Parallel promises in computing

Igor Aleksander

The Connection Machine. By W. Daniel Hillis. *MIT Press:1986. Pp.190. $22.50, £22.50.*

THE jacket of this book heralds "the connection machine" as "the second state in the evolution of digital computers"; the first sentence reads "Someday, perhaps soon, we will build a machine that will be able to perform the functions of the human mind, a thinking machine". It must be said immediately that while the concept of the connection machine is intriguing, and the book tells us a lot about the engineering behind it, any reference to "thinking machines" is thoroughly misleading.

The primary characteristic of the connection machine is parallelism. The prototype discussed by Hillis contains approximately 64,000 simple processors that may be linked in large parallel groups, and it is one of the strengths of the book that the author discusses his design decisions with great clarity. The main attraction of the machine is that it allows a user to research the behaviour of different "shapes" of parallel machines; for example, the properties of processors connected in a ring are very different from those of processors connected in a matrix, and the connection machine may be configured into these structures and many more. It almost becomes possible to believe the claim that the machine marks a new era in digital computing, as the possibility of working with different "shapes" of parallel structure is awkward (that is, slow) on conventional computers.

Hillis's main problem, in common with that of designers of similar machines, was to devise schemes that allow processors to be selected from a pool, and then to provide a sufficiently rich set of pre-defined connections so that the user's desired structure can be realized. Clearly, alternatives had to be found to the obvious solution of connecting every processor to every other one: that would require just too many connections, most of which would be unused during the execution of a particular task. Hillis lucidly reviews the various options and describes why he settled on a "Boolean *n*-cube topology" (the details of which might interest the topologist as an engineer's choice among theoretical alternatives). The end result of the parallel/serial compromise is an achievement of 1,000 million instructions per second (Mips) — the average modern machine in a research laboratory probably delivers about *one* Mip. Another achievement is that this machine may be programmed by means of an extension of LISP, easily mastered by those who use the language anyway. Hillis goes to some pains to argue that this is an arbitrary choice and advanced features of the machine could equally be applied to other languages.

But what scientific advances are likely to follow from the application of the connection machine? As is often the case with innovators in engineering, Hillis is less comfortable with the answer to this question than he is with the technical description of his machine. He points principally to questions arising in research on artificial intelligence, and suggests that image-processing could benefit by arranging the processors into two-dimensional arrays. Although this is true, and possibly useful, processor arrays have been built for some time; for example, Michael Duff's CLIP machines at University College London have been doing workmanlike tasks in practical image-processing for about 15 years, and those who have used them have probably discovered most of the principles that such methodology is likely to reveal.

Closer to the heart of AI, Hillis suggests that the connection machine might implement knowledge stores called "semantic networks". Here, two processors might have respective labels such as "lemon" and "sour", with a connection leading from the former to the latter and labelled as "tastes". The methodology was first introduced by Patrick Winston in 1970 and it is a matter of debate whether the lack of applications of such schemes in more recent research is solely due to the lack of a sufficiently parallel machine. Away from AI, Hillis proposes that the connection machine might be used to simulate silicon circuits (very large scale integrated systems), but at the rate of one processor (containing thousands of transistors) for each simulated transistor this has an air of overkill about it.

All this talk of application only occupies a few pages in the first chapter and is evidently not central to the author's interest. Hillis perhaps gets closer to the real potential for the connection machine in the last two pages of the book, where he suggests that through such "granular" computers of ever-decreasing grain-size, computational theory will extend its explanational powers to an increasing number of natural phenomena. This is where the payoff may really be found. The natural phenomena which, in my opinion, are most in need of attention are those concerning brain-like networks of cells, and the connection machine is an ideal simulation tool for such work. But despite the references to "thinking machines", Hillis does not actually discuss this possibility. The bibliography, although extensive, mostly bears little relation to the text and is notable for its omissions. The most serious of these is notice of the modelling of brain-like systems that has been pursued by Geoffrey Hinton at Carnegie Mellon University since the late 1970s. Exciting discoveries are being made about the way in which cellular learning systems can learn to "understand" concepts such as the shape of objects "seen" through a television camera, or build internal representations of "Italian name" against "English name" in a telephone directory database *without* having been programmed to do so directly. Investigations of this kind have much to do with what is meant by the word "thinking" and, if carried out through the connection machine, will contribute a great deal to understanding the functions of the brain. Indeed, simulations of Hinton's work are clumsy and counter-productive on conventional machines, and it is in this sense that the connection machine provides a useful tool with which "thinking" may be researched. To confuse this with the creation of a "thinking machine" is a category mistake of textbook quality. Moreover, despite the author's implication that the MIT-associated connection machine is unique, it must be seen as taking its place among a number of similar developments such as the "Non-Von" machine at Columbia University and John Darlington's ALICE machine at Imperial College, London.

The Connection Machine will be appreciated best by those who require a lucid description of a computer design that is truly ingenious and indicative of the trend in computer science hardware research towards massive parallelism. It is not for those with no knowledge of computer engineering at all, nor for others who hope to be enlightened on what a "thinking machine" might be. □

Igor Aleksander is Kobler Professor of Information Technology, Department of Computing, Imperial College, University of London, London SW7 2BZ, UK.

The real Frederick Soddy*

*As a number of people have pointed out, the photograph supporting the review of the collection of essays on Soddy (*Nature* **320,** 691; 1986) was in fact of Ernest Rutherford.