to the 'barrier' in the solid curve from the left and right correspond to structures having values for coordinate 2 of about 1 and 9, respectively. The dashed curve represents the higher energy structures that appear on the smooth, continuous, three-dimensional path. If the lower potential energy diagram were to be generated by driving coordinate 1, and care were not taken to note the discontinuity in coordinate 2, the barrier for interconversion of the two minima would be underestimated by a factor of 2 in this hypothetical example. (For an actual example of this phenomenon, see Cramer *et al*. 1994.)
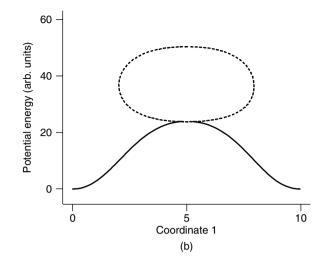
Figure 1.4   (*Continued*)

Finally, sometimes slices are chosen so that all structures in the slicing surface belong to a particular symmetry point group. The utility of symmetry will be illustrated in various situations throughout the text.

With the complete PES in hand (or, more typically, with the region of the PES that would be expected to be chemically accessible under the conditions of the experimental system being modeled), one can take advantage of standard precepts of statistical mechanics (see Chapter 10) to estimate equilibrium populations for situations involving multiple stable molecular structures and compute ensemble averages for physical observables.

1.3.3   Chemical Properties

One can arbitrarily divide the properties one might wish to estimate by computation into three classes. The first is 'single-molecule' properties, that is, properties that could in principle be measured from a single molecule, even though, in practice, use of a statistical ensemble may be required for practical reasons. Typical examples of such properties are spectral quantities. Thus, theory finds considerable modern application to predicting nuclear magnetic resonance (NMR) chemical shifts and coupling constants, electron paramagnetic resonance (EPR) hyperfine coupling constants, absorption maxima for rotational, vibrational, and electronic spectra (typically in the microwave, infrared, and ultraviolet/visible regions of the spectrum, respectively), and electron affinities and ionization potentials (see Chapter 9).

With respect to molecular energetics, one can, in principle, measure the total energy of a molecule (i.e., the energy required to separate it into its constituent nuclei and electrons all infinitely separated from one another and at rest). More typically, however, laboratory measurements focus on thermodynamic quantities such as enthalpy, free energy, etc., and

this is the second category into which predicted quantities fall. Theory is extensively used to estimate equilibrium constants, which are derived from free energy differences between minima on a PES, and rate constants, which, with certain assumptions (see Chapter 15), are derived from free energy differences between minima on a PES and connected transition-state structures. Thus, theory may be used to predict reaction thermochemistries, heats of formation and combustion, kinetic isotope effects, complexation energies (key to molecular recognition), acidity and basicity (e.g., pK_a values), 'stability', and hydrogen bond strengths, to name a few properties of special interest. With a sufficiently large collection of molecules being modeled, theory can also, in principle, compute bulk thermodynamic phenomena such as solvation effects, phase transitions, etc., although the complexity of the system may render such computations quite challenging.

Finally, there are computable 'properties' that do not correspond to physical observables. One may legitimately ask about the utility of such ontologically indefensible constructs! However, one should note that unmeasurable properties long predate computational chemistry – some examples include bond order, aromaticity, reaction concertedness, and isoelectronic, -steric, and -lobal behavior. These properties involve *conceptual* models that have proven sufficiently useful in furthering chemical understanding that they have overcome objections to their not being uniquely defined.

In cases where such models take measurable quantities as input (e.g., aromaticity models that consider heats of hydrogenation or bond-length alternation), clearly those measurable quantities are also computable. There are additional non-observables, however, that are unique to modeling, usually being tied to some aspect of the computational algorithm. A good example is atomic partial charge (see Chapter 9), which can be a very useful chemical concept for understanding molecular reactivity.

1.4   Cost and Efficiency

1.4.1   Intrinsic Value

Why has the practice of computational chemistry skyrocketed in the last few years? Try taking this short quiz: Chemical waste disposal and computational technology – which of these two keeps getting more and more expensive and which less and less? From an economic perspective, at least, theory is enormously attractive as a tool to reduce the costs of doing experiments.

Chemistry's impact on modern society is most readily perceived in the creation of materials, be they foods, textiles, circuit boards, fuels, drugs, packaging, etc. Thus, even the most ardent theoretician would be unlikely to suggest that theory could ever *supplant* experiment. Rather, most would opine that opportunities exist for *combining* theory with experiment so as to take advantage of synergies between them.

With that in mind, one can categorize efficient combinations of theory and experiment into three classes. In the first category, theory is applied *post facto* to a situation where some ambiguity exists in the interpretation of existing experimental results. For example, photolysis of a compound in an inert matrix may lead to a single product species as

analyzed by spectroscopy. However, the identity of this unique product may not be obvious given a number of plausible alternatives. A calculation of the energies and spectra for *all* of the postulated products provides an opportunity for comparison and may prove to be definitive.

In the second category, theory may be employed in a simultaneous fashion to optimize the design and progress of an experimental program. Continuing the above analogy, *a priori* calculation of spectra for plausible products may assist in choosing experimental parameters to permit the observation of minor components which might otherwise be missed in a complicated mixture (e.g., theory may allow the experimental instrument to be tuned properly to observe a signal whose location would not otherwise be predictable).

Finally, theory may be used to predict properties which might be especially difficult or dangerous (i.e., costly) to measure experimentally. In the difficult category are such data as rate constants for the reactions of trace, upper-atmospheric constituents that might play an important role in the ozone cycle. For sufficiently small systems, levels of quantum mechanical theory can now be brought to bear that have accuracies comparable to the best modern experimental techniques, and computationally derived rate constants may find use in complex kinetic models until such time as experimental data are available. As for dangerous experiments, theoretical pre-screening of a series of toxic or explosive compounds for desirable (or undesirable) properties may assist in prioritizing the order in which they are prepared, thereby increasing the probability that an acceptable product will be arrived at in a maximally efficient manner.

1.4.2 Hardware and Software

All of these points being made, even computational chemistry is not without cost. In general, the more sophisticated the computational model, the more expensive in terms of computational resources. The talent of the well-trained computational chemist is knowing how to maximize the accuracy of a prediction while minimizing the investment of such resources. A primary goal of this text is to render more clear the relationship between accuracy and cost for various levels of theory so that even relatively inexperienced users can make informed assessments of the likely utility (before the fact) or credibility (after the fact) of a given calculation.

To be more specific about computational resources, we may, without going into a great deal of engineering detail, identify three features of a modern digital computer that impact upon its utility as a platform for molecular modeling. The first feature is the speed with which it carries out mathematical operations. Various metrics are used when comparing the speed of 'chips', which are the fundamental processing units. One particularly useful one is the number of floating-point operations per second (FLOPS) that the chip can accomplish. That is, how many mathematical manipulations of decimally represented numbers can be carried out (the equivalent measure for integers is IPS). Various benchmark computer codes are available for comparing one chip to another, and one should always bear in mind that measured processor speeds are dependent on which code or set of codes was used. Different

kinds of mathematical operations or different orderings of operations can have effects as large
as an order of magnitude on individual machine speeds because of the way the processors
are designed and because of the way they interact with other features of the computational
hardware.

The second feature affecting performance is memory. In order to carry out a floating-point
operation, there must be floating-point numbers on which to operate. Numbers (or characters)
to be processed are stored in a magnetic medium referred to as memory. In a practical sense,
the size of the memory associated with a given processor sets the limit on the total amount
of information to which it has 'instant' access. In modern multiprocessor machines, this
definition has grown more fuzzy, as there tend to be multiple memory locations, and the
speed with which a given processor can access a given memory site varies depending upon
their physical locations with respect to one another. The somewhat unsurprising bottom
line is that more memory and shorter access times tend to lead to improved computational
performance.

The last feature is storage, typically referred to as disk since that has been the read/write
storage medium of choice for the last several years. Storage is exactly like memory, in the
sense that it holds number or character data, but it is accessible to the processing unit at a
much slower rate than is memory. It makes up for this by being much cheaper and being, in
principle, limitless and permanent. Calculations which need to read and/or write data to a disk
necessarily proceed more slowly than do calculations that can take place entirely in memory.
The difference is sufficiently large that there are situations where, rather than storing on disk
data that will be needed later, it is better to throw them away (because memory limits require
you to overwrite the locations in which they are stored), as subsequent recomputation of the
needed data is faster than reading it back from disk storage. Such a protocol is usually called
a 'direct' method (see Almlöf, Faegri, and Korsell 1982).

Processors, memory, and storage media are components of a computer referred to as 'hard-
ware'. However, the efficiency of a given computational task depends also on the nature of the
instructions informing the processor how to go about implementing that task. Those instruc-
tions are encoded in what is known as 'software'. In terms of computational chemistry, the
most obvious piece of software is the individual program or suite of programs with which the
chemist interfaces in order to carry out a computation. However, that is by no means the only
software involved. Most computational chemistry software consists of a large set of instruc-
tions written in a 'high-level' programming language (e.g., FORTRAN or C++), and choices
of the user dictate which sets of instructions are followed in which order. The collection of all
such instructions is usually called a 'code' (listings of various computational chemistry codes
can be found at websites such as http://cmm.info.nih.gov/modeling/software.html). But
the language of the code cannot be interpreted directly by the processor. Instead, a series
of other pieces of software (compilers, assemblers, etc.) translate the high-level language
instructions into the step-by-step operations that are carried out by the processing unit.
Understanding how to write code (in whatever language) that takes the best advantage of the
total hardware/software environment on a particular computer is a key aspect to the creation
of an efficient software package.

1.4.3   Algorithms

In a related sense, the manner in which mathematical equations are turned into computer instructions is also key to efficient software development. Operations like addition and subtraction do not allow for much in the way of innovation, needless to say, but operations like matrix diagonalization, numerical integration, etc., are sufficiently complicated that different algorithms leading to the same (correct) result can vary markedly in computational performance. A great deal of productive effort in the last decade has gone into the development of so-called 'linear-scaling' algorithms for various levels of theory. Such an algorithm is one that permits the cost of a computation to scale roughly linearly with the size of the system studied. At first, this may not sound terribly demanding, but a quick glance back at Coulomb's law [Eq. (1.2)] will help to set this in context. Coulomb's law states that the potential energy from the interaction of charged particles depends on the pairwise interaction of all such particles. Thus, one might expect any calculation of this quantity to scale as the *square* of the size of the system (there are $n(n - 1)/2$ such interactions where n is the number of particles). However, for sufficiently large systems, sophisticated mathematical 'tricks' permit the scaling to be brought down to linear.

In this text, we will not be particularly concerned with algorithms – not because they are not important but because such concerns are more properly addressed in advanced textbooks aimed at future practitioners of the art. Our focus will be primarily on the conceptual aspects of particular computational models, and not necessarily on the most efficient means for implementing them.

We close this section with one more note on careful nomenclature. A 'code' renders a 'model' into a set of instructions that can be understood by a digital computer. Thus, if one applies a particular model, let us say the molecular mechanics model called MM3 (which will be described in the next chapter) to a particular problem, say the energy of chair cyclohexane, the results should be completely independent of which code one employs to carry out the calculation. If two pieces of software (let's call them MYPROG and YOURPROG) differ by more than the numerical noise that can arise because of different round-off conventions with different computer chips (or having set different tolerances for what constitutes a converged calculation) then one (or both!) of those pieces of software is *incorrect*. In colloquial terms, there is a 'bug' in the incorrect code(s).

Furthermore, it is never correct to refer to the results of a calculation as deriving from the code, e.g., to talk about one's 'MYPROG structure'. Rather, the results derive from the model, and the structure is an 'MM3 structure'. It is not simply incorrect to refer to the results of the calculation by the name of the code, it is confusing: MYPROG may well contain code for several *different* molecular mechanics models, not just MM3, so simply naming the program is insufficiently descriptive.

It is regrettable, but must be acknowledged, that certain models found in the chemical literature are themselves not terribly well defined. This tends to happen when features or parameters of a model are updated without any change in the name of the model as assigned by the original authors. When this happens, codes implementing older versions of the model will disagree with codes implementing newer versions even though each uses the same name for the model. Obviously, developers should scrupulously avoid ever allowing this situation

Table 1.1   Useful quantities in atomic and other units

Physical quantity (unit name)	Symbol	Value in a.u.	Value in SI units	Value(s) in other units
Angular momentum	$\hbar$	1	1.055×10^{-34} J s	2.521×10^{-35} cal s
Mass	m_e	1	9.109×10^{-31} kg	
Charge	e	1	1.602×10^{-19} C	1.519×10^{-14} statC
Vacuum permittivity	$4\pi\varepsilon_0$	1	1.113×10^{-10} C^2 J^{-1} m^{-1}	2.660×10^{-21} C^2 cal^{-1} Å^{-1}
Length (bohr)	a_0	1	5.292×10^{-11} m	0.529 Å
				52.9 pm
Energy (hartree)	E_h	1	4.360×10^{-18} J	627.51 kcal mol^{-1}
				2.626×10^3 kJ mol^{-1}
				27.211 eV
				2.195×10^5 cm^{-1}
Electric dipole moment	ea_0	1	8.478×10^{-30} C m	2.542 D
Electric polarizability	$e^2 a_0^2 E_h^{-1}$	1	1.649×10^{-41} C^2 m^2 J^{-1}	
Planck's constant	h	2π	6.626×10^{-34} J s	
Speed of light	c	1.370×10^2	2.998×10^8 m s^{-1}	
Bohr magneton	μ_B	0.5	9.274×10^{-24} J T^{-1}	
Nuclear magneton	μ_N	2.723×10^{-4}	5.051×10^{-27} J T^{-1}	

to arise. To be safe, scientific publishing that includes computational results should always state what code or codes were used, *to include version numbers*, in obtaining particular model results (clearly version control of computer codes is thus just as critical as it is for models).

1.5   Note on Units

In describing a computational model, a clear equation can be worth 1000 words. One way to render equations more clear is to work in atomic (or theorist's) units. In a.u., the charge on the proton, e, the mass of the electron, m_e, and $\hbar$ (i.e., Planck's constant divided by 2π) are all defined to have magnitude 1. When converting equations expressed in SI units (as opposed to Gaussian units), $4\pi\varepsilon_0$, where ε_0 is the permittivity of the vacuum, is also defined to have magnitude 1. As the magnitude of these quantities is unity, they are dropped from relevant equations, thereby simplifying the notation. Other atomic units having magnitudes of unity can be derived from these three by dimensional analysis. For instance, $\hbar^2/m_e e^2$ has units of distance and is defined as 1 a.u.; this atomic unit of distance is also called the 'bohr' and symbolized by a_0. Similarly, e^2/a_0 has units of energy, and defines 1 a.u. for this quantity, also called 1 hartree and symbolized by E_h. Table 1.1 provides notation and values for several useful quantities in a.u. and also equivalent values in other commonly used units. Greater precision and additional data are available at http://www.physics.nist.gov/PhysRefData/.

Bibliography and Suggested Additional Reading

Cramer, C. J., Famini, G. R., and Lowrey, A. 1993. 'Use of Quantum Chemical Properties as Analogs for Solvatochromic Parameters in Structure–Activity Relationships', *Acc. Chem. Res.*, **26**, 599.

Irikura, K. K., Frurip, D. J., Eds. 1998. *Computational Thermochemistry*, American Chemical Society Symposium Series, Vol. **677**, American Chemical Society: Washington, DC.

Jensen, F. 1999. *Introduction to Computational Chemistry*, Wiley: Chichester.

Jorgensen, W. L. 2004. 'The Many Roles of Computation in Drug Discovery', *Science*, **303**, 1813.

Leach, A. R. 2001. *Molecular Modelling*, 2nd Edn., Prentice Hall: London.

Levine, I. N. 2000. *Quantum Chemistry*, 5th Edn., Prentice Hall: New York.

Truhlar, D. G. 2000. 'Perspective on "Principles for a direct SCF approach to LCAO-MO *ab initio* calculations"' *Theor. Chem. Acc.*, **103**, 349.

References

Almlöf, J., Faegri, K., Jr., and Korsell, K. 1982. *J. Comput. Chem.*, **3**, 385.

Cramer, C. J., Denmark, S. E., Miller, P. C., Dorow, R. L., Swiss, K. A., and Wilson, S. R. 1994. *J. Am. Chem. Soc.*, **116**, 2437.

Ghosh, A. 2003. *Curr. Opin. Chem. Biol.*, **7**, 110.