FROM: ANALOG VLSI AND NEURAL SYSTEMS.

BY CARVER MEAD.

ADDISON-WESLEY. 1989

CHAPTER 1

INTRODUCTION

Even the most naive observer can see that the nervous system is vastly different from a computer. Living systems are made from three-dimensional, squishy cells; computers are constructed of rigid inorganic matter in flat, two-dimensional sheets. Living systems are powered by metabolic biochemistry; computers are powered by transformers and rectifiers from the power mains. Living systems have approximately 100-millivolt nerve impulses lasting nearly a millisecond: computers have 5-volt signal levels switching at nanosecond intervals. The destruction of a few percent of the cells in a brain will cause no discernible degradation in performance; the loss of even a single transistor in a computer (save for in its memory) may cause complete loss of functionality. The average nerve cell dissipates power in the 10^{-12}-watt range; the average logic gate in a computer dissipates 10 million times as much.

Nonetheless, a more careful look reveals some underlying similarities between the two kinds of systems. Both process information. Signals are represented as differences in electrical potential, and are conveyed on "wires" formed by surrounding a conducting path with an excellent electrical insulator. Active devices cause electrical current to flow in a second "output" conductor due to the potential in a first "input" conductor. The "output" of an active device has more energy than was present in the "input" to that device; hence, the systems possess "gain"—the essential ingredient for unbounded information processing.

which is accompanied by an unavoidable dissipation of energy. A "power supply" maintains a near-constant average difference in electrochemical potential across the active devices. The active devices are formed of extremely thin *energy barriers* that prevent the flow of current between two electrical nodes. The passage of current is mediated by the potential on a third "control" electrical node. That current varies exponentially with the potential on the control node.

Heartened by these less obvious but deeper similarities between the two systems, we may be tempted to conclude that the brain is, indeed, a digital computer. Its nerve pulses encode information in much the same way as do pulses in a telephone exchange. Neurons perform Boolean AND and OR operations on the way to firing off a nerve pulse to the next stage of computation. Small neural memory elements store information in much the same way as computer memories do. Perhaps we are, after all, on the verge of discerning one of nature's most profound and best-kept secrets: the working of thought itself.

Speculations of this sort were rampant in the late 1940s. A film depicting the operation of the Whirlwind, an early vacuum-tube computer with magnetic-core memory, was called *Faster than Thought* [Bowden, 1953]. The field of artificial intelligence was born. Soon, however, signs of distress could be seen. By 1977, Marvin Minsky, one of the artificial-intelligence pioneers, opened a seminar at Caltech with the following observation:

> *Our first foray into Artificial Intelligence was a program that did a credible job of solving problems in college calculus. Armed with that success, we tackled high school algebra: we found, to our surprise, that it was much harder. Attempts at grade school arithmetic, involving the concept of number, etc., provide problems of current research interest. An exploration of the child's world of blocks proved insurmountable, except under the most rigidly constrained circumstances. It finally dawned on us that the overwhelming majority of what we call intelligence is developed by the end of the first year of life.* [Minsky, 1977]

The visual system of a single human being does more image processing than do the entire world's supply of supercomputers. The digital computer is extremely effective at producing precise answers to well-defined questions. The nervous system accepts fuzzy, poorly conditioned input, performs a computation that is ill-defined, and produces approximate output. The systems are thus different in essential and fundamentally irreconcilable ways. Our struggles with digital computers have taught us much about how neural computation is *not* done; unfortunately, they have taught us relatively little about how it *is* done. Part of the reason for this failure is that a large proportion of neural computation is done in an *analog* rather than in a digital manner.

Perhaps the most rewarding aspect of analog computation is the extent to which elementary computational primitives are a direct consequence of fundamental laws of physics. In Chapter 3, we will see that a single transistor can take at its gate a voltage-type signal and produces at its drain a current-type signal that is exponential in the input voltage. This exponential function is a direct

result of the Boltzmann distribution. We will see that addition and subtraction of currents follows directly from the conservation of charge. In subsequent chapters, we will encounter many examples of computations that follow directly from physical laws.

It is essential to recognize that neural systems evolved without the slightest notion of mathematics or engineering analysis. Nature knew nothing of bits, Boolean algebra, or linear system theory. But evolution had access to a vast array of physical phenomena that implemented important functions. It is evident that the resulting computational metaphor has a range of capabilities that exceeds by many orders of magnitude the capabilities of the most powerful digital computers.

It is the explicit mission of this book to explore the view of computation that emerges when we use this evolutionary approach in developing an integrated semiconductor technology to implement large-scale collective analog computation.

The biological questions asked about neural computation were for many years, and to a large extent still are, basically reductionist. It is tacitly assumed that, if we understand in detail the operation of each molecule in a nerve membrane, we will understand the operation of the brain. It is to this view that our knowledge of computers can bring some insight. A computer is built up of a completely known arrangement of devices; the operation of these devices is understood in minute detail. Yet it is often impossible to derive even a simple proof that a program that we ourselves write will compute the desired result or, for that matter, that the computation will even terminate!

The complexity of a computational system derives not from the complexity of its component parts, but rather from the multitude of ways in which a large collection of these components can interact. Even if we understand in elaborate detail the operation of every nerve channel and every synapse, we will not by so doing have understood the neural computation as a *system*. It is not the neural devices themselves that contain the secret of thought. It is, rather, the organizing principles by which vast numbers of these elementary devices work in concert. Neural computation is an emergent property of the system, which is only vaguely evident in any single component element.

Although study of the elements is an essential step in understanding the system organization, in and of themselves, the elements tell us very little. Furthermore, we have learned enough in recent years concerning the operation of nerves and synapses to know there is no mystery in them. In not a single instance is there a function done by a neural element that cannot, from the point of view of a system designer, be duplicated by electronic devices.

What then is to prevent us from creating a nervous system in silicon? Two barriers have historically blocked the way:

1. Neural systems have far greater connectivity than has been possible in standard computer hardware. Many early attempts to create neural systems failed simply because no workable technology existed for realizing systems of the requisite complexity.

2. Sufficient knowledge of the organizing principles involved in neural systems was not available.

The rapidly developing technology of very large scale integrated (VLSI) circuits has given us a medium in which it is presently possible to fabricate tens of millions of devices (transistors) interconnected on a single silicon wafer. This number will increase by two orders of magnitude before fundamental limitations are encountered [Hoeneisen et al., 1972a]. The densest and most widely available technology uses metal-oxide-silicon (MOS) transistors; it has been primarily conceived as a digital technology, and has been highly evolved for the production of microprocessors, memories, and other digital products. It might therefore be supposed that the most highly evolved fabrication process would not be suitable for the functions required in neural processing. The noise level of a typical device is higher, or the precision with which any two devices can be matched is lower, for example, than are those of technologies historically used for implementing analog functions. We observe, however, that the precision, reliability, and noise properties available in neural wetware fall short of those used in even the most rudimentary electronic systems. This lack of precision and reproducibility at the component level is more than offset by the redundancy introduced at the system level. Whether this property is the primary reason for the large connectivity in neural systems, or whether it is a byproduct of an organizing principle dictated by other system needs, is not a question open to us at the present time. We do know, however, that robustness under failure or imprecision of individual components is one important emergent property of neural systems. If we base our designs on the same organizing principles, we should not be concerned that individual devices will cause system malfunction. To the contrary, we can expect that systems with extraordinary reliability and robustness will result; so much so that useful integration at the scale of a complete wafer is feasible. In the chapters that follow, we will explore many ways in which an ordinary digital technology (complementary MOS, or CMOS) can be used to implement extraordinary systems based on neural paradigms.

In terms of discovering neural organizing principles, we are less well off. Although a great deal of progress has been made in recent years, there is still no global view of the principles and representations on which the nervous system is organized. There has been, however, a striking increase in knowledge of particular systems, due in large part to experimental techniques developed over the past decade. Detailed physiological studies have given us a picture of the mapping from the visual field onto the visual cortex [Schwartz, 1977], and similar information is available for many important auditory areas [Merzenich et al., 1977]. Several authors have put together a more unified view of these findings. Readable accounts of the gross connectivity among major areas of the brain have been given for the visual system [Van Essen, 1984; De Yoe et al., 1988] and the auditory system [Pickles, 1982; Kim, 1984]. A most notable account of the detailed synaptic circuits of each of several areas of the brain is given in *The Synaptic Organization of the Brain* by Gordon Shepherd of Yale University [Shepherd, 1979].

Many people have proposed hypotheses about the way computation is performed in these systems. To date, it has proved difficult if not impossible either to verify or to disprove any given hypothesis concerning the operating principles of even the simplest neural system. Major areas are so richly interconnected, and computation within a given area is so intertwined, that there exists no good way of separating one function from another. Our traditional scientific approach of studying the elements separately in order to understand the whole fails us completely. Simple neural systems based on clear, obvious principles may once have existed, but they are buried by the sands of time. Billions of years of evolution have presented us with highly efficient, highly integrated, and impossibly opaque systems.

A NEW APPROACH

Let us, then, undertake the following program. We have already noted that elementary operations found in the nervous system can be realized in silicon. We also note that many neural areas are thin sheets, and carry two-dimensional representations of their computational space. The retina is the most obvious example of this organization, which also occurs in the visual cortex and in several auditory areas. In both neural and silicon technologies, the active devices (synapses and transistors) occupy no more than 1 or 2 percent of the space—"wire" fills the entire remaining space. We can be confident, therefore, that the limitation of connectivity will force the solution into a very particular form. If the required functions could have been be implemented with less wire, nature would have evolved superior creatures with more computation per unit brain area, and they would have eaten the ones with less well-organized nervous systems.

We will therefore embark on a second evolutionary path—that of a silicon nervous system. As in any evolutionary endeavor, we must start at the beginning. Our first systems will be simple and stupid. But they, no doubt, will be smarter than the first animals were. We are, after all, endowed with the product of a few billion years of evolution with which to study these systems!

The constraints on our analog silicon systems are similar to those on neural systems: wire is limited, power is precious, robustness and reliability are essential. We therefore can expect that the results of our second evolution will bear fruits of biological relevance. The effectiveness of our approach will be in direct proportion to the attention we pay to the guiding biological metaphor. We use the term "metaphor" in a deliberate and well-defined way. We are in no better position to "copy" biological nervous systems than we are to create a flying machine with feathers and flapping wings. But we can use the organizing principles as a basis for our silicon systems in the same way that a glider is an excellent model of a soaring bird.

It is in that spirit, then, that we will proceed. First we will describe the relevant aspects of neural wetware at the level of abstraction where we will be working. We will then develop the operations that are natural to silicon, and examine how they can be used to implement certain known neural functions. Finally, we will show several examples of complete subsystems that have metaphors drawn, in one way or another, from biology.

It is the author's conviction that our ability to realize simple neural functions is strictly limited by our understanding of their organizing principles, and not by difficulties in implementation. If we *really* understand a system, we will be able to build it. Conversely, we can be sure that we do not fully understand a system until we have synthesized and demonstrated a working model.

The silicon medium can thus be seen to serve two complementary but inseparable roles:

1. To give computational neuroscience a synthetic element, allowing hypotheses concerning neural organization to be tested
2. To develop an engineering discipline by which collective systems can be designed for specific computations

The success of this venture will create a bridge between neurobiology and the information sciences, and will bring us a much deeper view of computation as a physical process. It also will bring us an entirely new view of information processing, and of the awesome power of collective systems to solve problems that are totally intractable by traditional computer techniques.

REFERENCES

Bowden, B.V. *Faster than Thought.* London: Sir Isaac Putnam & Sons Ltd., 1953.

De Yoe, E.A. and Van Essen, D.C. Concurrent processing streams in monkey visual cortex. *Trends in NeuroSciences.* 11:219, 1988.

Hoeneisen, B. and Mead, C.A. Fundamental limitations in microelectronics—I. MOS technology. *Solid-State Electronics,* 15:819, 1972a.

Kim, D.O. Functional roles of the inner– and outer–hair-cell subsystems in the cochlea and brainstem. In Berlin, C.I. (ed), *Hearing Science: Recent Advances.* San Diego, CA: College-Hill Press, 1984, p. 241.

Merzenich, M.M., Roth, G.L., Andersen, R.A., Knight, P.L., and Colwell, S.A. Some basic features of organization of the central auditory nervous system. In Evans, E.F. and Wilson, J. (eds), *Psychophysics and Physiology of Hearing.* New York: Academic Press, 1977, p. 485.

Minsky, M. Personal communication.

Pickles, J.O. *An Introduction to the Physiology of Hearing.* Orlando, FL: Academic Press, 1982.

Schwartz, E. Spatial mapping in the primate sensory projection: Analytic structure and relevance to perception. *Biological Cybernetics,* 3:181, 1977.

Shepherd, G.M. *The Synaptic Organization of the Brain.* 2nd ed., New York: Oxford University Press, 1979.

Van Essen, D.C. Functional organization of primate visual cortex. In Peters, A. and Jones, E.G. (eds), *Cerebral Cortex,* vol 3. New York: Plenum Press, 1984, p. 259.