

Questions and Answers about Molecular Dynamics on the DAP

by

David Fincham
DAP Support Unit, Queen Mary College

1. What is the DAP?

The Distributed Array Processor is a new type of computing device produced by ICL. It consists of 4096 processors arranged as a 64x64 array, each processor being capable of performing single bit arithmetic operations. The processors operate simultaneously, executing the same instruction on their own data. To provide more flexibility each processor has an activity register, under program control, which means that in effect it is either 'on' or 'off' for any particular instruction. The store associated with each processor consists of 4096 bits, making 2 Mbytes in total. (An 8 Mbyte version is also available). The whole device forms a specialised store module of a 2900 series ICL computer, so that the processing power is distributed into the store, avoiding problems of communication between main store and separate processing units which can degrade the performance of more conventional computers.

2. How is the DAP programmed?

The DAP is programmed in DAP Fortran, which is a version of Fortran with extensions for expressing parallel operations on matrices and vectors. For example, two matrices can be added with a single statement, rather than a double DO loop as in traditional Fortran. There are various methods of performing indexing operations, a particularly valuable one being the use of a logical expression in place of subscripts as a method of controlling the activity of the processors. These features result in a very simple, elegant and flexible language: most users describe it as 'fun'. Similar array-processing extensions will appear in the next ANSI Fortran standard.

3. What problems are suitable for the DAP?

To be suitable for the DAP a problem must satisfy two conditions. First, it must have a high degree of parallelism, so that many logically independent operations can be performed simultaneously. This sounds rather restrictive, but in fact parallelism is intrinsic to many problems, though the programmer may not be aware of it because he has previously been limited to algorithms expressed in strictly sequential languages. Second, for efficient utilisation of the DAP the problem must be mapped in some way onto a 64x64 array (which can also be regarded as a vector of 4096 elements). Thus the programmer must be aware of and able to exploit the architecture of the computer. It has always been true that users working with number-crunching problems at the limit of available computer power have had to do this to some extent.

4. Is the DAP suitable for MD calculations?

Yes. Each step of an MD simulation consists of two essential parts, the force evaluation and the integration of the equations of motion. The starting point of the step is a set of coordinates, three Cartesian components for each of the N particles, forming three 'vectors' of length N. In a simple MD program these are used to form the set of pair separations which can be regarded as an antisymmetric NxN matrix. From these are calculated the matrix of pair forces, and this evaluation is entirely parallel in nature, that is, it can be performed simultaneously for each pair. The pair forces are summed to give a vector of total forces, one for each of the N particles. The equations of motion of the particles can then be integrated, which again may be done in parallel. To use the DAP most efficiently the value of N should be a multiple of 64, 256 being a commonly used number for liquid simulations. If larger systems are to be studied some kind of list technique to avoid considering all pairs in the system, is advantageous as on other computers [1]. The neighbourhood list technique has been implemented on the DAP by Steve McQueen of ICL. He uses a logical mask rather than an actual list of indices; on the DAP logicals are single bit quantities so the storage problems which have made the technique unpopular on the 7600 do not arise. The in-range interactions then need to be 'packed' onto the DAP in 64x64 blocks. (A process akin to the GATHER operation on the Cray [2]). Other techniques are being investigated and I hope to report on these in a future article.

5. How does the DAP compare with the Cray?

Since some of the readers of this Newsletter are familiar with the Cray it is worth comparing the two machines. The Cray works by pipelining, that is overlapping, arithmetic operations in one floating point processor, rather than performing them simultaneously in many processors. However, to be overlapped the operations must be logically independent and so in principle capable of parallel processing, and thus algorithms for the two machines tend to be similar. On the Cray the performance is fairly independent of the vector length, whereas in the DAP blocks of data of size 4096 give maximum efficiency. In other respects the DAP is more flexible and easier to program. The indexing techniques can be used to handle conditionals, whereas these are difficult on the Cray, the only possibility being the choice of one out of two numerical values depending on the sign of a third. On the Cray also the common operation of summing the elements of a vector does not 'vectorise', while the bit-serial arithmetic on the DAP makes it a very rapid operation. (The bit serial arithmetic has other surprising effects, for example SQRT is faster than multiply!). Overall for floating point arithmetic the performance of the DAP lies between that of the CDC 7600 and the Cray. It should be remembered that the Cray-1 is a separate multi-million pound computer, whereas the DAP is an order of magnitude cheaper. In fact, since it provides additional storage on the ICL computer to which it is attached at little greater cost than straight storage alone it can be regarded as almost free by comparison with the Cray.

6. Can I use the DAP?

The first production DAP has been installed at Queen Mary College and is regarded by the Computer Board as a national facility for parallel processing. Communications with the 2980 computer which acts as 'host' to the DAP is possible via PSS, Metronet or SERC net. Anyone interested in using the DAP for MD calculations is invited to contact me in the first instance, though the use of substantial amounts of time requires a formal application to the SERC. The DAP Support Unit exists to give advice, assistance and education for users, and also to provide software. I plan to make available a fairly general purpose MD program for handling rigid polyatomic molecules with interactions of the site-site form; and perhaps to extend this to include constant pressure dynamics, point quadrupoles and fractional charges with Ewald sum.

References

- [1] D.M. Heyes, CCP5 Newsletter No.2.
- [2] D. Fincham and B.J. Ralston, Comp. Phys. Commun. 23 (1981) 127-134.