High-Performance Neural Networks*

W. B. Dress
Instrumentation and Controls Division
Oak Ridge National Laboratory
Oak Ridge, Tennessee 37831-6007

Introduction

The anticipated hardware and software needs of the 1990s and beyond will not be met by a continuation of the past evolutionary progress in computing capabilities. At a minimum, a dramatic acceleration over the 60% improvement each year (in processing power, memory capacity, software design, etc.) is needed. Better still, a revolution is called for, and indeed seems to be in the making. Even though the roots of this revolution extend far into the past, the underlying concepts are finding new and forceful expression today.

This revolution has various labels: connectionist models, biologically motivated computing, artificial neural networks, and so on. Some of these paradigms are direct extensions of classical computer science, while others are more radical in approach and extensive in scope. The more conservative approaches retain the preeminent position of the programmer, while the radical ones see eventual extinction of that breed. The conservative approach is extending the paradigms of artificial intelligence--an abstract parody of human behavior, while the radical is creating a synthetic, but nonetheless real, intelligence.

Synthetic Neural Systems

Simulating the massively parallel structure of the brain, even with greatly reduced and abstract neuronal models, requires enormous computational power on serial machines. The ideal machine for neural processing would have a large number of autonomous computing elements receiving, processing, and sending data streams, possibly communicating its end states to an expert system entrusted with the task of higher-level decisions. However, executing code for the network is not the most computationally demanding task to face designers of synthetic intelligent systems.

A synthetic neural system consists of a neural network interacting with an environment by means of sensors and effectors (Browning, 1964). In the system currently operating, the I/O (sensors and effectors) and the environment have been simulated on a Macintosh™ computer with MacForth™ Plus, while the network runs either on the Mac or on the Novix™ Beta Development System in polyForth.™

*Research performed at Oak Ridge National Laboratory, operated by Martin Marietta Energy Systems, Inc., for the U. S. Department of Energy under Contract No. DE-AC05-84OR21400.

So far, networks consisting of a few hundred nodes (neurons) and a few thousand connections (synapses) adequately support a dozen or so sensor inputs and drive the same number of simulated muscle cells. The resulting synthetic "insect" is able to "crawl" about the computer screen and learns to avoid the sides of the screen by feeling "pain," and is attracted to "food" by feeling "pleasure." It learns this behavior by modifying the synaptic weights (connection strengths from node to node) according to the information received by the sensor cells (a simulated eye, mouth, and feelers) and, in turn, acts on the environment by driving the muscle cells, producing motion.

Genetic Algorithms And Computational Needs

The system is thus far parameterized by about 50 "genes" which determine such properties as strength of muscular contractions, the spontaneous firing of the central network's cells, and so forth. This 50-dimensional, certainly nonorthogonal, system presents a seemingly impossible optimization problem to the system designer. (By way of contrast, the popular fruit fly has around 5000 genes, and so inhabits a successful region of a 5000-dimensional genetic space.) Taking a cue from biology, which may be credited with some reasonably successful systems, we can simulate a Darwinian (Darwin, 1859) type of selection, although artificial instead of natural. Thus the environment interacts with, and eventually determines, the genome describing the behavior of the class of synthetic neural systems observed in this work. The selection demanded by the Darwinian theory is provided by an algorithm which evaluates the performance of each synthetic creature and compares its performance to that of its parent. The one performing better according to the algorithm becomes a parent, passing all of its genes unmodified save for one to the offspring, while the other is killed off. This vicious but effective cycle repeats as long as the experimenter (no longer a programmer) has patience. A fascinating description of a similar genetic process, for evolving shape rather than behavior, has recently been given by Dawkins (1986).

Since a few thousand "ticks" of the system seem necessary to establish an identifiable behavioral pattern, and each tick may take up to one second in high-level Forth on the Mac in addition to the computations necessary for updating the sensors, muscles, and environment (an additional fraction of a second), the lifetime of a single creature tries the patience of most observers. When it is realized that many thousands of simulated life cycles must be evaluated to improve the synthetic species materially, the computational magnitude of the problem becomes apparent.

An interim solution (while awaiting the advent of affordable parallel supercomputers) is to run the network (brain part) on a NC4016-based coprocessor. One system tick--a complete update of all 200 nodes and 1760 connections (to pick a concrete example) takes 17 ms on an 8 MHz NC4016, equivalent to greater than 100,000 interconnect updates per second. This is greater than 4 times the performance of a commercially available neurocomputer consisting of 4 MC68010s operating in parallel. The table below shows times measured on a network and simulated environment as well as times estimated for other host configurations.

Synthetic Neural System On Various Hardware Configurations

Host Computer	Mac Plus [1]	Mac-II (est)	HP350 (est)	NC6016 (est)
Network	743 ms	250 ms	160 ms	7 ms
Environment	132 ms	45 ms	30 ms	1 ms
Communications [2,3]	20 ms	1 ms	——	——
Animation	20 ms	8 ms	4 ms	——
Totals [4]	895 (189)	304 (71)	194 (51)	8 (11)
Lifetimes/Day [5]				
Host	50	140	220	5400 [7]
NC4016 [6]	230	600	850	4000

Notes:

1. Current network, times measured for a 200-node network with 1760 connections.
2. Does not apply when host is running the network.
3. Communications are assumed to take place over a parallel, fast interface for for all but the Mac Plus as host.
4. The number in parentheses is the time with the NC4016 running the network.
5. Lifetime is 2000 passes through the network.
6. The time for one pass through the network is 17 ms.
7. Assumes one NC6016 is the host and another runs the network.

Table 1. Times and number of lifetimes per day for various configurations of the network and host computer systems. The times in the first column are measured values, while the others are estimated based on known clock rates and instruction-set properties.

Conclusions and Future Directions

The new Forth hardware architectures offer an intermediate solution to high-performance neural networks while the theory and programming details of neural networks for synthetic intelligence are developed. This approach has been used successfully to determine the parameters and run the resulting network for a synthetic insect consisting of a 200-node "brain" with 1760 interconnections. Both the insect's environment and its sensor input have thus far been simulated. However, the frequency-coded nature of the Browning network allows easy replacement of the simulated sensors by real-world counterparts.

The resulting synthetic entity, taking inputs from physical sensors, should organize its "brain" to function autonomously in the real world (we hope with reasonable success). At this point, a new order of hardware performance will be required both for the artificial Darwinian determination of the system parameters and running the evolved entity. One solution is to build special-purpose neural devices in silicon and configure an asynchronous nervous system with large numbers (several thousand) of parallel elements. Until this is feasible, we intend to run modular parts of the nervous system in parallel on a large number (hundreds) of Forth processors. Such a system will allow parallel evolution with competing synthetic creatures at a rate of about one million lifetimes per day as well as running an evolved system with sound, touch, and vision input at real-time (i.e., as fast as natural creatures respond to sight, touch, and sound) rates.

References

Iben Browning, 1964. "A Self-Organizing System Called 'Mickey Mouse,'" Panoramic Research, Inc., Palo Alto, Ca.

Charles Darwin, 1859. *On The Origin Of Species,* A Facsimile of the First Edition, Harvard University Press, Cambridge, Ma., 1966.

Richard Dawkins, 1986. *The Blind Watchmaker,* W. W. Norton, New York.