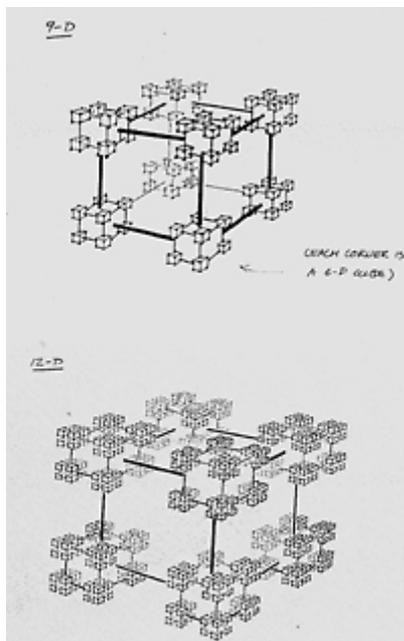
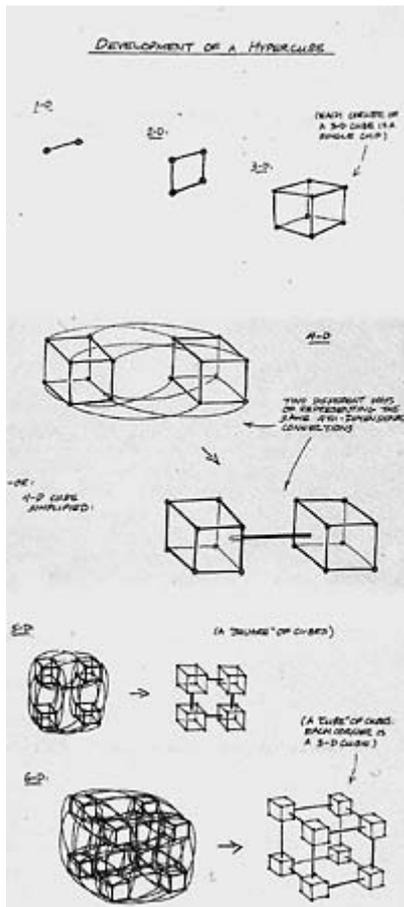


MACHINA COGITANS / THE VISUAL DESIGN OF THE CONNECTION MACHINES CM-1 AND CM-2 TAMIKO THIEL

Ten years ago, in June 1983, Thinking Machines Corporation started building the first commercial massively parallel supercomputer, the connection Machine CM-1. For myself and the young computer scientists designing the machine, it was to be the incarnation of our dreams of what a computer should be.

Today, parallel computing is the undisputed leading edge of computer technology. In 1983, however, it was still extremely controversial. At that time, high-speed sequential vector machines were regarded as the ultimate in computing power. They were up against nature's absolute limit, however; their signals could never exceed the speed of light. Theoretically, parallel machines could circumvent this barrier, gaining increased speed by having multiple processors execute calculations simultaneously, but existing machines were limited to special, small subsets of problems, and general purpose machines seemed impossible to program and impossible to build.

For Danny Hillis, working on problems in human cognition at the Massachusetts Institute of Technology's Artificial Intelligence Laboratory, existing sequential supercomputers were simply inadequate for the problems in which he was interested. He became convinced that it was necessary to design a parallel computer with a structure closer to that of a human brain, laying importance not so much on the processors themselves, but rather on the nature and mutability of the connections between them, hence the name "Connection Machine". In order to build the first of these new machines, he helped found Thinking Machines Corporation in 1983, which introduced the CM-1 in 1986 and the higher performance version, the CM-2, in 1987. Due to the controversial nature of the machine, the top management of Thinking Machines, especially Danny Hillis and the company president, Sheryl Handler, put a high priority on a package whose visual appearance could convince people that this was the first of a new generation, unlike any computer they had ever seen before. I was responsible for the mechanical and industrial design group that produced this package, which was used for both the CM-1 and the CM-2. Ted Bilodeau did the detailed mechanical design, Tom Callahan built the first prototype, and Allen Hawthorne and Gordon Bruce did the detailed industrial design. The CM-2 was superseded by the next generation Connection Machine CM-5 in October, 1991, which has a very different architecture and package, so this paper can also be seen as an epitaph for the original Connection Machine.



THE ARCHITECTURE OF A NEW MACHINE

The Connection Machines CM-1 and CM-2 were to have up to 65,536 simple 1-bit processors grouped 16 to a chip, for a total of 4,096 chips. Each processor could simultaneously perform the same calculation, each on its own separate data set. This kind of massively parallel architecture had been tried in the past, but what enabled the CM-1 and CM-2 to succeed where others had failed was an extremely flexible and fast hardware communications network

between the processors. If this network could operate fast enough, the software connections between processors could be arbitrarily configured to optimally fit the structure of any given problem, ignoring the physical connectivity of the hardware.

This hardware network was to be in the shape of a 12-dimensional hypercube. The term "12-D", far from having to do with warp drives and extraterrestrials, had the mundane but complicated meaning that each processor chip would be directly wired to 12 other chips in such a way that any two chips could communicate with each other in 12 or less steps. Overwhelmed by the effort to visualize a 12-D connection scheme for 4,096 chips, I expressed my bemusement to Richard Feynman, the Nobel-Prize winning physicist who was helping us design the network. (His son Carl was working for Thinking Machines, and Richard got so interested that he asked for a summer job too!) Feynman had the rare talent of being able to explain the most complicated things simply, and characteristically his reply to my complaint was, "Oh, that's easy!"

He connected two chips together to make a "1-D" cube, commonly called a "line". He joined two 1-D cubes together at their ends to make a "2-D" cube, commonly called a "square", and then two 2-D cubes together to make a "3-D" cube, which is the same as our usual understanding of a "cube". Then the step into the 4th dimension: He connected the two 3-D cubes together to make a "4-D" cube, the first stage of a hypercube with more than 3 dimensions.

It is difficult to represent these multidimensional structures in a 2-D drawing. I finally resorted to a radical graphic simplification: I represented all dimensions greater than 3 as thick "hyperlines", and drew cubes as somewhat solid objects in order to visually simplify the resultant structures. Using this simplification, it is easy to see that the structures always repeat themselves: a 4-D hypercube looks just like a 1-D line, except that it connects two cubes rather than two chips. A 5-D hypercube is a square of cubes, and a 6-D hypercube is a cube of cubes. Going further, a 9-D hypercube is a cube of 6-D hypercubes, and a 12-D hypercube is a cube of 9-D hypercubes.

This is the structure used in the CM-1 and CM-2, a cube with 2^{12} corners, or 4,096 chips, each connected to 12 other chips, each connection being one dimension of the cube. This structure was repeated throughout every level of the machine, in the traces connecting processor chips, in the connectors between the printed circuit boards, and in the 1000 feet of cable connecting the highest dimensions of the machine.

What is the "Function" of a Computer?

Now we could turn our attention to the aesthetic aspects of the Connection Machine package. The computer scientists whose visions of the machine influenced me the most, Danny Hillis, Brewster Kahle and Carl Feynman, talked of the machine as a cerebral starship, a vehicle that could open up boundless new frontiers, or as an immensely complex, constantly fluctuating electronic society — the image of an electronic brain. We didn't want to build a machine that looked just like a refrigerator or a washing machine, even if this was the most "practical" and "functional" way to package it. We wanted the design to express the excitement we felt about the machine and about its potential to revolutionize computer architecture.

We still felt, however, that "form should follow function": the form should have meaning and express the essence of the machine. The most obvious solution was to directly follow the lead of modern design and expose the physical structure, which for a computer is its boards and

cables. But how can you "see" the function of an electronic machine? Unlike a mechanical machine, a computer's real structure and function is buried in the invisible layers of silicon on a chip, or hidden in the tangle of cables and electronic circuits on a printed circuit board. It is illegible even to a vast majority of computer professionals, who are concerned exclusively with the software and often have not the slightest idea in which direction electricity flows.

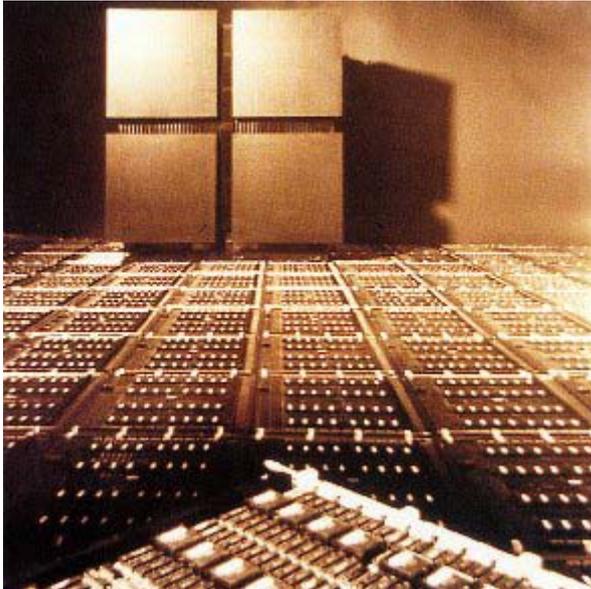
It became clear to me that for a computer, "function" went beyond physical structure into the abstract. To truly understand the function of a computer one has to look at the schematic representations that a computer scientist would use to talk about its architecture, the diagrams that describe the flow of data and commands within the system.

THE CUBE-OF-CUBES

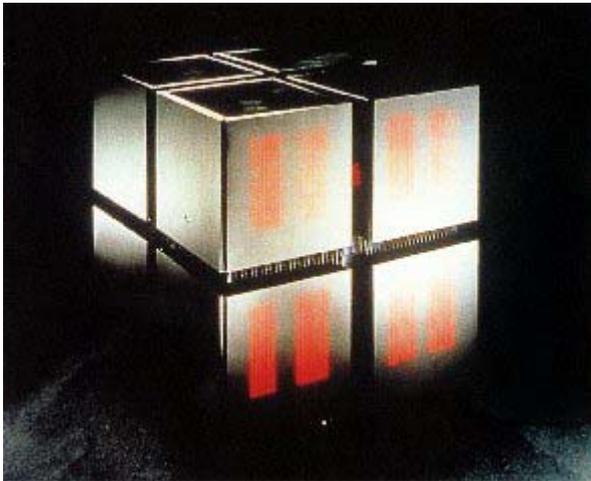
So if form should follow function, and function in a computer means the workings of the processors hidden in the silicon chips, then the real function of the CM-1 and CM-2 lay in the way the processors could communicate with each other, physically using the 12-D hypercube network hardware, but conceptually — that is to say in the software — independent of its structure. We had already developed a symbol that expressed these aspects of the machine: in a design I had done for the company T-shirt! (Richard Feynman liked the T-shirt enough to wear it on the cover of one of his popular books!)

This design showed a cube-of-cubes: a large cube built up out of 8 smaller cubes, representing the "hard" electrical connections of part of the 12-D network. Inside these hard rectangular boxes are the "fuzzy" software connections that can be changed independently of the physical wires and traces. Both were equally important aspects of the structure of the machine: the hard, physical wiring and the soft, programmable connections. How could we make something as abstract as software - which has the intangibility of a concept, a speech, a conversation — visible to the eye?

Carl Feynman had described a fantasy of the CM-1 as a vast cloud of lights that flickered as they sent their electronic messages back and forth, like the firing of neurons in a brain. Status lights are often installed on printed circuit boards to provide visual monitors of the current state of components — whether power is "on", or a chip is plugged in properly. We decided to use these lights to make the intangible and unseen activity of the processors visible to the naked eye.



Connection Machine mit 128 Prozessor-Boards



Connecton Machine CM2

A LIVING MACHINE

Thus, we chose to depict the hardware structure of the machine in the external form of the CM-1 and CM-2 package, and depicted the software connections within this hardware structure using the status lights of the chips: Eight cubes, each holding 9 dimensions of the hypercube, are visually plugged together to form the cube-of-cubes, just as the internal electronic components are physically plugged together to form the highest level of the machine, the 12-D hypercube. Through the doors of the machine glow the lights from 4,096 chips, flickering on and off as the processors work in parallel, each one computing its own part of the data. The microscopic elements of the machine, and the structure buried in the confusion of traces and cables thus become visible, to let the machine speak for itself.

This would communicate to the viewer the immense complexity hidden beneath the surface of the machine. A massive electronic brain, 1.5 meters in height, it is connected with cables to the data drives that feed its processors information and the workstations and monitors that let it communicate with its human users. A hard, geometric object, black, the non-color of sheer, static mass, it is filled with a soft, constantly changing cloud of lights, red, the color of life

and energy. This is how we felt a computer should look, a living, breathing machine, a machine that could be proud of us.

NOTES

The opinions expressed in this article are the author's alone, and do not necessarily represent Thinking Machines policy.

This article is an extract from a longer article to be published this fall in the journal "Design Issues", ed. Dick Buchanan, Chair, Department of Design, Carnegie-Mellon University, Pittsburgh, Pennsylvania, USA.

For more technical information on the Connection Machine CM-2 see "The Connection Machine", MIT Press, 1985, or "The Connection Machine", Scientific American, June 1987, vol. 256, pp. 108, both by Daniel W. Hillis.