Archives of Biochemistry and Biophysics xxx (2015) xxx–xxx

Contents lists available at ScienceDirect

Archives of Biochemistry and Biophysics

journal homepage: www.elsevier.com/locate/yabbi

Review

Technical advances in molecular simulation since the 1980s

Martin J. Field

Dynamo Team, DYNAMOP Group, UMR 5075, Université Grenoble 1, CNRS, CEA, Institut de Biologie Structurale, 71 Avenue des Martyrs, CS 10090, 38044 Grenoble Cedex 9, France

ARTICLE INFO

Article history:
Received 3 February 2015
and in revised form 5 March 2015
Available online xxxx

Keywords:
Algorithms
Hardware
Molecular simulation
Software
The last 30 years

ABSTRACT

This review describes how the theory and practice of molecular simulation have evolved since the beginning of the 1980s when the author started his career in this field. The account is of necessity brief and subjective and highlights the changes that the author considers have had significant impact on his research and mode of working.

Introduction

The first electronic computers were developed during and immediately after the Second World War. Their capability in tackling complex numerical problems was obvious and they rapidly gained an important foothold in scientific research. One of the areas that blossomed due to the advent of computing was that of molecular simulation in which the behavior of molecular systems was modeled at the atomic level. Previously calculations had to be done by hand or with primitive mechanical or electromechanical devices, but the new machines permitted the use of either novel or hitherto impractical techniques. Examples from the 1950s and 1960s include the use of Monte Carlo and molecular dynamics (MD)[1] methods to study gases and liquids [1], and the extension of molecular orbital (MO) Hartree–Fock (HF) and configuration interaction (CI) quantum chemical (QC) techniques to molecules containing more than a handful of atoms [2].

Progress since these early days has been rapid, and computer simulation is now an integral part of the scientific process. The most public indication of this in the area of molecular science has been the award of two recent Nobel Prizes in Chemistry. The first, in 1998, was given to Walter Kohn and John Pople for a mixture of theoretical and computational work, the former for his development of density functional theory (DFT), and the latter for his development of computational methods in quantum chemistry [3]. By contrast, the 2013 Chemistry Prize was given to three recipients, Martin Karplus, Michael Levitt and Arieh Warshel, all of whom are primarily computationalists. They were cited for their development of multiscale models for complex chemical systems in the late 1960s and 1970s [4].

This review provides an account of the changes in molecular simulation since the early 1980s when the author first started working in this area. The items that have been chosen are necessarily highly subjective and include technical advances that have impacted the practice of molecular simulation in addition to purely scientific developments.

Hardware

An autobiography

The world of scientific computing at the beginning of the 1980s was very different from that nowadays. Electronic hand-held calculators had only just become inexpensive enough to be widely used, and personal computing was in its infancy, with the release

E-mail address: martin.field@ibs.fr

[1] *Abbreviations used:* AMD, Advanced Micro Devices; CDC, Control Data Corporation; CI, configuration interaction; CM, Connection Machine; CPU, central processing unit; CRT, cathode ray tube; DEC, Digital Equipment Corporation; DFT, density functional theory; DIIS, direct inversion in the iterative subspace; Fortran, formula translation; FLOP, floating point operations per second (in the text double precision values are used where possible); FMM, fast multipole method; GB, generalized Born (method); GB, gigabyte (unit); GPU, graphics processing unit; HF, Hartree–Fock; HP, Hewlett–Packard; IBM, International Business Machines; IMSL, International Mathematics and Statistics Library; MB, megabyte; MD, molecular dynamics; MIT, Massachusetts Institute of Technology; MM, molecular mechanical; MO, molecular orbital; MPI, message passing interface; NAG, Numerical Algorithms Group; NIH, National Institutes of Health; OS, operating system; PB, Poisson–Boltzmann; PC, personal computer; PME, particle mesh Ewald; PS, picture system; QC, quantum chemical; SCF, self consistent field; TPS, transition path sampling; US, umbrella sampling; VAX, virtual address extension; VMS, virtual memory system.

http://dx.doi.org/10.1016/j.abb.2015.03.005
0003-9861/

of the Apple II at the end of the 1970s and that of the first IBM personal computer (PC) at the beginning of the 1980s. At this time, computational chemistry seemed to rely on large mainframe computers located in centralized centers, although it was being realized — in the US at least — that smaller so-called minisupercomputers within a research group could be just as productive and more cost-effective [5].

All the author's PhD work, in quantum chemistry, was done at the University of Manchester Regional Computer Center in the UK. The bulk of calculations was done on the two CDC 7600 computers at the center, although a more powerful CDC Cyber 205 computer became available towards the end of the author's thesis. Interaction with these mainframes was done via intermediate front-end computers, upon which files for job submission could be prepared and back to which results files would be transmitted once a job had finished. Initially communication to these front-ends was done via teleprinter, which had to be periodically reloaded with rolls of paper (!), but these were later superseded by CRT video display terminals that the author's host research group managed to acquire. However, this did not mean an end to the use of paper, as it was customary to print out all job output for perusal and for subsequent storage, leading to many boxes of results stacked around the office. Electronic backups, when performed, were carried out on 10.5 inch magnetic reel tapes which had to be done by physically visiting the computer center where the tape drives were to be found.

The CDC 7600s were notable machines. They were designed by Seymour Cray at CDC before he left to found his own company, Cray Research, which dominated supercomputer design in much of the 1970s and 1980s. The 7600s were actually quite old in the 1980s as they had appeared on the market at the beginning of the 1970s, and had been overtaken in terms of computing power by the Cray-1 which was released in the mid-1970s. The CDC 7600 used a word length of 60 bits, different from the 32 or 64 bits that are commonly used today, had an achievable top performance of approximately 10 MFLOP, and a dual memory system consisting of 64 Kword of small core memory and 192 Kword of large core memory (very roughly, 0.5 MB and 1.5 MB, respectively). This small memory presented quite a challenge when writing QC algorithms, as it limited the size of problems that could be handled and meant that there was much reading and writing of intermediate results to and from disk. As an example, an important step in HF calculations is the diagonalization of the Fock matrix to give the orbitals and their energies. If the number of basis functions in the calculation is N, the Fock matrix, which is symmetric, requires $N(N+1)/2$ words of memory, whereas the matrix of orbitals, which is square, requires N^2 words of memory. This implies that the maximum number of basis functions that could be handled in small core memory was approximately 200, given the assumption that these are the two principal arrays required for diagonalization.

After his thesis the author went as a postdoctoral fellow to the Department of Chemistry at Harvard University in Cambridge, Massachusetts. Here the principal computers for research were organized at a departmental level, in contrast to Manchester where they were in a multi-university regional center. The departmental machines were made up of mini-supercomputers, namely two DEC VAX 11/780s, supplemented later by a Convex C2. Although the VAXes were slower machines than the CDC 7600s that the author had used previously, with approximate speeds of 1 MFLOP, they were much more intuitive to use.

The Boston area was an exciting place to be in the mid to late 1980s as it was one of the centers of the computer industry. There were established companies, such as DEC, but also younger start-ups, including Alliant, Apollo and Thinking Machines. The latter was located at MIT and was one of the first manufacturers to produce computers with a massively parallel architecture. The author had the privilege of working on one of their earlier machines, the CM-2 Connection Machine, for a short while. The CM-2 consisted of a cube approximately 1.5 m per side and could house up to 65,536 (2^{16}) very simple single-bit processors. Numerical calculations were accelerated by adding floating-point units, typically one per group of 32 of the simpler processors. Although a full machine had, in principle, a performance in excess of 1 GFLOP, this was difficult to achieve for the type of molecular simulation algorithm that the author was interested in.

The trend to more localized computer power continued in the author's next major appointment at the National Institutes of Health (NIH) in Bethesda, Maryland. The host group, in the Division of Computer Research and Technologies, relied on their own Apollo workstations for most of their calculations, although there was also some access to the NIH's centralized IBM 3090 mainframes.

When the author left to set up his own group at the beginning of the 1990s in the Institute of Structural Biology in Grenoble, France, it was clearly preferable to purchase personal workstations rather than employ external resources. At the time, the most cost-efficient machines for calculation were HP 9000 700 series workstations, and it was on five of these that the group relied for the next several years. However, when it came time to replace these machines towards the end of 1990s, the choice was made to switch to commodity PCs, using Intel chips, as these had increased significantly in computational power during the decade and were unbeatable on cost.

Today the situation is very similar and Intel-based PCs, or their AMD equivalents, dominate computational research. Even the majority of the most powerful supercomputers consist of clusters of tightly-coupled nodes with Intel or AMD chips, although some other companies, including Fujitsu, IBM and ShenWei, also manufacture their own processors.

To terminate this section, it is worth reflecting on the progress made in general-purpose computational hardware in the last 30 years. The machines the author first used were room-size with approximate computational speeds and memory sizes of 1 MFLOP and 1 MB, respectively. Since then there has been at least four orders of magnitude improvement in performance as, currently, a high-end single processor PC, that comfortably fits onto a desktop, will have a speed of several GFLOPs and a memory of several GBs.

Specialized computational hardware

The previous section gave a brief and rather simplified autobiographical overview of the changes in general computing hardware over the last 30 years. This section describes some more specialized topics. One of these is the design of processors that are specific to certain types of calculation. In the molecular domain most work appears to have been done on chips for accelerating MD simulations with molecular mechanical (MM) force fields. Examples include the FASTRUN accelerator of Fine et al. [6], the Gemmstar project of Brooks et al. [7] that was being undertaken by the author's host group when he was at the NIH, and the MDGRAPE series of processors that are currently in their fourth incarnation [8,9]. The problem with these efforts is that the production of specialized chips is extremely demanding and so can be very slow in a niche market such as that for MD simulations unless substantial resources are provided. This has meant that MD-specific processors are often rendered obsolete by the newest generation of general-purpose processors because they cannot be developed rapidly enough.

Please cite this article in press as: M.J. Field, Arch. Biochem. Biophys. (2015), http://dx.doi.org/10.1016/j.abb.2015.03.005