PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Hardware and Quantum Mechanical Calculations [and Discussion]

E. Wimmer and A. M. Stoneham

Phil. Trans. R. Soc. Lond. A 1992 **341**, 361-371
doi: 10.1098/rsta.1992.0108

Hardware and quantum mechanical calculations

By E. Wimmer

BIOSYM Technologies Inc., 9685 Scranton Road, San Diego, California 92121-3752, U.S.A.

The remarkable progress in the architecture, speed and capacity of computer hardware continues to drive the development of quantum mechanical methods, thus allowing calculations on increasingly complex systems. Using high-end computers, accurate quantum mechanical all-electron studies are now possible for solids such as transition metal compounds containing about fifty atoms per unit cell. Pseudopotential plane-wave methods are being applied to unit cells with 400 silicon atoms, and organic molecules consisting of over 100 atoms have become tractable using *ab initio* methods. Smaller, yet still useful calculations can be carried out on workstations. The combination of graphics workstations and high-performance supercomputers, integrated in tightly coupled heterogeneous networks, has allowed the design of software systems with unprecedented convenience and visualization capabilities. Despite this progress, however, there is still an urgent need for new quantum mechanical methods which converge systematically to the exact solution of Schrödinger's equation while maintaining a reasonable scaling of the computational effort with the system size.

1. Introduction and historical perspective

During the past four decades the development of quantum mechanical methods for accurate calculations on solids and their surfaces has been strongly coupled to the progress in computer hardware. In fact, a variety of hardware aspects play a decisive role in this evolution, especially computer architectures, processor speed, memory size, and external storage devices. While operating systems, compilers, mathematical libraries, software analysis and optimization tools, and networking software represent other important aspects of a computing system, the focus of this contribution is on the interplay between hardware and theoretical methods.

To gain a perspective, the major milestones in the development of computer hardware are reviewed in the context of the development of computational solid state physics. In §2, the three dominant computer architectures are discussed, i.e. vector supercomputers, RISC workstations, and massively parallel machines. Examples of solid state as well as molecular cluster calculations illustrate the current capabilities. Additional remarks are then made on the role of three-dimensional graphics workstations. A section of possible future developments concludes this paper.

The first generation of viable computing machines for quantum mechanical calculations were vacuum-based electronic computers such as M.I.T.s Whirlwind I in the 1950s. The major application of this machine was fluid dynamics. Using spare computing time during the night, these machines were used for pioneering quantum mechanical calculations. The processor speed of this machine, which occupied two

Phil. Trans. R. Soc. Lond. A (1992) **341**, 361–371
Printed in Great Britain

floors of a building, was 0.02 MFlops and the memory size was initially 25000 16-bit words (Ralston & Meek 1976).

Although the invention of the transistor in 1948 by Bardeen, Brattain and Shockley initiated a major technological breakthrough for computer hardware, it was not until 1959 that the first transistor-based or second generation computers, such as the UNIVAC M460, the CDC 1604 and the IBM 7030, became available. With the introduction of the IBM 360 series in April 1964 and the CDC 6600 in the autumn of that year, two highly successful series of 'third generation' computers became more widely available to the growing community of computational solid state physicists. These third generation computers are characterized by the use of large scale integration (LSI) techniques for their semiconducting logic devices. Measured by today's standards, the capabilities of third generation computers were limited not only by slow processors, but also by their small and slow direct-access memories made of arrays of magnetic cores. For example, early versions in the IBM 360 series had only 128 kbytes of magnetic core memory, which is surpassed by more than an order of magnitude in today's lap-top computers. Nevertheless, this third generation of computer hardware, available in the late 1960s and into the 1970s enabled first-principles electronic structure calculations for compounds such as TiC in a sodium chloride structure containing two atoms per unit cell (see, for example, Neckel *et al.* 1976). These calculations provided a detailed understanding of the bonding mechanism and electronic charge distribution in these materials, allowing interpretations and predictions of properties such as X-ray absorption and photoemission spectra. In these early computations, full exploitation of symmetry was essential to keep the size of the hamiltonian matrices under the limit of 40×40.

Initially, these calculations were not self-consistent, solving the bandstructure problem for selected high-symmetry k-points for a given crystal structure and a given crystal potential, which typically was constructed from a superposition of atomic densities. In the late 1960s and throughout the 1970s, machines of increasing speed such as the Control Data CDC 7600 computers as well as further developments of the IBM 360 and 370 mainframes provided the hardware platforms for a large number of electronic structure calculations that were carried out self-consistently, but using a simplified shape for the electron density and potential in the form of the muffin-tin potential (Moruzzi *et al.* 1978). Pseudopotential theory evolved as a particularly successful approach for semiconductors (as reviewed by Cohen & Chelikowsky 1989) and a wealth of calculated data started to provide a systematic understanding of the electronic structures of bulk metals, semiconductors, and simple compounds.

Further progress was only partly due to improved hardware. Advances in theoretical methods and more efficient algorithmic implementations played an equally important role. For example, in the original augmented plane wave (APW) method (Slater 1937) the hamiltonian matrix elements depend on the energies. Consequently, the resulting eigenvalue problem is nonlinear and has to be solved by a tedious, discrete energy search. Similar computational complications occur in the Korringa–Kohn–Rostoker (KKR) method (Korringa 1947; Kohn & Rostoker 1954). The linearization of the APW and the KKR methods (Andersen 1975) constituted a significant step forward. Using a linearized approach, all eigenvalues and eigenvectors for a given k-point can be evaluated in one diagonalization step.

The computational effort in such a linearized method is spread between the evaluation of the hamiltonian and overlap matrix elements and the diagonalization.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

The size of these matrices is linearly proportional to the number of atoms, N. The diagonalization involves a step that scales like N^3. Memory requirements scale like N^2. While computing time presents a soft limit, the available memory size poses a hard constraint on the maximum number of atoms per unit cell, unless one finds 'out-of-core' solutions. For example, the memory of a CDC 6600-class machine could hold about 130000 floating point numbers, which allowed storage of two matrices (the hamiltonian matrix and the overlap matrix) of the size of approximately 200×200 each. Assuming that about 50 LAPW basis functions are needed per atom, only four inequivalent atoms per unit cell could be treated, if no symmetry was taken into account. The smaller number of basis functions needed in the KKR and linearized muffin-tin-orbital (LMTO) methods presented some advantages in this respect.

The unprecedented speed of the CRAY-1 vector supercomputer, which was introduced in 1976, represents another milestone for quantum mechanical calculations. While the CDC Star-100 preceded the CRAY-1 as vector processor, the success of the CRAY-1 was largely due to its balance between three aspects: high scalar performance, fast memory access, and vector registers. The CRAY-1 architecture is based on a fairly small instruction set and thus can be seen as 'vector RISC' (reduced instruction set computing) machine. Compared with previous mainframes, the CRAY-1 offered a fairly large memory of 8 Mbytes. With this new generation of computer hardware came three major advances in all-electron calculations: (i) surface calculations became possible through slab geometries (Krakauer *et al.* 1978); (ii) the muffin-tin approximation to the shape of the charge density and the potential was removed through the introduction of 'full-potential methods' (Wimmer *et al.* 1981); and (iii) total energy calculations became possible for challenging systems such as transition metal surfaces (Weinert *et al.* 1982). Together, these new capabilities opened the door to first-principles investigations of important phenomena such as adsorption (Wimmer *et al.* 1982), surface magnetism (Freeman *et al.* 1982) and surface reconstructions (Fu *et al.* 1984).

The success of the large-scale vector supercomputers encouraged several hardware manufacturers such as Control Data Corporation (through its subsidiary ETA Systems) and the Japanese manufacturers Fujitsu, NEC, and Hitachi to pursue the development of vector supercomputers. IBM offered vector facilities as addition to its 3090 mainframe series. Besides these mainframe machines, vector mini-supercomputers such as Convex and Alliant computers found rapid acceptance in many research departments in the second half of the 1980s.

2. Current hardware trends and capabilities

At present, three major computer architectures determine the hardware for quantum mechanical calculations: (i) shared memory vector machines with multiple processors; (ii) RISC workstations with high scalar speed; and (iii) massively parallel machines with distributed memory. An additional aspect is the availability of three-dimensional graphics capabilities which are integrated in workstations.

Shared memory vector machines with multiple processors such as the CRAY Y-MP, the Fujitsu VP series, the NEX SX series, and new generations of IBM mainframes are continuations of the mainframe developments of the past four decades. At present, NECs 4-processor SX-3 is the fastest shared-memory supercomputer, delivering 90% of its peak 25.6 GFlops for the LINPACK benchmark, whereas the 16-processor CRAY Y-MP/C90 with a theoretical peak

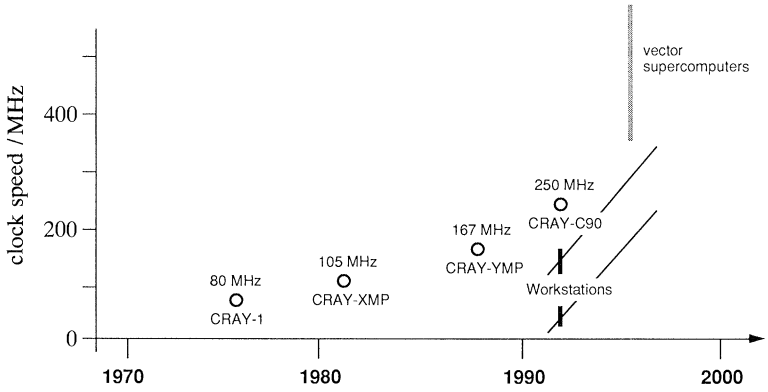

Figure 1. Development of processor speed shown for the example of CRAY supercomputers and typical workstations. Clock speeds are given for the first computer models in a series. Typically, the clock speed is increased during the lifetime of a series. The bar drawn for vector supercomputers around 1995 is an extrapolation based on current plans. Note that at present the processor speed of both supercomputers and workstations appears to increase at the same rate.

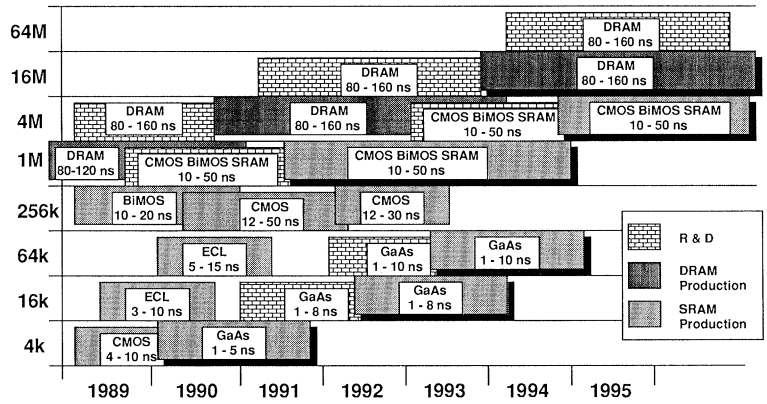

Figure 2. Development of size and access speed of memory components. SRAM stands for static random access memory and DRAM for dynamic random access memory. ECL refers to emitter coupled logic and CMOS is the abbreviation for complementary metal-oxide semiconductors. Note the trade-off between level of integration and access speed. (Courtesy of Cray Research, Inc.)

performance of 16 GFlops provides the greatest throughput for supercomputer work loads (Bell 1992). The maturity of the shared-memory architecture, the extremely high performance, and the familiarity with the programming model for these supercomputers makes these machines the work horses for large-scale quantum mechanical calculations as well as other applications involving solutions of linear algebraic equations.

RISC workstations such as the IBM RS/6000 are offering computing speed for individual researchers which is comparable with the performance of vector supercomputers a decade ago (such as the CRAY X-MP). While there is an impressive gain in speed and performance of workstations, the trend on the high-end supercomputers does not seem to be so different (cf. figure 1). Furthermore, increase in processor speed leads to higher system performance only if the memory access time can keep up. Figure 2 shows the trends in memory component technology. It can be seen that the development of the slower, but larger DRAM technology follows the

same trend as the faster, but smaller CMOS SRAM components. Thus, the issues for machine performance are perhaps better understood if one considers computers not only from the standpoint of processors, but from the viewpoint of memory technology. A critical trade-off exists between large/small and slow/fast memory. Large and slow memories can be combined with slower processors, whereas fast memory is required by faster processors, but then the memory cannot be made so big. Thus, until recently, CRAY supercomputers had a relatively small, but very fast memory using emitter coupled logic (ECL) memory technology (cf. figure 2). For these reasons, the first CRAY Y-MP/8 machines had only 32 Mword of memory. Because of the improvements in the speed of CMOS memory components (cf. figure 2), today's CRAY supercomputers rely on larger CMOS memory technology. The CRAY-2 used a very large DRAM (and later SRAM) memory with only one memory port to the processors. Thus, the performance of the CRAY-2 is best when memory references are infrequent compared with the rate of floating point operations, but is relatively slow in the opposite case. Typical quantum mechanical solid state calculations represent the first case and thus the CRAY-2 continues to be a good, albeit somewhat unique architecture.

The actual improvements in performance can be seen from the following example. The CRAY Y-MP/C90, which was introduced in 1991, has a theoretical peak performance of 16 GFlops. Compared with the CRAY-1, which offered 160 MFlops, this represents an increase of computing speed of two orders of magnitude within 16 years. Significantly, during the same time, the memory size has increased by almost three orders of magnitude. Because of the maturity of this architecture and the existing experience with its programming model, a sustained performance of 10 GFlops for quantum mechanical calculations is now achievable on a machine such as the CRAY Y-MP/C90 without major changes of existing FORTRAN programs. Assuming a scaling of the computational effort with a third power as a function of system size, today's supercomputers would allow calculations for systems about five times as big as ten years ago. Alternatively, one could carry out calculations for 100 different geometries to map the energy hypersurface of a system where electronic structure calculations for only a single geometry were possible ten years ago. The actual progress is even larger because of significant improvements in computational methods and algorithms.

An example of current capabilities of quantum mechanical calculations is the determination of the electronic structure of an organometallic complex containing over one hundred atoms and a total of 1031 basis functions (Mertz & Andzelm 1991). Using a gaussian-based density function approach (Andzelm & Wimmer 1992) on an 8-processor CRAY Y-MP, a total speedup of 7.4 over a single processor was achieved. Table 1 provides an overview of the level of parallelization of the various steps in the calculation. It can be seen that already eight processors lead to a significant deviation from the ideal parallelization. Contrary to intuition, the analytic part of integral generation (three-centre integrals over gaussian functions) has the highest level of parallelization whereas the numerical grid generation shows relatively the lowest speedup. This may indicate that grid based methods are not necessarily 'naturally parallel'.

Workstations such as Apollo, SUN, Hewlett-Packard, and the IBM RS/6000 series introduced a new level of impressive performance at low cost. The key for this technology is a very high level of integration on the processor chip, combined with large and economic memory chips. Current IBM RS/6000 processors operate at a

Table 1. *Speedup due to parallelization of a density functional gaussian-type-orbital calculation of an organometallic Zn complex containing 112 atoms (1031 basis functions)*

(The results were obtained on an 8-processor CRAY Y-MP supercomputer using the DGauss program. The total CPU time to reach SCF convergence was approximately 40 minutes. XC fit refers to the numerical fitting to the exchange-correlation potential.)

	processors		
task	2	4	8
integrals	2	3.9	7.7
density fit	2	3.8	7.0
grid generation	2	3.7	6.2
XC fit	2	3.9	7.4
total	2	3.8	7.1

rate of about 40 MHz. The emerging generation of RISC processors such as the Digital Equipment Alpha chip will operate at about 150 MHz, thus surpassing the clock speed of the original CRAY-1 by almost a factor of two. As mentioned above, the speed of processors is just one of many hardware aspects. In fact, the development of RISC processors has shifted the bottleneck in the design of computer hardware even more from the processor to the issues of memory speed and communication (cf. figure 2).

The capabilities of workstations for quantum mechanical calculations are remarkable. For example, the self-consistent evaluation of the electronic structure for a molecule containing 52 atoms (560 basis functions) requires approximately 10 h of wall-clock time on a cluster of eight IBM RS/6000-320 workstations (S. Brode, personal communication). In these calculations, the workstations are coupled to form a parallel system. Compared with a single processor, a speed-up by a factor of 6.7 was achieved using an eight-processor cluster. In a similar calculation on an organometallic complex with 1017 basis functions, a speed-up of a factor of 10 was observed using a cluster of 14 processors (S. Brode, personal communication).

Another example of the computational capabilities of workstations is provided by a quantum mechanical calculation on a organometallic titanium complex modelling a catalytic reaction (M. Wrinn, unpublished results). The system, which has C_2 symmetry, contains 57 atoms and is described by 502 basis functions. Using a Silicon Graphics Personal Iris 35, the electronic structure and total energy has been evaluated with a density functional approach using numerical basis functions as implemented in the DMol program (Delley 1990) within about 10 h.

Massively parallel machines have captured the minds of many researchers in the field of computational quantum mechanics as well as other scientific disciplines, and there is no doubt that this architecture will play an increasingly important role in computational solid state physics and computational chemistry. The idea of parallel machines is by no means new. For example, the ILLIAC IV was built in 1967 with 64 processors. The design goal was 1 GFlop with 256 processors, but actually only 4 Mflop were achieved with 64 processors (Almasi & Gottlieb 1989). The reason was the use of pre-VLSI technology causing manufacturing and reliability problems. Another parallel architecture was implemented in the ICL DAP as described by Hockney & Jesshope (1981). Interestingly, before starting Cray Research, Seymour Cray worked on a shared memory scalar parallel machine as a successor for the CDC 7600. This parallel machine, which had a round shape and was internally called CDC

8600, was never completed. Instead, Seymour Cray designed the CRAY-1 with one processor, but vector registers.

Today's generation of parallel machines are based on VLSI technology and RISC processors. Currently there are a number of companies competing in this promising field including Intel, Thinking Machines, nCUBE, Kendall-Square, MasPar, Cray Research, IBM, and Meiko. It is now generally recognized that the key issue is not just processor speed, but rather the communication between various processing elements. The single instruction multiple data (SIMD) architecture turns out to be optimal only for specific applications whereas the multiple instruction multiple data (MIMD) architecture is more flexible and thus is emerging as the preferred approach. Memory topology and access speed across processing elements will be the decisive factors for success, as will be the convenience of the programming model and the corresponding system software.

Compared with today's shared memory multiprocessor vector machines and the RISC workstations, the architecture and usability of massively parallel machines is in an early stage. Nevertheless, despite the communication bottleneck, impressive results have already been obtained. The communication bottleneck can be seen in the following example. In a recent calculation of the 7×7 reconstruction of the Si(111) surface, Brommer *et al.* (1992) used a CM-2 parallel machine. In their calculations 25% of the total time was spent in a routine which sums numbers from all processors, in other words, gathering single numbers from all processors and adding them up. Independently, Stich *et al.* (1992) performed a first-principles structural and energetic determination using a similar approach using an Intel iPSC/860 hypercube and a 64-node Meiko i860 Computing Surface. In their work, both Brommer *et al.* and Stick *et al.* could explain the structure and energetics of this technologically important surface. Clearly, both calculations represent major milestones in electronic structure calculations.

3. Heterogeneous networks

In the previous section, three different architectures were discussed, shared memory parallel vector supercomputers, RISC workstations, and massively parallel machines. There are two other hardware aspects that need to be considered in the context of quantum mechanical calculations: the availability of three-dimensional (3D) graphics workstations and high-speed computer networks. The aim of the quantum mechanical calculations is not only a quantitative prediction, expressed in a few numbers such as a binding energy, lattice constants, and bulk modulus, but also physical insight into the atomic-scale phenomena of matter. To this end, visualization has become a significant component. In the early 1980s, graphics displays were often done on separate hardware devices such as the Evans & Sutherland Picture Systems and the transfer of data from the computational hardware to the graphics device was often cumbersome and slow. The introduction of 3D graphics workstations in the early to mid 1980s such as the Silicon Graphics, Stellar, and Ardent computers provided a new hardware concept by offering both impressive computational speed as well as convenient 3D-graphics capabilities. Compared with molecular biology, the use of 3D graphics and animation in computational condensed matter physics and chemistry is not yet fully developed and exploited, although impressive results have already been shown such as oxygen diffusion in bulk silicon (Joannopoulos 1992).

3D graphical user interfaces also allow a unique and unprecedented way to interact

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

with the set-up and the execution of calculations. Most of the quantum mechanical calculations performed so far for solids and surfaces were done in batch jobs using text input files to direct the calculation. This is a tedious process which is quite susceptible to mistakes. Thus, large amounts of precious time is spent in tasks which are quite unrelated to the scientific investigation. This issue of scientific productivity is becoming more urgent as more large-scale numerical simulations are carried out in industrial research laboratories.

These factors have stimulated the development of integrated software systems combining the convenience of graphical user interfaces with the speed and computational capabilities of supercomputers. One such system, called UniChem, was developed at Cray Research with the aim to provide such a fully integrated quantum mechanical environment for molecular and cluster calculations. In this system, it is possible to build a molecular structure interactively on the workstation screen, then select a quantum mechanical method such as semi-empirical, density functional theory, or Hartree–Fock theory. The selection of the computational parameters is done in easily understandable pull-down menus and dialogue boxes. Launching a calculation across the network is literally one click on the mouse button. The progress of the calculation, for example a geometry optimization, can then be monitored on the workstation. Upon completion of the job on the supercomputer, the results can be immediately visualized including molecular structures, molecular orbitals, electron densities, spin densities, and electrostatic potentials. Normal vibrational modes can be visualized through animation of the molecule by selecting lines in the calculated vibrational spectrum.

While this system increases the productivity of expert users, it also allows – for better or worse – the non-expert to begin using these tools. From a user's point of view, the entire system behaves as a single entity, yet a number of different computer architectures are actually involved including the special-purpose graphics processors of the workstation, the vector capabilities of the shared-memory supercomputer, and the scalar processing power of the workstation for post-processing. Perhaps this type of transparently integrated system points into a direction of more generally integrated heterogeneous systems combining the three types of architectures discussed in the previous section.

4. Future

To predict future developments, it may be useful to state again the goals of quantum mechanical calculations of solids and surfaces. The aim could be described as the qualitative understanding and quantitative prediction of atomic-scale phenomena including geometric structures to within about ± 0.001 Å†, energy changes to within ± 0.01 eV for structural changes and ± 0.00001 eV for magnetic changes, the description of dynamic behaviour including chemical reactions with the ability to derive thermodynamic quantities such as free energies, and responses to external forces and electromagnetic fields. This should be possible for ordered as well as disordered systems. Full quantum mechanical descriptions of about 1000 atoms per system or unit cell are meaningful. Beyond that, it should be possible to embed the domain of the quantum mechanical description into the effective field created by the environment, where the environment could be treated by a force-field approach or even a continuum model.

† $1\ \text{Å} = 10^{-10}\ \text{m} = 10^{-1}\ \text{nm}$.

While current methods and hardware clearly do not meet the above goals, there is reason for optimism. For example, ground state structures such as lattice constants can be predicted to within about ± 0.02 Å, systems containing 400 atoms have been treated, and relative energies of the Si(111) 3×3, 5×5, and 7×7 reconstructed surfaces have been predicted from first principles calculations with energy differences of only about 0.02 eV per surface atom (Stich *et al.* 1992; Brommer *et al.* 1992). Perhaps the biggest challenges of quantum mechanical calculation of atomic assemblies is the formulation of theoretical methods that allow systematic convergence to the exact many-body result while maintaining a reasonable scaling of the computational effort. Traditional quantum chemical methods fulfil the first criterion, but not the second. Density functional fulfils rather the second than the first criterion. Definitely, more theoretical investigations are urgently needed. In the exploration of new ideas, the hardware has to be flexible, easy to use, and fast enough to enable the testing of methods and approaches.

Since the pioneering work of Slater (1951), the use of concepts from the homogeneous interacting electron gas as formalized in density functional theory has been the predominant many-body approach in quantum mechanical calculations for solids. The new implementation of gradient corrections to the exchange and correlation term (Becke 1988; Perdew 1986) has been shown to improve binding energies, yet there is no systematic way to improve density functional theory beyond this approximation. Nevertheless, density functional theory lends itself to computational implementation that could scale linearly in the number of atoms. Furthermore, the simplicity of the basic density functional equations offers great freedom in the choice of the most efficient computational implementation. In contrast, for over four decades *ab initio* quantum chemistry has been dominated by the use of gaussian-type basis functions. The major reason is the necessity to evaluate four-index two-electron integrals in Hartree–Fock theory. This requirement does not exist in density functional theory and thus it can be anticipated that a number of different numerical approaches will continue to emerge, making best use of new hardware architectures. Promising developments in this direction include pure plane-wave based methods with appropriate pseudopotentials (see, for example, Teter *et al.* 1989), non-homogeneous plane wave basis sets (F. Gigy, personal communication), and multigrid numerical approaches (J. Bernholc, personal communication).

An exciting development is the combination of quantum mechanics and molecular dynamics as demonstrated by Car & Parrinello (1985). This type of calculation has become possible through a unique combination of an elegant theoretical and computational approach and the speed of hardware. This approach sets the stage for first-principles simulation of phenomena such as adsorption, diffusion, as well as chemical reactions on surfaces, interfaces, and the bulk materials as found in epitaxial growth, chemical vapour deposition, catalysis, electrochemistry, and corrosion. An important aspect in this work is the unification of hitherto different theoretical disciplines, namely electronic structure theory and molecular dynamics. It can be anticipated that in the future this unification of different sub-disciplines will continue thereby leading to sophisticated integrated approaches.

It is reasonable to assume that in such integrated approaches, different methods and algorithms will have to be applied, each requiring different characteristics of the hardware. Thus, integrated heterogeneous hardware architecture may well turn out to be the most appropriate platform for such approaches. Specifically, there will be

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

massively parallel components, shared memory vector components, and fast scalar processors, together with sophisticated 3D visualization capabilities and access to large external storage devices, possibly using optical storage technology. It is quite conceivable that such heterogeneous compute environments do not have to be physically at one location, but could be spread out even over different continents. From this perspective, the band widths of computer networks becomes an equally important hardware component together with the computer architecture, processor speed and memory size.

Virtual reality technologies as applied to scientific problems are creating great excitement, especially in the non-scientific community. While this is indeed a fascinating development, one should not get carried away by impressive pictures which may or may not represent reality. In fact, one could argue that 'one good number is worth more than a thousand pictures' (D. A. Dixon, personal communication). Despite the impressive progress made over the past decades in the theoretical description and computational predictions of phenomena in solids and on their surfaces, we should be aware of the great complexity and subtle intricacies of even relatively simple systems such as pure silicon or carbon. We have to realize that our tools are still very crude, capturing only small bits of real systems. It can be hoped that as we improve our theoretical methods and the hardware tools, we also increase our respect for the beautiful architecture and marvellous properties of condensed matter.

The author thanks his former colleagues at Cray Research for many fruitful and stimulating discussions, and especially Charles Grassl for the information on memory component developments. Furthermore, the communications with Stephan Brode (BASF) and Michael Wrinn (BIOSYM) have been extremely useful and are gratefully acknowledged.

References

Almasi, G. S. & Gottlieb, A. 1989 *Highly parallel computing*. Redwood City, California: The Benjamin/Cummings Publishing Company.

Andersen, O. K. 1975 *Phys. Rev.* B **12**, 3060.

Andzelm, J. & Wimmer, E. 1992 *J. chem. Phys.* **96**, 1280.

Becke, A. D. 1988 *Phys. Rev.* A **38**, 3098.

Bell, G. 1992 *Science, Wash.* **256**, 64.

Brommer, K., Needles, M., Larson, B. E. & Joannopoulos, J. D. 1992 *Phys. Rev. Lett.* **68**, 1355.

Car, R. & Parrinello, M. 1985 *Phys. Rev. Lett.* **55**, 2471.

Cohen, M. L. & Chelikowsky, J. R. 1989 *Electronic structure and optical properties of semiconductors*, 2nd edn. Berlin, New York: Springer-Verlag.

Delley, B. 1990 *J. chem. Phys.* **92**, 508.

Freeman, A. J., Krakauer, H., Ohnishi, S., Wang, D. S., Weinert, M. & Wimmer, E. 1982 *J. Phys., Paris* **43**, C7, 167–176.

Fu, C. L., Ohnishi, S., Wimmer, E. & Freeman, A. J. 1984 *Phys. Rev. Lett.* **53**, 675.

Hockney, R. W. & Jesshope, C. M. 1981 *Parallel Computers*. Adam Hilger.

Joannopoulos, J. D. 1992 *Science, Wash.* **256**, 44.

Koelling, D. D. & Arbman, G. O. 1975 *J. Phys.* F **5**, 2041.

Kohn, W. & Rostoker, N. 1954 *Phys. Rev.* **94**, 1111.

Korringa, J. 1947 *Physica* **13**, 392.

Krakauer, H., Posternak, M. & Freeman, A. J. 1978 *Phys. Rev. Lett.* **41**, 1072.

Mertz, J. E. & Andzelm, J. 1991 *Cray Channels, Fall 1991*, Cray Research, Eagan, p. 10.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Moruzzi, V. L., Janak, J. F. & Williams, A. R. 1978 *Calculated electronic properties of metals*. New York: Pergamon.

Neckel, A., Rastl, P., Eibler, R., Weinberger, P. & Schwarz, K. 1976 *J. Phys.* C **9**, 579.

Perdew, J. P. 1986 *Phys. Rev.* B **33**, 8822.

Ralston, A. & Meek, C. L. (ed.) 1976 *Encyclopedia of computer science*, p. 1456. New York: Petrocelli/Charter.

Slater, J. C. 1937 *Phys. Rev.* **51**, 846.

Slater, J. C. 1951 *Phys. Rev.* **81**, 385.

Stich, I., Payne, M. C., Kingsmith, R. D., Liu, J.-S. & Clarke, L. J. 1992 *Phys. Rev. Lett.* **68**, 1351.

Teter, M. P., Payne, M. C. & Allan, D. C. 1989 *Phys. Rev.* B **40**, 12255.

Weinert, M., Wimmer, E. & Freeman, A. J. 1982 *Phys. Rev.* B **26**, 4571.

Wimmer, E., Krakauer, H., Weinert, M. & Freeman, A. J. 1981 *Phys. Rev.* B **24**, 864.

Wimmer, E., Freeman, A. J., Weinert, M., Krakauer, H., Hiskes, J. R. & Karo, A. M. 1982 *Phys. Rev. Lett.* **48**, 1128.

Discussion

A. M. Stoneham (*Harwell Laboratory, Didcot, U.K.*): In your excellent user-friendly, general-purpose code you still had a theorist running it for an experimenter. Is the theorist needed in this role? Surely the experimenter wants to do it all.

E. Wimmer: Eventually, experimenters should be able to carry out more and more computations by themselves. However, theorists will always be needed to handle difficult problems beyond routine calculations. In fact, user-friendly computer codes and interfaces will free the theorists to focus on challenging problems thus advancing the capabilities and usefulness of computer simulations while non-experts can accomplish many routine calculations on their own. Clearly, the boundary between advanced and routine work will shift with time.

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES