DR. JACK SCHWARTZ
Courant Institute, New York University

The Supercomputer Situation

As the U. S. moves to meet the very ambitious supercomputer plans announced by Japan, the general level of architectural activity in the supercomputer area has been rising rapidly. Many universities have become involved; over fifty designs for parallel computers of various types have been proposed and more are coming. This creates a substantial problem of choice for the administrative agencies (principally DARPA, DOE, and NSF) that will have to set the main directions of research funding in this area.

Figure 1 gives a rough taxonomy of one major subclass of the supercomputers that have been proposed. It shows the parallel machines that are based on substantial individual processing elements where "substantial" means at least a high-performance microprocessor. These machines are to be contrasted with the other main class, shown in Fig. 2 -- machines that are composed of minimal processing elements, e.g., at an extreme, single bit processors. The first, "substantial processor" class of machines tend to use "universal" interconnections; machines of the second class tend to be more severely constrained in their choice of interconnection scheme by silicon layout considerations.

Figure 1 shows the substructure in the "substantial processor" machine subfamily.

A subfamily of these consist of packet-switching machines which use various types of optimal communication nets for coupling many microcomputers very efficiently and tightly. Among the machines of this subclass, there is a significant group of machines which are designed to be programmed in a fairly conventional "procedural" style -- one in Illinois, one at N.Y.U., one being developed commercially by Sullivan Associates, and lately one at Cal Tech having a slightly different, message-passing rather than shared memory design. Down the next branch of the taxonomic tree shown in Fig. 1, we find a class of data-flow machines distinguished by a different sort of programming paradigm; these will be discussed in more detail by Professor Arvind. Finally, the "tightly coupled" family of machines shown in Fig. 1 includes another branch on which appears the circuit switching, optimal communication net, TRAC machine developed by the University of Texas. Finally, getting further away from the ULTRA class of machine shown in the lower left hand of Fig. 1, one begins to find computing devices that from the point of view of the relatively tightly coupled "ULTRA" or "TRAC" machines are more esoteric; these use various types of supplemented nearest-network communication nets.

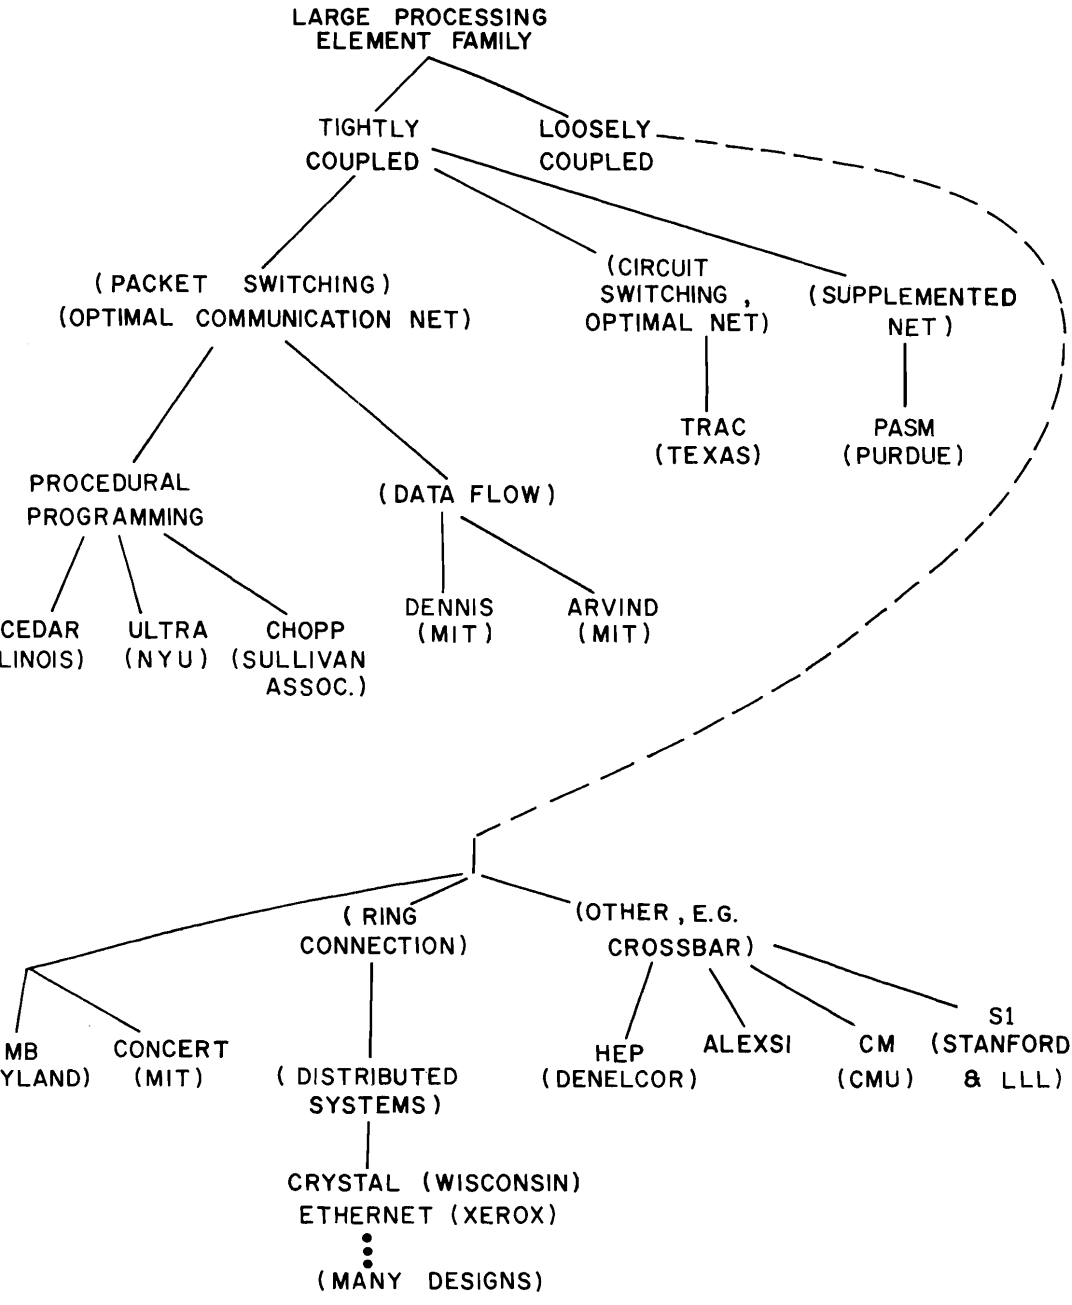

Fig. 1 Machines composed of substantial processors.

In a short talk like this I clearly cannot review too many machines in detail. By now there are at least 50, perhaps as many as 150, university supercomputer designs that have been proposed. The total number is continuing to expand rapidly as universities continue to get excited about this area.

Next I turn to the other part of our taxonomic diagram, Fig. 2, which shows machines composed of minimal processing elements.

These ultra-small-individual-processor designs tend to be constrained (though they are not invariably constrained) in their communication pattern; since designers of machines of this class are trying to optimize the use of silicon area, they ordinarily opt for simplified communication designs which lay out well in two dimensions. (However, there is a special subclass of these machines, including the so-called MIT "connection" machine, currently under active development, that use a more universal logarithmic communication network.) Typical of this class are the tree machines, which use a logarithmic but severely bandwidth-limited communication net; also the class of image processing machines exemplified by the ICL DAP. Finally, we have H. T. Kung's class of systolic array machines within which data flows through an "assembly line," with operations being done as the data moves, until finally results emerge.

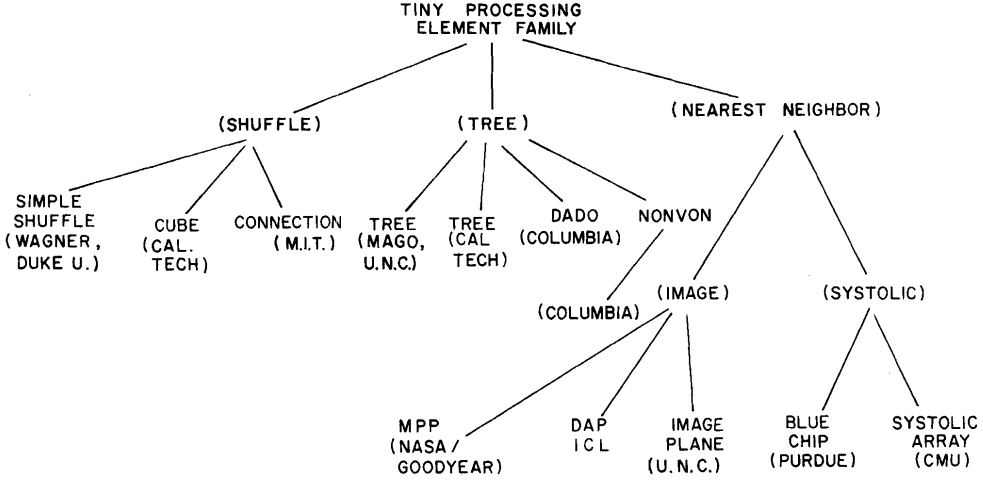

Fig. 2 Machines composed of "minimal" processing elements.

In connection with this general survey of architectural proposals, I cannot resist the temptation to say a few more words about NYU's own "ULTRACOMPUTER" proposal; this is shown in Fig. 3. The advantage claimed for this machine is a particularly "vanilla," general purpose design. A programmer would simply see it as a large collection of processors, each having a certain limited amount of private memory, but all connected to what the programmer would see (on the other side of the data communication switch shown in Fig. 3) as a giant, entirely homogenous, shared global memory. Relative to some of the more highly optimized, but also more special purpose machines that use powerful data communication schemes, the ultracomputer's reliance on shared global memory implies acceptance of a (hopefully slight) memory access inefficiency in order to increase the generality and easy use of this machine. However, this design decision does increase the weight of the hardware substantially, because of the necessity of accelerating memory communication as much as possible.

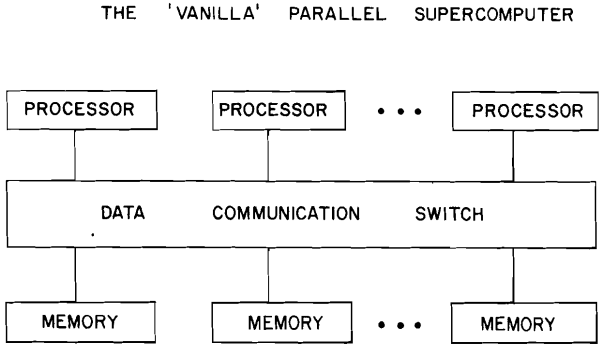

Fig. 3. NYU ultracomputer. Note that fairly substantial individual processors are used.

Additional details concerning the physical structure of the "omega network" that supports memory-to-processor communication in this machine are shown in Fig. 4.

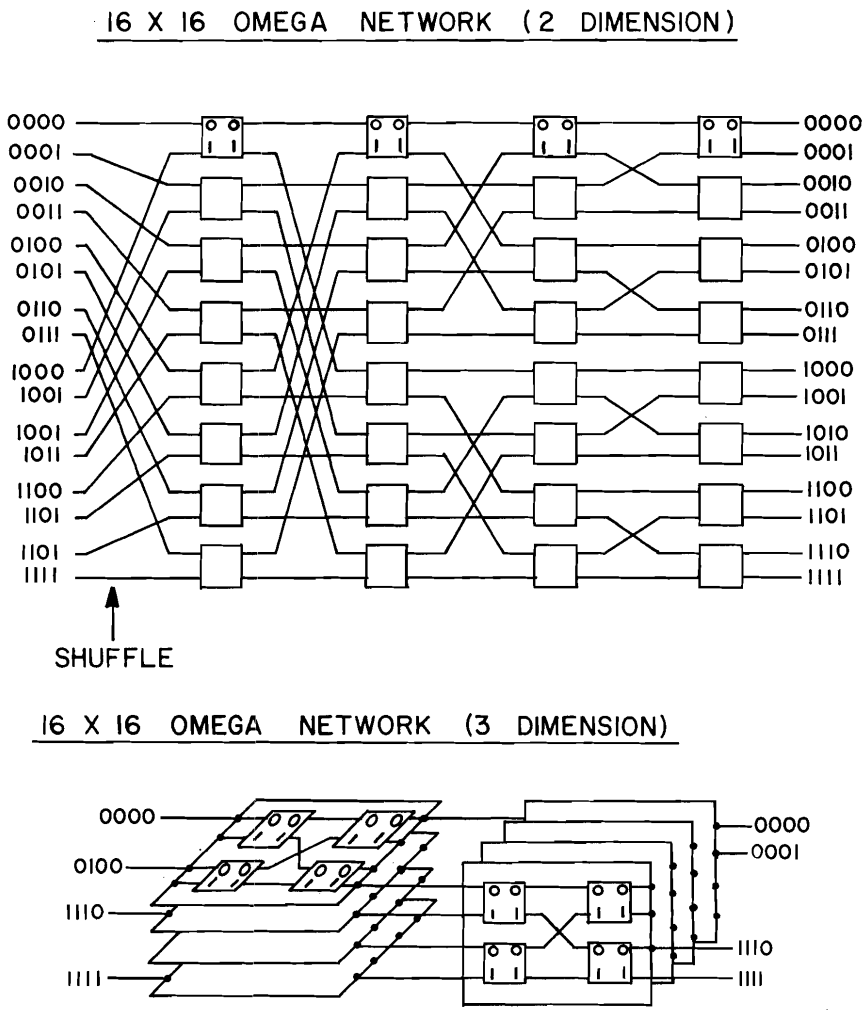

Fig. 4 Details of the physical structure of the "Omega network" that supports memory-to-memory processor communication in the NYU ultracomputer.

So much by the way of a quick survey of U. S. offerings in the supercomputer area. Next I would like to make some prognostications about the developing Japanese - U. S. competitive relationship in regard to supercomputers, which is one of the factors motivating activity within the United States now. It is easy to predict the Japanese supercomputer effort, like all their efforts, will be very well managed technically. Whatever they seek to do and are capable of defining precisely, they do very rapidly and well. On the other hand, I would say that the present conceptual basis for their fifth-generation machine is weak. Nevertheless, since Japan does represent competition that is very strong technically, the success of the developing U. S. response will depend on our ability to follow a better strategy. This will in turn depend on funding agency realism and will also require the effective involvement of industry: if only U. S. universities are involved, and a well-organized industrial participation able to move forward quickly from the university work is absent, it is easy to predict that U. S. universities will innovate very successfully, but only for the benefit of Hitachi and Fujitsu.

The administrators responsible for shaping the U. S. program-to-be in the supercomputer area therefore need to discern the strongest designs, the likeliest winners, from within a growing crowd which already includes many vocal contenders. Already something like 150 universities are each loudly proclaiming that their machine is best. How then should the funding agencies proceed? Concerning this difficult question, I

have time for only one comment. I believe it is important to avoid too heavy a concentration on artificial intelligence longshots. What one wants to do is fund a balanced set of architectural alternatives which can serve to explore the whole of the taxonomic spectrum set forth above; but one must also try to concentrate on those classes of machines most likely to be capable of serving a variety of purposes.

Next I would like to make several longer-term prognostications. I believe that the wave of design innovation represented by the best of the machines appearing in Figs. 1 and 2 will be successful, and that immense parallel machines, presently entirely hypothetical, will become everyday realities to which computing centers will become accustomed. No more than a few years hence, I expect these to be commercially available as the "Cray IV," the "IBM 5999," or what have you.

There is no secret in the construction of these parallel machines. Once one has perceived the new possibilities that large-scale parallelism opens up, the lines of design, especially of general purpose parallel machine design, are fairly obvious. I believe that the U. S. and Japanese large parallel machines will come on the market within a few years of each other. Thus the present race is for a quite temporary advantage.

Once this race has come to its natural end, i.e., once the first few of the new generation of superspeed parallel machines are around and computing centers start ordering them, the ensuing competition will take on a normal commercial character. Competition will then be a matter of quality of software

supplied, speed reached, features available, and price-performance. The crucial factor will simply be the level of corporate commitment, here in the U. S. and in Japan, to maintain a strong position in the large computer area.

A final technical comment. I believe that future large scientific applications systems will become partly hybrid. The pure "vanilla" machine appearing in Fig. 2 is a reasonable first supercomputer, but I expect that eventually one will have various types of special processors attached to this massive general purpose parallel computer base. Certainly in an environment like Fermilab, where there are many major computations that can be greatly accelerated by special- purpose devices, such an admixture of special and general purpose computing devices can have real advantage.

It is easy to surmise from what is already happening that future supercomputers will include attached image-processing machines like the Goodyear Aerospace, MPP, various signal processing devices, graphics chips, etc. Some of these attachments will have large enough markets to become regular market offerings of vendors concerned to furnish a rounded line of special-purpose devices supplementing their basic computer line.