

Peter M. Asaro

1 Introduction

I was first introduced to cybernetics in 1995. I came to it via a strange route through a philosophical concern with Artificial Intelligence (AI), and an attempt to historicize AI as a science in order to alleviate the philosophical problems AI had generated. As I read the cybernetic literature, I became intrigued that as an approach to the mind which was often described as a predecessor to AI, cybernetics had a much more sophisticated approach to mind than its purported successor. I was soon led to Prof. Herbert Brun's seminar in experimental composition, and to the archives of the Biological Computer Laboratory (BCL) in the basement of the University of Illinois library. Since then, I have been trying to come to terms with what it was that was so special about the BCL, what allowed it to produce such interesting ideas and projects which seem alien and exotic in comparison to what mainstream AI and Cognitive Science produced in the same era. And yet, despite its appealing philosophical depth and technological novelty, it seems to have been largely ignored or forgotten by mainstream research in these areas. I believe that these are the same concerns that many of the authors of the recent issue of *Cybernetics and Human Knowing* (Brier & Glanville, 2003) express in regard to the legacy of von Foerster and the BCL. How could such an interesting place, full of interesting things and ideas have just disappeared and been largely forgotten, even in its own home town?

Stuart Umpleby, in his (2001) paper, identifies several of the key reasons why he believes that the BCL was such a unique place. At the top of this list he puts the fact that Heinz von Foerster was preoccupied with epistemic issues in a way that most scientists were not. Consequently, von Foerster directed his lab in a way that illuminated those issues, unlike most scientific and engineering labs, a fact that begins to explain the philosophical sophistication of their ideas and projects. Bernard Scott (2003) has shown in his analysis of the historical development of his thought that von Foerster held a concern with the position of the observer beginning in his early work in quantum theory, and continuing throughout his career. At the time when he was directing the BCL, he was finally in a position to develop projects that directly addressed these issues. This active period for the BCL also happened to coincide with several other developments in which cybernetics and systems sciences began to take reflexive turns, ultimately leading to what came to be called Second Order Cybernetics (Scott, 2004). Ranulph Glanville (2003) has identified three "machines" which served as recurrent themes in von Foerster's work – Maxwell's Demon, Eigen Forms, and Non-Trivial Machines – though they are conceptual gedanken experiments rather than technological artifacts. I

wish to augment these histories with a more careful examination of the early BCL machines and their relation to von Foerster's conceptual framework, and contemporaneous scientific movements.

The Self-Organizing Systems and Bionics movements interest me both because of their historical origins in carrying forward some of the most interesting insights and ideas from the cybernetics movement of the 1940s and 50s, and because of their odd position in light of later historical developments between the fields of AI and Connectionism of the 1970s, 80s and 90s, especially the rivalries and debates between them. It is interesting to see the stature of the cyberneticians within these communities. The popular histories of Connectionism and AI mention the inspirational ideas of feedback control and goal-directed behavior as a means for producing mechanistic descriptions of intelligence, but these histories rarely acknowledge the extent to which cyberneticians mentored the young founders of AI, nor do they mention the continuing influence of cybernetic work on AI and Connectionism well into the 1960s, after each field of research was well under way. Indeed, it has only been recently that histories of neural networks have begun to acknowledge that research in this field had been undertaken in a serious manner before 1970 (Anderson & Rosenfeld, 1998).

A brief survey of the scientists present at the first Self-Organizing Systems and Bionics conferences reads like a "who's who" of Cybernetics, AI, and early neural network research. In attendance at the first Self-Organizing Systems conference, for instance, were AI luminaries Marvin Minsky, John McCarthy, Herbert Simon, Allen Newell, and John Shaw, neural network pioneer Frank Rosenblatt, and the cyberneticians Warren McCulloch and Gordon Pask, in addition to von Foerster. By the end of the 1960s, AI had firmly set itself apart from cybernetics and neural networks. Partly this was due to political struggles over funding, and partly due to processes of institutionalization and disciplinarity. So it is rather interesting to see the founding fathers of each of these paradigms describing their work alongside one another as being part of a common scientific movement some years before sharp distinctions had been drawn between them. The other intriguing feature that draws my attention to these movements is the leading role played by Heinz von Foerster and the unique projects of the BCL.

The BCL was created and directed by von Foerster just one year before the Self-Organizing Systems and Bionics conferences began, in 1959. The BCL is itself a fascinating place which seems to be an historical anomaly according to the widely received history of this era. Unfortunately, the Biological Computer Lab has never been given the historical attention that the Digital Computer Lab, run by the same department at the University of

Illinois, has received. The Digital Computer Lab has been praised for creating ILLIAC, the first non-military high speed digital computer, and ILLIAC IV, arguably the first super-computer (Goldstine, 1972). Indeed, histories of AI and Cognitive Science have focused almost exclusively on digital computers and their programs as being the very essence of these movements. Thus the BCL, with its staggering array of analog machines built over the course of the 1960s has been largely overlooked. More importantly, the role of these analog machines in articulating the Self-Organizing Systems and Bionics movements has been completely ignored.

The Self-Organizing Systems and Bionics movements were motivated by highly abstract theoretical principles, much like cybernetics. However, there is more to science than just ideas and concepts. Crucial to understanding scientific progress is to look also at both the social networks that coordinate the activities of scientists, and the material engagements of scientists and engineers with the technologies and material phenomena that constitutes their work. Recent developments in the sociology of science, such as Actor-Network Theory (Latour, 1987), has provided a useful framework for thinking about socio-technical networks, while the Mangle of Practice (Pickering, 1995) has stressed the importance of looking at material engagements with the world, and how these unfold in real time. Pickering (2004) has also considered cybernetics explicitly, and the ways in which work in early cybernetics revolved around specific material apparatus and phenomena. He calls these material apparatuses the 'gallery of monsters,' and includes Norbert Wiener's anti-aircraft predictor, W. Ross Ashby's Homeostat, and Ilya Prigogine's Belousov-Zhabotinsky reaction in his gallery. Even though the cybernetic sub-disciplines of the 1960s, Self-Organizing Systems and Bionics, were highly theoretical, they too still had a material grounding in specific technologies and natural phenomena. This essay hopes to show that an important part of why the BCL was able to make such significant contributions to cybernetics in the 1960s was the instantiation of these ideas in specific technologies.

Rather than dwell only on concepts in this paper, I wish to discuss some of the machines built at the BCL in its first few years. In doing this, I want to show how von Foerster sought to position his lab and its work into the neo-cybernetic disciplines of Self-Organizing Systems and Bionics. While much of the scholarship on von Foerster and the BCL to date has focused primarily on the concepts and essays produced there, I wish to show that the machines were a significant element of what made the lab so interesting. As embodiments of these disciplines, understanding these machines can add new depth and appreciation to our understanding of von Foerster's ideas. In particular, the Adaptive Reorganizing Automaton demonstrates the concepts

of self-organization, while the Numa-Rete illuminates the concepts of pre-organization and property filters, and the Dynamic Signal Analyzer embodies neighborhood logics, and bionic engineering.

To make the point in ontological terms, the idea is that science is not constituted by ideas alone, and technology is not constituted by material things alone. And then, of course, there are the observers. When we consider science, it is imperative to consider each of these three aspects – ideas, materials and observers – as agents in the construction of the networks we call knowledge and technology. When viewed in this way, we can see a scientific technology as an interface between the dynamic conceptual, material and social networks that constitute scientific practice and around which the activities in each realm can be coordinated. These machines are metaphysical conduits between theories and natural phenomena in virtue of being embodiments of scientific theory. They are in this sense physical metaphors. And because of this embodiment, they can be shared objects of reference between observers, and are subject to certain canonical techniques of engagement and manipulation. It is these techniques that define the style of work, or disciplinary matrix, within the paradigm by providing scientists with ways of proceeding, both theoretical and technical, when they meet resistance in trying to extend and materialize concepts, or encounter the emergent properties of the system.

While it took philosophers, historians and sociologists of science until the late 1990s to come to these realizations (though perhaps not in the terms I have just presented them in), it seems that von Foerster was well aware of this triadic ontology of ideas, things and observers. In fact, it is made quite explicit in the summary report of the ONR grant upon which the BCL was founded. “Toward the Realization of Biological Computers” was submitted on December 31, 1963, when the ONR funding to the lab was suddenly and unexpectedly revoked. In that document, von Foerster outlined the main lines of approach taken by the BCL in its attempt to “realize” biological computers: Theoretical, Experimental, and Symposia. I will return to the details of each later, but I want to point out here that these correspond exactly to the development of theoretical concepts, experimental artifacts, and the social coordination of observers – the scientists attending the Symposia.

The role that von Foerster played in the early formation of the Self-Organizing Systems and Bionics movements might best be described as that of the “visionary.” On the one hand, he is very unlike the traditional textbook notions of a “great scientist” in that he is not really credited for a great and heroic theory, discovery or accomplishment. Yet he was enormously influential on the scientific community, both in terms of personal influence and charisma, and in the shaping of scientific movements, and disciplines.

I think that he achieved this through being a superlative spokesman for the movements, and thus leading them conceptually in ways that promoted their interdisciplinary strengths. He also did this by supervising the construction of several significant biological computers at the BCL, which served as exemplars of the goals and practice of the fledgling movements, and also served to promote the scientific communities engaged in those movements. Of course, von Foerster did not work alone. He was aided by a collection of bright and capable graduate students, including Murray Lewis Babcock (who was working towards his Ph.D. in electrical engineering), Paul Weston and George W. Zopf (both in electrical engineering as well), Crayton Cann Walker (who was working toward a masters degree in psychology), and Albert A. Mullin (who was studying mathematics and philosophy). It was this handful of individuals who would produce the theories and the analytic and synthetic tools that would be used in realizing and elaborating the first biological computers in the early days of the BCL.

2 The Mission of Self-Organizing Systems

The Self-Organizing Systems movement was organized around a rather esoteric idea, stemming from the consequences of classical thermodynamics and the Second Law of Thermodynamics that states that in all systems, entropy tends to increase over time. While this simply means that differences in heat or energy tend to equalize over time, it is also understood to mean that order breaks down over time, and that organization decays into chaos within both individual systems and the whole of the universe. Given that this is one of the central laws of physics, how can it be that there is any order at all in the universe? Then there is the even more perplexing matter of biological evolution. Commonly cited in this movement is the book *What is Life* (1948) by the physicist Erwin Schrodinger. He frames the problem as being that biological systems actually increase in complexity over evolutionary time, and thus appear to break the second law. Similar arguments can be made for social organizations, neural systems, and other complex systems. This presents the two-fold question of how it is possible for these “self-organizing systems” to run upstream in the river running from order to chaos, and at the same time avoid violating an apparently central law of physics.

The theoretical answer, given by von Foerster at the very first of the Self-Organizing Systems conferences (von Foerster, 1960), is that biological organisms and other complex systems consume energy and order from their environments. And so, while entropy will steadily increase globally,

locally organisms can capture and transform energy and produce islands of increasing order. This resolution of the apparent paradoxes provided the basis for thinking that there was indeed a coherent set of phenomena and principles regarding self-organization, such that they could be fruitfully studied by a new discipline. Von Foerster thus offered a fundamental element of theory to the new field, a justification that there really was something to be studied here. What seems less obvious is what sort of research agenda should follow from such an esoteric theoretical concern, and how it should proceed. What did it mean in terms of experiments, instrumentation, and the control of natural phenomena?

3 Adaptive Reorganizing Automaton

At the time of the first conference on Self-Organizing systems, Murray Babcock was working on his doctoral research under the direction of von Foerster at the BCL. Babcock was in fact completing the instrumental embodiment of a Self-Organizing System – the Adaptive Reorganizing Automaton, a truly unique device (see Figures 1 & 2 from Babcock, 1960.)




Figure 20. Front View of Adaptive Reorganizing Automaton




Figure 21. Rear View of Adaptive Reorganizing Automaton

In his dissertation, Babcock appeals to a great deal of neurophysiology in his justification of the design for its electronic circuits (Babcock, 1960). The

analog circuits were modular, and designed to be connected together into numerous different network configurations. They were of three types: an Energy Transducer Module, an Autonomous Component Module, and a Facilitator Module. The energy transducer module was based on the McCulloch and Pitts neuron model. The autonomous component was designed to offer random inputs to the rest of the system when the system's activity became low, and was random to prevent the other components in the system from adapting to it.

The basic idea was that the adaptive automaton would be a continuous and potentially recurrent network that would “evolve” certain patterns of behavior due to its specific topology, environmental input conditions, and the system's own input history:

The automaton will be composed of many elementary components, customarily known as artificial neurons, which are all essentially alike functionally. These elementary components are connected to each other in a “very general manner” – to be explained later – such that the information flow between elementary components is through a variable conductance path called a facilitator. The conductance of the facilitator is time and use dependent. That is, the more a particular facilitator is used with respect to the total information transferred by it in a given time, the greater is the effect of that information upon the receiving elementary component. The effect of the facilitator may be considered as a variable gain whose value is dependent upon the time integral of the information flow through it. Thus the gain of the facilitator is the ratio of the magnitude of the output energy of the facilitator to the input energy to the facilitator where both energies are those connected with the coded information being transferred through the facilitator – i.e. the signal energy.

As a result of the facilitators, preferred paths of information flow within the machine will be established, these paths being dependent upon the information content, its code, its source in the sense of its input location in the machine, and the total information stored or in transit within the machine. Thus the state of the automaton and its changes of state will be dependent upon the stimulus history of the machine (Babcock, 1960, pp. 45–46).

Because components can take the outputs of other components as inputs, one adaptive automaton could be made to “observe” another and hence they could together achieve self-recognition. The self-recognizing system resulting from the combination of the two systems could observe the effects of the outputs from its adaptive subsystem and modify them – thus becoming a truly self-organizing system. Considered from a perspective outside of the Self-Organizing Systems research agenda, the Adaptive Reorganizing Automaton is a very strange machine indeed. As an instrument, it does not measure anything. As an engineering accomplishment, it does not perform

any needed tasks, or serve any obvious functions. It is only when the device is viewed within the context of its disciplinary matrix that we can see its true purpose: it is a perfectly concrete embodiment of the abstract concept of a Self-Organizing System according to von Foerster's theory. That is, it takes energy and information from its environment, and transforms this into an internal order of connections. And yet, it manages to avoid being designed to do this in a pre-determined way. If it were to do this, it would be "other-organized" instead of "self-organized."

The device becomes an instrument to the extent that it allows researchers to make direct observations of the processes of self-organization. Unlike biological organisms that experience these processes too quickly, too slowly, or obscure them in microstructures, the Adaptive Reorganizing Automaton is designed to make its processes observable. And it is unlike a computer program that strictly determines the transformations of a computer memory and therefore has an observed behavior which can only depart catastrophically and uselessly from specified behavior. While the behavior of its units are specified, the behavior of the whole assembly and its interactions and transactions with its environment are not. Thus, the Adaptive Reorganizing Automaton constituted an embodiment of the theory of Self-Organizing Systems. It was not the only machine that explored the ideas of self-organization, however, and through the development of the early bio-computers of the BCL the ideas of self-organization were extended to the notion of preorganization.

4 Principles of Preorganization and the Numa-rete

Another significant element in von Foerster's development of the fundamental theory underlying Self-Organizing Systems was something he called "Preorganization". This notion is first mentioned in a BCL report "Some Principles of Preorganization in Self-Organizing Systems" (Babcock, 1960b), just after the Allerton Conference on "The Principles of Self-Organization" which was organized by the members of the BCL in June of 1960. The basic idea of "preorganization" was that systems, organisms or machines, did not deal with the totality of the universe, but only dealt with certain aspects of it and filtered the rest out. As von Foerster states in his preface to the conference proceedings:


The main theme of this report is a particular facet of the general problem of pre-organization in self-organizing systems, namely, the theory and circuitry of information processing networks.

One may consider these networks as a special type of parallel computation

channels which extract from the set of all possible inputs a particular subset which is defined by the internal structure of the network. The advantage of such operationally deterministic networks in connection with adaptive systems is the obvious reduction in channel capacity of the adaptors, if it is possible to predetermine classes of inputs which are supposed to be meaningful for those interacting with the automaton (von Foerster, 1963, p. ii).

The main principles of preorganization investigated at the BCL were property filters, periodic functions and neighborhood logics. These latter two are foundational principles for parallel computation. Periodic functions are simple functions that can be repeated numerous times fairly “cheaply” in engineering terms, yet collectively they perform a sophisticated process. We will see an example of this later with the Dynamic Signal Analyzer. Neighborhood logics are similarly repeated parallel functions, but these interact with topologically close neighbors. A classic example of this is a neural phenomenon called “lateral inhibition.” This theoretical principle was exemplified in a parallel computer known as the “Numa-Rete” (see Figures 3 & 4 from Halacy, 1965).





Dr. Heinz von Foerster and *New Scientist*

The *a-seer* reports the number of top and bottom edges of each object that it "sees" in each strip. The difference between the sum of edges from each two adjacent strips totals 16 here; one quarter of this is the number of objects.

While not nearly so ambitious as the Adaptive Reorganizing Automaton, the Numa-Rete was probably the most conceptually accessible of the BCL projects due to its simplicity and efficacy. "Numa" refers to number, and "Rete" refers to an anatomical mesh or network, etymologically applied to veins and arteries. The simple device had an input field of photocells arranged in a 12 x 12 matrix, and a two-layered electrical circuit which computed the number of opaque objects placed on the input field and displayed the result on a digital counter. The device cost \$1,200 to build in 1961 and was described as a "toy" in many of the presentations made by von Foerster and others from the lab (Weston, 1961). A larger 20 x 20 device was also constructed, and plans were made for building a much more sophisticated device capable of several more topological functions and consisting of a 70 x 70 or even a 400 x 400 input matrix, but further development was halted after the ONR canceled their funding to the BCL. Because of its straightforward demonstration of crucial theoretical aspects of parallel computation, the device was frequently described and displayed in public presentations and popular publications like *Electronics* (Weston, 1961) and *New Scientist* (von Foerster, 1962b) and even made its way into a popular book on *Bionics* (Halacy, 1965).

Designed by Paul Weston, then a graduate student in electrical engineering, and built by him with the assistance of M. Knott, G. Goodall, and G. Gunsalus, the device effectively computed the number of objects on its input surface in

parallel by utilizing the homeostatic properties of feedback circuits involving the photocells. The Numa-Rete actually processed the stimuli in a sequential fashion – from the moment it was turned on each row of photocells was activated in sequence. If light struck the photocell (no object was present at that location), then the computational circuits attached to the cell would receive an appropriate current. If no light fell on the photocell (meaning that the field was obstructed by an opaque object at that point), then the input circuit would not complete and no current would be directed to the attached computation circuits. The device's ingenuity was to be found in its computation circuits – all one of two types arranged in a checkerboard pattern – each of which performed a simple summation over input currents from three adjacent photocells at a time (Inselberg, et al., 1960, p. 2–41).

While summation is a simple function which seems to carry less information than the photocell circuit itself due the loss of the specific configuration of the matrix (in the sense of atomic impressions) in computation, when appropriately arranged as a network the computational circuits perform as edge detectors, or more precisely as edge counters. If all three of the photocells attached to the circuit are receiving light (or all are in the dark since the edges detected are simply contrasts and “light” and “dark” are not representational values in the network – an image and its negative constitute the same number of objects and the machine “sees” these as the same) the summation is 0 ($2 + -1 + -1 = 0$). If one of the adjacent photocells receives light, but not the other two, the summation is -1 ($0 + 0 + -1 = -1$), and if only one of the adjacent photocells is dark and the other two receive light the summation is 1 ($2 + 0 + -1 = 1$). Given the topological constraint that every discernable object must be finite and therefore have two edges within the field, one can calculate the number of objects intersecting each of the one dimensional rows of the field by adding their respective summations and dividing by 2. This also works when an object is so small as to cover only one photocell or row and not any of the adjacent ones – the summation results in two edges ($2 + 0 + 0 = 2$). By then performing a summation and division by 2 of the columns in another layer of circuits, it is possible to count the total number of objects on the input field regardless of topology – objects could have holes, or objects inside of holes – and regardless of the location on the field where those objects are placed (provided that there is at least one photocell between two objects – a limitation of resolution).

The Numa-Rete was designed to be a network capable of computing the “*n*-ness” impending upon its sensory field. It did not matter where the object was on the field, it was capable of extracting the invariant property of “number”. This was argued by von Foerster to be a product and virtue of the periodicity

of the computation – the idea that the computations (in this case summation) themselves are dreadfully simple but when repeated in periodic intervals, as in a parallel network, it was possible to extract invariant properties from the patterns of inputs.¹ Other connections between periodicity and complexity would come later from the study of chaos, fractals and automata (Stevens, 1974, Mandelbrot, 1977, Wolfram, 1994). Warren McCulloch had long been interested in finding “universals” in the generalizing and abstractive properties of networks (Pitts & McCulloch, 1947), and von Foerster argued that the invariants extracted by a network of this kind were the “Platonic Ideas” of the Twentieth Century (von Foerster, 1962c, 1962d).

In the true form of a biological computer, there remains some ambiguity as to whether the Numa-Rete device is “truly” a parallel or a sequential computer. As noted earlier, it operates sequentially, though technical reports claim that this is merely an engineering solution aimed at building the device inexpensively (Babcock, 1960c, pp. 76–77). Also, the device is only able to accept one set of inputs in a single operation, and must be switched “off,” or reset, in order to recalculate the number of objects on its input field, so it could not operate with dynamic inputs. Still, every step in the sequence is a parallel performance of 12 calculations in the input layer and even more in the calculation layers below; the entire calculation is completed in just 12 steps. Moreover, it does all of its calculation for enumerating discrete objects in analog rather than digital form – since numbers are represented by physical quantities of current rather than “digits” (except for the conversion at the digital display device for final output) – and there is nothing analogous to a recursive serial “counting” procedure (the computation is achieved purely by exploiting the properties and constraints of the network and its problem domain). Thus, the computation was performed by both the 12 x 12 and 20 x 20 devices in 0.20 of a second, comparable to a typical human performance of the same task, given the resolution of the input field. However, this implies

1 “I have first made mention of networks in general and have later introduced the restriction of periodicity. It is significant that this restriction still leaves us with a host of different kinds of property filters, because repetition of a structure is a simple task compared to altering the structure itself. This has of course, been realized by nature again and again, because a genetic command “repeat structure X so and so many times” is a simple and reliable operation, particularly after X has proven itself to be an evolutionary success. However, the command “change structure X into Y” is difficult and risky. Y may turn out to be a flop – on the other hand, it may turn out to be a great success. Only the epigones can tell, and they will always tell a success story, for in the other case there are no epigones!” (von Foerster, 1962d, p. 36)

a “counting” speed of 20,000 objects a second for the proposed 400 x 400 photocell device, which would put its performance well ahead of any human performance for the counting task.

While the Adaptive Reorganizing Automaton sought to demonstrate the fundamental principles of self-organization, the Numa-Rete sought to illustrate some extensions to the basic theory. Preorganization as a concept is already deeply tied to engineering concerns for how to implement and realize technological systems. Its discussion presumably arose out of discussion of how to make devices like the Adaptive Reorganizing Automaton more adept at dealing with real environments, and with explaining how self-organizing brains might do this – all in the conceptual framework of self-organization. That is, they could have simply talked about perception or specialized sensors, but instead sought to devise systems that achieved similar goals within the theoretical framework of self-organization. The engineering work of the BCL was thus conceptually-driven in this sense. We will see more of these extended principles – periodicity and property filters – in the Dynamic Signal Analyzer, but first we will consider the Bionics movement.

5 Mission of Bionics

Bionics was conceived of as an interdisciplinary venture of a very particular sort by the Air Force flight surgeon, psychiatrist and neurophysiologist Major Jack Steele. He coined the term in August of 1958 for the new science of “systems whose function is based on living systems, or which have characteristics of living systems, or which resemble these” (Gérardin, 1968, p. 11). Its research agenda was soon sharpened to solving various sorts of technical problems based on knowledge of living systems through a radical mix of techniques taken from biological, mathematical, and engineering expertise. While the typical relationship between specialists from these areas might be for the engineer to design tools, for the biologist to inquire into life, and for the mathematician to provide tools for the analysis of the data obtained from such inquiry, this was explicitly *not* Bionics as Steele envisioned it:

The biologists are not helping the engineers with their problems. The attitude of the biologists – that they know nothing of value, and that only at some future date, after having received still more assistance from the other two disciplines, will they acquire marvelously useful knowledge – is one of the greatest impediments to successful collaboration.


This situation is the consequence of natural selection. A man with a primary passion for synthesis or creative design simply does not become a biologist. The problems and attractions of biology are not those of synthesis and design, but of

observation and analysis. Lest you think I am over-critical of the biologist I shall remind you of the problems brought to the collaboration by the engineer and the mathematician.

The engineer typically resents the sloppy amorphous quality of biological knowledge, its lack of precision and its multivariant complexity. The engineers often feel they have nothing to learn from this messy science and that mathematics is adequate to their needs. This means that someone like Boole must convert biological principles into nice neat equations so the engineers will not be offended.

Finally to the mathematician. He enjoys manipulating symbols. If these symbols stand for nothing recognizable to anyone else, biologist or engineer, then the mathematician refers to the manipulation as Pure Mathematics. The mathematician prefers abstractions, the less he abstracts from, the more he can reject of reality, the happier he is. The mathematician's "neurone" is stark indeed. He quickly strips it of all its biologically interesting features to get on more quickly with the more enjoyable occupation of symbol manipulation. He strips his problems of these features which would make them interesting to others to make them rigorously solvable and therefore of more interest to himself (Steele, 1960b).

For the three fields to combine their talents and techniques, a more innovative form of collaboration would be needed, and this is precisely what Bionics proposed to do. The interdisciplinary mix was captured in the symbol used by the annual symposia, first held in September of 1960 (see Figure 5 from the cover of Steele, 1960).



In his preface to the conference proceedings, von Foerster describes the relation between the three fields constituting Bionics and its heritage in the earlier interdisciplinary movement of Cybernetics:

All this [the early excitement of Cybernetics] is now history, and in the decade which elapsed since these early baby steps of interdisciplinary communication, many more threads were picked up and interwoven into a remarkable tapestry of knowledge and endeavour: Bionics. It is good omen that at the right time the right name was found. For, bionics extends a great invitation to all who are willing not to stop at the investigation of a particular function or its realization, but to go on and to seek the universal significance of these functions in living or artificial organisms.



The reader who goes through the following papers which constitute the transactions of the first symposium held under the name Bionics will be surprised by the multitude of astonishing and unforeseen connections between concepts he believed to be familiar with. For instance, a couple of years ago, who would have thought to relate the reliability problem to multi-valued logics; or, who would have thought that integral or differential geometry would serve as an adequate tool in the theory of abstraction? It is hard to say in all these cases who was teaching whom: The life-sciences the engineering sciences, or vice versa? And rightly so, for it guarantees optimal information flow, and everybody gains.

Hence, there could not have been found a better symbol to represent this situation than the one to be seen at the cover of this volume: A scalpel and a soldering iron connected by an integral sign. Is it to mean that Medicine and Electronics have joined forces for a better end? No, this would not be in the right spirit of bionics. The proper meaning of this symbol must be: "The Allegorical Unification of Analysis and Synthesis." The newborn child has started life with a generous endowment of productive and creative genes. May it use them well (von Foerster, 1960b).

Bionics is best understood not as a science, but rather as a technological design practice, a way of doing engineering that emphasizes utilizing knowledge of biological systems. Steele's conception of it is remarkably comprehensive in demanding interdisciplinary education, multiple specialization and a consciousness of one's own organizational and financial situation on the part of its practitioners, in addition to the practice of technical design. But it was von Foerster who would realize these demands in the production of a Bionic machine.

6 Dynamic Signal Analyzer

The Dynamic Signal Analyzer offers what is perhaps the most convincing early articulation of what Bionics could, or should, do – it was the material realization of a computational model of a specific biological function (see Figures 6 & 7 from Babcock et. al, 1962).



The function in question is one that was believed to be performed by the basilar membrane of the inner ear in mammals. The project's objective as stated in its Air Force Aeronautical Systems Division grant for May 1, 1959 to October 1, 1961 was:

to conduct research on the analyzing principles of the mammalian auditory system through the construction of a system which shall embody auditory analyzing principles as they are known in man, and permit in its design sufficient latitude as

to check the validity of present day theories of the functioning of the auditory system (von Foerster, 1962, p. 1).

Research into the functioning of the human ear had revealed some remarkable features of the basilar membrane, which lines the fluid-filled nautilus-shaped chamber of the inner ear and is covered with fine hairs. Namely that, due to the specific curvature of its physical structure and the lining surface's relative distances from the source of pressure waves emanating from the ear-drum and small bone assembly at one end of the inner ear, the arrangement of sensory hairs on the membrane effectively performed a known signal differentiation function on sound waves entering the ear.² This function was argued to be computationally equivalent to a Fourier transform being performed on the signal – an interesting result because this was the process commonly used to “sharpen” a signal in electrical engineering.


What Babcock, and his collaborators R. J. Erikson and D. M. Neill, had seen in the empirical studies of the inner ear was not simply the performance of a single Fourier transform, but instead that, due to the rows upon rows of tiny hairs, a vast array of transforms were being performed in parallel and in real-time by the physical structure of the basilar membrane. Thus, they sought to duplicate this parallel architecture in an analog computational device of their own construction.

The Dynamic Signal Analyzer was essentially an early, if somewhat glorified, comb filter. The output from the Dynamic Signal Analyzer was displayed on three oscilloscopes: one for each of the second and fourth partial derivatives of the Fourier transform, and one which could be connected to monitor any combination of the output coming directly from selected spectrum filters (see Figure 8 from Babcock, et al., 1962).

The device functioned quite well at discriminating and sharpening a sound's fundamental through a continuous real-time operation. The research intent was to utilize the Dynamic Signal Analyzer in finding correspondence relationships between the input of phonemes and other speech patterns and the oscilloscope's processed output. That is, the goal of the Dynamic Signal Analyzer was not simply to mimic a biological computation in a special

2 I am told it is now believed that the cochlea actually achieves this in virtue of the temporal extension that takes place as the sound waves spiral down the nautilus. This implies that the basilar membrane is actually sensing fine-scale disturbances over time, rather than scanning for a range of specific properties in parallel as the research with the DSA posits. It is not necessarily fair to hold historical knowledge up to current knowledge, but the comparison invites inspection.

purpose electronic computer, but to develop a scientific research tool for sophisticated sound analysis.³



Beyond exemplifying the design practice of Bionics, the Dynamic Signal Analyzer also served as the basis for extensible research and the exploration of other key theoretical concepts in the movement: property filters, periodic functions and pattern recognition. More specifically, it was a property filter that operated in parallel. It performed a kind of pattern analysis, namely analyzing patterns with respect to variations in a single dimension – time. Another intriguing feature of the device was its dynamic analysis of properties that were time dependent, and hence necessarily required parallel processing to be analyzed robustly in real-time. The spectrum filters of the Dynamic Signal Analyzer and the hairs of the basilar membrane also correspond to the principle of preorganization called a periodic function, in which a simple function is duplicated in periodic intervals to achieve a more complicated overall function. In this case identical filters tuned to periodic frequencies and hairs positioned at different distances from the signal source result in both mechanisms responding quickly and naturally to a broad range of different frequencies.

Unlike the Adaptive Reorganizing Automaton that was meant to be the object of investigation, the Dynamic Signal Analyzer's primary aim was the computational analysis of speech and sound patterns. Thus, it matched the

3 According to Paul Weston that the Dynamic Signal Analyser was designed and built under contract to the Air Force, and upon its completion and testing it was sent to Wright-Patterson Air Force Base. He does not know what experiments they actually conducted with it (personal communication, March, 2004).

research goals of Bionics exactly. It used biological knowledge, under an appropriate mathematical formulation, to address an important problem in engineering. The Dynamic Signal Analyzer was truly a Bionic Ear, insofar as it realized in electronic form an actual known biological computation and processed an input signal analogously to the basilar membrane of a human ear, producing comparable output in the form of visualized mathematical data. This was the essence of Bionics. By instantiating the theoretical framework of Bionics, and serving as a basis for further experimentation to extend the mathematical understanding of more advanced sound processing of the ear, the Dynamic Signal Analyzer embodied the principles of Bionics laid out by Steele and von Foerster.

7 Conclusions

Self-Organizing Systems and Bionics were both supplied with potent material instantiations of their ideas by von Foerster's lab. This being the case, what happened to these apparently forgotten fields? And what happened to von Foerster's BCL? Ultimately, the BCL was to shut down in the mid-1970s, and by that time, researchers were rarely using the phrases "self-organization" or "bionics" and there were no longer any conferences or symposia focused on these topics. In that sense, the movements were over. Which is not to say that the ideas and technologies developed during the 1960s did not have an influential legacy. Rather, I believe, the technologies of this early work were eclipsed by the technology of digital computer simulations, which overtook these movements by the late 1960s.

After von Foerster left the BCL, his work entered another phase. His primary concern with the construction of machines was replaced with a primary concern for language. His interest in Second-Order Cybernetics and Radical Constructivism, which had been developing throughout the 1960s, came to dominate his work. A good history of this development is given by Scott (2004). I believe that the highly linguistic and social focus of his writings in the two decades which followed the BCL are indicative of his departure from engineering concerns which had grounded his theory, and the fields of Self-Organizing Systems and Bionics, during the BCL years. While he remained sensitive to the issues of technology, engineering and design, he was no longer concerned with maintaining a lab, and the securing of funding and construction of devices which that entailed. So for his own part, von Foerster abandoned the analog machines of the BCL as well, but rather than trading them for digital computers, traded them for the abstract theorization of observing systems.

A small number of digital computer simulations were being conducted by the early AI researchers, as well as some of the neural network researchers as early as the late 1950s. By the mid- to late-1960s, these types of simulations were dominating the presentations at the Self-Organizing Systems and Bionics conferences, and the mainstream had chosen the digital computer program as its paradigm for understanding the mind and complex systems. Certainly, there were computers in the late 1950s, though access to them was difficult to obtain. Due to their value to business and government, however, there was a huge investment in computer development, making them cheaper and easier to program. By the end of the 1960s, most researchers in academia and industry would have access to some kind of digital computer and training in how to program them with the required simulations. So while Bionics and Self-Organizing Systems conferences supported neural network research, they soon excluded or ignored work in analog networks and focused exclusively on computer simulations of neural networks.

The analog machines of the BCL had the advantage of being clear and direct analogies of the theoretical systems they sought to instantiate. But the digital simulations were only at a slight disadvantage, requiring a further level of numerical abstraction. The greatest disadvantage to digital simulations is that even though they are physically realized within the computer they are not directly manipulable in the ways that the experimental analog apparatus were. This advantage, apparently, was not enough to save the practice of building analog machines.

The digital computer simulation had several advantages over specialized analog machines. The main advantages were the speed and ease with which simulations could be replicated and distributed. Whereas a lab desiring to extend the work on the Dynamic Signal Analyzer, Numa-Rete or Adaptive Reorganizing Automaton would first have to spend several months building a machine from scratch, a computer simulation of a neural net or logic theorem prover could be copied to paper or magnetic tape and run on a similar computer with little or no modification. Moreover, one hand-built copy of an analog computer such as the Automatic Reorganizing Automaton might behave very differently than another – and this reliability issue has a huge impact on the replicability of observations and experiments. Reliability was a major issue in early computing, but after a huge investment of time and resources, and the refinement of manufacturing processes, digital computers became remarkably reliable in performing calculations, and replication of a program's performance came to be almost taken for granted. Apart from the practical concerns, the computer simulations also supported the symbol manipulation so desired by mathematical theories.

While a great deal of attention on von Foerster's work focuses on his writings, I believe that his contributions to the technologies of cybernetics are as deserving of this attention. The early machines of the BCL clearly contributed to Self-Organizing Systems and Bionics as "concepts made flesh" and provided insights to the researchers in those fields that could never be achieved by reflection alone. Indeed, one of the features that made the BCL so unique was that it did build numerous machines using a conceptual framework and design methodology derived from cybernetic ideas, and the role of observing systems. Thus the very notion of "observing systems" for von Foerster meant not only sitting back and observing, but actually building systems to be observed, and systems which could themselves observe their environments. This was what the early bio-computers of the BCL aimed to do.