CHARACTERIZATION OF CYBORGED ECOSYSTEMS

by

O. GRANT CLARK

Department of Agricultural & Biosystems Engineering
McGill University
Montreal

August, 1999

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

0-612-64539-8

ABSTRACT

In this thesis, a philosophy and lexicon for the engineering of biosystems are established. The focus is on a specific class of biosystems (*ecocyborgs*) created by combining ecosystems and technological components. This work is part of the EcoCyborg Project, a highly interdisciplinary research program which concerns the development of a general theory for biosystems engineering, with an emphasis on system autonomy as a design goal. In the short term, the objective is to develop computational models and simulations for use in the study of ecocyborgs as representative instances of substantially autonomous biosystems. Accordingly, in this thesis an explicit conceptual basis is established for the EcoCyborg Project, as well as for biosystems engineering in general.

First, in the body of the thesis, a biosystem is defined as a coherent assemblage of entities that is alive to some degree as a whole. The sole criterion for life is considered to be comportment that is somewhat *autopoietic*, whereby local interactions among the components combine to continually renew the overall system. Next, concepts related to *autonomy*, or the formulation and pursuit of proprietary goals, are elaborated. The degree of autonomy of a system is seen to depend on its *consciousness*, or ability to reason using a model of itself. Hence, a substantially autonomous system requires an ensemble of information storage and processing devices (*mind*) of the type and sophistication (*intelligence*) appropriate for this. The approach that is taken here to the creation of ecocyborgs with such minds is described, and a specific mental architecture is delineated, comprising functionally semidifferentiated, intermediate-scale components arranged according to a semihierarchical control organization. Finally, the characterization of such systems is scrutinized as an epistemic process in which knowledge is generated by an observer, but in which only a limited degree of objectivity is possible. A paradigm appropriate to the engineering of ecocyborgs is defined as an illustration, and associated archetypal concepts and descriptive procedures (such as measures) are given that are useful in this context. Such tools are required by significantly autonomous ecocyborgs because they must characterize themselves. They are also necessary to observers with scientific and engineering agendas.

RÉSUMÉ

Dans cette thèse sont établis une philosophie et un lexique pour l'ingénierie de biosystèmes. Les biosystèmes considérablement autonomes créés par une combinaison d'écosystèmes avec des systèmes de contrôle (*écocyborgs*) sont ici d'un intérêt particulier. Ce travail fait partie du projet ÉcoCyborg qui, à long terme, concerne le développement d'une théorie générale des biosystèmes, avec l'accent sur l'autonomie substantielle de ceux-ci. À court terme, l'objectif consiste en une étude de l'utilité de différents écocyborgs pour un tel projet d'ingénierie. Dans la thèse une base conceptuelle explicite est proposée pour le projet ÉcoCyborg et pour l'ingénierie des biosystèmes en général. Ce travail est fortement interdisciplinaire, englobant l'étude de l'ingénierie, des sciences de la vie, des sciences cognitives et des systèmes complexes.

Dès le départ un *biosystème* est défini comme un assemblage d'entités qui est, à un certain degré, vivant dans l'ensemble. Le critère unique pour être vivant est le comportement, qui se distingue par une certaine *autopoïèse*, de telle manière que les interactions locales entre certains composants se combinent pour continuellement renouveler le système entier. Par la suite sont élaborés des concepts liés à l'*autonomie*, c'est-à-dire à la formulation et la poursuite d'objectifs privés. Le degré d'autonomie d'un système dépend de sa *conscience*, voire de sa capacité de raisonner en utilisant un modèle de lui-même. Il lui faut donc un ensemble d'information et des dispositifs de traitement de l'information (*traitement mental*) d'un type et d'un raffinement (*intelligence*) appropriés. Une explication est alors donnée de l'approche ici adoptée pour la création des écocyborgs avec de telles capacités mentales. Ensuite, une architecture mentale est décrite sous l'angle des composants fonctionnellement semidifferenciés et d'échelle intermédiaire, disposés selon une organisation semihiérarchique de contrôle. En conclusion, la caractérisation de tels systèmes est présentée comme un processus épistémique où un observateur produit la connaissance, mais où seulement un degré limité d'objectivité est possible. Un paradigme approprié à l'ingénierie des écocyborgs est défini en tant qu'illustration, et des archétypes conceptuels ainsi que quelques procédures descriptives utiles dans ce contexte (telles que des mesures) sont donnés. De tels outils sont cruciaux pour les écocyborgs considérablement autonomes parce qu'ils doivent

s'auto-caractériser. Ils sont également requis par tout autre observateur ayant des visées scientifiques (descriptives) et d'ingénierie (prescriptives).

ACKNOWLEDGEMENTS

I am indebted first to my thesis supervisor, Professor Robert Kok. I am sincerely grateful to him for his tireless effort, enormous patience, and unfailing support. He walked with me not only along the high, clear ridges but also through the tangled thickets. Without him, my journey might have been easier, but it would certainly have been vastly less rewarding.

I am thankful for the support, advice, and friendship of the faculty, staff, and students of the Agricultural and Biosystems Engineering Department, and sundry other denizens of the Macdonald Campus of McGill University. Among these people, Drs. Pierre Dutilleul, Eric Norris, René Lacroix, Jacques-André Landry, and Vijaya Raghavan have my gratitude for serving as members of my advisory committee. I owe special thanks to the past and current members of the EcoCyborg Project Group: Dr. Steven Hall, Petra Kalshoven, Jennifer Karsten, Dr. René Lacroix, Dr. Robert Molenaar, and Lael Parrott. Ms. Parrott, Dr. Lacroix, and Dr. Molenaar contributed substantially to the writing of parts of this thesis, and all of the members of the group were involved both directly and indirectly in the work and play that led to this document's completion. Dr. Lacroix, Dr. Valérie Orsat, and Ms. Kalshoven reviewed various versions of the résumé. Dr. Stephen Murphy served as the external examiner, and provided the helpful suggestions that are addressed in the Addenda of the thesis.

I gratefully acknowledge the financial support of the Natural Sciences and Engineering Research Council of Canada, the J.W. McConnell Foundation, the Keith Gilmore Foundation, the President's Office of McGill University, and the Department of Agricultural and Biosystems Engineering.

During my studies, I had the good fortune to encounter many kind souls among the people of Ste-Anne-de-Bellevue and greater Montreal in general. Dr. Darwin Lyew has been my guardian angel and steadfast companion throughout my stay at Macdonald Campus, and Hilda Lyew has worried over me as if I was her own son. I am blessed that they chose to take me in. In other venues, the Seidokwan Judo Club and the choir of the Union United Church have always welcomed my sometimes fickle presence without question or criticism.

My being is rooted in the family that gave me life and nurtured my growth, both of which they continue to do despite the intervening distances. I thank them for their unconditional love. My sister Christine Stacey and my brother Peter Clark also took the time to read drafts of several chapters of this thesis.

Finally, I am filled with joyous wonder by the turns of fate that led to the union of my destiny with that of my wife, Norma Bautista-López. Our love for one another transforms even the most ordinary of days into uncommon adventure. Our little child, though still unborn at the completion of this thesis, is nevertheless a source of joy and inspiration.

TABLE OF CONTENTS

LIST OF FIGURES

CONTRIBUTIONS OF AUTHORS

This thesis has been prepared in accordance with the February, 1999 revision of the *Guidelines for Thesis Preparation* (Faculty of Graduate Studies and Research, McGill University). It is stated therein that:

> "As an alternative to the traditional thesis format, the dissertation can consist of a collection of papers that have a cohesive, unitary character making them a report of a single program of research. [...] The thesis must be more than a collection of manuscripts. All components must be integrated into a cohesive unit with a logical progression from one chapter to the next. In order to ensure that the thesis has continuity, connecting texts that provide logical bridges between the different papers are mandatory.
>
> "In general, when co-authored papers are included in a thesis the candidate must have made a substantial contribution to all papers included in the thesis. In addition, the candidate is required to make an explicit statement in the thesis as to who contributed to such work and to what extent. This statement should appear in a single section entitled 'Contributions of Authors' as a preface to the thesis."

In accordance with these guidelines, the contributions made by the various authors to each of the journal articles included in this thesis are indicated here. It should be noted that many of the concepts described in these particular chapters and in the rest of the thesis were developed over a period of several years, with contributions being made by all members of the EcoCyborg Project Group. It is, therefore, sometimes difficult to attribute the origin of a particular idea to any one person. Those people who made substantial contributions to each of the chapters are, however, indicated in the following paragraph and in the citation information in the connecting texts.

Chapters 1, 2, 8, 9, and 10 of this thesis, as well as all prefacing and connecting texts, were written entirely and solely by O.G. Clark. As indicated in the acknowledgement sections of the thesis overall and of some of the individual chapters, R. Kok extensively reviewed and corrected the entire thesis, and several other people also reviewed drafts of some parts. Chapters 3, 6 and 7 were co-authored by O.G. Clark and R. Kok, who together developed the bulk of the content of these chapters during frequent meetings that occurred from September, 1993 to August, 1999. Some of the ideas presented in Chapter 4 are based on work originally done in cooperation with R. Lacroix during the period of 1990 to 1994, material that was later expanded into its final form by O.G. Clark and R. Kok. Most of the themes presented in Chapter 5 were originally conceived by R. Kok and some of his previous students (G. Desmarais, L. Gauthier, and R. Lacroix; see the literature review in Chapter 2). All members of the group working on the EcoCyborg Project during the writing of this thesis (O.G. Clark, R. Kok, R. Molenaar, and L. Parrott) advanced the original concepts and worked toward their implementation. The text of Chapter 5 is based partly on sections of O.G. Clark's thesis proposal, the forerunner of this dissertation, although those sections were extensively rewritten and expanded by the authors: L. Parrott, O.G. Clark, and R. Kok.

CHAPTER 1. GENERAL INTRODUCTION

1.1 Biosystems engineering

This doctoral dissertation is written from an engineering perspective, as opposed to a scientific one, a distinction that is not always made clear elsewhere. Although engineering may involve the use of scientific methodology and the application of knowledge acquired through such means, its underlying philosophy differs fundamentally from that of science. Science is an explanatory enterprise, while the focus of engineering is, ultimately, not so much on explanation but on creation. The scientist is concerned with observing, understanding, and describing that which already exists, whereas the engineer is concerned with imagining and then bringing into being something that has never existed before. Engineering activities include, therefore, the design, creation, operation, maintenance, repair, modification, and upgrading of systems, usually with the intent of achieving certain predetermined objectives. The work described here was conceived in this spirit; although it is focused on the characterization of certain kinds of systems, the intent is to specify these in a prescriptive manner. Thus, the observation, understanding, and description of existing systems are seen as a means to the end of creating new ones. Such effort is facilitated by an intellectual framework that explicitly emphasizes these creative intentions. The exploration and development of an appropriate engineering philosophy is, therefore, an important aspect of this thesis.

Engineering is an extremely broad discipline that encompasses and draws upon many overlapping fields. One of these is biosystems engineering, itself a very broad area, and one that has yet to be clearly defined. The reason for this lack of clarity is that, to this point, there has been no concise definition of the class of systems that are of interest (i.e., biosystems). This illustrates that in any field of endeavor a suitable lexicon is of fundamental importance. Hence, the development of a lexicon appropriate for biosystems engineering is another significant aspect of this thesis, part of which centers around the definition of the term *biosystem* and related concepts. The biosystem class as defined here (Chapter 3) is quite diverse, with members ranging in organizational scale, physical size, and type from molecular to planetary, from natural to artificial, and from imaginary to real. Thus, living things are not considered as being necessarily biological, but life is instead interpreted as an essentially informational (*virtual*) phenomenon that can reside on a wide variety of substrates, biological as well as non-biological. Biosystems engineering is, therefore, a correspondingly diverse field, dealing with ways in which such systems can be created or manipulated, usually with the intent that they fulfill some predetermined objectives, however general.

Much of the work described in this thesis is generally applicable to all biosystems, but there is an emphasis on a particular class of these (*ecocyborgs*) that are organized at the ecosystem scale and that have been augmented with technological components. These are dealt with as consisting of a collection of biological organisms and their abiotic surroundings (i.e., an ecosystem), together with a set of components that, in this work, are added with the intent of guiding the comportment of the overall system. Equivalently, they may also be systems that are organized in a fashion similar to a biological ecosystem, but that comprise altogether different kinds of components. The ecosystem part of an ecocyborg may range from purely natural, to somewhat modified, to entirely artificial. Ecocyborg engineering reduces to pure ecosystem engineering if the technological components are insignificant, but this would be an extreme case. Most current engineering practice involves some kind of technology, so that the resulting systems are usually cyborged to some extent. Even the rehabilitation of a natural habitat, for instance, frequently involves the installation of technological components like monitoring equipment and pumps. Thus, the engineering of ecocyborgs is important for a large number of applications not only in research and industry, but also, for example, in education, recreation, and conservation. Accordingly, the systems that are created can fulfill a broad spectrum of possible objectives related to productivity, environmental remediation and enhancement, exploration, housing, etc. Thus, ecocyborgs can range from production greenhouses in which food crops or ornamental plants are grown, to fermentation vats that produce foodstuffs (such as cheese and yogurt) or chemicals (like pharmaceuticals and alcohol). As well, they include tropical aquariums, botanical gardens, and managed natural habitats. They might even be submarine or polar living quarters, or orbital space stations containing human crews, with the life support system for the International Space Station being an instance of the latter.

1.2 Mind and autonomy

Every stage in the engineering of a system, from design to upgrading, involves a great deal of informational activity. Humans have for millennia augmented their physical capabilities with animals and machines, and have done the same, to some extent, with their informational capabilities. Technologies such as writing, for example, are important aspects of many cultures. Recently, however, the possibility of shifting mental tasks from humans to other entities has increased enormously, due to advancement in the cognitive sciences and the development of technologies such as computational electronics. It is now feasible to consider the engineering of systems that can independently perform many kinds of informational tasks that could once only be done by humans. Systems with such abilities are able to act in a much more autonomous manner than was previously possible. This involves not only automation, but also the capacity to formulate goals, as well as to invent and execute strategies for attaining them.

Significant autonomy is desirable whenever a system must fulfill particular objectives in an unpredictable environment (including basic objectives such as persistence). This is the case whenever human guidance is rendered impossible or impractical by expense, distance, danger to human operators, or by the extreme complication of the system. For example, the effectiveness of direct human guidance of many natural and modified natural ecosystems is limited by these systems' intricacy. Making them significantly autonomous could improve their ability to respond to unpredictable, otherwise disruptive changes in their surroundings and in their own constitutions. Substantial autonomy is also desirable for ecocyborgs, such as goal-oriented production facilities or space habitats, whose ecosystems are entirely artificial, because it can make them more robust and self-reliant.

Although autonomy has, in various guises, been discussed for centuries, the engineering of systems so that they are substantially autonomous is, as mentioned, a relatively novel pursuit. This is especially true with respect to biosystems in general, and ecocyborgs in particular. A lexicon suitable for the coherent discussion of such an engineering exercise has therefore been lacking. Another important theme in this thesis is, therefore, the definition of a vocabulary that is appropriate for this purpose (Chapter

4). Along with autonomy, other related concepts, such as consciousness, intelligence, and mind, are also explored.

1.3 The EcoCyborg Project

Two of the principal themes that are dealt with in this thesis, namely biosystems engineering (especially of ecocyborgs) and the engineering of systems for substantial autonomy, come together in the guise of the EcoCyborg Project. This research program is being conducted in the Department of Agricultural and Biosystems Engineering of McGill University. The long-term goal of the EcoCyborg Project, as originated by Professor Robert Kok and his students (Chapter 5), is to develop a general theory of biosystems engineering, with emphasis on the design goal of substantial system autonomy. The engineering philosophy that underlies the project, already mentioned as a principal theme of this thesis, is explained in some detail in Chapter 5. In light of the long-term goal of the EcoCyborg Project, the short-term goal is to develop computational models and simulations for use in the study of ecocyborgs as candidates for the engineering of systems that are substantially autonomous. Thus, in Chapter 5, there is also an examination of the enhancement of ecosystems with technological components as an approach to making them more independent in their comportment.

Kok and his students are the first to have formalized and researched the idea of cyborging ecosystems in order to enhance their autonomy, but activities are already underway in other venues that can be considered as the cyborging of both natural and artificial ecosystems with this result. For instance, the rapid development of telecommunications and remote-sensing technology and its deployment on a planetary scale can be considered as the cyborging of the terrestrial biosphere in a way that might conceivably make it more autonomous. On a more modest and immediate scale, many greenhouses, industrial fermentation facilities, and animal housing installations are being made increasingly autonomous due to the growing sophistication of their extrinsic control systems.

The particular case that is currently being studied in this project (i.e., the EcoCyborg itself) is a hypothetical orbital space platform. The initial phases of this project are focused on the development of computer-based modeling, simulation, and characterization tools for the study of different configurations of the EcoCyborg. These are being used to describe the interaction of its constitution (composition and structure), initial state, and comportment (dynamic changes in state) with forcing functions (such as weather). In the future, the focus of the project will shift toward understanding how to engineer the various parts of the system so that it will possess particular design features, including substantial autonomy. This will require the engineering not only of the ecosystem, but also the configuration of the control system so that it hosts a mind of the appropriate type and sophistication (Chapter 6).

1.4 The characterization of cyborged ecosystems

The philosophical and lexical themes that are presented here are all related to the characterization of biosystems generally and ecocyborgs particularly, especially those of substantial autonomy. These topics are presented, as mentioned, from an engineering perspective, with the motivation of contributing toward the development of a general engineering theory for such systems. Hence, this thesis can, in a larger sense, be considered as an exercise in characterization.

A comprehensive approach to the characterization of large-scale biosystems, such as ecocyborgs, has been lacking to this point. Methods of characterization are necessary in the EcoCyborg Project for three reasons. First, any biosystem that is substantially autonomous must be capable not only of observing and responding to its surroundings, but it must also be able to observe and control itself and therefore requires effective characterization methods. Second, such methods are required by the observers of these systems in order to generate scientific descriptions and, third, they are necessary for creating prescriptive, engineering specifications. The latter part of this thesis is devoted, therefore, to an examination of the epistemics of characterization (Chapter 7). This is illustrated with suggestions for characterization methods that are appropriate for use in the engineering of ecocyborgs.

1.5 Objectives

Thus, the principal objectives underlying the work described in this thesis are:

1) **To describe an engineering philosophy** that facilitates the understanding and creation of substantially autonomous biosystems.

2) **To develop a coherent lexicon for use in the characterization of biosystems**, and to explain this in a systems-theoretic context that is appropriate for the engineering of such entities.

3) **To develop a coherent lexicon of the concepts related to autonomy**, and to describe how this lexicon can be employed in the (descriptive and prescriptive) characterization of substantially autonomous systems.

4) **To describe a viable approach to the engineering** of novel biosystems, especially those of the ecosystem scale, through the combination (cyborging) of biological and technological components.

5) **To outline the mental architecture** required for substantial autonomy in cyborged biosystems.

6) **To examine and illustrate the characterization** of substantially autonomous biosystems as an epistemic process.

The achievement of these objectives to any extent will contribute toward the engineering of ecocyborgs, as studied under the auspices of the EcoCyborg Project. This work also has much wider applicability, corresponding to the long-term goal of developing a general theory for biosystems engineering, with an emphasis on substantial system autonomy.

CONNECTING TEXT

Chapter 2 has been published as: **Clark, O.G., and R. Kok. 1998. Engineering of highly autonomous biosystems: review of the relevant literature.** ***International Journal of Intelligent Systems*** **13(8):749-783**, (Copyright © 1998, John Wiley & Sons, Inc.).

This chapter is the general literature review for the thesis. Having been published in 1998, it does not include any references to literature published in after that date. Moreover, due to the multidisciplinary nature of this project, the bibliography was not intended to be comprehensive, but rather to present a general overview of literature associated with the relevant themes. The reference sections of the other chapters should, therefore, be consulted for more recent and specific citations relating to the corresponding topics.

CHAPTER 2. ENGINEERING OF HIGHLY AUTONOMOUS BIOSYSTEMS: REVIEW OF THE RELEVANT LITERATURE

Abstract

This article is a general guide to the literature associated with the development of highly autonomous biosystems. The specific context of the article is the EcoCyborg Project, in which computer models are used to investigate the engineering of ecosystems combined with artificially intelligent control networks. The project exists at the nexus of several expansive fields of research, and the review therefore is not comprehensive. Instead, it is a general guide to the literature associated with the relevant themes. First we give the definition of a biosystem as an adaptive, complex, dynamic system that is alive to some degree. A brief overview is given of the historical development of holistic ecology, followed by a discussion of what it means for a system to be "alive". Second we review the engineering of natural, modified, and entirely artificial biosystems for various purposes. The next section is on the engineering of biosystems for autonomy, including the characterization of mind, artificial intelligence, the implementation of mind in biosystems, and the history and current nature of the EcoCyborg Project. Finally, mention is made of techniques for the characterization and comparison of highly autonomous biosystems, since these techniques are necessary both for the objective study of such systems and for their own self-examination and control.

2.1 Introduction

This paper is a review of literature that is relevant to the engineering of biological systems. It was written in the context of the EcoCyborg Project, a research program with the long-term goal of learning how to engineer highly autonomous biosystems. The case studies that are currently under way in the EcoCyborg Project are intended to investigate the possibility of creating large-scale autonomous biosystems containing both ecological and technological components. Although this is the kind of system that serves as the focus for this review, the ideas that are dealt with here are extensible to other kinds of biosystems. The central themes are the nature of biosystems, the nature of mind, and the engineering of biosystems. There is also some discussion of methods for the characterization and comparison of biosystems. These methods serve two purposes in this project: first, they serve as tools for the objective scientific investigation of the systems of interest, as well as for the development of a theory of engineering such systems. Second, such methods must be available to the biosystems themselves if they are to be autonomous to any degree. This is so that they can be made aware of their own state in relation to their environment, and of changes in that state, enabling them to react in an intelligent manner.

The literature cited here is drawn from a number of different fields of research, reflecting the interdisciplinary nature of the EcoCyborg Project. However, there are extensive bodies of literature directly and indirectly associated with each of the themes mentioned above, and to attempt a comprehensive review of all of them would exceed the bounds of this article. Therefore, most of the references that are presented are overviews, works of a general philosophical nature, or representative samples of the current state of knowledge in the relevant fields. Also cited are works that have had a particularly significant influence on the evolution of the EcoCyborg Project.

2.2 The nature of biosystems

The engineering of biosystems forms the context for the EcoCyborg Project. Since biosystems engineering is a relatively new field, the associated terminology is in a state of flux. It is therefore necessary to define the paradigm and lexicon that will be adopted in this article. To this end, biosystems are defined as adaptive, complex, dynamic systems that are alive to some degree. A review is presented below of the background literature on which the definition is based. The underlying concepts are generally applicable to biosystems of all scales, but in accordance with the focus of this article, the literature of most direct importance is that which treats ecological systems theory. Nonlinear dynamics and complex systems theory are also very applicable to the current understanding of large-scale biosystems, and so these topics are also briefly touched upon. Finally, the definition of biosystems used here specifies that they are living entities in their own right, and this requires a brief review of the literature that treats the nature of life.

2.2.1 Ecological systems

Ecology has been a strongly integrative discipline since its inception, placing importance not only on the properties of the individual system components, but also on those of the overall system. This approach is extensible to the study of biosystems at all scales, but was originally developed in the context of ecosystems. This is the stream of literature that will be traced here.

A holistic view of nature has roots in ancient religious beliefs, and has occasionally been popular during the history of Western scientific thought. Plato, for example, espoused the idea of pantheism, of the universe as a god, or a single living thing synonymous with the creator. The concept was revived during the Renaissance, when Greek and Roman ideas were reintroduced to Western Europe, and promulgated by scholars such as Bruno, Spinoza, and Goethe (Margulis and Sagan 1995). More recently, Smuts wrote about holism early in the century, and was a direct influence on the founders of modern ecology (Tansley 1935).

Clements was an influential ecologist at the beginning of the twentieth century who espoused a holistic view of natural systems. He spent his professional career studying the vegetation of western North America, and described his theory of succession in a famous monograph, *Plant Succession: An Analysis of the Development of Vegetation* (Clements 1916). He asserted that vegetative communities develop toward a particular stable configuration, the character of which is dependent primarily on the local climate and physiography. The member populations of this *climatic climax* are so functionally integrated that the system as a whole can be considered to be a "complex organism" or "superorganism" in its own right.

Tansley, a contemporary of Clements, disagreed with the use of the term "complex organism" in this context. He instead coined the term *ecosystem* to describe a community of plants together with the associated animal community and all of the physical factors forming their environment (Tansley 1935). This term is derived from the Greek *oikos*, meaning "household", and the root *systema*, which denotes a whole compounded of several parts. In his classic article on "vegetational concepts and terms", Tansley states that although our human prejudices often cause us to place importance on those parts of an ecological system that are individual living plants and animals,

> "...certainly the inorganic 'factors' are also parts – there could be no system without them, and there is constant interchange of the most various kinds within each system, not only between the organisms but between the organic and the inorganic. These *ecosystems*, as we may call them, are of the most various kinds and sizes. They form one category of the multitudinous physical systems of the universe, which range from the universe as a whole down to the atom. The whole method of science [...] is to isolate systems mentally for the purposes of study, so that the series of *isolates* we make become the actual objects of our study, whether the isolate be a solar system, a planet, a climatic region, a plant or animal community, an individual organism, an organic molecule or an atom. Actually the systems we isolate mentally are not only included as parts of larger ones, but they also overlap, interlock and interact with one another. The isolation is partly artificial, but is the only possible way in which we can proceed."

The currently accepted hierarchy of ecological systems, as listed by Odum (1993), includes, from least to most complex: organism, population, biotic community, ecosystem, landscape, biome, biogeographic region, and the biosphere. There is some qualitative distinction between these designations, since abiotic components only begin to be considered at the level of the ecosystem.

The holistic approach to ecology spread and strengthened during the middle part of the twentieth century, as typified by the work of the Odum family. This family of American ecologists emphasized the understanding of ecosystems first as functional wholes, and then through the investigation of the synthetic interaction of their biotic and abiotic components (Odum 1997). To accomplish this, they made explanatory use of formal systems theory. From this perspective, an ecosystem can be described using a set of state variables corresponding to: its principal properties; forcing functions, which are outside energy sources or causal forces that drive the system; flow pathways that connect properties with one another and with forces; interactions of forces and properties by which the flows are modified; and feedback loops by which a flow pathway will have an influence on an "upstream" component or flow. This kind of formalized analysis helps to makes clear the organizational patterns of ecosystems, the nature of the mass and energy flows within and through them (e.g., nutrient cycles), the dynamics of the constituent populations, and the way in which the system as a whole changes with time (e.g., the succession of plant communities as described by Clements).

Ecologists such as the Odums promoted a holistic, synthetic approach to all of science, to complement the reductionist, analytic approach that had been prevalent. As well, they advocated greater awareness of ecological interdependence in the formulation of economic and political policy. Partly as a result of the efforts of ecological proponents such as the Odums, the middle decades of the twentieth century saw an awakening of scientific and public concern regarding the impact of humanity on the state of the environment. Incidents such as Love Canal contributed to the widespread realization that humans are not isolated from nature, but are an integral part of it. Actions that significantly changed the environment could have a dramatic long-term impact on peoples' health and living standards. This concept continued to be popularized through publications such as *Silent Spring* by Carson (1962). The space programs of the Soviet Union and the United States also had an important but more subtle impact on the environmental awareness of the public. Images from space enabled people to see the entire planet for the first time, to perceive it as a single isolated entity in the immense emptiness of space, and to gain some appreciation of the unity and fragility of the biosphere. Ironically, the driving force behind the space race, the development of long-range nuclear weapons, made even more real the threat of global ecological destruction (Sagan 1994).

With this increased global awareness, the perception of the biota and the abiotic environment as an integrated unit expanded from the level of the ecosystem to include the entire planetary biosphere. Margulis, a microbiologist, described the "microcosm" as a single planetary network that includes all living things, and on which humans are integrally dependent for their existence (Margulis and Sagan 1986). She considered the evolution of the biota of the world as being based primarily on the cooperative coexistence and coevolution of microbial populations. According to Margulis, for the largest part of the history of life on Earth, the biota consisted exclusively of unicellular bacteria, and when eukaryotic life finally arose, it did so through the symbiosis of communities of prokaryotes. In turn, multicellular organisms, which even today compose only a relatively insignificant part of the biota, could be thought of as intricate symbiotic networks of nucleated cells: walking communities of bacteria. She considered the whole of the global biota as a single, integrated bacterial community, which is manifested in some instances as tightly integrated symbiotic colonies that we perceive as multicellular organisms.

Lovelock carried this idea further in his Gaia hypothesis (Margulis and Lovelock 1974). He emphasized the dramatic impact of the biota on its abiotic surroundings, and the resultant difficulty in distinguishing between the living and the nonliving aspects of the biosphere. Lovelock first elaborated the Gaia hypothesis in the early 1970s, while employed by the National Aeronautics and Space Administration (NASA) to develop methods for detecting life on Mars. He realized that the constituent gases of the Earth's atmosphere are in gross chemical disequilibrium, whereas those of the Martian atmosphere are not. For example, the highly reactive gases of free oxygen and methane coexist in large proportions in the Earth's atmosphere. Lovelock's explanation was that the biota of the Earth regulates the balance of the atmosphere so that it remains in a state that is favorable for life. He proposed that over millennia, intricate feedback loops have developed that involve the Earth's vast microbial populations, atmospheric and geochemical nutrient cycles, and to some extent, populations of macroscopic organisms. Most importantly, these feedback loops maintain terrestrial conditions so that they remain favorable for the existence of life. Mars lacks such a regulating influence, and its atmosphere reflects this. Lovelock, following Margulis, went so far as to consider the Earth's biosphere as the analog of a single, enormous, self-regulating organism.

Concern for ecological issues continued to increase through the 1980s. Demographers warned of the dramatic potential consequences of the exponential growth of the world population. The "revolution of rising expectations" continued to increase the pressure on natural resources throughout the world, as burgeoning populations struggled to improve their living standards following the historical model of natural resource exploitation for short-term gain. Alarms were raised about deforestation in both developing tropical and industrialized temperate nations, the resulting extinction of unique species of organisms, and the possible effects of these trends on the integrity of the biosphere (Schultz and Mooney 1994; Wilson 1992). Meanwhile, the continued evolution of satellite technology strengthened global ecological awareness. For the first time it became possible to continuously monitor, in real time and on a planetary scale, trends such as the progressive destruction of forests; the extent and health of food crops; the encroachment of urban areas on agricultural land; and the changing composition and temperature of the atmosphere. With regard to the latter, tentative links were proposed between industrial activities and newly discovered global atmospheric phenomena such as polar "holes" in the ozone layer and global warming. At the same time, improvements in communication and transportation technology and infrastructure made the global culture stronger and more integrated. More people than ever before gained a knowledge and appreciation of other cultures and of the natural world as a whole.

This widespread adoption of a more holistic perspective has begun to color the whole of Western society, including scientific thought. Practitioners of the "hard" sciences have begun to approach the study of natural phenomena from the viewpoint of synthesis as well as from a purely analytical perspective, and to take an interdisciplinary interest in the problems of the social and life sciences. In turn, students of the latter disciplines are applying methodology from other fields to their work. This trend is deemed by many to be valuable for all of science, and beneficial to ecology in particular by increasing the depth and scope of the field and strengthening the body of mathematical theory available for the description of ecological systems (Patten et al. 1995). The broader study of holistic systems theory, abstracted largely from ecology, has taken the form of nonlinear dynamics and complex systems theory. Because of the current importance of these approaches to the study of living systems (biosystems in general and ecological systems in particular), a brief review of the development of these fields is presented below.

2.2.2 Nonlinear dynamics and complex systems theory

As mentioned above, the holistic perspective of the world is an ancient one. Capra (1996) presents a brief history of this philosophy up to the present day. However, reductionism has prevailed in Western science since the time of Descartes, and was especially strong from the middle of the nineteenth century until the middle of the twentieth. During this time the scientific disciplines were somewhat isolated from one another, and new ideas did not spread rapidly among them. This was probably due largely to the prominence of reductionism itself, which favored the detailed study of isolated aspects of phenomena and was not conducive to interdisciplinary efforts. Another factor contributing to this insularity between disciplines may have been that biological and sociological systems are difficult to describe using conventional mathematical tools. The mathematical methods available before the middle of the twentieth century were more suited to the analysis of linear systems, such as those dealt with in Newtonian physics. Although nonlinear behavior is commonplace in nature, it is generally intractable to classical methodology. It was therefore often considered to be the result of "random noise" or was neglected entirely. The illusion of the physical world as a linear one became pervasive, in the spirit of the adage, "when the only tool one has is a hammer, every problem begins to resemble a nail." Since biological and social systems are very rarely linear, a division grew between the so-called hard and soft sciences.

The aforementioned current of holistic thought began to stir in ecology at the end of the nineteenth century. Concurrently, radical changes in theoretical physics began with the formulation of relativity theory and then of quantum mechanical theory, which at very large and very small scales seemed to provide more appropriate models of the universe than did traditional Newtonian physics. The rediscovery that the physical universe is not always adequately describable in the linear, reductionist Newtonian paradigm, but is instead often exceedingly nonlinear, gained increasingly widespread acceptance. Digital computer technology and new mathematical techniques, such as improved numerical analysis methods, provided new means of studying complicated nonlinear dynamics.

In the 1970s there was a surge of interdisciplinary interest in *chaos theory*, the investigation of the general principles governing structurally simple deterministic systems that demonstrate unpredictably complicated behavior. Chaotic systems were described in fields such as biology, meteorology, fluid dynamics, structural mechanics, electronics, and economics, to name only a few (Gleick 1988). The field of chaos or nonlinear dynamics now boasts a well-established body of theory and an extensive technical literature, including many dedicated journals, textbooks (Scheinerman 1996; Peitgen et al. 1992; Thompson and Stewart 1986) and trade books (Hall 1991; Schröeder

1990; Gleick 1988). As often happens, it was realized that previously developed, relatively arcane bodies of mathematical theory were appropriate for the description of the new phenomena. Chaotic dynamics, as well as the geometry of many of the patterns that they generate in nature, could be described using mathematical objects called *fractals*. Historically regarded as pathological exceptions, or "mathematical monsters", the utility of fractals in the description of natural systems was pioneered by Mandelbröt (1983). Many of them proved to be exceedingly beautiful, and this helped to fuel popular interest in the field of chaos.

This interest in structurally simple but highly nonlinear systems paved the way for the study of larger, more complicated nonlinear phenomena. The late twentieth century has seen a widespread interest develop in the kind of complicated nonlinear systems that biologists, ecologists, and social scientists have been studying for decades. This has given rise to yet another new field of study, known as *complex systems theory* (Casti 1995; Waldrop 1992). As with chaos theory, complex systems theory (i.e., *complexity*, which is an umbrella term for this field of research) deals with the search for unifying principles that govern seemingly disparate systems. It therefore attracts a very interdisciplinary group of researchers. As a result, a strong spirit of cooperation has developed between different scientific disciplines. This is typified by the research and publications affiliated with the Santa Fe Institute, a privately funded institute that promotes the "transdisciplinary" study of new disciplines such as complexity, complex systems, and complex adaptive systems.

Many of the terms used to describe the phenomena in these new disciplines stem from ecology. For example, the terms *complex system* and *emergence* were both used by ecologists at the turn of this century. Clements (1916) referred to climax plant communities as "complex organisms", and Tansley (1935) attributed the idea of "emergent evolution" to Smuts, who, as mentioned previously, wrote about holism in the early twentieth century. According to this concept, the juxtaposition and interaction of a collection of components can give rise to a new entity that can be regarded as an independent and integral unit in its own right. In a hierarchical manner, the interaction of collections of entities at one level give rise to new entities at the next. Causal feedback loops are deemed to be the root of such emergent structures, whereby the communal association of the components creates a unique milieu which in turn affects the behavior of the individual components. Thus, when one's viewpoint is shifted to include the totality of the actions of the individual constituents, the system can in a sense be said to be the cause of its own activities. Later on, this kind of system property was described by Fuller (1969) as *synergy*, i.e., the "behavior of whole systems unpredicted by the separately observed behaviors of any of the systems' separate parts or any subassembly of the systems' parts".

The modern science of complexity is still very much in its formative stages, and researchers are grappling with the relevant lexicon. Notable among them is Rosen (1988), but there are numerous others (Edmonds 1998; Mikulecky 1995; Silvert 1995). These particular authors follow the ecologists mentioned previously, defining a *complex system* as an integrated assembly of many components that interact at a local level, resulting in global characteristics that are not predictable based on the analysis of the components in isolation. They also use the adjective *emergent* to describe these global characteristics. An example of an emergent property that has recently received much interest is *self-organized criticality* (Bak and Chen 1991). A critically self-organized system is one which, as a result of many small perturbations over an extended period, evolves to a state in which a given input may result in a small, local disturbance or a very large disturbance of system-wide consequence. There is no consistent relation between the scale of a disturbance and that of the perturbation that initiated it. However, the magnitudes of the disturbances that occur over a period of time are distributed according to an inverse power function, i.e., large disturbances happen much less frequently than small ones.

It must be noted that various authors believe that the definition of a complex system given in the paragraph above is deficient. For instance, Rosen (1988) points out that *complexity* (referring in this context to a particular characteristic) is not solely a system property per se, but is also dependent on the sophistication of the observer. Thus complexity is better described as a property of the relationship between observed (natural) and observing (formal) systems. Accordingly, Rosen suggests that complexity be measured by the class of inequivalent models (formal descriptions) that could be made of the system. Silvert (1995) agrees with this, but proposes that the definition can be made more useful if cast in terms of the amount of information that the model system can process, as indicated by a measure such as the Shannon information index (Shannon 1948). These refined definitions may be more theoretically sound than that carried forward from Clements and his contemporaries, but they are difficult to apply practically. Therefore, keeping in mind the deficiencies mentioned here, the prior definition of a complex system is used in the discussion of biosystems and biosystems engineering in this article.

2.2.3 The nature of life

Biosystems are defined in this article as composite entities that are alive at the system level. In traditional biology, it is well accepted that structures at the cellular and organism scales of organization are alive. Following the tradition of Clements (1916), many ecologists assert that biosystems of the ecosystem scale, and even of the biosphere scale, are also living entities. These claims require that the commonly accepted notion of "life" be reexamined to accommodate this broader perception. The following literature forms the background for this discussion.

Margulis and Sagan (1995) present a brief history of the ideas about life that have prevailed in Western societies, highlighting some of the major personages and events that have influenced these ideas over the centuries. This history is one of the transformation of the idea of life from a metaphysical concept to a physical one. According to these authors, animism was probably a predominant feature of belief systems from prehistory until the time of the Greek and Roman empires. Animism is the belief that not only are animals and plants alive, but that all other objects and phenomena in the world, such as rocks, rivers, and storms, also possess, or are inhabited by, spirits. Historically, there has often been accompanying dogma about the nature of the life force that animated these entities. For example, it was commonly believed that breath, fire, or some invisible fluid was the animating substance.

The animist view of the world gradually changed to one of polytheism. Examples of such belief systems are portrayed by Greek and Roman mythology, in which a limited number of spirits or gods of varying power were believed to inhabit and influence the world. The cohort of the living included these supernatural beings as well as animals and plants. However, objects such as rocks and sticks were generally not considered to be alive at all. The metaphysical realm had begun its retreat from the natural world. The rift between the metaphysical and the physical was deepened with the rise of the monotheistic Abrahamic religions. These religions went so far as to state that there was only one god, who was the source of all life, and that in the natural world only humans contained any spark of divinity.

The division between the metaphysical and the physical was carried to an extreme during the Renaissance by dualist thinkers such as Descartes. He held the then-common conviction that the soul and the flesh were distinct, (although Descartes suggested that they were connected through the pineal gland). The soul was considered to be divine, while the body was a base part of nature. The natural world possessed no inkling of active divinity or animism. The physical universe was like a gigantic and very complex clockwork that had been created by God and left to function on its own, unfolding according to a divine plan. Margulis and Sagan (1995) assert that this perspective opened the door to experimentation upon, and exploitation of, the natural world. This mechanistic view of nature was important for the development of modern scientific understanding. Influential natural philosophers such as Laplace and Newton bolstered this world view by formulating deterministic mathematical laws that successfully described many natural phenomena.

Kauffman (1993) also offers some historical insight on the development of popular Western beliefs about life. His interpretation agrees with that of Margulis and Sagan, describing how Darwin's theory of evolution pushed back the metaphysical by eroding the idea of the divine origin of humanity. If one accepted Darwin's claim of a common ancestry shared by humans and other primates, and ultimately by all creatures, then humans also became part of the clockwork mechanism of the natural world. This view was bolstered and transformed into "neodarwinism" with Gregor Mendel's discovery of discrete "atoms of heredity" and the widespread acceptance of Weismann's belief in the existence of a "germ plasm" that was the directing agency of morphogenesis. Schrödinger (1955) advocated the search for a physicochemical agent, an aperiodic crystal, to fill these roles. Francis and Crick succeeded in this quest when they identified deoxyribonucleic acid (DNA) as the principle molecular carrier of heritable information (Watson 1968). This discovery and the subsequent determination of the mechanisms of

DNA replication and protein synthesis were the crowning successes of the mechanistic interpretation of life. As a result, current dogma places life firmly in the realm of physical, and presumably understandable, phenomena.

Margulis and Sagan (1995) believe that although molecular biology is correct in that it provides valuable insight into what it means to be alive, this alone cannot provide a complete definition of life. In its extreme form, a purely mechanistic viewpoint holds that the universe is a vast machine devoid of self-awareness and self-determination, unwinding according to the laws of physics. This perspective is still rooted in the religious assumption that the universe was created and is unfolding according to some divine plan. Moreover, it is blindly reductionist, and does not address the phenomenon of life at larger scales. Margulis and Sagan argue that there must be room in the world-view for conscious decision and free will, and that there must be a more scientific explanation for the intricacies of the universe, which appear to be so improbable when interpreted from a mechanistic viewpoint. They therefore reject the extreme mechanistic view of life, advocating a broader understanding which draws upon the systems philosophy described previously.

The systems perspective of life which is advocated by Margulis is based upon her previously mentioned belief in the importance of microorganisms in the origin and continuance of life. She regards all modern organisms as having descended from bacterial ancestors (Margulis and Sagan 1986). She proposes that eukaryotic cells may have had their origins in the symbiotic merger of prokaryotes. In turn, multicellular organisms are tightly integrated cooperative colonies of the resulting eukaryotes, together with associated symbiotic populations of prokaryotes. Finally, the biosphere is a living whole composed of myriad individual organisms. This theory is one of *synergy*, a term coined by Fuller (1969) to describe entities that behave as more than the sum of their parts. Margulis also draws on the ideas of Koestler when she describes terrestrial life as a *holarchy* based on the coexistence of smaller beings, or *holons*, in larger assemblages (Margulis and Sagan 1995). This view of life is more expansive than the traditional one in which only individual animals or plants are considered to be alive. It includes biosystems of all scales, from the simplest unicellular bacteria, to the unified biotic network of the entire planet. In this aspect it recaptures, in a more sophisticated form, some of the animistic philosophy of ancient times.

As mentioned above, other twentieth century scientists have carried the renewal of holistic animism even further than Margulis. The Russian scientist Vernadsky (1863-1945) described the Earth as a single unified whole; in fact, he was responsible for popularizing the term *biosphere*, which was originally coined by Austrian geologist Seuss (Margulis and Sagan 1995). However, Vernadsky did not describe the Earth as a living being. Instead he described organisms as "living matter," and life as the greatest of all geological forces, transforming and transporting the elements of the planet's crust. Lovelock's Gaia hypothesis, described previously, also blurs the distinction between the animate and inanimate. However, his approach is the opposite of Vernadsky's, considering all of the Earth's crust, atmosphere, and biota to be a single, living, self-regulating entity (Margulis and Lovelock 1974).

A common essence uniting life at all of these scales, from the cellular to the biospheric, was identified by Schrödinger (1955) as being thermodynamic in nature. He was the first to popularize the idea of living systems as organizing themselves into a far-from-equilibrium steady state by rejecting entropy to their environment. Thus biosystems appear to act contrary to the Second Law of Thermodynamics, in that they become more orderly with time. However, a larger frame of reference reveals that the total entropy balance, including all input and output streams, is indeed positive.

Prigogine (1980) called such self-organizing phenomena "dissipative systems". Dissipative systems comprise a large class that includes not only living organisms, but also inanimate structures such as vortices and waves. However, the idea of life as a dissipative process has been further refined, for instance in Maturana and Varela's (1980) concept of *autopoeisis*. An autopoietic or "self-making" system is a causally closed network of processes in which each component serves to produce or transform other components in the network, so that the overall assemblage constantly regenerates itself. The components of a living system are continually renewed, and as described by the Second Law of Thermodynamics, the system must be open to a flow of nutrients and energy in order to maintain its constituent processes. However, the system as a whole is "organizationally" closed. No external controlling agent or imported information is required to affect its continuance. Another definitive feature of an autopoietic network is the possession of a selectively permeable boundary (e.g., a cell membrane) that is part of the self-making network and serves to contain and distinguish it from its surroundings.

The "self-making" aspect of life extends beyond the persistence and growth of the individual entity to reproduction. Life has been shaped by natural selection so as to maximize its own survival and reproduction. Those living things that exist today do so because their forebears were successful in reproducing. Their success was not necessarily due to any grand design or conscious effort, but to historical happenstance and the inherent characteristics of the successful organisms. According to Margulis and Sagan (1995), the "exuberance" of life tends to lead to crises of population and pollution, critical junctures at which life is forced to adapt, and the resolution of which has often resulted in increased overall complexity in the biosphere.

Autopoietic networks can themselves be components in larger networks, forming a hierarchical structure of emergent entities. This view of life accommodates biosystems of all scales. For instance, cells can be thought of as autopoietic metabolic networks that are cooperatively engaged in a larger organization, the organism. Odum and Odum (1955), following Clements (1916), describe how the organisms living in coral reefs together display the synergy of a single integrated living creature. At the extreme, all living things can be considered as part of a single extended network, some parts of which are more densely connected than others, making the distinction between them somewhat subjective (Margulis and Sagan 1995).

Maturana and Varela (1980) carry the idea of autopoiesis further, stating that autopoietic systems are engaged in another kind of continual change. In what is known as the "Santiago theory", they assert that living systems constantly respond to their environment in a process called *structural coupling*. Environmental stimuli provoke nonlinear responses that lead to structural changes in the system. These responses are not externally directed, but depend on the internal nature of the network. The Santiago theory is the controversial claim that life and the process of knowing are one and the same. Thus all living things are cognitive, regardless of whether or not they possess a nervous system, just as the humblest bacterium is continually perceiving and reacting to its surroundings. This view of life will be reintroduced later in the discussion about mind.

Whether or not one agrees with Maturana and Varela in equating life with mind, the definition of life as an autopoeitic network is expansive enough to accommodate biosystems of every scale, from the cellular to the biospheric, and also includes the traditional view of individual organisms as living things. Moreover, it can also accommodate systems that are combinations of biotic and abiotic components. This is not so unconventional as it first seems, since at the smallest scale, prokaryotic cells are themselves composed of nonliving organic molecules. Margulis and Sagan (1995) point out that just as the distinction between individual living things is often arbitrary, so is the division between the living and the nonliving. For example, most woody tissue is composed of "dead" substance, as is epidermal and exoskeletal material. All of these, however, are considered to be part of the living organism. Vernadsky (Margulis and Sagan 1995) and Lovelock (Margulis and Lovelock 1974) consider the same to be true at coarser scales. They propose that the entire biosphere of the Earth, including the atmosphere and lithosphere, could be considered as a single living entity. For example, in Lovelock's Gaia theory, he asserts that the community of organisms on Earth has coevolved in conjunction with the oceans, atmosphere, and crust of the planet to the point where the entire ensemble can be thought of as a single autopoietic network (Margulis and Sagan 1995). According to this hypothesis, planetary nutrient and energy cycles involving biological, meteorological, and geological processes form a complex, persistent, self-regulated pattern of organization.

Finally, Margulis and Sagan (1995) propose that technology can also be considered as an extension of the global organism, and speculate that symbioses between organic creatures and technological artifacts may one day result in even more complex forms of life. Just as the emergence of eukaryotic life incorporated but did not displace prokaryotes, and multicellular creatures incorporated but did not displace unicellular organisms, so cyborged entities of various scales will incorporate, but not displace, biological creatures. Humans, or their evolutionary successors, may constitute essential elements of these more complex living entities.

Thus the definition of life described by Margulis and Sagan (1995), although not exact and rigorous, suits the theme of the EcoCyborg Project. The definition comfortably accommodates biosystems of a variety of scales, and includes hybridized biological and technological entities such as ecocyborgs. Since the aim of the EcoCyborg Project is to learn about the engineering of such systems, the next section is a review of literature that deals with engineering in this context.

2.3 Engineering biosystems

Engineering in general includes the design, construction, operation, maintenance, repair, and upgrading of a system, usually in order to achieve a particular goal in the face of a set of constraints. Historically, the engineering of large-scale biosystems stems from agricultural, silvicultural, and aquacultural activities, which have been practiced in various forms since prehistoric times. Indeed, proficiency in these kinds of practices was probably a primary reason for the current success of the human species. These practices can all be considered as instances of the more general area of biosystems engineering, in which biosystems, as previously defined, are considered to be adaptive, complex, dynamic systems that are alive to some degree. This broader perspective of biosystems engineering allows one to envision other applications of the discipline, some of which may even be purely hypothetical at this time. These might range from medical applications, to the design of androids, to altering the surfaces of other planets so that they are suitable for terrestrial life. In keeping with the theme of this article, the following review will concentrate on biosystems of the ecosystem scale. Instances of biosystems engineering at this scale fall into three general categories: the management and repair of natural systems; the modification of natural systems; and the creation of artificial systems. The literature reviewed below is that most relevant to each of these categories.

2.3.1 Management and repair of natural biosystems

Historically, the management and repair of natural biosystems was not an issue in most cultures. Wilderness in sparsely settled areas was generally thought to be so vast as to be inexhaustible, and even in populated regions of the world, little value was placed on unaltered wild biosystems. Only in the past century, with exponential increases in population and resource exploitation, together with the rise of ecological awareness, has the understanding and conservation of wilderness biosystems been considered important in Western cultures. This new emphasis on the conservation of wilderness has made necessary the development of effective methods of managing and repairing natural biosystems.

Responsibility for the management and repair of large-scale biosystems has fallen largely to government. This is because natural biosystems are generally very geographically extensive, and their management often concerns large and diverse groups of people. Moreover, public lands are often included in these areas. As a result, the most extensive areas of literature relevant to the management and repair of natural systems are probably those relating to governmental policy for wildlife and natural resource management. This includes the studies and recommendations of advisory bodies, documents defining policy and strategic initiatives, as well as evaluations of the implementation and impact of policy. Examples of government activity in the sphere of natural biosystem management and repair include the regulation of recreational hunting and fishing; the establishment and administration of public park systems; the regulation of primary industries such as mining, petroleum extraction, forestry, and commercial fishing; the establishment and enforcement of guidelines for reclamation of disturbed areas after the abandonment or exhaustion of a particular resource; and the monitoring and regulation of the environmental impacts of other industries, (e.g., atmospheric and watershed pollution).

The effective management of any large-scale biosystem requires information about its current state and about changes in that state. It also requires a sound understanding of how the system functions and responds to a given intervention. Obtaining this information and understanding is equivalent to the characterization and comparison of the static and dynamic aspects of biosystems. Since the final section of this article treats the literature written about these subjects in particular, they will not be discussed in detail here. It is sufficient to say that efforts in this area have been limited until recently by the difficulty and expense of working at the large scales of most natural biosystems. In the past few decades, however, new technology, such as remote sensing and very large database management, have made such activities more practical. Currently, a number of cooperative international programs are aimed at the observation and understanding of the biosphere and its constituent biogeographical regions.

Humanity's capacity to observe, and therefore better understand, the environment at this scale is thus improving rapidly. Unfortunately, there is still a dearth of basic knowledge about how natural biosystems function. Moreover, the romance of the aforementioned high technology, such as satellite observation, tends to obscure the value of basic ecological research on the ground. Such work continues to be done, but the vastness and complexity of the systems of interest are immense, making their observation and understanding a difficult, expensive, and time-consuming task. Furthermore, in nearly every area of the world, ecologists find themselves in a race against industrial exploitation, trying to learn as much as possible about a particular natural system before it is polluted, destroyed in the process of resource extraction, or altered for agricultural or recreational purposes (Wilson 1992).

Finally, given information about the systems in question and an understanding of their workings, strategies must be formulated and implemented for their management or repair. With some exceptions, this too is left to government, which usually operates through the legislated regulation of private sector activities, or through financing initiatives such as the Superfund (i.e., the U.S. Comprehensive Environmental Response, Compensation, and Liability Act of 1980) (Reisch 1983). The efficacy of such management efforts is often compromised by the difficulty of monitoring and enforcement, as well as political factors such as lobbying by diverse interest groups, and the lack of long-term continuity of policy due to changes of government.

Politics are further compounded when we deal with biosystems that transcend national boundaries. Little in the way of actual management or restoration has been done at these scales, although there have been some attempts to tie environmental issues to monetary loans given to developing nations, for instance. Some international accords have also resulted in partial success stories. For example, nations that signed the 1987 Montreal Protocol on Substances that Deplete the Ozone Layer agreed in principle, and to a large extent have been successful in their efforts, to stop the production of chlorofluorocarbons (United Nations 1995). This family of chemicals is believed to be active in the upper reaches of the atmosphere, causing polar "holes" in the ozone layer. The limited overall effort at management of large-scale natural systems probably is due largely to public perception; environmental issues are still not deemed to be urgent or important enough to warrant more effort in this area.

As previously discussed, most large-scale environmental initiatives take place in the sphere of government because of the complicated public issues involved and the extensive resources required. Some large international conservation groups such as Greenpeace and the Sierra Club have been organized, but their activities are limited primarily to lobbying governments in an attempt to influence legislation. However, there are some privately funded, grass-roots management and restoration programs underway in various parts of the world. For example, in the United Kingdom, the Trees for Life group has established as its long-term objective the reconstruction of a large tract of the Caledonian Forest (McPhillimy 1997). This forest covered much of Scotland in prehistoric times, but, like most of the primal European forest, only the tiny remnants remain that are either difficult to access or are otherwise unsuitable for industrial exploitation.

The activities and literature discussed in this section are oriented toward one aspect of biosystems engineering: the management and repair of natural biosystems in their wild state. As mentioned, the advent of such activities is fairly recent. Far more ancient and widespread is the practice of altering natural biosystems in order to achieve particular goals. This is the topic of the next section.

2.3.2 Modification of natural biosystems

The modification of natural biosystems is a consequence of most intensive human activities. If it is intended to achieve a particular goal, the alteration of a biosystem may be considered an instance of biosystems engineering. The industries of agriculture, aquaculture, and silviculture are general classes of these kinds of activities. These industries have very ancient histories, originating from the first attempts to alter the environment so as to produce more of the biomass necessary for human survival. For example, agriculture has developed from the hunting and gathering of animals and plants to the actual planned management of biosystems. Moreover, like all industries, agriculture is still being continually transformed by the development of new technologies. Aquaculture and silviculture have followed the same trend, but the management of aquatic and forest biosystems historically has not been as sophisticated as agriculture, perhaps because the scale and abundance of these resources have not made the need for advanced management very pressing until recently. This situation is now changing, as wild resources dwindle ever more rapidly and more powerful technologies become available.

Aside from agricultural, silvicultural, and aquacultural applications, natural systems are also being modified for less traditional industrial purposes. For example, wetlands are being used to treat the effluent from municipalities and industrial installations before it is discharged into the environment. Large-scale biosystems are also being transformed to serve recreational purposes. Lakes are routinely stabilized to improve boating and swimming, and to minimize damage to waterfront developments due to fluctuations in water levels. Golf courses are constructed in all manner of settings, and mountainsides are cleared for ski slopes.

There is a great deal of literature associated with each of the industries mentioned above, especially with agriculture. Included in this literature, and particularly relevant to the EcoCyborg Project, is work on the application of systems analysis to these kinds of biosystems (Liao 1991; Spedding 1988). Where development is taking place in previously wild areas, the documentation generated is similar to that concerning the repair and management of natural biosystems: i.e., impact studies and recommendations to government; policy and strategic documents; and implementation and assessment reports.

The transition from the maintenance and repair of unaltered natural biosystems to their modification can be extended to a third area of biosystems engineering: the creation of entirely new systems. Again, because of the extent and complexity of these systems, and a lack of knowledge about how they function, the creation of biosystems at the ecosystem scale or larger is, with a few exceptions, a relatively new practice. This is the area of biosystems engineering that most directly concerns the current case studies in the EcoCyborg Project.

2.3.3 Creation of artificial biosystems

There is usually a particular goal behind the construction of an artificial biosystem. Artificial biosystems of the ecosystem scale can be created for a variety of purposes, including education, research, industry, and, most recently, for the support of life in space. There is a great deal of literature associated with each of these pursuits. Although the ideas described in this article are extensible to all of these areas, the current case studies in the EcoCyborg are envisioned as being housed in a space station. The engineering of space life support systems is therefore most immediately relevant to the project, and will be emphasized in this review.

Artificial biosystems are often intended for use in recreation, education, conservation, or research. Examples include household aquariums and terrariums, botanical gardens, conservatories, zoos, game parks, and aviaries. For the most part, these are meant to preserve collections of one or a few organisms of each of a number of different species. With some exceptions, they have not generally been designed to be well integrated at the system level, nor are they sustainable without an input of new organisms. Conditions are often inappropriate for the reproduction of the specimen organisms, and the captive populations are usually not large enough to be viable. This is changing, however, as more is learned about the organisms, the wild systems of which they are naturally a part, and the construction of artificial ecosystems that emulate natural conditions.

The purpose of most industrial biosystems is to produce a particular kind of biomass, whether it be food, pharmaceuticals, or flowers. Examples include: fermentation facilities in which the diverse metabolic abilities of bacteria and fungi are harnessed for the production of compounds such as alcohol and acids; facilities for the microbial production of pharmaceuticals; and greenhouses and phytotrons for the growth of a diversity of plants such as geraniums, tissue-cultured pineapples, and pine seedlings for reforestation programs. Some of these are batch systems, in which a single population of organisms is brought to a certain point in their development, harvested, and the system is then "reset" to its initial conditions. Many, however, are continuous-flow systems that are expected to transform a steady inflow of nutrients into an uninterrupted outflow of product, ready for sale or further processing. The latter kind of system has more in common with the autopoietic networks that were previously defined as being alive. Many disciplines contribute to the knowledge base required for the creation of these kinds of systems, including industrial microbiology and horticulture, as well as biosystems, agricultural, chemical, food, and process engineering.

The most recent and romantic motive for the creation of artificial biosystems is to support the exploration and colonization of space. This research has been driven primarily by the space programs of the United States and the former Soviet Union, Russia being the primary inheritor of the latter. Work is also being carried out in Japan and Europe (Nitta et al. 1990; Tamponnet 1992; Tsiolovsky 1960). The objective of this research is to create ecological systems that provide a comfortable, sustainable environment for humans in extraterrestrial colonies or during space flight. The systems provide all of the services supplied on Earth by the microbially based planetary biosphere, and should be as materially closed as possible. The required services include the recycling of wastes (solid, liquid, and gaseous) into usable forms (food, drinking water, and breathable air), and the maintenance environmental factors such as temperature, pressure, and humidity within a range comfortable for humans.

Russian scientists were among the first to think seriously about the idea of creating closed ecological systems in space. In the late nineteenth century, while writing about the principles of space flight, Tsiolkovsky discussed the need for "space greenhouses" (Tsiolovsky 1960). In the early nineteen-sixties, scientists working in the Soviet space program were inspired by Tsiolkovsky's work, as well as Vernadsky's previously mentioned writings about the biosphere (Margulis and Sagan 1995). Gitelson, Shepelev, and Meleshka, of the Moscow Institute of Biomedical Problems, experimented with a small sealed chamber called the *Siren*. A human occupant and tanks of green algae lived in balance in the *Siren* for periods of several days. From the mid-1960s on, a succession of larger facilities, *Bios*-1, *Bios*-2, and *Bios*-3, were constructed. The Bios chambers included higher plants as well as algae, and were much more comfortable for the human occupants than was the cramped *Siren*. *Bios*-3 was capable of sustaining a human crew of three for up to 6 months (Gitelson et al. 1989). The air was completely recycled, as was 95% of the water and 50% of the required food. However, little of the waste was recycled, and its removal resulted in a depletion of trace elements. There were also problems with the build-up of trace organic gases, which perhaps was due to a paucity of microbes in the soil-less hydroponics system. Nevertheless, the Bios program was a ground-breaking investigation of the practicality of closed ecological systems.

In the United States, major contributions to the understanding of closed ecological systems were made by Folsome, a microbiologist at the University of Hawaii (Folsome and Hanson 1986). Beginning in 1968, he began experimenting with randomly collected marine microbial communities in small sealed flasks. Given a high enough initial diversity, he discovered that these closed communities would reach a viable steady state, each stabilizing with a unique balance of organisms and atmospheric gases. Some have remained viable for over twenty years, and may persist indefinitely. His work was corroborated by Maguire, at the University of Texas, and Hanson, at the California Institute of Technology, who had both independently established similar "ecosystems in a bottle" (Folsome and Hanson 1986).

In the 1970s, the National Aeronautics and Space Administration (NASA) of the United States also began to experiment with closed ecological systems as a means of sustaining humans on extended space missions. They encountered problems similar to those that the Soviet researchers had experienced previously. This line of research led in 1977 to the initiation of NASA's Advanced Life Support (ALS) program (Volk 1996; Allen 1991), an umbrella program that currently funds research related to the development of Controlled Ecological Life Support Systems (CELSS), or Bioregenerative Life Support Systems (BLSS). ALS research is being conducted at several NASA installations, principally the Johnson Space Center, the Ames Research Center, and the Kennedy Space Center. A number of university-based organizations have also been affiliated with the ALS program, notably the NASA Specialized Center of Research and Training (NSCORT) at Purdue, which was active from 1990 until 1995 (Mitchell 1994), and the New Jersey NSCORT at Rutgers University and the Stevens Institute of Technology, which was established in 1996 (Ting et al. 1997).

Privately funded groups have also been involved in researching closed ecological systems. In 1984 Space Biosphere Ventures (SBV) was initiated (Nelson et al. 1994). The purpose of this private company was to better understand the terrestrial biosystem by engineering closed ecological systems, and to learn how to construct biosystems for space exploration and settlement. SBV constructed a large experimental facility called *Biosphere* 2. Located in the Arizona desert, the entire complex covered 3.15 acres (13,000 m^2), and was designed to operate as a completely mass-closed system (less than 10% air exchange per year) open only to energy and information exchange. Its technological systems were intended to work in conjunction with a number of artificially assembled ecological communities, which together contained over 3000 species of organisms, including human occupants. These communities were patterned after rain forest, ocean, marsh, savannah, thorn scrub, and desert biomes. There was also an intensive agriculture module and living quarters for human occupants.

The first major human enclosure trial of *Biosphere* 2 extended from 1991 to 1993, when eight people were sealed inside (Cohen and Tilman 1996). The second was in 1994 for 6.5 months with seven people. The researchers encountered unexpected difficulties, including imbalances in atmospheric gases, nutrient contamination of the water supply, the loss of many critical species, and domination of the ecology by others. Because of questionable management and public relations practices, and an apparent misunderstanding of the objectives of the program on the part of the public and the media, the *Biosphere* 2 project suffered a great deal of criticism and a loss of credibility as center for scientific research (Kaiser 1994; Macilwain 1996; Beardsley 1995; Kaiser 1996). However, many recognized the value of the project and the importance of the *Biosphere* 2 facility as a research tool, and in 1996 it was brought under the management of the Earth Institute of Columbia University (Dempster et al. 1997; Cohen and Tilman 1996; Odum 1996). The facility is currently being renovated to serve as an educational and research facility, but with a shift in focus. Ecological experiments will continue, but without long-term human occupants.

In the instances of biosystems engineering described above, there are examples in which the engineered systems consist completely of natural biotic and abiotic components. There are also examples of systems that include not only natural components, but technological components as well. This is true in the case of the modified natural systems such as farms, and artificial biosystems such as space habitats. In all cases, however, the biosystems in question include both "living" and "nonliving" components. As mentioned previously, authors such as Vernadsky have considered this distinction to be an arbitrary and artificial one. In fact, if one considers the system to be alive in its own right at a higher "holarchical" level (to use Koestler's terminology) this renders moot the question of whether the individual "holons" are themselves alive. As many authors propound, it is the network of relationships between the components that defines the biosystem, not the nature of the components themselves (Logofet 1993). As mentioned previously in the discussion of complexity, this perspective brings the study of ecological systems together with the study of systems that are not necessarily composed of biological organisms, but which nevertheless possess similarly structured networks of interrelationships. These include economies and societies, for instance, as well as computer-based systems.

The field of *artificial life* involves the study of systems that demonstrate the presumed formal qualities of living biological systems, but which are based on substrates which need not be biological themselves (Langton et al. 1991). Electronic, computer-based systems are of primary interest, encompassing, for instance, genetic algorithms, cellular automata, and computational ecologies. However, the field also extends to robotics and biochemistry. These systems are considered to emulate various characteristics of living systems, or in fact to be alive in their own right. If one adopts the latter, "hard" approach to artificial life, then these systems would be considered biosystems per the definition adopted in this article. Whether or not this is deemed to be true, many researchers believe that their study has important ramifications for the understanding of economic, sociological, and natural biological systems (Hoffmeyer 1997; McGlade 1993). Among the numerous popular books now available about artificial life is an overview written by Levy (1992). In terms of scientific publications, a series of important conferences were hosted at the Santa Fe Institute, resulting in several volumes of proceedings (Langton 1994; Langton et al. 1991; Langton 1989). A number of refereed journals are also being published that relate to this field.

Another factor that varies between the systems described previously is the kind of management, or control strategy, that is associated with each (Kok and Lacroix 1993). Unmodified natural biosystems have evolved over millennia so that those extant today are functionally integrated networks that are viable over long periods of time. Their viability is partly due to inherent homeostatic mechanisms, or *intrinsic control* mechanisms. At the other extreme, a highly modified ecosystem, such as a mechanized farm, requires extensive management, or *extrinsic control*, on the part of its human operators in order to remain productive. (Of course, this interpretation depends on whether or not one chooses to include the human operators within the system boundaries.) In many modified or artificial biosystems, such as greenhouses, extrinsic control is to some degree managed by technological control systems. This kind of hybridization is the approach proposed in the EcoCyborg Project as a means of increasing the autonomy of biosystems. The sum of all of control mechanisms, both intrinsic and extrinsic, forms the basis of the *mind* of the biosystem in question. In order to provide some background information on this topic, the next section is a brief review of some literature related to the theory of mind and its implementation in artificial entities.

2.4 Engineering mind in biosystems

2.4.1 Characterization of mind

The autonomy of artificial biosystems implies that they possess some degree of mental capacity, and this makes relevant the literature of the cognitive sciences. These include artificial intelligence, psychology, and neurophysiology. The former is defined by Chalmers (1994a) as the production of (usually computational) models that cohere to some extent with human behavioral or neurological data. Since the thrust of the field of artificial intelligence is based on the attempt to emulate human thought, it is engaged in a constant comparative dialogue with the more traditional branches of cognitive science: psychology, which is the study of human behavior, and neurophysiology, the study of how the brain supports cognition. Of course, it would be impossible within the limits of this article to give a comprehensive review of all of the literature associated with these fields. However, beginning with psychology, a few works will be mentioned that were particularly influential in the development of the ideas presented in this article.

Howard Gardner (1993), a prominent educational psychologist, is the proponent of a modular theory of intelligence. This view of human psychology runs contrary to the characterization of the mind as a single, integrated intelligence. Instead, the mind is described as possessing a variety of mental faculties, or collections of abilities specialized in the performance of particular kinds of mental tasks. The mind of an intelligent entity, be it human or not, can therefore be characterized by a unique profile indicating its proficiency in performing each kind of task. Pinker (1994) discusses the modular description of mind from the perspective of a linguist in the tradition of Chomsky. This modularity forms the basis for the proposed structure of the computational component of cyborged biosystems.

Ideas about the neurological basis of the mind are explored by Susan Greenfield (1995). She gives an overview of the current state of neurophysiology as it relates to the question of how conscious intelligence arises from the brain. The central thesis of her theory is that mental activities occur as "gestalts", or groups of neurons that are coordinated in their activity. The strength and inclusiveness of these neuronal gestalts change continuously, corresponding to the transient prominence of ideas or mental events in the consciousness. This theory is based on the perception of the human brain as a self-configuring, massively parallel, sparsely connected network of computational agents (neurons). The view of the mind as a network structure is an important part of the approach taken in this article.

An aspect of cognition that is of particular relevance to this article is consciousness, since in the adopted framework, consciousness forms the foundation for the desired characteristic of autonomy. Chalmers (1994b), as well as Hofstadter and Dennett (1982), have written on the nature of consciousness, the latter two having also edited popular collections of essays about issues central to the cognitive sciences. These authors have reviewed and contributed to the ongoing debates about the philosophy of mind, consciousness, and artificial intelligence. They concern themselves primarily with metaphysical or philosophical arguments that surround the study of mind and consciousness, but are also prominent in the debate over the feasibility of attempting to create artificial intelligence.

2.4.2 Artificial intelligence

The EcoCyborg Project, which is the context of this article, is concerned with the engineering of autonomous biosystems. This work involves the engineering of mind in biosystems; therefore the field of artificial intelligence (AI) has a strong bearing on this project. Chalmers (1994a) has defined artificial intelligence as being based on the production (and implementation) of computational models of the mind. Computational models of the mind were popularized in the 1940s by the originators of cybernetics. These were the participants of the Macy Conferences, including Wiener, Bateson, Shannon, and von Neumann, the originator of the digital electronic computer (Capra 1996). This group interpreted human thought from the perspective of communications and control theory. Subsequently, important contributions to the computational modeling of human mental processes were made by the likes of Marr (1982). Although Marr unfortunately died at a young age, he performed pioneering work in the field of machine vision. He broke down his approach into three steps: the formulation of a computational theory of the mental process in question; the algorithmic representation of the theory; and the hardware implementation of the algorithm.

The field of artificial intelligence has, since its inception, been divided roughly into two philosophical camps: "hard" AI and "soft" AI. Proponents of hard AI believe that it will one day be possible to construct machines that think as a human does. Researchers of soft AI stop short of claiming that their creations will be capable of thought per se, but will merely emulate certain aspects of it. Hofstadter and Dennett (1982), mentioned previously, have contributed to and edited excellent collections about the philosophical arguments associated with both of these perspectives. Many popular overviews of the field have been authored from an optimistic standpoint (Waldrop 1987). From a more conservative perspective, Schank (1984) might be considered a cautious proponent of hard AI. He argues that machines will one day be constructed that can think as humans do, but that the practical difficulties that must be overcome will delay their development until well into the future. Penrose (1989), on the other hand, refutes even this reserved perspective, claiming that thought is a unique activity of which only the human brain is capable. He asserts that machines will never be anything more than glorified calculators which cannot be compared with the human mind.

2.4.3 Mind in biosystems

The idea of actually engineering a biosystem of the ecosystem scale so that it possesses a mind of a certain character is a relatively recent idea. However, the concept that biosystems of the ecosystem scale or larger might possess or somehow develop mental capabilities is not new. In fact, this is a concept inherent in animist theologies. Theology aside, scientists of the nineteenth and twentieth centuries conceived of the planetary biosphere as possessing, or having the potential to develop, a global mind. For example, Teilhard de Chardin, a Jesuit priest, paleontologist, and philosopher, referred to this global mind as the *nöosphere* (Margulis and Sagan 1995). More recently, Maturana and Varela (1980), in their Santiago theory, as well as Bateman (Capra 1996), have proposed that cognition and life are one and the same phenomenon. By extension, if one adheres to the idea that the biosphere is alive in its own right, it must then by definition also be considered cognitive, or possessed of a mind.

A cybernetic interpretation of ecosystem theory leads to a similar conclusion. Wiener and his associates (which, as mentioned previously, included Bateman) developed cybernetics as a perspective from which to analyze computational processes, with reference to human thought (von Neumann 1963). Not long after this, Patten (1959) applied the cybernetic perspective specifically to ecosystem theory. Its applicability to both underlines the parallels between the human mind and ecological systems. Each can be interpreted as a communications-based, control-oriented structure, or information-processing system. If one adheres to the idea that thought is a computational phenomenon, then the cybernetic properties of ecosystems are equivalent to mind. The cybernetics of ecosystems continues to have proponents in the field of ecology (Haug 1983; McNaughton and Coughenour 1981; Patten and Odum 1981).

Others have discussed the role of anthropogenic technology in endowing large biosystems with mental abilities, or enhancing those that already exist. In the nineteenth century Butler discussed the evolution of machines, and the possibility of world dominance by a civilization of machines (Dyson 1997). Fuller (1969), in the tradition of Butler, suggested that humans will be replaced by machines in roles of specialization, while continuing to contribute the capacity for forward-thinking, integrative, generalized intelligence. Margulis and Sagan (1995) also write of technology as another step in the development of the biosphere, as the emergence of a new level of complexity in living systems, based on the integration of technological and biological components. Dyson (1997) elaborates the idea of hybridized biological and technological life of planetary scale, describing how world-encompassing technological systems are evolving as the infrastructure of a single global intelligence, the equivalent of a planetary ecocyborg.

At the same time that the aforementioned authors have been writing of their grand visions of a cognitive planet, the creation of intelligent biosystems has been proceeding behind the scenes, as it were. Engineers are interested in building useful systems that operate with limited supervision, and these include highly autonomous biosystems. The advent of electronics has spurred the rapid advancement of control technology, and highly automated industrial installations are often created that include biological components. Examples are greenhouses, phytotrons, fermentation facilities, and pharmaceutical bioreactors. The electronic control networks of these systems often possess considerable computational capabilities, resulting in a significant degree of automatic behavior. Artificial intelligence software has made these systems increasingly capable of dealing with variable environments, thereby increasing their autonomy. Thus, hearkening back to cybernetics, the engineering of mind can be seen as an extension of process control technology.

Of particular relevance to the EcoCyborg Project, NASA is developing intelligent software for the control of autonomous systems (Williams and Nayak 1996). The kinds of systems that the developers of this software have in mind are large, complicated, and essentially immobile. Their combination with intelligent software would result in immobile robots, or *immobots*. These are autonomous systems with minds (or regulatory components) that are primarily occupied with the robust control of internal functions. The autonomy of this kind of system is based on its capacity to model itself, which is similar to the approach taken in the EcoCyborg Project (Kok and Lacroix 1991). Williams and Nayak (1996) mention the possible use of these autonomous software agents as the mental component of the controlled ecological life support systems that NASA is developing. The resulting systems would resemble the ecocyborg currently envisioned as the case study of the EcoCyborg Project.

2.4.4 Evolution of the EcoCyborg Project

The evolution of the EcoCyborg Project began with the work of Kok and Desmarais (1985), who considered how a technological system might be implemented in a greenhouse so as to endow it with intelligence. Based on an analysis of the structure of hierarchical control systems in biological entities, a four-level hierarchy of computer-based controllers was envisioned. The control levels were designated as *physical*, *instinctual*, *Pavlovian*, and *intelligent*. Each of these levels was more flexible but less rapid in its actions than the previous one. An example of a physical level control is a propane pressure regulator, in which the functions of sensor, controller, and final control element are integrated in a single device, the actions of which result from its inherent physical structure. In instinctual devices, the functions are performed by distinct elements, but the control loop is straightforward and nonprogrammable. For instance, a safety loop closes the propane supply valve to a furnace if the pilot light is extinguished. Pavlovian devices are programmable to some extent. For instance, the setpoints of a proportional-integral-derivative (PID) controller can be changed by another control device. Finally, intelligent controllers were envisioned as being capable of abstract activities such as memory, reasoning, decision-making, and communication. They would generate and maintain models of various aspects of the greenhouse system and its environment for use in predictive control and in adaptively reprogramming themselves for optimal performance. The vision of an intelligent "ecological cyborg", or "eco-cyborg", was also conceived at this time. This was envisioned as a composite entity consisting of biological and technological components that might take the form of an intelligent greenhouse, for example (Kok and Desmarais 1985).

To develop their ideas, Kok and Desmarais (1988, 1987, 1985) outfitted a physical greenhouse with a system that included sensors, effectors, and controllers resident on digital computers. The implementation of the Pavlovian and intelligent levels of the control hierarchy proved to be more difficult than originally hoped. This was reflective of the widespread realization on the part of artificial intelligence researchers in the mid-1980s that the emulation of human intelligence was not a trivial matter of merely building faster computers and compiling larger rule sets. The research community began to take stock of what had been accomplished, and to revise their expectations in light of a better understanding of the challenges they faced. Kok and his group did the same, articulating the need for a better understanding of the concept of intelligent control, for the development of an appropriate symbolic language with which it could be described, and for the determination of a detailed conceptual design before its physical implementation (Kok et al. 1986). Work on an automated greenhouse control system continued, but the development of higher-level control devices was primarily limited to conceptualization and definition (Kok and Desmarais 1988, 1987).

Kok and Gauthier continued to develop the idea of an autonomous biosystem, but in the larger context of an entire farm. Gauthier (1987), as part of his doctoral work, developed and tested a "prototype integrated program package" for farm production management, which was subsequently used for several years on the research farm of the Macdonald Campus of McGill University. It consisted primarily of a database management system for recording histories and characteristics of the farming operation. The package relied entirely on human operators for higher-level control functions such as decision-making, and for formulating and effecting all control strategies. However, in his thesis, Gauthier did briefly discuss the development of more advanced mechanisms for decision support and control activities. He listed some functional requirements for a farm operating/management system in order of increasing intelligence levels, from data entry through to imagination and creation, (e.g., design of equipment, processes, models, etc.), and gave a brief discussion of how a cognitive control system might be structured.

The idea of an intelligent farm management system was given more detailed treatment in later publications. Kok and Gauthier (1989) discussed the general design considerations and construction requirements for "integrated farm control software" (IFCS). The IFCS was envisioned as offering decision support as well as assuming some autonomy in making and implementing farm management decisions. However, it was acknowledged that the software available at that time was inadequate for the actual creation of such a system.

At the same time, Kok and Gauthier (1988, 1987) developed in more detail the concept of the ecocyborg: a large, integrated system of many biological, technological, physical, and virtual components. They emphasized that this kind of system might be considered a living entity in its own right, capable of intelligent, conscious, independent behavior, and interaction in a social context. These entities were proposed as possible terrestrial production units or space stations. The fusion of artificial intelligence and ecology was also suggested as a general approach to the engineering of more sustainable, ecologically sound agricultural practices. It was speculated that artificial intelligences could have biases different from those of humans, and that they might serve to balance the often short-sighted, anthropocentric inclinations of humans, that frequently result in unsustainable management practices. The design of an experimental ecocyborg was outlined, the principal physical components of which included underground chambers connected by tunnels; a control network similar to those described above; and biological components, including cultures of plants, insects, and fish, as well as occasional humans (Kok and Gauthier 1988). The physical infrastructure of a prototype system was largely completed on the Macdonald Campus of McGill University, but was never made fully functional. Work on this facility was suspended in 1989 due to lack of funding.

After suspension of work on the physical ecocyborg, Kok and Lacroix turned to computer modeling and simulation to further develop the idea of an artificially intelligent agricultural system. Lacroix (1994) defined the main objectives of his doctoral work as the development of tools to aid in the design of enclosed agroecosystems, and the use of these tools to create a prototype simulation-based control system. Three tools were developed: a conceptual framework for the creation of such systems (Kok and Lacroix 1993; Lacroix and Kok 1994, 1991a, 1999), a simulation approach utilizing the environment of a multitasking operating system (Lacroix et al. 1996; Lacroix and Kok 1991b), and a virtual greenhouse system. The latter included a control system which, in step with advancing artificial intelligence techniques and computer control technologies, came closer to the originally conceived goal of intelligent control, as opposed to that of mere data management or decision support. A prototype simulation-based controller was produced that attempted to minimize the heating load for the greenhouse in light of forecast meteorological conditions. This work made use of "primary consciousness", which took the form of predictive modeling abilities implemented using neural networks (Lacroix et al. 1996; Kok et al. 1994; Kok and Lacroix 1993; Lacroix et al. 1993; Kok et al. 1991; Lacroix and Kok 1991a).

2.4.5 Current EcoCyborg Project

A brief summary of the current EcoCyborg Project is presented here to set the context for this section of the article. The project carries forward in the spirit of the work of Kok, Desmarais, Gauthier, and Lacroix, as described previously. The long-term objective is the creation of large cyborged ecological systems that are independent, intelligent entities in their own right. For practical and financial reasons, work is not currently proceeding in the physical realm. Instead, the project is following a trend prevalent in much of the academic and industrial world; computer technology is being used to create models and simulations of the system of interest. This system is currently envisioned as a hypothetical ecosystem and control network enclosed in an orbital space platform. The exercise is proceeding in three parallel streams: the development of the ecocyborg control structure and the simulation framework that will be used to implement the models; the creation of the models of the ecosystem and forcing functions; and the development of techniques for the characterization and comparison of different configurations of the system. A review of the literature recently published and presented by members of the EcoCyborg Project Group is given here.

General overviews of the EcoCyborg Project have been published that include descriptions of the philosophy and underlying approach (Parrott et al. 1996; Kok et al. 1995). More detailed descriptions of some components of the modeling and simulation software have also been written. Molenaar et al. (1995) have described the structure of the overall simulation framework created as part of his doctoral research, which was based on the previous work by Lacroix and Kok. The conceptual design of the ecosystem model has been described by Parrott, and will be implemented during her doctoral work (Parrott 1995; Parrott and Kok 1995). Clark et al. (1997, 1995) have described the software used to generate the terrain in the ecosystem model. Parrott et al. (1995) have outlined the generation of temperature values, one of the dynamic forcing functions that will drive the system. The control network of the EcoCyborg has been discussed in a number of articles that extend the previously described work on intelligent control systems. These articles include discussions about the conceptual perspectives from which cognition can be considered, the associated lexicon, and the general approach to the implementation of cognition in engineered biosystems (Clark et al. 1997, 1996). In a more specific article about the implementation of the intelligent control components, Molenaar and Kok (1995) have discussed a Pavlovian control mechanism that is being developed.

As mentioned, one aspect of the EcoCyborg Project is to compile methodologies for the characterization and comparison of complex biosystems, especially those of the ecosystem scale. The section that follows is not intended to give a detailed description of the methodology and techniques available for the treatment of this problem. This will be discussed in subsequent articles. Instead, an overview is given of some salient research programs that are currently underway, and the general areas of literature that are associated with them.

2.5 Characterization of biosystems

As described above, the widespread adoption of an integrated systems perspective of ecology is relatively recent. Most early work in ecology focused on isolated aspects of ecological systems, such as the dynamics of one or a few populations, or the cycling of a particular nutrient. However, one important theme of this review is the evolution of a holistic understanding of ecosystems. The advancement of this understanding is contingent on, or perhaps equivalent to, the ability to characterize and compare large-scale biosystems as wholes. Research programs have been established that are based on this new holistic perspective, and which therefore generate literature about the problem of characterizing biosystems in their entirety. Moreover, from the standpoint of artificial intelligence, the creation of highly autonomous systems requires that the systems possess the capacity to observe, model, and reason about themselves. The ability to characterize and compare biosystems, or different states of the same biosystem, is a necessary ability of systems like those being researched in the EcoCyborg Project.

It was mentioned previously that difficulties inherent in the observation and understanding of large-scale natural biosystems have historically caused these activities to remain largely in the public sphere. Even given the resources of governments, the difficulty and expense of such activities has imposed practical limits on their extent. Monitoring has usually been restricted to specific aspects of environmental quality (e.g., acid rain, pesticide contamination), to limited geographical regions, and to issues of industrial or military importance. Until recently, comprehensive studies have usually been possible only in the case of relatively small systems or parts of larger systems, while holistic studies of large systems have generally been limited in their scope.

An early example of the study of a large-scale system is the work on coral reef communities that so influenced the ideas of Odum (Odum and Odum 1955). This relatively comprehensive study was done on only a small part of a reef at Eniwetok Atoll, in anticipation of monitoring the effects on coral reef ecosystems of nuclear testing in that area of the Pacific. The results of the study seemed to confirm that coral reefs and, by extension, other ecosystems, displayed system characteristics that were different from what could be expected from the study of their components in isolation. This study had a great influence on the research community, and led to many studies of similar kind. Odum (1977) makes reference to some other early studies of similar scale.

The first notable effort to gain an understanding of entire biogeographic regions was the International Biological Program (Van Dyne 1995; Loucks 1986). Many nations initiated a variety of projects under the auspices of this program in an effort to better understand planetary ecology. The United States' contribution to the International Biological Program (US/IBP), initiated in 1964, had lofty objectives regarding the development of a comprehensive understanding of global ecological systems and their management, (i.e., predictive modeling capacity). The results did not fully meet these ambitious initial expectations, proving the difficulty of developing and maintaining an understanding and awareness of biosystems of this scale (Odum 1977). However, the IBP established a precedent in that it led to the adoption of an "ecosystem processes paradigm" for large-scale ecological research, and many subsequent programs have adopted similar approaches (Van Dyne 1996). The program could be interpreted as humanity's first attempt to directly perceive the biosphere at a planetary scale. Blair, before the U.S. House Subcommittee on Science, Research, and Development, Committee on Science and Astronautics, testified that "for the first time in the history of Man on this planet the pertinent scientists of nearly all countries are joining together for a unified and coordinated look at Man and his environment on a world-wide basis" (Van Dyne 1996).

As mentioned previously, large-scale ecological research is being facilitated by advances in technology, such as satellite monitoring, global positioning systems (GPS), geographical information systems (GIS), computer modeling of large-scale processes, and very large database (VLDB) management technology. The availability of these new technologies is making it easier for humanity to perceive the environment on a global scale. Moreover, the rise in public environmental consciousness is encouraging governments, conservation groups, and scientific organizations at the national and international levels to use these new technologies to monitor regional and global ecological systems. Cooperative research programs are currently under way that are the successors to the International Biological Program mentioned above. These include the World Climate Research Programme (WCRP), which is part of the World Climate Programme, and the International Geosphere-Biosphere Programme (IGBP). The objective of these programs is to better understand global change, including the ecological aspects thereof. Agencies such as the International Council of Scientific Unions (ICSU) and the International Group of Funding Agencies for Global Change Research (IGFA) have been created to provide leadership for the coordination and scientific planning of these kinds of international programs (Global Change Research Information Office 1996).

The challenge of characterizing and comparing large-scale biosystems has been broached from the two camps of applied and theoretical ecology. The new, advanced technology tools and high-profile international programs are frequently associated most strongly with the latter. The aim of theoretical ecological research is to gain a formal understanding of the general governing principles that are common to all ecosystems, and that are perhaps extensible to the larger set of biosystems. This is the realm of mathematical ecologists and biologists, whose pursuits overlap with the disciplines of nonlinear dynamics, complexity, and artificial life. Theoretical ecologists often work with idealized abstractions and computational models, and tend toward the use of overarching, generalized characteristics. Characteristics that are deemed to be important by this group are those that capture the "essence" of the biosystem of interest, and that can serve as a basis for formal models thereof. However, these attributes may be impractical or impossible to quantify directly in a large, physical system.

Applied ecologists tend to identify ecosystem characteristics that can be quantified or estimated through field measurements. These measures follow in the tradition of work done by early ecologists like Clements (1916), whose monograph about plant succession included an exhaustive analysis of the vegetative communities in question, and the environmental factors that shaped them. Researchers in the area of restoration ecology, for example, have defined sets of attributes that serve to characterize ecosystems (Aronson et al. 1993a, 1993b). These attributes range from perennial species richness and soil biota diversity, to rain use efficiency and microsymbiont effectiveness. Similar characteristics have been identified for use during environmental assessment studies by numerous government agencies and conservation groups concerned with natural resource management and industrial development. The community of applied ecologists can also be considered to include managers of the various kinds of artificial and modified natural ecosystems mentioned above. The effective control of agricultural, silvicultural, aquacultural, industrial fermentation, and other goal-oriented systems requires that they be monitored, and this necessitates measures by which to characterize and compare their states. These measures and methodologies differ, depending on the emphasis placed on the management of the system: soil nutrient analysis on a wheat farm; timber assaying in a commercial forest; or the temperature in a trout tank.

Since the characterization and comparison of biosystems involve the sampling and summarization of the attribute values of populations of systems components, statistical methods often play an important role. Physical biosystems of the ecosystem scale are often spatially extensive, and spatial statistics are therefore useful in their analysis, as well as in the analysis of data from models meant to represent such systems. Recalling the definition of biosystems as living, dynamic systems, it is apparent that temporal statistics are also frequently called for. In fact, statistical inference can serve as the link between measurements performed by the applied ecologist, and the generalized formal characteristics proposed by theoretical ecologists. Evidence of the underlying principles of organization may be abstracted from raw field data through the judicious application of statistical methods.

2.6 Conclusions

The EcoCyborg Project is an investigation of possible approaches to engineering highly autonomous biosystems. As such, it exists at the nexus of several broad areas of research: ecological systems theory, including complex systems theory and artificial life; the engineering of natural, modified, and completely artificial biosystems; the cognitive sciences of psychology, neurophysiology, and artificial intelligence; and the characterization of ecosystems, cognitive systems, and hybrids thereof. Because the extent of each of these research areas prohibits a comprehensive literature review, the aim of this article is to serve as a general guide to the salient domains.

The approach that is being followed in the EcoCyborg Project is to increase the autonomy of large-scale biosystems using an artificially intelligent control network. The details of the project as it is currently progressing may eventually be incorporated in the design of such systems. It is highly probable that many future biosystems will be tightly integrated with advanced control subsystems that include artificially intelligent control components. However, whether or not the EcoCyborg is physically realized in detail is not important. It is believed that exploratory research of this kind is fundamental to the development of biosystems engineering. The ongoing development of a philosophical framework and lexicon, as well as modeling and simulation methods for the creation of speculative design approaches, are all necessary to the advancement of the field.

2.7 Acknowledgments

The authors wish to acknowledge Christine Stacey and Lael Parrott for reviewing initial drafts of the article, and the Natural Sciences and Engineering Research Council of Canada and McGill University for providing financial support for this research.

2.8 References

Allen, J. 1991. *Biosphere 2: The Human Experiment*. New York, NY: Penguin.

Aronson, J., C. Floret, E. Le Floc'h, C. Ovalle and R. Pontanier. 1993b. Restoration and rehabilitation of degraded ecosystems in arid and semi-arid lands. II. Case studies in Southern Tunisia, Central Chile, and Northern Cameroon. *Restoration Ecology* 1(3):168-187.

Aronson, J., C. Floret, E. Le Floc'h, C. Ovalle and R. Pontanier. 1993a. Restoration and rehabilitation of degraded ecosystems in arid and semi-arid lands. I. A view from the South. *Restoration Ecology* 1(1):8-17.

Bak, P., and K. Chen. 1991. Self-organized criticality. *Scientific American* January:46-53.

Beardsley, T. 1995. Down to Earth: Biosphere 2 tries to get real. *Scientific American* August:24-26.

Capra, F. 1996. *The Web of Life*. New York, NY: Doubleday.

Carson, R. 1962. *Silent Spring*. Boston, MA: Houghton Mifflin.

Casti, J.L. 1995. *Complexification*. New York, NY: HarperCollins.

Chalmers, D.J. 1994a. A computational foundation for the study of cognition. *The AI CD-ROM*, Revision 2. Irving, TX: Network Cybernetics Corporation.

Chalmers, D.J. 1994b. Consciousness and cognition. *The AI CD-ROM*, Revision 2. Irving, TX: Network Cybernetics Corporation.

Clark, O.G., R. Kok and P. Champigny. 1995. Generation of a virtual terrain. NABEC Paper No. 9510. St. Joseph, MI: ASAE.

Clark, O.G., R. Kok and P. Champigny. 1997. Generation of a virtual terrain. *Environmental Software* 12(2/3):143-149.

Clark, O.G., R. Kok and R. Lacroix. 1996. Engineering intelligent biosystems. NABEC Paper No. 9662. St. Joseph, MI: ASAE.

Clark, O.G., R. Kok, R. Molenaar and L. Parrott. 1997. The implementation of mind in engineered biosystems. In *Proceedings of the 1997 Conference of the Canadian Society for Civil Engineering, Sherbrooke, Québec, May 27-30, 1997, Volume A: Agricultural, Food, and Biosystems Engineering*, 582-589. Montreal, QC: CSCE.

Clements, F.E. 1916. *Plant Succession: An Analysis of the Development of Vegetation*. Pub. 242. Washington, DC: Carnegie Institute of Washington.

Cohen, J.E., and D. Tilman. 1996. Biosphere 2 and biodiversity: the lessons so far. *Science* 274:1150.

Dempster, W.F., M. Nelson and J. Allen. 1997. Biospherian viewpoints. *Science* 275:1247-1249.

Dyson, G. 1997. *Darwin Among the Machines*. Redwood City, CA: Addison-Wesley.

Edmonds, B. 1998. What is Complexity?: the philosophy of Complexity per se with application to some examples in evolution. In *The Evolution of Complexity*, ed. F. Heylighen and D. Aerts. Dordrecht, the Netherlands: Kluwer Academic Publishers.

Folsome, C.E., and J.A. Hanson. 1986. The emergence of materially-closed-system ecology. In *Ecosystem Theory and Application*, ed. N. Polunin, 269-288. New York, NY: Wiley.

Fuller, R.B. 1969. *Operating Manual for Spaceship Earth*. Carbondale, IL: Southern Illinois University Press.

Gardner, H. 1993. *Frames of Mind*, 2nd ed. New York, NY: HarperCollins.

Gauthier, L. 1987. Development and use of a database and program package for farm management. Unpublished Ph.D. thesis. Department of Agricultural Engineering, McGill University, Montreal, QC.

Gauthier, L., and R. Kok. 1989. Integrated farm control software: I. Functional requirements and basic design criteria. *AI Applications* 3(1):27-37.

Gitelson, I.I., I.A. Terskov, B.G. Kovrov, G.M. Lisovskii, Y.N. Okladnikov, L.S. Sid'ko, I.N. Trubachev, M.P. Shilenko, S.S. Alekseev, I.M. Pan'kova and L.S. Tirranen. 1989. Long-term experiments on Man's stay in biological life-support system. *Advanced Space Research* 9(8):65-71.

Gleick, J. 1988. *Chaos*. New York, NY: Penguin.

Global Change Research Information Office. 1996. Our Changing Planet: The FY 1996 U.S. Global Change Research Program. Global Change Research Information Office, University Center, MI.

Greenfield, S.A. 1995. *Journey to the Centers of the Mind*. New York, NY: Freeman.

Hall, N. (ed.). 1991. *Exploring Chaos*. New York, NY: New Scientist.

Haug, P.T. 1983. Living systems theory: conceptual basis for ecosystem modelling. In *Analysis of Ecological Systems: State-of-the-art in Ecological Modeling*, ed. W.K. Lauenroth, G.V. Skogerboe and M. Flug, 45-56. Amsterdam, the Netherlands: Elsevier.

Hoffmeyer, J. 1997. Biosemiotics: Towards a new synthesis in biology. *European Journal for Semiotic Studies* 9(2): 355-376.

Hofstadter, D.R., and D.C. Dennett. 1982. *The Mind's I*, 2nd ed. Harmondsworth, Middlesex, UK: Penguin.

Kaiser, J. 1994. Wiping the slate clean at Biosphere 2. *Science* 265:1027.

Kaiser, J. 1996. New manager for Biosphere 2. *Science* 273:1327.

Kauffman, S.A. 1993. *The Origins of Order*. New York, NY: Oxford University Press.

Kok, R., and G. Desmarais. 1985. An integrated hierarchical control system for an intelligent greenhouse. ASAE Paper No. NAR85-403. St. Joseph, MI: ASAE.

Kok, R., and G. Desmarais. 1987. Structure d'un systeme de controle pour une serre intelligente. CSAE Paper No. 87-103. Saskatoon, SK: CSAE.

Kok, R., and G. Desmarais. 1988. A hierarchical control system for an intelligent greenhouse. EurAgEng Paper No. 88.249. Silsoe, UK: European Society of Agricultural Engineers.

Kok, R., and L. Gauthier. 1987. Design considerations for an agricultural eco-cyborg. CSAE Paper No. 87-105. Saskatoon, SK: CSAE.

Kok, R., and L. Gauthier. 1988. Design considerations for an agricultural eco-cyborg. EurAgEng Paper No. 88.247. Silsoe, UK: European Society of Agricultural Engineers.

Kok, R., and R. Lacroix. 1991. An analytical framework for autonomous and enclosed agro-ecosystems. A case study: the greenhouse. In *2° Convegno Nazionale "Il vivaismo oriticolo"*, 83-95. Foggia, Italy, May 2-4.

Kok, R., and R. Lacroix. 1993. An analytical framework for the design of autonomous, enclosed agroecosystems. *Agricultural Systems* 43:235-260.

Kok, R., and R. Lacroix. 1993. Consciousness and the agro-ecosystem. NABEC Paper No.93-106. St. Joseph, MI: ASAE.

Kok, R., G. Desmarais and J. Larouche. 1986. A control/operating system for an intelligent greenhouse. CSAE Paper No. 86-602. Saskatoon, SK: CSAE.

Kok, R., L. Parrott, G. Clark, R. Molenaar and R. Lacroix. 1995. The engineering of autonomous biosystems: the EcoCyborg Project. CSAE Paper No. 95-602. Saskatoon, SK: CSAE.

Kok, R., R. Lacroix and E. Taillefer. 1991. Greenhouse climate modelling with an artificial neural network. CSAE Paper No. 91-212. Saskatoon, SK: CSAE.

Kok, R., R. Lacroix, O.G. Clark and E. Taillefer. 1994. Imitation of a procedural greenhouse model with an artificial neural network. *Canadian Agricultural Engineering* 36(2):117-126.

Lacroix, R. 1994. A framework for the design of simulation-based greenhouse control. Unpublished Ph.D. thesis. Department of Agricultural Engineering, McGill University, Montreal, QC.

Lacroix, R., and R. Kok. 1991a. A cognitive controller with recourse to simulation. Unpublished paper presented at the Automated Agriculture for the 21st Century ASAE Symposium. Chicago, IL, December 16-17.

Lacroix, R., and R. Kok. 1991b. Greenhouse simulation in a multitasking environment using OS/2. Unpublished paper presented at the Annual Meeting of NABEC. Ste-Anne-de-Bellevue, QC, July 22-23.

Lacroix, R., and R. Kok. 1994. Un cadre conceptuel pour le controle des ecosystemes. Unpublished paper presented at the 62nd Congress of L'Association Canadienne-Française pour L'Avancement des Sciences. Montreal, QC, May, 18.

Lacroix, R., and R. Kok.1999. Simulation-based control of enclosed ecosystems – a case study: determination of greenhouse heating setpoints. *Canadian Agricultural Engineering* 41(3):175-184.

Lacroix, R., R. Kok and L. Gauthier. 1993. Simulation-based determination of greenhouse temperature setpoints. ASAE Paper No. 934043. St. Joseph, MI: ASAE.

Lacroix, R., R. Kok and O.G. Clark. 1996. Use of a multitasking operating system as a setting for the simulation of an enclosed agro-ecosystem under cognitive control. *Canadian Agricultural Engineering* 38(2):129-138.

Langton, C.G. (ed.). 1989. *Artificial Life, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume VI.* Redwood City, CA: Addison-Wesley.

Langton, C.G. (ed.). 1994. *Artificial Life III, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume XVII.* Redwood City, CA: Addison-Wesley.

Langton, C.G., C. Taylor, J.D. Farmer and S. Rasmussen (eds.). 1991. *Artificial Life II, Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume X.* Redwood City, CA: Addison-Wesley.

Levy, S. 1992. *Artificial Life*. New York, NY: Random House.

Liao, C.-M. 1991. An analysis of the dynamics of agroecosystems via the viewpoint of engineering cybernetics. *Journal of Chinese Agricultural Engineering* 37(2):14-25.

Logofet, D.O. 1993. *Graphs and Matrices*. Boca Raton, FL: CRC Press.

Loucks, O.L. 1986. The United States' IBP: An ecosystems perspective after fifteen years. In *Ecosystem Theory and Application*, ed. N. Polunin, 390-405. New York, NY: Wiley.

Macilwain. 1996. Biosphere 2 begins fight for credibility. *Nature* 380(28 March):275.

Mandelbröt, B.B. 1983. *The Fractal Geometry of Nature*. New York, NY: Freeman.

Margulis, L., and D. Sagan. 1995. *What is Life?* New York, NY: Simon and Schuster.

Margulis, L., and D. Sagan. 1986. *Microcosmos*. New York, NY: Simon and Schuster.

Margulis, L., and J.E. Lovelock. 1974. Biological modulation of the Earth's atmosphere. *Icarus* 21:471-489.

Marr, D. 1982. *Vision*. San Francisco, CA: Freeman.

Maturana, H.R., and F.J. Varela. 1980. *Autopoiesis: the Organization of the Living*. Dordrecht, the Netherlands: D. Reidel Publishing Company.

McGlade, J. 1993. Alternative ecologies. *New Scientist* 13(February):14-16.

McNaughton, S.J., and M.B. Coughenour. 1981. The cybernetic nature of ecosystems. *The American Naturalist* 117:985-990.

McPhillimy, D. 1997. A forest for Scotland. In *Ecoforestry*, ed. A. Drengson and D. Taylor, 159-166. Gabriola Island, BC: New Society Publishers.

Mikulecky, D.C. 1995. Life, complexity and the edge of chaos: cognitive aspects of communication between cells and other components of living systems. *Centre for Policy Modelling, Manchester Metropolitan University*. <http://griffin.vcu.edu/complex/mikulecky/rev.htm> [Accessed 11 August, 1999].

Mitchell, C.A. 1994. Bioregenerative life-support systems. *American Journal of Clinical Nutrition* 60:820s-824s.

Molenaar, R., and R. Kok. 1995. Development of a Pavlovian controller for the guidance of autonomous biosystems. NABEC Paper No. 9523. St. Joseph, MI: ASAE.

Molenaar, R., R. Kok and R. Lacroix. 1995. Structure of the EcoCyborg simulation under OS/2. CSAE Paper No. 95-604. Saskatoon, SK: CSAE.

Nelson, M., T.L. Burgess, A. Alling, N. Alvarez-Romo, W.F. Dempster, R.L. Walford and J.P. Allen. 1994. Using a closed ecological system to study Earth's biosphere: Initial results from Biosphere 2. *BioScience* 43(4):225-236.

Nitta, K., K. Otsubo, S. Kibe and H. Ohja. 1990. Japanese research activities of life support system. SAE Paper No. 901205.Warrendale, PA: SAE.

Odum, E.P. 1977. The emergence of ecology as a new integrative discipline. *Science* 195:4284.

Odum, E.P. 1993. *Ecology and our Endangered Life Support Systems*, 2nd ed. Sunderland, MA: Sinauer.

Odum, E.P. 1996. Cost of living in domed cities. *Nature* 382(July 4):18.

Odum, H.T., and E.P. Odum. 1955. Trophic structure and productivity of a windward coral reef community on Entiwetok Atoll. *Ecological Monographs* 25(3):291-320.

Parrott, L. 1995. The EcoCyborg Project: a model of an artificial ecosystem. Unpublished M.Sc. thesis. Department of Agricultural Engineering, McGill University, Montreal, QC.

Parrott, L., and R. Kok. 1995. Creation of an ecosystem model: structure and composition. CSAE Paper No. 95-609. Saskatoon, SK: CSAE.

Parrott, L., R. Kok and R. Lacroix. 1995. Daily average temperatures: modeling and generation with a Fourier transform approach. NABEC Paper No. 9522. St. Joseph, MI: ASAE.

Parrott, L., R. Kok and R. Lacroix. 1996. Daily average temperatures: modeling and generation with a Fourier transform approach. *Transactions of the ASAE* 39(5):1911-1922.

Parrott, L., R. Kok, G. Clark and R. Molenaar. 1996. Crafting New Life-Engines. *Resource* 3(3):8-10.

Patten, B.C. 1959. An introduction to the cybernetics of the ecosystem: the trophic-dynamic aspect. *Ecology* 40(2):221-231.

Patten, B.C., and E.P. Odum. 1981. The cybernetic nature of ecosystems. *The American Naturalist* 118:86-895.

Patten, B.C., S.E. Jorgensen and S.I. Auerbach (eds.). 1995. *Complex Ecology*. Englewood Cliffs, NJ: Prentice Hall.

Peitgen, H.-O., H. Jürgens and D. Saupe. 1992. *Chaos and Fractals*. New York, NY: Springer-Verlag.

Penrose, R. 1989. *The Emperor's New Mind*. Oxford, UK: Oxford University Press.

Pinker, S. 1994. Mind design. In *The Language Instinct*, 1st ed, 404-430. New York, NY: William Morrow.

Prigogine, I. 1980. *From Being to Becoming*. San Francisco, CA: Freeman.

Reisch, M.E.A. 1983. A legislative history of the Comprehensive Environmental Response, Compensation, and Liability Act of 1980 (Superfund), Public Law 96-510: together with a section-by-section index. Prepared by the Environmental and Natural Resources Policy Division of the Congressional Research Service of the Library of Congress, for the Committee on Environment and Public Works, U.S. Senate, U.S. G.P.O, Washington, DC.

Rosen, R. 1988. The epistemology of complexity. In *Dynamic Patterns in Complex Systems*, ed. J.A.S. Kelso, A.J. Mandell, and M.F. Shlesinger, 7-30. Singapore: World Scientific.

Sagan, C. 1994. *Pale Blue Dot*. New York, NY: Random House.

Schank, R.C., and P. Childers. 1984. *The Cognitive Computer*. Sidney, ON: Addisson-Wesley.

Scheinerman, E.R. 1996. *Invitation to Dynamical Systems*. Upper Saddle River, NJ: Prentice Hall.

Schrödinger, E. 1955. *What is Life?* Cambridge, MA: Cambridge University Press.

Schröeder, M.R. 1990. *Fractals, Chaos, Power Laws*. New York, NY: Freeman.

Schultz, E.-D., and H.A. Mooney (eds.). 1994. *Biodiversity and Ecosystem Function*. Berlin, Germany: Springer-Verlag.

Shannon, C.E. 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27(July-October):379-423, 623-656.

Silvert, W. 1995. Complexity. *Journal of Biological Systems* 4:585-591.

Spedding, C.R.W. 1988. *An Introduction to Agricultural Systems*, 2nd ed. New York, NY: Elsevier.

Tamponnet, C., M. Colasson and M. Pastor. 1992. European involvement in CELSS: definition of a closed ecological systems test bed. SAE Paper No. 921200. Warrendale, PA: SAE.

Tansley, A.G. 1935. The use and abuse of vegetational concepts and terms. *Ecology* 16(3): 284-307.

Thompson, J.M.T., and H.B. Stewart. 1986. *Nonlinear Dynamics and Chaos*. Toronto, ON: Wiley.

Ting, K.C., P.P. Ling and G.A. Giacomelli. 1997. Sustaining human lives in outer space. *Resource* 4(3):7-8.

Tsiolkovsky, K. 1960. *The Call of the Cosmos*. Moscow, Russia: Foreign Languages Publishing House.

United Nations. 1995. *Montreal Protocol on Substances that Deplete the Ozone Layer: 1994 Report of the Economics*. New York, NY: United Nations.

Van Dyne, G.M. 1995. Ecosystems, systems ecology, and systems ecologists. In *Complex Ecology*, ed. B.C. Patten, S.E. Jorgensen and S.I. Auerbach. Englewood Cliffs, NJ: Prentice Hall.

Volk, T. 1996. Considerations in miniaturizing simplified agroecosystems for advanced life support. *The Journal of Ecotechnology* 6:99-108.

von Neumann, J. 1963. The general and logical theory of automata. In *John von Neumann: Collected Works*, vol. 5., ed. A.H. Taub, 288-328. Oxford, UK: Pergamon.

Waldrop, M.M. 1987. *Man-Made Minds*. Rexdale, ON: Wiley.

Waldrop, M.M. 1992. *Complexity*. New York, NY: Touchstone.

Watson, J.D. 1968. *The Double Helix*. New York, NY: Atheneum.

Williams, B.C., and P.P. Nayak. 1996. Immobile robots: AI in the New Millennium. *AI Magazine* 17(3):16-35.

Wilson, E.O. 1992. *The Diversity of Life*. Cambridge, MA: Belknap Press.

CONNECTING TEXT

Chapter 3, **Characterizing biosystems as autopoietic entities**, was authored by O.G. Clark and R. Kok and, at the time this thesis was submitted, had been sent for review to the editors of the journal *Oikos*.

This chapter contains a detailed discussion of the definition of *biosystem* as the class of all systems that are alive to some degree. Aliveness, or vitality, is one of the two principal system characteristics that are examined in this thesis, the other being autonomy, which is introduced in Chapter 4. In this chapter, therefore, one of the primary conceptual streams of the thesis is introduced and developed, setting the general context for the rest of the thesis. The related lexicon is defined and presented from a systems-theoretic perspective, which is useful in the characterizion of systems for scientific and engineering purposes.

CHAPTER 3. CHARACTERIZING BIOSYSTEMS AS AUTOPOIETIC ENTITIES

Abstract

This article is an exposition of the term *biosystem*, which is defined as a collection of entities forming a coherent assemblage that is alive to some degree, where *autopoiesis* is proposed as the sole criterion for being alive. Autopoiesis is a special type of comportment in which local interactions among individual components combine to continually renew the overall system of which they are a part. This does not imply that the system will grow or replicate itself, although these activities may both arise as the result of autopoiesis. The latter is driven by a gradient in order, either between inputs and outputs, or between sources and sinks within the system. The article opens with a brief discussion of some epistemological issues surrounding the definition of the term *system*, the specification of *system types*, and how particular instances of these types might be described. Since there currently are no direct measures of autopoiesis, measures of related characteristics, such as order, complexity, and emergence, are discussed in some detail. Next, state space is introduced as a paradigm in which system comportment can be framed. Static, periodic, and chaotic behavior are reviewed as distinct modes of comportment that may correspond to attractors in an appropriately defined state space, and autopoiesis is also presented in this way. Finally, the term *biosystem* is re-examined in light of this interpretation to show how it can include a broader range of systems than are conventionally considered to be alive. It is shown to accommodate systems that range in scale from the subcellular to the biospheric, that are of organic or inorganic composition, and are of either natural or artificial origins.

3.1 Introduction

In this article the term *biosystem* is defined so that naturally occurring biological phenomena of different scales, as well as systems that are of nonbiological composition, and constructs that are partly or completely artificial in origin, are all encompassed. Thus, a *biosystem* is considered to be a collection of entities forming a coherent assemblage that is alive to some degree, where *alive* is taken to mean *autopoietic*. Autopoiesis is a dynamical mode wherein the local interactions of the components of a system combine to continually renew that system, which is therefore "self-producing" when considered as a whole (Maturana and Varela 1980). Self-production is both a necessary and sufficient condition for autopoiesis; thus, it is not implied that the system will grow or replicate itself, although these activities may both arise as the result of autopoiesis. This definition is explored and developed in the article from an engineering perspective, where the intent is to create new systems. An example of such a system is an *ecocyborg*, which is a biological ecosystem hybridized with a technological control network. Ecocyborgs are of special interest because they are currently the subjects of the research program that is the context of this article (Clark et al. 1999).

The body of the article opens with a discussion of the concept of *system*. *Order* and *disorder*, *complexity*, and *emergence* are then examined in some detail. High measurements of these system characteristics accompany, but are not sufficient to guarantee, a high degree of autopoiesis in the comportment of a system. Since, however, there is currently no means to directly measure autopoiesis, these characteristics are useful in the study of living things. A general discussion follows, therefore, of how measures of order and disorder, complexity, and emergence can be used in the characterization of biosystems. *State space* is next introduced as a framework for representing change in a system, and three qualitative classes of system comportment are described: stationarity, periodicity, and chaos. *Autopoiesis* is then described as another qualitatively distinct mode of comportment that, when portrayed within an appropriate state space, tends toward a special class of attractor. Finally, the definition of *biosystem* is re-examined in this light to show how it can include a variety of systems of different scales, compositions, and origins, including large-scale, artificial constructs like ecocyborgs.

3.2 System characterization

This section is a treatment of some epistemological issues related to the concept of *system*, from which *biosystem* is derived. A system is a set, i.e., a number of component entities that are considered together as a unit. Any system is an instance of a *system type* which may be specified by an observer in one of a number of ways. First, the observer might explicitly specify membership criteria for system components, often including the requirement for some kind of mutual interaction. The observer can then identify instances of that system type (i.e., particular systems) by evaluating the extent to which these criteria are satisfied. Second, a reverse approach can be used, whereby the observer identifies a system by declaring a boundary, all entities within which are considered to be components of that system. The identification of a particular system in this way also implies the specification of a set of membership criteria (i.e., a system type) that conforms to the salient attributes of the system components, rather than vice versa. Given only one instance of this system type, however, these criteria are likely to be highly indeterminate. Only as more instances of that type are identified, perhaps by declaring similar boundaries around other sets of entities, will the criteria become more clearly and precisely specified. This method is often used to identify as a single, continuous system a set whose membership is somewhat variable. This is a reasonable approach so long as the system does not exhibit a dramatic shift in its qualitative nature. Third, an observer might identify a set of entities as a system while neither evaluating their concordance with any criteria, nor declaring an explicit boundary around them. This approach can be considered as equivalent to the declaration of a (possibly highly disjoint) boundary, or to the implicit specification of a (perhaps very indeterminate) system type based on a set of membership criteria that are tailored to match the attribute values of the included entities. Although it might be possible to relate these different methods by formulating corresponding sets of explicit membership criteria, boundary types, and lists of components, this is not always a trivial task, nor a necessary one. In all cases, the system type specification is entirely dependent on an observer, whose approach might change depending on the context. The identity of any individual system is therefore always subjective.

When dealing with living things, which are more or less integrated with one another and with their environment, the discrete identification of a particular system is often problematic. This is true whether the type of the system in question is based on the concordance of its components with a set of membership criteria, the declaration of a system boundary, or the selection of a set of entities. For example, a human being might be identified as an instance of *biosystem* by declaring a system boundary that is coextensive with the skin. Such pat distinctions can often prove to be awkward, however, since transplanted organs, prostheses, ingested materials, clothes, or even the car that the human might be driving, for example, are frequently considered as part of the human. An alternative to the discrete approach is to specify system types, or classes, in which entities can have a variable degree of membership. This approach is reflected in natural language, whereby all people can, for instance, be considered "short" to some degree, even if some are "not at all short". The equivalent formal concept is the *fuzzy set*, in which every item is assigned a membership value that can range from zero to unity (Kosko 1993). Although the analysis of some types of systems is made easier if component membership values are discretized by imposing threshold values, the use of fuzzy sets constitutes a more realistic and convenient approach for many other system types. This becomes apparent, for example, in the discussion of coupled autopoietic networks, later in this article.

Once an observer has identified a particular system, it can then be characterized, and once a number of systems have been identified and characterized, they can then be compared. As with the specification of a system type, the approach to characterization and comparison is also determined by the observer. The focus might, for instance, be on system *constitution*. This includes *composition*, i.e., the number and kinds of system components, and *structure*, referring to the overall pattern of functional relationships between them. The constitution of a system serves as the framework for its *state*, which comprises the values of all of the attributes of its components, as well as for the way in which the state changes with time, which is the system's *comportment*. Characterization and comparison can be based not only on constitution, but also on a system's state or comportment. Consideration of comportment is especially important in the understanding of biosystems, because their defining feature, autopoiesis, is a special dynamical mode. In characterizing a system, not only can the various features of the system itself be taken into consideration, but so can the way in which it interacts with its surroundings. Such interaction can affect any or all of the composition, structure, state, or comportment and if this is so, then the system is said to be *open* with respect to that particular aspect; otherwise, the system is said to be *closed*. Closure in the first case means that components neither enter nor leave the system; in the second case, the relationships between the components are not changed by external influences such as *forcing* *functions*; and in the last two cases, the values of the components' attributes are independent of external influences. The Earth, for instance, can be considered as approximately closed with respect to both composition and structure, since the amount of mass that enters from or escapes to space is negligible, and the relationships between its physical components are relatively constant. It is, however, open with respect to its state and comportment, because the attributes of its components are significantly influenced by radiation from the Sun and the gravitational effects of the Sun and Moon.

For a system that is subject to regular, sustained external influences, it might be convenient to internalize these by respecifying the system boundary. The forcing functions are then considered to be internal *sources*, and their influence is deemed to be an inherent property of the system. Thus, one might expand the specification of the Earth system so that the effects of the Sun and Moon are considered to stem from internal sources. Similarly, *sinks* can be included that diminish the magnitude of some observable quantity. A system without sources or sinks is *conservative* (as opposed to *nonconservative*). Any real, physical system is conservative with respect to mass and energy, for instance, although, when the availability of these quantities is affected by some kind of transformation, it is sometimes useful to identify such a system as being of a nonconservative type. The description of an automobile, for example, can be simplified by considering the motor as an internal source of energy, although a different description might not acknowledge an energy source, but rather a transformation of energy from a chemical, potential form to a mechanical, kinetic one.

Whether a system is closed or open, and conservative or nonconservative, has important consequences for its possible comportment. Due to dissipative effects such as the degradation of energy, most completely closed, conservative, physical systems eventually settle to an *equilibrium* state in which no more change occurs, even if their comportment might appear to be approximately nondissipative in the short term. In general, change can actually persist in a conservative, physical system only if it is open in some respect (some exceptions are noted below). On the other hand, an open system may be changing constantly, but this might not be apparent to the casual observer, if components enter and exit at approximately equal rates. Such comportment is called *steady state*. Organisms, for instance, often exhibit a kind of steady state, since they are open to flows of mass (composition) and energy (state) which are accommodated in the maintenance of autopoietic comportment. Persistent change can also occur in a system that is completely closed if it is nonconservative and contains sources and sinks that balance each other over the long term. If the biosphere, for instance, is considered to be part of a nonconservative system that includes the Earth, sources of gravitational force and radiant energy (the Moon and Sun), and a radiation sink (outer space), then its comportment can be considered to persist indefinitely, despite the complete closure of the system as it has been identified. Finally, virtual systems (and some physical systems outside of everyday experience, involving quantum mechanical phenomena such as superconductivity) can be nondissipative, and thereby display bounded comportment that continues to change indefinitely despite the absence of inputs, outputs, sources, or sinks.

System constitution, state, and comportment, including any interaction with the environment, can be characterized and compared with differing emphasis on the resolution of observation. In a reductionist approach, for example, the study of local phenomena is emphasized. The effectiveness of this approach depends on the validity of two underlying assumptions, the first being that a system which has a simple constitution will demonstrate correspondingly simple comportment. If this is true, then the phenomenon that is being studied can be divided into successively smaller parts until components are discovered that are easily understandable. The second assumption is that the objects of study combine in a tractable manner. In such circumstances, larger-scale phenomena are easily interpreted as combinations of local processes, and the explicit study of these can therefore be neglected. These two assumptions, however, do not always hold true for biosystems. Living things do not necessarily reduce to simple, easily understandable components that combine in a tractable fashion. The cells that make up biological organisms, for instance, are themselves sophisticated living things, an example of what Stewart (1995) calls the "reductionist nightmare." Moreover, when biological components interact as parts of a larger system, their combined behavior often differs surprisingly from that which might be expected from the study of isolated components. It is therefore useful to complement reductionism with approaches that emphasize the study of large-scale or multiscale features.

As pointed out above, various approaches can be used to characterize and compare systems, some pertaining to their constitution, and others related to their state, or to their comportment. All of these approaches require, however, measurement of the target characteristics. Accordingly, measures must be defined and employed so that they are appropriate to the situation in which the system is observed, and function at the range of resolution of interest to the observer. The measures that are discussed in this paper are relatively abstract, meaning either that they are of rather low resolution, or that they are based on other, more direct measures. Abstraction has both disadvantages and advantages. For instance, the further removed a measure is from direct observation, the more detail is lost. On the other hand, more abstract measures are of greater value when dissimilar systems are compared. For biosystems, relatively abstract measures often highlight the importance of features at both small and large scales, as well as across scales. In biosystems, of course, the most important multiscale feature is autopoiesis, and ideally, abstract measures would be available for its direct evaluation. No such measures currently exist, however, and so measures of other, closely related characteristics, are employed instead. These are described below.

Measures of *order* quantify the degree of correlation between features of a system. Measures of *disorder*, on the other hand, quantify the variation in a system. This variation might be random, or it might be associated with pattern, such as autopoiesis. Measures of *complexity* quantify the variation associated with pattern, and thus can be used to gauge the difficulty of describing such pattern. Measures of *emergence* quantify the degree to which global phenomena are influenced by local structure. They can therefore be used to evaluate the difference between observed phenomena and those that might be expected to occur if multiscale relationships were not taken into account. Multiscale interaction among system components is an essential aspect of autopoiesis (Maturana and Varela 1980) and therefore biosystems are highly emergent. In this regard, Logofet (1993) has expressed the opinion that, for example, in ecosystems "it is mainly the interactions themselves among constituent species or ecosystem components that form the structure of the system." Measures of order and disorder, complexity, and emergence are discussed in more detail in the following sections.

3.3 Measures of order and disorder

Both order and disorder are important concepts for the characterization of biosystems, and measures of these can be employed for the analysis of state and comportment, as well as constitution. *Order* is the degree of correlation between comparable features of a system; conversely, *disorder* is the lack of correlation, or the degree of difference, between these. Measures of order and disorder can, in some cases, be constructed as complements, but this need not always be so. Although absolute disorder implies an absence of order, and vice versa, for intermediate ranges the sum of the two quantities need not always be a constant value. Whereas the values obtained for order and disorder are reflective of the objective properties of the system being examined, i.e., they represent some real aspect of the system as it exists, they are strongly affected by the nature of the specific measure being employed. In this regard, the features that are compared, as well as the number of meaningfully different states that each of the particular features of interest might attain, influence the result. This means that, although there might actually be order and/or disorder in a system, they are not necessarily detectable with a particular approach to observation.

When the order or disorder of the state of a system is measured, the features that are compared are the attributes of components, at a given instant. Order, in this case, is the degree of correlation between the values of the attributes, and disorder is the degree of difference between these. As pointed out above, the resolution of observation is specified by the observer, so that any order or disorder actually present might very well go undetected at a particular resolution. Very similar measures can also be applied to other aspects of the system, such as comportment, which is describable with a set of parallel time series (histories) of the attribute values. Hence, for this second case, order can be defined in terms of the degree and type of autocorrelation and cross-correlation of the attribute values as they change with time. Conversely, disorder is the degree of change in these values over time, a concept that is particularly appropriate to the study of biosystems because of their inherent dynamics. An analysis of order and disorder might also be conducted with respect to the constitution (composition and structure) of a system. In this third case the number and kinds of components, or their interrelationships, would be considered. Although according to the strict definition of *system* that was presented previously, the observer would be dealing with a series of different systems if the constitution were to change over a period of time, it is sometimes convenient to relax this criterion and characterize the order or disorder inherent in the resulting history as if it pertained to a single system.

Measures of disorder and of order both bear a relationship to *information* and, in fact, the same unit (bit) is often employed for all of them. When particular features of a system are observed at a given resolution, only a certain number of distinctions, or quantity of information, is required to generate as complete a description as possible. For instance, for a particular system that is observed at a given resolution and for which a certain set of features are considered, in the circumstance where these are found to be maximally uncorrelated there would be as much variation among them as possible, and the disorder of the system would be at its theoretical maximum. In all other circumstances, however, the disorder would be less than this, so that a complete description could be formulated based on fewer distinctions (i.e., less information). Given such a theoretical maximum value for disorder, order might then be defined as a complementary measure, equal to the difference between the actual disorder and the theoretical maximum. If this system now were observed at a finer resolution, more possible values would be distinguishable for the features of interest, and more information would therefore be required to completely describe the maximally disordered situation. A larger measurement of system disorder might therefore be obtained for the system than when it was observed at lower resolution, and a larger value might be found for its order as well.

Order and disorder can result either from influences that are internal to a system, such as structural constraints and sources, or from external ones, like forcing functions. Overall, ordering influences tend to increase the similarity between the features of a system and thus decrease the variation, whereas disordering influences have the opposite effect. For instance, when internal structural constraints are present, the features of the system that are affected will be more correlated than when these constraints are not present, and there are fewer degrees of freedom. The effect is, of course, similar when such structural constraints are externally induced (the system must be open with respect to structure for this to be possible). The number of degrees of freedom can also be altered, and the system order and disorder affected, by the addition or removal of components. This is possible if the system is open with respect to composition, or it might be due to the activity of an internal source or sink. The influence on system order and disorder would then be related to the effect on the correlation between the features which depend on those components that were added or removed. Forcing functions may also affect the order and disorder of a system, by influencing what region of the state space is reachable and thus changing the number of possible microstates (this is described in more detail immediately below and in Section 3.6). As well, they may influence the probability with which the system inhabits its microstates. Thus, their effect might be to narrow the distribution of the components' attribute values, which would correspond to greater order. On the other hand, they might broaden the distribution, in which case the disorder would increase. In the case of comportment, disordering influences increase the variation among attribute values with respect to time, possibly causing temporal features such as cyclicity or chaos to appear in previously more correlated data.

Order and disorder can be quantified with many different measures. Of these, entropy was one of the first to be defined. It is a measure of disorder that was devised by Boltzmann in 1877 as a statistical thermodynamic quantity linking the molecular theory of matter and the concept of unavailable thermodynamic work in an ideal gas system (Broda 1983). It quantifies the degree of uniformity of the distribution of attribute values. Essentially, entropy describes an observer's ignorance of the actual microstate (molecular configuration) of a closed ideal gas system that is in a particular macrostate (global state) (Gell-Mann and Lloyd 1996). The observer is ignorant of exactly which microstate such a system is in, and can only describe the probability associated with the system being in each of them. A system in a lower-entropy macrostate inhabits one of fewer possible microstates (exactly which one being unknown) than it would if it were in a given reference state. Conversely, a system in a higher-entropy macrostate inhabits one of a greater number of possible microstates than it would if it were in the reference state. In other words, the probability distribution associated with the microstate population of a system in a higher-entropy macrostate is broader than that associated with the microstate population of the same system when it is in a lower-entropy macrostate. The maximum value for the entropy measure is reached when the probability distribution is completely uniform. Since there is a theoretical maximum value for such a measure of disorder, a complementary measure of order can be defined as the difference between the actual value and the theoretical maximum. Instead of comparing two states of the same system, two different systems can also be compared. These might have different numbers of microstates due to different constitutions. The entropy of the system with the larger number of microstates is, however, not necessarily greater than that of the other, since the entropy value depends also on the probability distribution associated with the inhabitation of the microstates.

Thermodynamic entropy is just one of a family of measures that can be defined to quantify disorder. Similar measures can be defined based on other macrostates and microstates that reflect system features and a resolution of observation that are useful in a given scenario. Shannon (1948) became the first to recognize the general utility of such measures. He applied them in communication theory, defining a measure that is now often referred to as *Shannon information*. Whereas for thermodynamic entropy a macrostate was originally defined as the global state of an ideal gas system, in the case of Shannon information it refers to a given set of messages that might be transmitted during communication. Accordingly, *microstate* refers in this instance to one of the possible messages that might be transmitted. The value of the measurement is therefore zero if the observer is certain beforehand which message will be received, whereas it is maximized if all messages are equally likely to be received (the uniform distribution case). Thermodynamic entropy, Shannon information, and other, similar measures have become standard for many different contexts.

Measures of order and disorder are especially relevant to the study of living things. This is because biosystems are usually dissipative, meaning that there is a tendency for their disorder to increase. In the case of physical systems, this disorder is ultimately manifested as a degradation of energy. Biosystems must export this disorder which, in the physical case, requires a flow of a substrate, such as matter or energy, through the system. Boltzmann, in his 1886 monograph entitled *The second law of the mechanical theory of heat* first expressed this idea when he defined entropy flow as necessary for life (Broda 1983, pp.79-80):

> "The general struggle for existence of living beings is therefore not a fight for the elements – the elements of all organisms are available in abundance in air, water, and soil –, nor for energy, which is plentiful in the form of heat, unfortunately untransformably, in every body. Rather, it is a struggle for entropy that becomes available through the flow of energy from the hot Sun to the cold Earth."

Schrödinger (1955) popularized the equivalent idea that living things require a source of order, or *negative entropy*, that they can degrade so as to maintain themselves in an ordered state. This implies that physical biosystems are conservative and must be open to their environments so as to allow an exchange of some medium that can transport this negative entropy. Alternatively, biosystems might be specified as nonconservative, with order originating from an internal source. Even in this case, however, order is degraded as the biosystem maintains its structure. The attendant autopoietic comportment is itself moderately ordered, as evidenced by the variety of temporal features that can be seen when it is appropriately observed and analyzed. Life cycles of organisms are an example of such temporal features that can be characterized with appropriate measures of order and disorder.

As described above, measures of disorder quantify the variation in a system's constitution, state or comportment, and measures of order quantify the correlation within these. These measures alone, however, cannot be used to determine how much of the variation and correlation is associated with randomness and how much with pattern. *Pattern* comprises the relationships between features of the system, involving change under some kinds of transformations, and invariance under others. The ability to distinguish between variation that is associated with pattern and that which is random depends on the observer's understanding of the system, i.e., the possession of appropriate models of the relationships between its features. Understanding can be derived from, but can also lead to, the perception of pattern. Accordingly, an unsophisticated observer, i.e., one with a limited understanding of a system, might identify only few patterns in it, whereas a more sophisticated observer might perceive many different patterns. Both pattern and understanding have a bearing on how predictable an observer finds a system to be. Given perfect understanding of the pattern in the system, the observer has a maximal ability to predict the features of interest, with respect to some independent variable such as time or distance. It must be noted, however, that *maximal* predictability does not imply *complete* predictability, since some patterns are inherently unpredictable in detail, and because some features are unpatterned. Measures of complexity can be used to distinguish between variation that is due to recognized pattern and that which is due to features that are poorly understood or unpatterned. In the study of biosystems, the first kind of variation is of primary importance, and it is therefore useful to make a distinction between these two cases.

3.4 Measures of complexity

In the most general sense, *complexity* is the difficulty of performing a given task (Li 1997). In the more narrow sense relevant to this article, complexity is the difficulty of describing the patterns that exist in a system or system type, and that a particular observer is capable of recognizing. As with measures of order and disorder, therefore, the magnitude of a given complexity measurement depends not only on the observed entity, but also on the observer. This is because the recognition of pattern is a function of perceptual ability, understanding (as described above), and the chosen approach to observation. In other terms, complexity measures quantify the variation associated with the perceivable and understandable pattern in a particular system or system type. This pattern, and therefore the complexity of the system, might, as with order and disorder, arise from either internal or external influences.

Complexity, order and disorder share other general commonalties. For instance, measures of complexity bear a relationship to information, similar to the way that measures of order and disorder do. This is because a certain, minimal number of distinctions are required to describe the pattern present in a system, and so complexity measures, in quantifying this pattern, yield measurements of a size (possibly measured in bits of information) that is characteristic of that particular system. The description of the pattern in that system requires more information than the description of the pattern in the system type to which it belongs, because the former possesses all of the patterns pertaining to the type, as well as those that are unique to the system as an individual. Any system is therefore minimally as complex as any larger class to which it belongs.

As with order and disorder, any of a variety of complexity measures can be employed to characterize and compare a given set of entities based on various different aspects. The latter might include, for instance, the composition, structure, state, or comportment of the entities. Gunther et al. (1994) indicated that any such measures should yield values that: *i*) are zero for both the cases of total order and of total system disorder, but positive in between; *ii*) do not increase for both of two independent systems as the consequence of the direct interaction of those systems; *iii*) do not increase as a result of simple enlargement of a system; and *iv*) have values that depend on the method used to describe the system. In accordance with the first characteristic, complexity measures are likely to produce high values for entities that are dominated neither by ordering nor disordering influences. In this regard, intermediacy between complete order and extreme disorder is sometimes illustrated with reference to a physical system at the point of a phase transition: one phase (e.g., solid) is a more ordered state, whereas the other phase (e.g., gas) is less ordered (Langton 1990). Sophisticated patterns often appear in such systems. The second characteristic is a principle of conservation, and the third is a statement that complexity should be independent of the size of a system. The description of the fourth characteristic underscores the importance of how a measure is defined (Silvert 1995; Crutchfield 1994a, 1994b). For instance, just as temperature can be measured with a wet-bulb thermometer or a dry-bulb one, different complexity measures can be used for characterizing biosystems in different ways.

Numerous complexity measures have been defined, each relating to particular features of the entity of interest (e.g., Li 1997; Gell-Mann and Lloyd 1996; Wackerbauer et al. 1994; Gunther et al. 1994; Crutchfield 1994a, 1994b). None, however, have been universally accepted, nor is it necessary to specify any for the purposes of this article. Nevertheless, one measure that conforms particularly well to the usage adopted here is *effective complexity*, which is roughly equivalent to the length of a compact description of the patterns that are recognized by a given observer as being present in a particular entity (Gell-Mann and Lloyd 1996). It is complemented by a quantification of the information required to describe the remaining, apparently unpatterned, and perhaps random variation. These two measures together quantify all of the detectable variation in an entity.

3.5 Measures of emergence

Biosystems can be characterized not only with measures of order, disorder and complexity, but also with measures of *emergence*. These quantify the influence that local interactions have on global phenomena, revealing the effect of pattern that extends across different scales. One could observe, for instance, a set of entities that were isolated from one another, and compare the features of this set with those of a system comprised of similar entities but that were engaged in mutual interaction. Any observed differences would be the result of the multiscale structure of the system, and this difference is quantifiable with a measure of emergence. Such measures, as with those discussed previously, depend not only on the objective properties of the observed system, but also on the abilities, understanding, and approach of the observer. Hence, they reflect the difference between the features of a system as predicted by an observer with a limited understanding of its structure, and the features of the system as predicted by an observer with greatly superior understanding. A large measure of emergence reflects a substantial difference between the complexity of the system, as it is understood by the first observer, and the complexity as measured by the second observer. Emergence, then, bears a relationship to complexity, although it is not necessarily a straightforward one. Simple interactions at the local scale sometimes belie great overall complexity, but it is also possible for complication at the local scale to underlie simplicity at the system scale, this being the gist of the aforementioned "reductionist nightmare" (Stewart 1995).

Since measures of emergence, like those of complexity, are based on other, more direct measures, they are relatively abstract. Moreover, because emergence measures quantify features across scales, the more direct measures on which they are based must be meaningful at all of the scales in question. This can be illustrated with reference to the way in which the autopoietic networks of biosystems arise from the local interactions of their components (Capra 1996). If, for example, one were to measure emergence in the flight of a set of birds, then population would not be an appropriate variable to use as a basis, because, although useful at the community scale, it is meaningless with reference to a single animal. Position, on the other hand, is meaningful at both the global and local scales, and could serve as a basis for a measure of emergence, as follows. First, the paths of a set of isolated birds could be recorded and compared with their average path. From this, a correlation value could be found for each bird and a mean correlation coefficient calculated for the set. Second, similar measurements could be recorded for the members of a flock as they interacted with one another. The ratio of the two mean correlation coefficients would indicate the degree of emergence of the flock as compared to the isolated birds. If the pattern of interactions between the birds was not understood, then their predicted behavior would be very different from that actually observed for the flock.

Numerous other emergence measures could also be defined to quantify the relationship between fine- and coarse-scale patterns in a system, and might apply variously to its composition, structure, state or comportment. As with the measures discussed previously, the manner in which these might be used and the results that they yield will be highly dependent on how they are defined and employed, and on the abilities of the observer that employs them. If used appropriately, however, they can be powerful complements to measures of order, disorder, and complexity in the study of biosystems. The next section is a discussion of *state space*, a framework in which such measures can be applied.

3.6 State space

The use of state space for system characterization is founded on the idea of a *state vector* that is composed of a number of measures. The suitability of these measures depends on how the observer specifies the system type, and then identifies the particular system in question. Once the measures that are to compose the state vector have been chosen, the accuracy and precision of the associated values will be influenced by the observer's understanding of the system and his powers of perception. Ideally, the chosen measures will correspond to variables that are mutually orthogonal, with each one corresponding to a degree of freedom of the system. The value of the vector therefore uniquely represents the state of the system at any given time (Casti 1992). A system that has been identified as an ideal oscillator in one spatial dimension, for example, can be characterized as having two degrees of freedom. Its state can therefore be completely described with two appropriately chosen measures, such as angular displacement and angular velocity, or energy and momentum. A system that has been characterized as more complicated than this will require more measures for a complete description. Generally, a system comprising *N* components, each with *d* independent attributes, has $N \cdot d$ possible degrees of freedom, and an equal number of measures must therefore be chosen to compose the state vector. The range of each chosen measure can be envisioned as extending along an axis that defines one dimension of an abstract region called a *state space*. As discussed before, constitution is, strictly speaking, constant for a given system, and so state space provides a "static backcloth" for the characterization of the system's actual and potential comportment. It must be noted that not all parts of the state space are necessarily ***reachable***, since for many system types, such as biosystems, the number of possible microstates is limited by strong interactions between the components (Casti 1992), i.e., the chosen variables are not necessarily independent. In such cases the state space might be respecified based on truly independent variables, so as to encompass only reachable states, but this is generally not convenient, and this approach is not pursued further herein.

State space can be used to characterize a system in several ways. First, a series of values of the state vector can be represented in state space as a trajectory. Every observable history of the system can thereby be uniquely represented. Secondly, the structure of a system can be characterized in state space. This can be done by superimposing the *vector field* of the system on the static backcloth. The vector field represents the potential comportment of the system as it will develop from any given conditions, as a result of the interactions between the system components. It is isomorphic to a set of differential equations, or a set of rules, that relate the state variables. Various features of the vector field correspond to characteristic kinds of behavior. These might, for instance, represent the inherent tendency of the system comportment to move toward relatively stable, persistent, dynamical modes. This particular kind of abstract object in state space is called an ***attractor***. Features such as attractors might be revealed with analytical mathematical methods, or by repeatedly reinitializing the system and mapping the trajectories that result. The latter is an indirect approach to exploring the structure of the system, but, although it is sometimes useful when working with laboratory apparatus or computer simulations, it is mostly impractical for physical biosystems. If extensive series of observational data are available, the simultaneous mapping of data segments that begin with similar values might, however, also serve to reveal the vector field.

A detailed state-space representation of biological phenomena is usually difficult to visualize, simply because a very high-dimensional state space will be required to represent a biosystem's many degrees of freedom. It is often appropriate, therefore, to collapse the state space by using more abstract measures to characterize the collective features of the components. Instead of using a separate measure to denote the water content of each of many components, for example, a single measure might be used to describe the average water content of all of them. This kind of reduction in resolution may very well reveal emergent patterns that are not apparent in the features of the individual components themselves. This is similar to the way in which classical laws of physics, such as those of Newtonian mechanics, describe the behavior of extremely large numbers of molecules. Collectively, these dynamics are easily described with classical mathematics, but resolve into complicated and unpredictable behavior when observed at very fine scales.

It is sometimes useful, as mentioned earlier, to relax the requirement of strict compositional closure and to identify a changing set of components as a system. The resolution of observation can be of significance when characterizing such a compositionally open system. For example, even if the flow of components into and out of such a system were balanced so that the overall composition was qualitatively constant, the state space would not appear to be static if defined with measures of fine resolution. It might, on the other hand, appear to be approximately stable if defined with measures of coarser resolution. Organisms, for instance, are materially and energetically open, and their local composition therefore changes continuously as components are replaced by new inputs and rejected to the environment. This means that if the biosystem were characterized in a high-resolution state space, axes would continually appear and disappear as components were assimilated and rejected. More abstract measures, however, might be used to filter out these immensely numerous local changes and render global patterns more easily detectable. Accordingly, when characterizing a human, a state space might be founded on axes corresponding to measures such as blood pressure, pulse, and electrical potentials in the brain. For an ecosystem, species populations, nutrient concentrations, and the partial pressures of atmospheric gases might be employed. More abstract measures can be used to characterize features at larger scales, and may reveal the presence of multiscale patterns, like autopoiesis, that are not entirely evident at smaller scales.

For system types that are specified as being structurally open, the relationships between the components of a system are considered to be modifiable by outside influences. Of course, structural alterations can also accompany shifts in composition, if a system is open in this latter respect. If a system is characterized by means of a state space, then such changes in structure are reflected by corresponding shifts in the vector field. Generally, a gradual drift in structure corresponds to a gradual change in the vector field; accordingly, features will grow or shrink while the overall topology remains similar (Capra 1996). Occasionally, however, sudden, qualitative shifts called *bifurcations* can occur in the nature of a system as some threshold is crossed. In biosystems, such changes occur, for example, during morphogenesis, wherein gradual development leads to sudden events like cell division. Bifurcations are reflected by catastrophic changes in the vector field. About twenty different types of these have been catalogued, and when they occur, various features in the state space, such as attractors, suddenly appear, disappear, or change dramatically in quality (Capra 1996).

3.7 Comportment and attractors

Attractors are abstract objects that exist in appropriately defined state spaces of systems which are capable of qualitatively robust comportment in the absence of external interactions. For such systems, there exist ranges of initial conditions that result in these characteristic dynamical modes. Trajectories that begin in the regions of state space which correspond to these initial conditions will, without the action of forcing functions, inputs or specific outputs, converge to bounded surfaces. These surfaces are the attractors; the ranges of initial conditions of the convergent trajectories define *basins of attraction*; and a system with such features is called *self-ordering*.

For a self-ordering system that is capable of several independent, persistent modes of comportment, an equal number of attractors will coexist in the state space. These can be envisioned by thinking of the state space as an undulating landscape. The depressions in the landscape then represent basins of attraction and the lowest point of each depression corresponds to an attractor (Figure 3.1). A marble dropped onto the landscape may end up at the bottom of any one of the depressions, which one depending on where it falls initially. External perturbations might subsequently jostle the marble without removing it from a particular depression, or they might be sufficient to knock it from one depression into another, corresponding to a shift between modes of comportment. This illustrates how comportment evolves as the convolution of external and internal influences, the latter being due to structural features which may be represented by attractors.

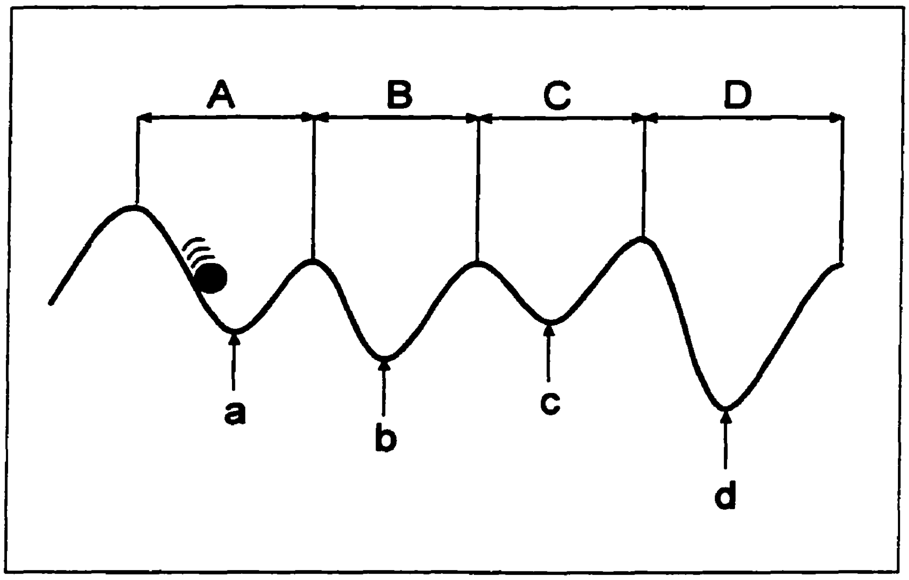

Figure 3.1. A marble rolling on an undulating landscape that models a state space with several basins of attraction, (A, B, C, D...), and corresponding attractors (a, b, c, d...).

Although conveniently described above as depressions in a landscape, attractors are more accurately thought of as surfaces because, although possibly very convoluted, they have a lower dimensionality than the overall state space. The number of points on an attractor, therefore, is infinitely less than the number of points contained in the overall volume of the state space. As a consequence of this, it is extremely unlikely that a particular trajectory will actually lie on an attractor, although there is a remote possibility that a set of initial conditions might coincide exactly with an attractor's surface. In the overwhelming majority of cases, however, a trajectory will only approach the attractor asymptotically. In fact, for a self-ordering system that is subject to external influences the trajectory may be constantly perturbed so that it never actually comes close to settling on an attractor.

Given appropriate measures and an observer who is capable of resolving these in sufficient detail, a system might be seen to display any of a number of different types of comportment, including *stasis*, *periodicity*, and *chaos* (Wolfram 1984). For a system to be considered as strictly self-ordering, however, its type must be specified so that its characteristic modes of comportment result from internal structural constraints alone (Casti 1992). Only when the comportment persists in the absence of external influences (forcing functions, inputs, and outputs) can it be said to tend toward an attractor. The kinds of attractors corresponding to the modes of comportment mentioned above are, respectively, *point attractors*, *limit cycles*, and *strange attractors*. In this paper, autopoiesis is also considered to be a characteristic mode of comportment that may correspond to a special kind of attractor.

As mentioned above, the structural features that are represented by an attractor are likely to significantly affect the comportment of a system, even in the presence of external influences. It is also possible for an open system to demonstrate the aforementioned modes of comportment as a result of the interaction between inherent structural constraints and external influences, even though it may be that no attractor actually exists. Only if the system type is respecified to include these external influences as internal sources and sinks can the comportment be considered as tending toward an attractor. Only if the necessary gradient of order were internalized, for example, could the autopoietic comportment of a biosystem be considered as tending toward an autopoiesis attractor.

The identification of attractors in the state space of a biosystem provides a convenient way of characterizing the biosystem's structure and potential comportment. It must be remembered, however, that the detection of attractors depends not only on the nature of the system, but also on how the system type is specified, on the measures that are chosen to construct the state space, and on the ability of the observer to resolve the measurements. Accordingly, comportment that appears to be qualitatively persistent at one resolution might be revealed, upon examination at a different resolution, to shift between modes. Furthermore, comportment that appears to be of a certain type may, in fact, prove to be of another. For instance, a dynamic that is believed to be chaotic might eventually repeat itself over longer time scales, and thereby prove to be periodic. There is, in fact, some degree of arbitrariness in the identification of any comportment or attractor as being of a given type. As with systems, the specification of types, the identification of particular instances of these, and the formulation of the state spaces in which they can be represented, all depend on the observer. These kinds of considerations are ubiquitous, however, and do not negate the usefulness and necessity of a characterization approach. The taxonomy of comportment and of attractors, therefore, bears closer examination.

The simplest of the dynamical modes mentioned above is static comportment, or the complete absence of change with respect to a given reference frame. It is denoted in state space by a *stationary point*, of which there are three kinds (Scheinerman 1996). First, an *unstable* stationary point represents static comportment that persists only in the complete absence of perturbations, like a marble balanced on the point of a cone. Second, a stationary point might be *neutrally* or *marginally stable*, so that state space trajectories cycle indefinitely around it with frequencies and amplitudes determined by the initial system state. Finally, a *stable stationary point* is one surrounded by a basin of attraction. Stable, static comportment may result from a combination of external influences and internal structural constraints but this kind of comportment corresponds to an attractor only if it converges even in the absence of external influences. If the type of a biosystem is specified appropriately and certain of its features are observed with very coarse and abstract measures, then point attractors might be identified in the vector field. For instance, mammals are homeostatic for very constrained ranges of comportment with respect to measures such as body temperature.

The second kind of comportment discussed here is periodicity. This dynamical mode results from countervailing influences, some of which pull the comportment toward a focus, while others drive it away. In cases where such opposing influences are regular, the comportment may tend toward a cycle, so that the corresponding state space trajectory approaches a closed curve. This curve can be a simple circle, or might be exceedingly complicated (Seborg et al. 1989). If the opposing influences are entirely internal to the system then the curve is an attractor called a *limit cycle*. In this case, when the comportment is moderately perturbed, it will tend to revert to its original periodic behavior, so long as the state of the system is not forced entirely out of the basin of attraction of the limit cycle. If two or more such regular cycles combine in a non-rational frequency ratio, then the resulting comportment is *quasi-periodic*. In this case, the state space trajectory is no longer a closed curve, but spirals around the surface of a toroid without crossing itself (Casti 1992). Although the same state never recurs exactly, in this case the trajectory is, nevertheless, precisely predictable given knowledge of its contributing cycles and initial conditions. Approximate periodicity is often observed for biosystems, examples in organisms being the rhythmic contractions associated with the circulatory and respiratory systems.

System comportment can also be more complicated than periodicity and quasi-periodicity. One example of this is *chaos*, which can arise in simply structured, completely deterministic systems, but which is nevertheless unpredictable in detail beyond the short term. The reason for this unpredictability is a sensitive dependence on arbitrarily small variations in initial conditions. Chaotic comportment diverges exponentially with respect to at least some quantities, and immeasurably tiny fluctuations in these are magnified to such an extent that they have a dramatic impact at macroscopic scales. No model of the system will be sufficiently precise for use in predicting local features of its comportment over the long term. Feigenbaum, one of the key figures in the development of chaos theory, described such comportment as consisting of "Nonrandom complicated motions that exhibit a very rapid growth of errors that, despite perfect determinism, inhibits any pragmatic ability to render accurate long-term prediction" (Peitgen et al. 1992, p.6). Although divergent with respect to some quantities, chaotic comportment is also bounded. In this case also, it is the presence of external influences or internal sources, combined with compensating constraints, like friction in physical systems, that yields such comportment. If both the driving and constraining influences are internal, then the trajectories approach a surface in state space that can be described as a *strange attractor* (Casti 1992).

The apparently repetitive comportment of many natural systems is often simplistically interpreted as being periodic when, in fact, chaos would be a more appropriate model. The climate, for example, is strongly seasonal due to the regular motions of the Earth, Moon, and Sun, but there is also a more complicated component that often is not recognized. Thus, dismay is frequently expressed that there have been three successive winters of very heavy snowfall, for instance, or that the summer was unusually hot, when in fact such unpredictable climatic variations are the norm, and not the exception. The realization that deterministic systems can be locally unpredictable, even in principle, has led to a shift from the formulation of precise, predictive models to the development of more general, explanatory ones. Although chaos precludes the possibility of foreseeing the occurrence of specific events, one can sometimes devise a useful, probabilistic description of such modes of comportment.

Whereas chaos is a useful model, only a limited number of complicated, natural phenomena can be adequately described in such terms. If a system is chaotic then, with the passage of time, the probability becomes uniform that its trajectory will pass through any arbitrarily small segment of the bounded region of state space in which it evolves (Peitgen et al. 1992, p. 554). Many systems, however, display aperiodic comportment in which sophisticated patterns are retained indefinitely and which is therefore more complex than chaos. Like periodicity and chaos, such comportment can only persist in the presence of both driving and constraining influences. If there are constraints but no driving influences, the system will tend toward equilibrium and, contrariwise, it will diverge. If a system does display the aforementioned comportment, and if it is specified so as to include both the driving and constraining influences then, if characterized with appropriate measures, it can be considered as self-ordering. Thus, the vector field representing the system structure will include one or more corresponding attractors. If a biosystem is characterized in this way, its vector field will contain one or more autopoiesis attractors.

3.8 Autopoiesis and the conventional interpretation of life

Autopoiesis has been defined as a special type of comportment in which local interactions among individual components combine to continually renew the overall system of which they are part (Maturana and Varela 1980). This definition might seem to imply that the characterization of a given system as being autopoietic is a binary distinction, but this need not be so; rather, a system may be autopoietic to any degree, and autopoiesis can therefore be regarded as a fuzzy characteristic. Although the degree of autopoiesis of a system is often intuitively evident, there are, as pointed out above, no formal, direct measures with which this characteristic can be evaluated. It is closely related, however, to the order and disorder, complexity, and emergence of a system, and these can therefore be used as indicators of autopoiesis.

Biosystems are generally highly complex and emergent, and as such they occupy an intermediate zone between absolute order and complete disorder with respect to constitution and comportment, as well as state. In terms of their composition, for example, they are not extremely ordered since, at the local level, there is continual change as components enter and leave the system (Maturana and Varela 1980). Neither are significantly autopoietic systems very disordered; for instance, in terms of structure there is always some permanence of pattern. In a dynamic situation, which is a prerequisite for autopoiesis, this requires both the availability of order from a source or external influence, and the continual reduction of disorder through its elimination, either in a sink or by rejection to the surroundings. With respect to complexity, biosystems possess sophisticated patterns of constitution, state, and comportment that can be fully characterized only with extensive descriptions, resulting in high values of such measures. As well, they display highly emergent comportment arising from the way in which their structure relates phenomena across scales. For instance, the features of a rabbit could be compared with those of a trillion independent rabbit cells scattered on the forest floor; the characteristic features of a rabbit obviously depend heavily on the way in which its many cells interact with one another as a coherent system. It can be surmised, therefore, that systems which exhibit intermediate degrees of order and disorder, together with high degrees of complexity and emergence, are likely to be substantially autopoietic.

A single cell, such as a bacterium, is an instance of a relatively simple system type that is generally considered to be alive in the conventional sense. In terms of the perspective developed here, a cell is indeed significantly autopoietic: the metabolic cycles within its cytoplasm form a complicated web of mutually interactive processes, the overall result of which is the renewal of the constitution of the cell (Maturana and Varela 1980). The corresponding comportment is autopoietic, as well as homeostatic for autopoiesis. As discussed above, an inflow of ordered matter and energy (e.g., nutrients) is always required for the renewal of the components of such a dissipative, conservative system, as is a more disordered outflow to the environment. In physical, thermodynamic terms, this order and disorder are manifested as negentropy and entropy that are associated with flows of matter or energy. These environmental interactions are necessary for the cell to maintain itself in a far-from-equilibrium state (Prigogine 1980). Without them, the cell will eventually die as the result of disordering processes such as the generation of heat and metabolic byproducts.

The failure to recognize the dependence of physical, living things on external interactions sometimes leads to the mistaken belief that they spontaneously decrease in entropy and thus violate the Second Law of Thermodynamics. This apparent conundrum is resolved when living creatures are viewed from a broader perspective and are identified as part of a larger system that, on average, tends toward increasingly probable states (Layzer 1988). In fact, cells and larger-scale organisms, if characterized in the way that such biosystems are described conventionally, cannot be considered to be truly "free-living" in themselves, although they are often referred to in this manner. This is because a coherent description of them must include reference to ordered inflows and disordered outflows or, in other words, to their integration into a more extensive system. Thus, their comportment, although perhaps highly autopoietic, does not correspond to an autopoiesis attractor unless the sources of order and sinks for disorder are specified as internal to the cell (or organism). This is, however, not the way that such systems are conventionally characterized.

Associated with cells and multicellular organisms are a number of features that, although they are not central to the definition of autopoiesis, are common to most natural biosystems. For instance, the internal structure of a cell is delineated and protected by a selectively permeable membrane that attenuates external influences. This barrier is generated in the course of the autopoietic comportment of the cell. It isolates the cell from its surroundings sufficiently to safeguard the internal metabolism, yet allows interaction between the cytoplasm and the environment. It may be that the conduct of the cell and the conduct of its neighbours become functions of one another, a situation that Maturana and Varela (1980) have referred to as *coupling*. This can be considered as an extension of autopoiesis beyond the confines of the cell. Maturana and Varela (1980) have asserted that this is, in fact, the essence of communication and cooperation, and underlies the existence of emergent, composite living entities. An extreme example of this is described in the endosymbiont theory, wherein it is proposed that some cell organelles originated from parasitic bacteria which invaded their hosts and then coevolved with them until, today, host and invader are characterized together as a single entity (Margulis and Sagan 1995). At the intercellular level, this kind of cooperative comportment has progressed to the point where coordinated collectives have emerged. Bacterial plaques, tissues, organs, and multicellular organisms are examples of these. Organisms, in turn, couple with their external environment and can thus be considered part of even larger collectives. In all these cases, an extended *autopoietic network* is created that can itself be identified as a system that is alive to some degree. As before, measures can be defined as the basis of a state space in which the comportment of such larger-scale biosystems will display autopoiesis.

Up to this point, biosystems have been described as instances of a special system type having features that can be characterized with measures such as those of order and disorder, complexity, and emergence. Biosystems are neither extremely ordered nor extremely disordered. They are highly complex, meaning that the patterns of their constitution, state, and comportment are relatively difficult to describe. Biosystems are also highly emergent entities whose features are correlated across scales. This multiscalar structure contributes to autopoiesis, such that the components of the biosystem are continually renewed. If the biosystem is characterized appropriately, then autopoiesis corresponds to a bounded surface in state space. Furthermore, if the system is characterized so that the flux of order on which this internal organization depends occurs between internal sources and sinks, then autopoiesis can be described in terms of an attractor. This homeostatic behavior is usually reinforced in natural biosystems by the generation of a proprietary barrier, and is often extended through coupling with the surroundings. In fact, the coupling of biosystems that are in proximity to one another leads to the formation of larger-scale, composite entities whose comportment may, in turn, be characterized as autopoietic.

3.9 A broader perspective of life

The definition of *biosystem* presented above encompasses the widely accepted notion of living things as comprised by the various kingdoms, such as animals and plants. In this article autopoiesis has been presented, however, as the sole defining feature of biosystems and, when framed in this way, the definition accommodates a much broader class of entities, not all of which are necessarily considered to be substantially alive in the conventional sense. Thus, the type *biosystem* may include entities that are: incapable of growth or reproduction; either smaller or larger in scale than a single organism; composed of constituent substances other than organic compounds; derived from the artifacts of humans; or that are virtual, instead of physical, in essence. Some of these systems might be more strongly autopoietic than others, and so a fuzzy interpretation of *biosystem* is appropriate.

The acceptance of autopoiesis as the sole criterion for membership in the set *biosystem* leads to the rejection of a number of other conventional criteria, such as the ability to grow or reproduce, for instance. Although evolutionary dynamics have resulted in natural biosystems that are robustly capable of growth and reproduction under certain circumstances, and albeit that these capabilities may stem from autopoiesis, an entity need not actually grow nor reproduce to be considered substantially alive. A virus, for example, cannot grow, and neither can many adult, multicellular creatures. As well, organisms such as many domestic plants, although certainly deemed to be very much alive, have come to depend on humans for propagation. Moreover, the reproduction of most sexual creatures depends on their being part of a larger-scale, collective organization, just as does autopoiesis itself when an entity is characterized as being conservative and dissipative. Thus, a system may be substantially alive, yet neither grow nor reproduce; it is sufficient that a system demonstrates multiscale, internal processes, such as metabolism, and possibly engages in external interactions, so that it remakes itself.

Scale is another criterion that is not included in the wider definition of life that is presented here. Systems other than those of cellular and multicellular scales can be obviously alive. At the subcellular level, closed catalytic networks have been shown to arise in chemical solutions that contain a suitably large number of interacting molecular species (Kauffman 1993). These networks are not biological organisms in the conventional sense: they do not have distinct boundaries, nor do they reproduce or grow. They need not even involve naturally-occurring organic molecules, nor do the processes that constitute their comportment necessarily resemble the protein and ribonucleic acid metabolism of biological cells. Given a flow of suitable compounds (nutrients), however, the various chemical species in these solutions affect one another in a complex web of interactions so as to constrain their relative proportions. Then, when the concentrations of these species are taken as state variables, the state space trajectories originating from a range of initial values tend toward a bounded surface in the state space. In the face of external forcing functions, such as fluctuations in the composition of the incoming flow, the relative proportions of the chemicals respond in a manner that maintains the catalytic network. Such autocatalytic solutions are therefore autopoietic to some degree and, although some may have a stronger autopoietic tendency than others, all have some membership in the set *biosystem*. These autocatalytic systems are, in fact, thought by some to resemble the prebiotic antecedents of terrestrial life (Margulis and Sagan 1995).

Entities larger than single organisms are also alive to varying degrees as systems. These most often arise from the integration of smaller entities, some of which might be alive in themselves, to form composite biosystems. Although the components might be living things in their own right, their integration often causes the distinction between them to become somewhat arbitrary. As a case in point, slime mold cells spend part of their life cycle as individuals but they may, in response to environmental stimuli, gather together to form a collective entity that resembles a multicellular organism. At a larger scale, social insects like ants, termites, and bees can survive in the long term only as communities, and the structure and comportment of these communities result in the continual renewal of themselves. Here, the individuals are so interdependent that they may be considered as being substantially alive both by themselves and as a population. Extreme interdependence can occur also with populations of different kinds of organisms. For example, the algae and fungi that compose lichens are so tightly associated with one another that they essentially share a single metabolism. Less integrated systems have also been accorded a substantial degree of vitality. For instance, some founding ecologists have described climax communities of vegetation as "super-organisms" (Clements 1916), a perspective that becomes more convincing when it is considered that the roots of trees in groves are often extensively interconnected by networks of fungi. In a marine context, Odum and Odum (1955) have found that the functional interactions among the member populations of a coral reef community result in more effective nutrient use than could be expected of the constituent populations in isolation. This seems to indicate that reefs are emergent entities that are somewhat alive in their own right. Finally, at the extreme, in the Gaia hypothesis all living things on Earth are considered to be part of an extended network in which some subsets of components, such as organisms, have more intimate internal connections than others (Figure 3.2) (Margulis and Sagan 1986). Since mutual links to phenomena beyond the Earth's biosystem are apparently insignificant, this is presently the upper limit to which one can reasonably extend the idea of integration. All of the systems mentioned above qualify as biosystems; their essence arises from the coupling of components which, as discussed previously, is an extension of autopoiesis from the local level upward in scale to encompass the system as a whole.

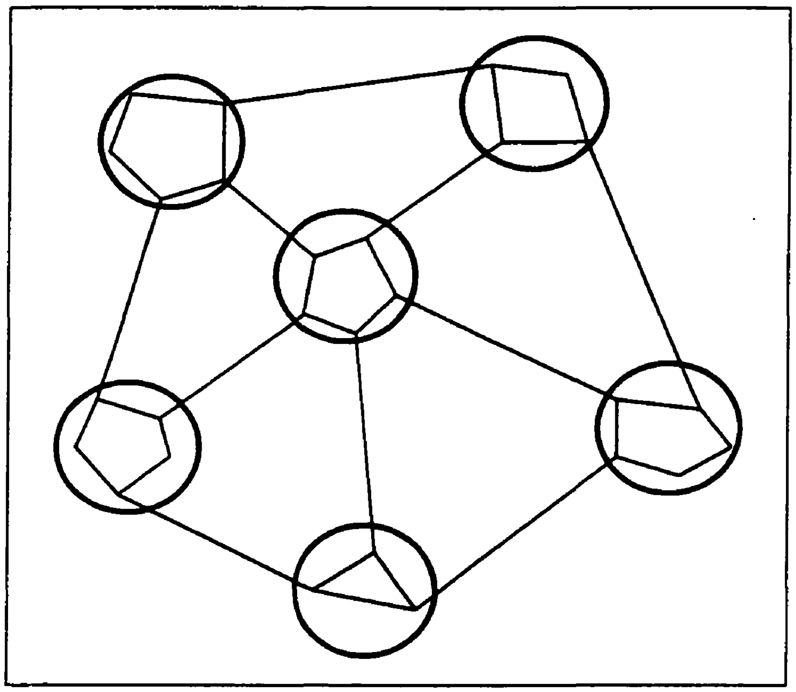

Figure 3.2 An autopoietic network decomposed into smaller networks.

Another criterion that is unnecessary for the definition of life is that the constituent components of a system be of a particular kind. Entities that are conventionally considered to be substantially alive, although they consist primarily of naturally-occurring organic materials, frequently include other kinds of components as well. As an extrapolation of this, it is even conceivable that biosystems could be manifested on a physical substrate completely devoid of organic material. Systems similar to the autocatalytic sets mentioned previously, for example, might be composed entirely of inorganic molecules. At an extremely macroscopic scale, it is speculated by some that terrestrial life has coevolved in conjunction with the atmosphere, hydrosphere and lithosphere of the planet to the point where the entire ensemble can be thought of as part of a single biosystem, of which only a tiny fraction is actually organic (Margulis and Lovelock 1974). Autopoiesis, as illustrated above, is not subject to any limitations with respect to the kinds of components that it might involve.

Just as life need not be limited to systems composed of specific kinds of components, neither do the components need to be of natural origin. In fact, the distinction between natural and artificial is questionable; human technology can be considered as equivalent to the artifacts of other organisms, such as spider webs, bird nests, and the mounds and hives of social insects, all of which are normally considered to be natural (Margulis and Sagan 1995). Many times during the evolution of life on Earth, organisms have become increasingly dependent on such artifacts, to the point that now they are so integrated with them that they are considered to be one and the same. The skeletons of many animals are composed of crystalline calcium compounds and may have evolved from deposits of a previously useless and polluting metabolic byproduct (Margulis and Sagan 1995). This kind of assimilation of artifice is presently occurring with humans; most people wear clothes, for example; some extend their memories by writing and reading; others ride bicycles, wear glasses, have prosthetic hip joints, and drive cars. Some humans are now so tightly integrated with technological components, like pacemakers, that the latter are considered part of the person, and may in fact be vital to their survival. This is to say that their continued autopoietic comportment now depends on these artificial components. Moreover, there is no reason that a system could not be considered alive to some extent even if it were composed entirely of artificial components, a simple example being the autocatalytic networks discussed previously. These autopoietic systems could be composed entirely of chemicals that are not only inorganic, as mentioned, but also of completely artificial origin. They would, nevertheless, exhibit the autopoietic comportment characteristic of significantly alive systems.

Finally, it can be argued that life need not even be confined in its essence exclusively to the physical realm. Minds, economies, societies, and some computer programs can be viewed, for instance, as primarily virtual systems that are alive to some degree. Certainly these are all based at some level on physical substrates, but it is their comportment, not their composition, that is of the essence. For example, the models and simulations that are being developed as part of the aforementioned EcoCyborg Project can be considered, from this viewpoint, as marginally alive (Clark et al. 1999). Societies, too, warrant some degree of membership in the set *biosystem*. They have long been compared with natural organisms in the way they are organized so as to perpetuate themselves, as Thomas Hobbes discussed in his famous work, *Leviathan* (e.g., Hobbes 1982). If one adopts the definition of life that is presented here, these similarities are, however, more than conveniently explanatory. They demonstrate that informational constructs such as computer programs and well-established societies can often be considered as being alive to some degree, and the essence of this vitality is their autopoietic comportment.

It has been argued in this section that life, as characterized by autopoiesis, can occur in a very broad range of system types. Accordingly, the characterization and comparison of systems of such a varied, inclusive class require rather abstract measures. These must be based on variables that are appropriate to the description of the local features of the individual systems, and will also vary according to the interests of the observer. A slime mold might, for instance, be characterized on the basis of the interactions of individual cells; a coral reef on the basis of the interactions of many species of organisms; and the biosphere on the basis of the interactions of entire biogeographical regions. In each case, however, measures can be chosen as the basis of a state space in which the system comportment can be described as more or less autopoietic, and, if the necessary conditions of characterization are met, as tending toward an autopoiesis attractor.

A discipline that is very much contingent upon an adequate understanding of autopoietic networks is the intentional engineering of such systems. It is still in a very early stage of its development, but as the understanding of autopoietic systems improves it may very well become possible to engineer new forms of life. Given a sufficiently broad specification of *biosystem*, such as the one presented here, the variety of these will not be restricted to the conventional conception of life as limited to the sphere of biological organisms. The single criterion by which a system will be identified as alive is its degree of autopoiesis. Although the challenge of manipulating and creating such systems appears to be daunting, related work is already progressing in a variety of fields. With genetic engineering technologies, for example, organisms can now be shaped much more rapidly than through older practices such as selective breeding. Equally dramatic possibilities exist through the combination of biotic and abiotic components. An illustration of this is the cyborging of ecosystems with cybernetic control mechanisms. Modest examples of this already exist in the form of automated greenhouses and similar structures (Linker et al. 1998). Some authors have gone so far as to interpret the development of global telecommunication and computational networks in terms of the cyborging of the entire planet, an exercise that might enhance the biosphere's potential for abstract mentation, which might in turn make it more autopoietic (Dyson 1997; Vernadsky 1945). Even at fairly modest scales, this kind of activity might make the comportment of the biosphere more robust in the face of increasingly disruptive interactions between its components. It may also be possible to engineer biosystems without including any biological components at all. This is the vision of researchers working in the field of *artificial life*, whose creations range from communities of semi-autonomous agents, to inorganic chemical systems, to entirely virtual ecosystems that reside, like the EcoCyborg models and simulations, on digital computers. Every one of these systems, albeit they lack biological or even physical components, can be shown, when appropriately characterized, to exhibit the hallmarks of autopoiesis: intermediate order and disorder, great complexity, and significant emergence.

It is interesting to speculate on the ultimate origin of autopoiesis. There are, of course, myriad hypotheses that address the question of *how* living things first arose, as well as *why* they did so. A partial response to the latter question is the proposal that autopoiesis is the manifestation of a universal tendency toward maximum entropy generation, and that phenomena such as the biosphere arise spontaneously because they are capable of generating a higher entropy flux than would occur in their absence (Swenson and Turvey 1991). If this is so, it might make sense to characterize and compare biosystems specifically in terms of their pattern of entropy production. Authors such as Johnson (1995) have observed that, although some living systems of natural origin do tend to maximize the rate of entropy production per unit energy (this can be seen, conversely, as a tendency toward least energy flow, or "least attainable dissipation"), others seem instead to tend to maximize the throughput of energy ("greatest attainable dissipation"). The former tendency appears to dominate over shorter (evolutionary) time frames and in the presence of strongly cyclical forcing functions, such as those experienced by isolated arctic lakes. On the other hand, the latter tendency seems to dominate over longer time frames and in the presence of less variable forcing functions, as is characteristic of many tropical ecosystems. This area of thought remains richly controversial.

3.10 Summary

The term *biosystem* has been defined in this article as a system that is alive to some degree. The body of the article is an exploration of this definition and its inclusiveness. The definition rests on the characterization of biosystems as instances of a system type that is specified with reference to certain membership criteria, kinds of boundaries, or particular sets of entities. Although these are often specified in a discrete manner, it is acknowledged that a fuzzy approach is also useful in many circumstances. The single necessary and sufficient criterion for life is proposed in this article to be autopoiesis, or comportment in which the components of a system are continually renewed as a result of the overall web of interactions between the components themselves. Having specified a certain system type, particular instances of biosystems can be identified, characterized and compared by means of any number of different measures. No direct measures are currently available for autopoiesis, so related measures must be used instead to evaluate the degree of vitality of a system. Various measures of order and disorder, complexity, and emergence can be used for this. Entities that are significantly alive, or autopoietic, are of moderate order, high complexity, and high emergence. These and other chosen measures might be used as the basis for a state space in which the autopoietic comportment of a biosystem could be represented. Autopoietic comportment must be driven away from equilibrium by internal or external influences, and simultaneously constrained, as the biosystem exploits sources of order and rejects disorder. In the conventional view of life as a conservative, dissipative phenomenon, autopoiesis can only persist as the result of integration into a larger system from which this order is derived, and to which the disorder is rejected. In a different view, the required influences might be included as internal to the biosystem. The trajectories corresponding to autopoiesis could then be considered to tend toward an attractor in an appropriately constructed state space. As the set of all systems in which this occurs to some degree, *biosystem* can include entities that are not necessarily capable of growth or reproduction, that range in scale from the viral to the biospheric, that are partially or completely abiotic, constructed wholly or in part from human artifacts, and that might even be completely virtual in essence.

3.11 Acknowledgments

This work was made possible by the gracious support of the Natural Sciences and Engineering Research Council of Canada, the J.W. McConnell Foundation, and McGill University. The authors also owe their gratitude to the members of the EcoCyborg Project Group of the Department of Agricultural and Biosystems Engineering, McGill University, for discussion of the concepts presented in this article, as well as to M.C. Stacey and P.C. Clark for reviewing initial version of the manuscript.

3.12 References

Broda, E. 1983. *Ludwig Boltzmann: Man, Physicist, Philosopher*. Woodbridge, CT: Ox Bow Press.

Capra, F. 1996. *The Web of Life*. New York, NY: Doubleday.

Casti, J.L. 1992. *Reality Rules*, vol. 2. Toronto, ON: Wiley-Interscience.

Clark, O.G., R. Kok and R. Lacroix. 1999. Mind and autonomy in engineered biosystems. *Engineering Applications of Artificial Intelligence* 12(3):389-399.

Clements, F.E. 1916. *Plant Succession: An Analysis of the Development of Vegetation*. Pub. 242. Washington: Carnegie Institute of Washington.

Crutchfield, J.P. 1994a. The calculi of emergence: computation, dynamics, and induction. *Physica D* 75:11-54.

Crutchfield, J.P. 1994b. Observing complexity and the complexity of observation. In *Inside Versus Outside*, ed. H. Atmanspacher and G.J. Dalenoort, 235-272. Berlin, Germany: Springer-Verlag.

Dyson, G. 1997. *Darwin Among the Machines*. Redwood City, CA: Addison-Wesley.

Gell-Mann, M., and S. Lloyd. 1996. Information measures, effective complexity, and total information. *Complexity* 2:44-52.

Gunther, R., B. Schapiro and P. Wagner. 1994. Complex systems, complexity measures, grammars and model-inferring. *Chaos, Solitons & Fractals* 4:635-651.

Hobbes, T., and C. B. MacPherson (eds.). 1982. *Leviathan*. New York, NY: Penguin Classics

Johnson, L. 1995. The far-from-equilibrium ecological hinterlands. In *Complex Ecology*, ed. B.C. Patten, S. E. Jorgensen and S. I. Auerbach, 51-103. Englewood Cliffs, NJ: Prentice Hall.

Kauffman, S.A. 1993. *The Origins of Order*. New York, NY: Oxford University Press.

Kosko, B. 1993. *Fuzzy Thinking: the New Science of Fuzzy Logic*. New York, NY: Hyperion.

Langton, C.G. 1990. Computation at the edge of chaos: phase transitions and emergent computation. *Physica D* 42:12-37.

Layzer, D. 1988. Growth of order in the universe. In *Entropy, Information, and Evolution*, ed. B.H. Weber, D.J. Depew and J.D. Smith. Cambridge, MA: MIT Press.

Li, W. 1997. The complexity of DNA. *Complexity* 3:33-37.

Linker, R., I. Seginer and P.O. Gutman. 1998. Optimal CO_2 control in a greenhouse modeled with neural networks. *Computers and Electronics in Agriculture* 19:289-310.

Logofet, D.O. 1993. *Graphs and Matrices: Stability Problems in Ecosystem Models*. Boca Raton, FL: CRC Press.

Margulis, L., and J.E. Lovelock. 1974. Biological modulation of the Earth's atmosphere. *Icarus* 21:471-489.

Margulis, L., and D. Sagan. 1995. *What is Life?* New York, NY: Simon and Schuster.

Margulis, L., and D. Sagan. 1986. *Microcosmos: Four Billion Years of Evolution from Our Microbial Ancestors.* New York, NY: Simon and Schuster.

Maturana, H.R., and F.J. Varela. 1980. *Autopoiesis and Cognition: the Realization of the Living.* Dordrecht, the Netherlands: D. Reidel Publishing Company.

Odum, H.T., and E.P. Odum. 1955. Trophic structure and productivity of a windward coral reef community on Entiwetok Atoll. *Ecological Monographs* 25:291-320.

Peitgen, H.-O., H. Jurgens and D. Saupe. 1992. *Chaos and Fractals: New Frontiers of Science.* New York, NY: Springer-Verlag.

Prigogine, I. 1980. *From Being to Becoming: Time and Complexity in the Physical Sciences.* San Francisco, CA: Freeman.

Scheinerman, E.R. 1996. *Invitation to Dynamical Systems.* Upper Saddle River, NJ: Prentice Hall.

Schrödinger, E. 1955. *What is Life? The Physical Aspect of the Living Cell.* Cambridge, UK: Cambridge University Press.

Seborg, D.E., T.F. Edgar and D.A. Mellichamp. 1989. *Process Dynamics and Control.* New York, NY: Wiley.

Shannon, C.E. 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27(July-October, 1948): 379-423, 623-656.

Silvert, W. 1995. Complexity. *Journal of Biological Systems* 4:585-591.

Stewart, I. 1995. *Nature's Numbers.* New York, NY: BasicBooks.

Swenson, R., and M.T. Turvey. 1991. Thermodynamic reasons for perception-action cycles. *Ecological Psychology* 3(4):317-348.

Vernadsky, W.I. 1945. The biosphere and the nöosphere. *American Scientist* 33:1-12.

Wackerbauer, R., A. Witt, H. Atmanspacher, J. Kurths and H. Scheingraber. 1994. A comparative classification of complexity measures. *Chaos, Solitons & Fractals* 4:133-173.

Wolfram, S. 1984. Universality and complexity in cellular automata. *Physica D* 10:1-35.

CONNECTING TEXT

Chapter 4 has been published as: **Clark, O.G., R. Kok, and R. Lacroix. 1999. Mind and autonomy in engineered biosystems.** ***Engineering Applications of Artificial Intelligence*** **12(3):389-399**, (Copyright © 1999 Elsevier Science Ltd.).

In this chapter, *autonomy* and related concepts are explored in some detail. Autonomy is one of the two principal system characteristics that are examined in this thesis, the other being aliveness, or vitality, which was introduced in Chapter 3. Hence, in this chapter, the second major conceptual stream of the thesis is established. A lexicon related to the concept of system autonomy is defined and elaborated in a very general sense, although it is illustrated with reference specifically to cyborged ecosystems.

CHAPTER 4. MIND AND AUTONOMY IN ENGINEERED BIOSYSTEMS

Abstract

Biosystems are unitary entities that are alive to some degree as a system. They occur at scales ranging from the molecular to the biospheric, and can be of natural, artificial or combined origin. The engineering of biosystems involves one or more of the activities of design, construction, operation, maintenance, repair, and upgrading. Engineering is usually done in order to achieve certain preconceived objectives by ensuring that the resultant systems possess particular features. This article concerns the engineering of biosystems so that they will be somewhat autonomous, or able to pursue their own goals in a dynamic environment. Central themes include: the computational abilities of a system; the virtual machinery, such as algorithms, that underlie these abilities (mind); and the actual computation that is performed (mentation). A significantly autonomous biosystem must be engineered to possess particular sets of computational abilities (faculties). These must be of sufficient sophistication (intelligence) to support the maintenance and use of a self-referencing internal model (consciousness), thereby increasing the potential for autonomy. Examples refer primarily to engineered ecosystems combined with technological control networks (ecocyborgs). The discussion is focused on clear working definitions of these concepts, and their integration into a coherent lexicon, which has been lacking until now, and the exposition of an accompanying philosophy that is relevant to the engineering of the virtual aspects of biosystems.

4.1 Introduction

This paper comprises a philosophical and lexical basis for engineering the minds of highly autonomous biosystems. Biosystems are collections of physical and virtual components that perform together as integrated, living units. They range in organizational scale from the molecular to the biospheric and, as well, vary greatly in their degree of autonomy. The discussion presented here is general, but is illustrated with reference to a particular kind of biosystem called an ecocyborg (Parrott et al. 1996). Ecocyborgs consist of both biological and technological components that interact at the scale of an ecosystem, where the latter is defined as a community of organisms, together with their abiotic surroundings. Biosystems of this type can be engineered for a variety of purposes, which may be best served by tailoring their computational abilities; i.e., their capacity to transform input signals from their surroundings into outputs signals.

Currently, ecocyborgs are usually artificial in origin, or are derived from natural ecosystems by human modification. Humans have historically modified ecosystems to favor their own survival, and this has in part allowed them to expand their range outside of the ancestral environment to which they are evolutionarily adapted. They have accomplished this by introducing and extirpating species, supplementing soil nutrients, and altering the hydrological properties of watersheds, for instance. Such activities form the basis of modern industries like agriculture, aquaculture, and silviculture. Insofar as these practices involve the modification of biosystems in pursuit of particular objectives, they can be considered as examples of biosystems engineering at the ecosystem scale.

The modification of ecosystems, as it has been practiced, is a primary reason for the rapid growth of the human population that has occurred during recent millennia. Human activities are, in turn, altering the Earth's ecosystems more rapidly and on a larger scale than ever before. The magnitude of these alterations is such that ignorance or carelessness could potentially affect the integrity of the biosphere. The changes that could result from further human activity should therefore be carefully considered, as should the ongoing impact of changes that have already taken place. The modification of ecosystems on such vast scales must proceed with attention to design, construction, operation, maintenance, and repair considerations, aspects of engineering practice that have until now been largely neglected when dealing with biosystems of this class.

In the short term, biosystems engineering principles could be used to moderate environmental crises by making ecosystems more *autonomous*, or independent in the establishment and pursuit of their own goals. This would increase their persistence in the face of external perturbations or the self-serving activities of component species such as humans. The idea of engineering ecosystems in this way is new, and until now has not been framed in the context of biosystems. It implies modifying their computational abilities, or the manner in which the pattern of interactions between their components transforms input signals into outputs. All biosystems have some ability for computation, but natural ecosystems are incapable of the abstract mentation necessary for significant autonomy. Ecosystems might, however, be endowed with the required abilities by transforming them into ecocyborgs through the addition of technological components. An ecocyborg could be engineered to have computational abilities of the appropriate type and sophistication for *consciousness*, meaning that it would be aware of itself to some degree in the context of its environment. This in turn would increase its potential for autonomy. This approach, and the lexicon that is developed here, could prove to be valuable in the engineering and sustainable management of Earth's ecosystems.

In the long term, the engineering of biosystems at the ecosystem scale not only could help to safeguard against environmental crises, but might also provide for the continued growth and survival of the human species. Expansion into space, for example, will be necessary if humanity is to continue to increase, simply because the Earth's finite resources cannot sustain perpetual growth. Moreover, planet-bound life is vulnerable in the face of planetary events such as collisions between asteroids and the Earth (Sagan 1994). The establishment of self-sustaining colonies in space would provide practically limitless room for growth, and would better ensure the security of the species. Since people can only exist in an appropriate environment, extraterrestrial expansion will require the creation of artificial ecosystems that include humans. These will undoubtedly include many technological components, making them ecocyborgs. Moreover, since they will have to be self-sustaining in the isolation of space, they will have to be engineered to be highly autonomous. The survival of space-borne ecosystems would be more secure if their autonomy were independent of humans, since the ecocyborgs would then be able to function even if human guidance became impossible or ineffective. This might occur if the occupants were incapacitated or neglectful; it is also entirely possible that such ecocyborgs would simply be too large and complicated to be effectively controlled entirely by humans. The International Space Station (ISS), of which construction began in 1998, is an example of such a space-bound ecocyborg. The philosophy and lexicon presented in this article could be useful conceptual aids to engineering the ISS and its successors as viable, integrated, goal-oriented biosystems that include humans as components.

The cyborging of ecosystems illustrates how one class of biosystems might be engineered to be highly autonomous, but many of the concepts related to such an exercise are also relevant to the engineering of a much broader class of biosystems. Until now, these concepts have not been clearly defined as part of a coherent and useful lexicon such as the one presented here. In this paper, each concept is first discussed in broad terms, and then illustrated with examples. Frequent reference is made to animals, especially humans, since they are the most accessible and intensively studied autonomous entities. The ideas are then expounded in the specific context of ecocyborgs, and integrated into a conceptual framework that facilitates the engineering of these and other kinds of significantly autonomous systems. Although the framework presented here is loosely based on traditional human psychology, it is certainly not the only approach that might be appropriate. Since large-scale biosystems such as ecocyborgs are often composed of semi-independent agents, a serviceable framework might also be developed, for instance, from a sociological perspective.

4.2 Implementing mind in biosystems

In this paper, the *mind* of a biosystem is defined as the virtual machinery, including algorithms, that make possible all of its computational abilities. All biosystems possess some computational abilities, but these abilities, the virtual machinery that gives rise to them, and the physical substrates in which that machinery is embodied, can all differ greatly from one biosystem to another. Humans, for instance, possess a nervous system comprising specialized organs that embody highly adapted virtual machinery. This machinery gives rise to very specific computational abilities that make possible some degree of consciousness and autonomy. Natural ecosystems do not have such specialized structures, and so do not possess the kind of minds that humans do. Instead, their computational abilities reside in the way that input signals are transformed into outputs through processes such as interactions between the populations of their constituent species, the cycling of nutrients, and subtle phenomena like the transport of biologically active trace chemicals (McNaughton and Coughenour 1981; Patten and Odum 1981). The cumulative result of these processes is certainly computationally complex, but it does not make natural ecosystems conscious or autonomous in the sense that a human is. The virtual machines corresponding to these processes are more analogous to those embodied in the workings of the human digestive and circulatory systems than to those of the nervous system. Thus, this virtual machinery all gives rise to computational abilities, but is not generally considered to contribute to the capability of natural ecosystems to model or reason about themselves in the context of their surroundings (Engelberg and Boyarsky 1979). They cannot, therefore, establish and work towards their own goals. They can, however, serve as a basis for engineered biosystems of greater consciousness and autonomy.

Biosystems can be engineered to have minds similar to those of humans. Ecosystems, for instance, might be endowed with an infrastructure to support the computational abilities required for high degrees of consciousness and autonomy. This can be done by including components that are not native to a naturally occurring ecosystem. The resulting comportment is then a consequence of both the inherent dynamics of the natural ecosystem segment and the influence of the additional computational components. If the latter are added to a biosystem with the express intent of regulating its comportment so as to achieve particular goals, then the exercise is one of guidance or control. Control can be intrinsic or extrinsic, depending on the conceptual boundaries that are defined. If the guidance components are considered to be internal to the biosystem, then the control is *intrinsic*, whereas if they are considered to exercise a controlling influence from outside the system boundary, then the control is *extrinsic*. Components called *perceptors* sense signals in their surroundings, and create information corresponding to values of the *observed* variables. *Control mechanisms* structure this information and devise strategies to keep the values of certain *controlled* variables within a particular range. Lastly, *effectors* implement these strategies by parsing them into the values of *manipulated* variables, or directives that induce final control elements to generate output signals (Kok and Lacroix 1993). In expansive systems these components are often arranged in distributed networks, being widely separated in space but still linked together by communications channels so as to influence each other's activities (Kok and Lacroix 1993).

The control of large-scale biosystems can be illustrated with the example of human intervention in ecosystems. As discussed previously, humans habitually exercise control over ecosystems in order to improve their own welfare in the short term. It is believed that modern humans are relative newcomers to most parts of the world, having spread from the African continent only during the last two-hundred-thousand years or less (Vigilant et al. 1991). They have inserted themselves into a variety of ecosystems to which they were not originally native, and now regulate these in order to meet their own needs. Humans thereby guide the ecosystems by acting as perceptors, control mechanisms, and effectors. Whether this guidance is considered to be either intrinsic or extrinsic depends on whether or not the humans are included in the ecosystem definition.

Ecosystems might also be engineered to be highly autonomous by cyborging them with technological devices. Computer control systems already endow some greenhouses and industrial fermentation units, for instance, with a modest degree of autonomy. This approach might also be applied to other ecosystems, with machinery replacing human muscles as effectors, electronic instrumentation performing sensory tasks, and computers acting as control mechanisms. These technological components would endow the ecosystems with minds that would increase the independence of their comportment, enabling them to guide themselves toward particular goals. In the future, cyborged ecosystems might serve as habitats in the human colonization of space, and the entire biosphere of the Earth might someday gain greater autonomy through cyborging with sensory and communications networks (Dyson 1997).

Two general examples have been given of how the minds of biosystems might be engineered by including some kind of control system. Many other methods might be described that apply to different kinds of biosystems, and the advent of new technologies and novel applications in the future will make possible the creation of biosystems that cannot be foreseen today. It is therefore important to be able to discuss the concepts associated with mind in biosystems in a manner that is relatively context-independent. The adoption of several complementary perspectives on the computational essence of mind can facilitate this objectivity.

4.3 Perspectives on computation

Three alternative perspectives are presented here that can be employed when discussing computation in biosystems (Figure 4.1). The first refers to the *virtual machinery* that endows a biosystem with its computational abilities. This is the *mind* of the biosystem, and includes formal methods such as algorithms, although it may not be limited to these. The mind is referred to as being *virtual* because it creates, communicates, and manipulates information, but it must, nevertheless, reside on a physical substrate such as a brain or a landscape. This substrate might affect the performance of the mind, but in theory this is of only incidental importance. For example, an abacus, a Babbage engine and an electronic calculator can all potentially be used to perform the same mathematical operations. Although the speed of the operation might vary according to the instrument, the formal methods that are used can be qualitatively the same in each case. In an ecological context, for example, a rainfall event in a watershed might be transformed into discharge. The transfer function that relates the rainfall and the discharge might be the same for two different watersheds, but be mediated by different physical structures. In a completely natural ecosystem the transfer function might depend on topology and the hydrological characteristics of the soil, whereas in an ecocyborg it might result from the actions of a computerized network of drainage canals and hydraulic control structures that might be considered as extrinsic to the ecosystem. What is the same in both cases is the formal method, or virtual machine, that generates the output from the input.

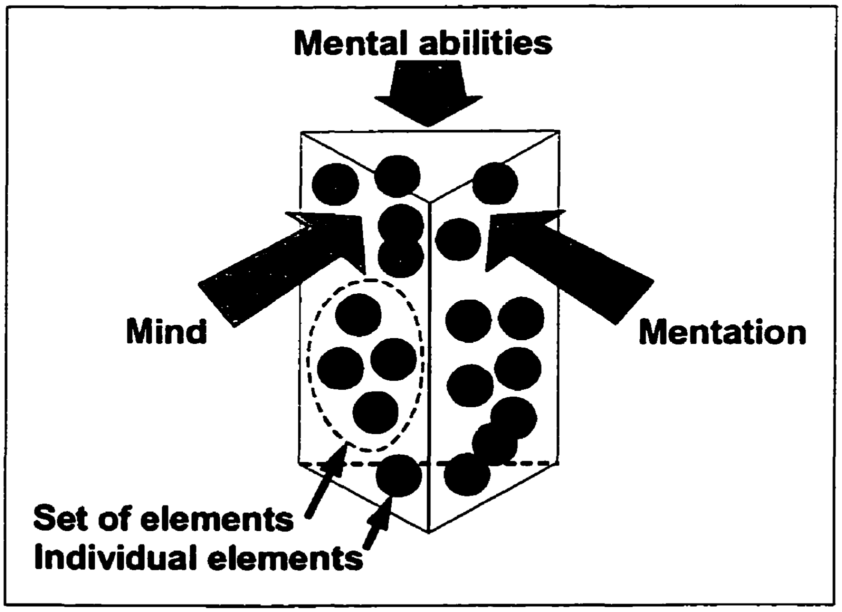

Figure 4.1 Three-fold perspective of a computational entity.

Virtual machinery can be grouped into sets and supersets on the basis of the functions to which they give rise. Instructions, for example, are the most basic embodiment of virtual machinery in the context of the digital computers that constitute the computational infrastructure of current ecocyborgs (e.g., automated greenhouses). Instructions can be grouped into subroutines, and the subroutines into programs that can perform particular tasks. The boundaries of these sets and supersets are, however, arbitrary and can overlap. The same subroutine, for instance, might be used in several different programs. The virtual machinery of future ecocyborgs might be organized less like the linearly structured program code that is currently common, and more like natural biological mechanisms. As a case in point, artificial neural networks already exist that are modeled after biological nervous systems. As well, evolutionary programming techniques have been developed, based on the principles of natural selection, and are used to create virtual entities that are specialized for a particular task. Sets of these sorts of virtual machines might be more appropriately referred to as communities and populations, rather than as subroutines and programs.

The second of the three perspectives discussed here refers to the *computational abilities* of a biosystem. These arise from the operation of the virtual machinery described previously, and can be envisioned as forming an epistemic space of potential computational activities to which the mind is limited. Like virtual machines, computational abilities can be grouped by function into sets. Many researchers have proposed lists of candidate sets, or *faculties*, in order to delineate the mental architecture of naturally-occurring intelligent entities such as humans and other animals (Pinker 1994; Gardner 1993; Goldman 1986). A similar taxonomy is proposed for the faculties of ecocyborgs, and is discussed later in this paper.

Finally, computation in biosystems can also be characterized by the information-processing activities that are actually performed. This movement through the space of potential computation is the dynamical manifestation of the computational abilities of a biosystem, and is referred to as *mentation*. (This is a general term that describes the activities of any biosystem; the term *thought* is used with reference to humans and similar animals.) Mentation can differ greatly between individual biosystems, in accordance with their goals, constraints, and unique experiences, even though their minds and computational abilities might be similar. For instance, two identical greenhouses might maintain entirely different internal climates in order to grow different species of plants.

The perspectives described here are useful when comparing the computational characteristics of biosystems that might be radically different in their physical structure and in their histories of experience and mentation. For instance, two ecocyborgs might differ enormously in their structure, the computational abilities of the first being based largely on virtual mechanisms that are intrinsic to its ecosystem component, and extrinsic components forming the foundation for the mind of the second. The two entities could nevertheless have the same capacity to regulate their internal temperature in the face of climatic fluctuation. In the first ecocyborg the temperature regulation might be mediated by the thermal mass of a pond, whereas in the second this might be accomplished by a technological control network including thermostats, digital controllers, and propane heaters that are extrinsic to the ecosystem.

4.4 Intelligence

Intelligence measures are useful for comparing the computational abilities of different biosystems, and a variety of intelligence indices have been devised for use in various applications. In the past, for example, the mental ability of a human has often been viewed as a cohesive phenomenon, and has been characterized accordingly with a single-valued Intelligence Quotient. This is more informative than a binary distinction between *intelligent* and *not intelligent*, but an even more detailed description can be provided by evaluating a number of characteristics on continuous scales and then collecting their values into a vector. Strengths and weaknesses can then be compared among different biosystems if the scales are calibrated with standard points. Minsky (1985) suggested a scale of intelligence normalized in this way, for instance, with the mental ability of an average human defined as unity (Figure 4.2). The adoption of such a scheme would be useful in the engineering of ecocyborgs with particular computational abilities, such as those required for autonomy. One basis for such a vectorized intelligence measure is the grouping of computational abilities into faculties. Accordingly, a set of faculties is proposed below for the particular case of ecocyborgs.

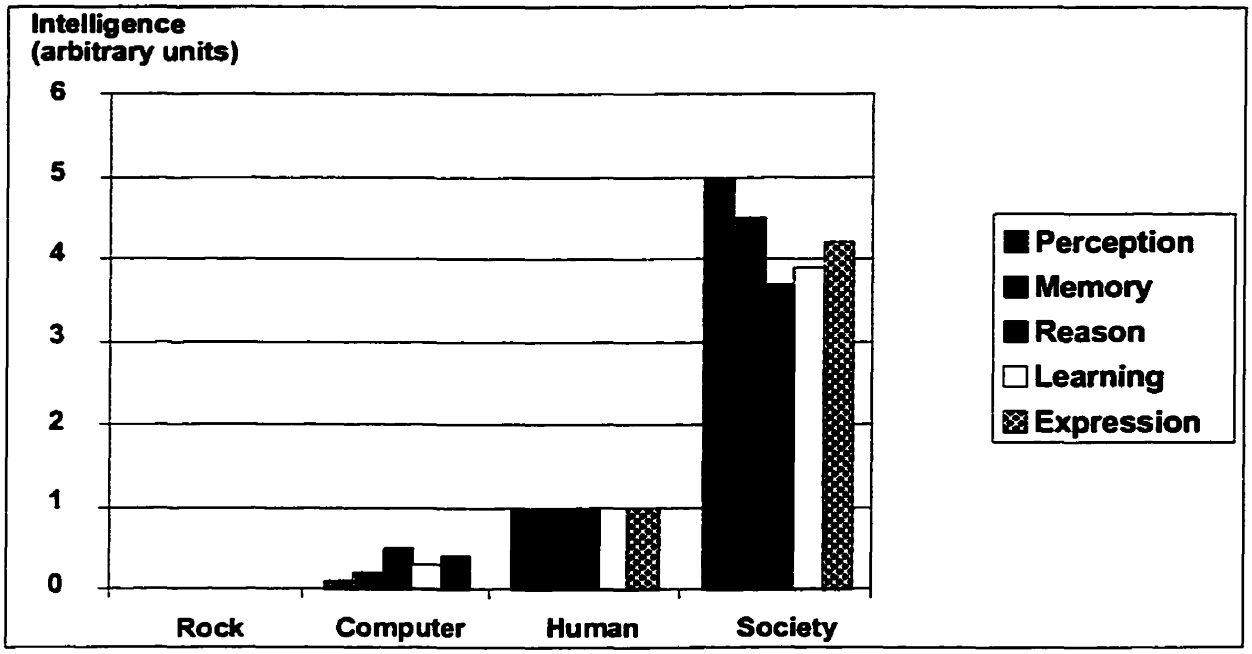

Figure 4.2 Vectorized measures of intelligence, using an average human as the standard.

4.5 Mental faculties

In an extreme interpretation, the whole causal network that connects input with output can be considered as one, unified transfer function. Alternatively, an interpretation can be employed that distinguishes between types of computational abilities. Such a scheme inevitably results in indistinct categories that overlap to a degree, since in any taxonomy the manner in which computational abilities are grouped together is somewhat arbitrary. Some taxonomy must nevertheless be imposed in order to proceed with an analysis. Here, a scheme is presented that categorizes computational abilities into five groups: the faculties of perception, reason, memory, learning, and expression (Figure 4.3). For each faculty there is a general discussion, which is illustrated with reference first to animals and then to ecocyborgs. The mental faculties of an ecocyborg might arise from either biological or technological components that could be either intrinsic or extrinsic to the ecosystem.

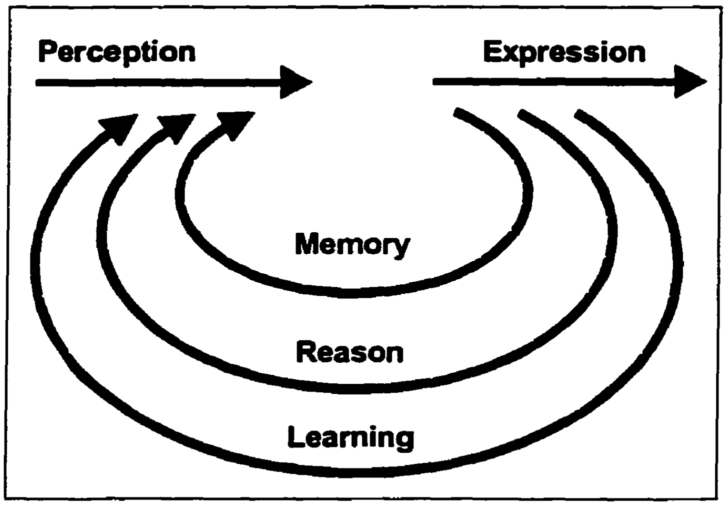

Figure 4.3 The five mental faculties of an intelligent system.

4.5.1 Perception

The faculty of *perception* encompasses the ability to create information from signals. These signals may be of external origin, but if the biosystem is capable of self-observation then some may also originate internally. The abilities included in this faculty arise partly from virtual machinery embodied in an array of sensory devices (*perceptors*). The physical embodiment of the perceptors is of special relevance, since they form the interface with the physical surroundings. As well as creating information from incoming signals, the virtual machinery also transforms it so that it is accessible to other parts of the mind. If the information created by certain perceptors is always structured in a particular way, then the associated virtual machinery might be highly optimized for the specific tasks that are involved, as reflected by the intransigence of the physical substrate. Flexibility is sacrificed in this case, since the specialized configuration that results serves as a base-level filter for the information that is created.

In biological organisms such as mammals, the perception of external signals depends largely on massively parallel arrays of specialized sensory neurons in the epidermis, like the retinas of the eyes, the cochlea of the ears, and the olfactory buds in the nose and mouth. Specialized sensory neurons throughout the body also interpret internal signals. Highly adapted computational abilities are associated with each of these specialized arrays, which create information based on particular kinds of input signals. These kinds of abilities can dramatically impact the whole physiology and mode of existence of an entity. Bats, for example, have evolved to be extremely dependent on their ability for acoustic imaging, and the physiology of temperate plants is centered on the way that this type of vegetation perceives sunlight and seasonal changes in the environment.

In an ecocyborg, the ecosystem segment would have the inherent ability to perceive and respond to signals like solar radiation flux, rainfall, and the partial pressures of atmospheric gases. Technological mechanisms could also track these variables, as well as others that would not normally be perceived by a natural ecosystem, such as the unit cost of heating fuel. In the case of ecocyborgs with extensive ecosystem segments (intended for human habitation, for instance), it seems appropriate that any technological perceptor arrays should be massively parallel and highly distributed. This would result in the generation of large amounts of information, but because an ecocyborg would likely be immobile or primarily focused on managing its internal state, the task of perception would be somewhat simplified as compared to the case of an animal. Perception could be simplified even further if the internal sensors were immutable and immobile with respect to the rest of the biosystem.

4.5.2 Memory

Memory includes all of the abilities required to index, retain, and retrieve information. This can be interpreted as the ability to create or perceive patterns in information, or to create deeper semantic structures based on information generated through the faculty of perception. When new information is acquired, it is subsumed into the mind so that the structure of the constituent virtual machinery is contingent on the history of its mentation. This retained information is indexed by detecting any similarities to previous information. These relationships are made explicit through the creation of links between informational constructs, or equivalently, by grouping the constructs. This process is equivalent to the creation of a semantically deeper layer of information that can be described as *meta-information*. The indexing process can be iterated to create a richly structured network. The associative patterns within the network then serve to index the information and recall it in the appropriate context. Memory is therefore dependent on the capacity to detect, create, and compare patterns.

As with all mental phenomena, memory in animals arises from virtual machinery whose functioning corresponds to the physical activity of neurons. Animals with more developed nervous systems have correspondingly sophisticated memories that appear to correspond to the synchronized firing of many neurons (Greenfield 1995). Patterns of relationships in retained information, i.e., associated memories, might correspond to the firing of subsets of neurons that are shared among various synchronized populations. Because of the vast numbers of neurons involved, it is possible to represent relatively large informational structures. The physiology of animal nervous systems has inspired the creation of similarly structured artificial neural networks. These have proven to be eminently capable of retaining, processing, and recalling patterns of information such as those that might be created by a biological sensory array.

In order to be significantly intelligent, ecocyborgs must retain, structure, and recall large amounts of information, just as animals do. The manner in which the required pattern-processing abilities will be implemented in ecocyborgs will depend on the underlying virtual machinery and the corresponding physical substrate in which it is embodied. Biological systems demonstrate an approach that involves massively parallel networks of information storage devices. In artificial systems, these devices might be packaged in a single structure, such as a silicon computer chip, but their basis will be ultimately reducible to large numbers of distinct components such as transistors. In order to support the required virtual machinery, these must be able to change state, and it should be possible to make their state dependent on that of other devices, so that they can be used to encode the sophisticated networks of information described previously. Finally, in order for this information to be kept current and accessible, there must be an interface with the other mental faculties of the ecocyborg.

4.5.3 Reason

It is speculated that increased autonomy improves the viability of an entity by heightening its ability to respond independently to an unpredictable environment. This implies the flexible and sophisticated formulation of appropriate responses to unforeseen stimuli. *Reason* is the faculty that encompasses the computational abilities required for this. It is bracketed by perception and expression, the faculties by which signals are translated into information and vice versa. Reasoning transforms the pool of information retained in the mind into mental products that potentially have an impact on the surrounding environment, or on the internal structure of the entity itself. These mental products include judgments, decisions, inferences, conclusions, and solutions to problems.

Human reason is the epitome of flexibility and sophistication, as evidenced by the great variety of behavior that it engenders. It is therefore difficult to completely catalog the abilities that it comprises, and it often seems that new ones become apparent in every scenario. A number of qualities can be used to characterize these abilities, corresponding qualities being definable to characterize the virtual machinery from which the abilities arise and the mentation that they enable. Of these, the qualities of mentation are the most commonly referred to. *Depth* and *breadth* are two of these: depth refers to the length of the chain of mediating events leading from inputs to outputs, whereas breadth indicates the number of alternative paths that are explored. Thus, reasoning might be fairly narrow and shallow, or it might simultaneously involve a vast array of different mechanisms in parallel and/or in series, each influencing the outcome to some degree. In the former case, the reasoning process might be sufficiently transparent so that the mentating system itself can observe, understand, and explain it; in the latter, it might be so diffuse and convoluted, with various branches reinforcing and inhibiting one another, that the process becomes entirely intractable. This is often referred to as *intuitive* reasoning. Reasoning may also be either *deductive* or *inductive*. Deduction moves from general premises to logical conclusions, and is supported by theoretical understanding, whereas induction is the inference of general principles from particular instances and relies on experience. Overall, the relationships between inputs and outputs can be extremely complicated, with many inputs taken into consideration and the activities of various reasoning mechanisms interacting with one another. The end result is often uncertain and multivalent.

If an ecocyborg is to have a high degree of autonomy, its mind must possess a wide variety of reasoning mechanisms that can interact flexibly with one another. The faculty of reason should therefore be composed of many semi-independent abilities that arise from such mechanisms, a scheme that is similar to some current interpretations of how the human mind functions (Pinker 1997; Minsky 1985). Each of these abilities could involve a different combination of the qualities mentioned above. The virtual machine that gives rise to each could operate on different kinds of information that might originate externally or be generated by other virtual mechanisms. The activity of this machinery might modify the internal state of the ecocyborg through the creation of mental products such as those mentioned above, and some of these mental products could also stimulate the faculty of expression to radiate signals into the surrounding environment.

4.5.4 Expression

The faculty of *expression* is the complement of perception. It encompasses the computational abilities required for the transposition of mental products into output signals. These signals can propagate outward to affect the external surroundings, or they can influence the internal state of the system. In a physical context, this involves the manipulation of material objects, whereas in a virtual setting it entails the manipulation of information, and can also include communication with other entities. As with perception, there can be one or more adjunct abilities permanently associated with each effector to enable the rapid and effective execution of habitual tasks, such as the parsing of directives intended for the effector.

As with the faculty of perception, some of the virtual machinery that underlies expressive ability forms an interface between the mind and the physical world, and so the physical embodiment of these virtual machines is again of particular relevance. In animals, effectors that impact the external surroundings are generally fewer in number and more localized than the vast arrays of perceptors described earlier. This is perhaps due to the tendency of a signal to disperse as it radiates from its source through an unconfined environment. The bulk of many animals is, nevertheless, made up of effectors and associated devices, through which physical signals are generated. For instance, the arms and legs of a human constitute effectors that interface with the external environment. There are also effector arrays, such as the peristaltic musculature, that influence the internal state of the body. Other expressive abilities, however, are oriented more toward the virtual rather than the physical realm, and so are not necessarily as directly dependent on the configuration of the material substrate in which they might be embodied.

As mentioned, most future ecocyborgs will probably be immobile, and therefore will not require the kinds of effectors that animals need for locomotion. External effectors will more likely be associated with activities such as maintaining a selectively permeable barrier between the ecocyborg and its surroundings, and with virtually-oriented tasks such as communication. Following the biological pattern, the internal effectors of an ecocyborg should be of a parallel and distributed nature, so that effects can be visited upon the entire extent of the system. Their type could vary greatly, depending on the nature of the ecocyborg; if it included a large ecosystem segment, the internal effectors could be as diverse as irrigation networks, air conditioning systems, or troops of pruning robots.

4.5.5 Learning

Learning includes the abilities that enable a mind to restructure itself adaptively. The idea of adaptation implies the improvement of performance, or increased viability in a particular context. Effective learning makes the mind of a biosystem more adept at interpreting the stimuli it encounters, and at responding in a manner that has favorable results. This requires that the biosystem be able to recognize in perceived information patterns that correspond to frequently encountered and exceptionally important environmental situations. The biosystem must also be capable of identifying associated patterns of mental activity that result in desirable outcomes in particular circumstances, and of generating new ones if the old ones are ineffective. In learning, important patterns are retained so that they can be quickly identified (in the case of perceived patterns) or reproduced (in the case of mental activities). The effectiveness of learning therefore depends on the ability to acquire or create new patterns and to retain those that are most useful. In a stable environment, this should make a biosystem increasingly successful, by whatever means this is measured. A changing environment could, however, require that the biosystem continuously restructure itself in order to deal with new situations. Depending on how challenging the environment is, a biosystem might not be able to keep pace, and it might become relatively less suited to its surroundings. There is more of an advantage if the faculty of learning is recursive, and can operate on itself to acquire better ways of learning. In a highly unpredictable environment, therefore, the autonomy of a biosystem is very dependent on its ability to learn, and on its ability to learn about learning.

Of all the biosystems that have been observed, humans are probably the most effective and versatile learners. Their ability to adapt to a wide variety of different environments is evidence of this. As suggested, the human faculty of learning encompasses the ability to adapt to significant environmental scenarios, and to determine which new scenarios are, in fact, significant. Humans can also reproduce courses of action that were successful in past circumstances, improve upon past actions, and, if necessary, even formulate entirely new strategies. Finally, humans can learn new ways of learning, indicating that this faculty can operate recursively on itself. For instance, a linguist who has learned several languages can draw upon past experience to acquire another one more quickly than someone who is unilingual.

In order to learn, an ecocyborg must be capable of recognizing, generating, evaluating, comparing, and reproducing patterns. The apparent ease with which the human mind accomplishes these tasks can be somewhat misleading. Cognitive scientists attempting to simulate these abilities on computers are discovering how difficult it is to reproduce them (Pinker 1997). Nevertheless, methods have been developed that emulate some aspects of human learning, and that might also endow an ecocyborg with rudimentary learning abilities. One example is the training of artificial neural networks by back-propagation of error. An ecocyborg can only learn effectively, however, if it has the creative capacity to discover or invent new patterns of relationships. Creation in this context can involve *optimization*, whereby existing patterns are varied according to some scheme and the results are evaluated. More dramatic creative efforts are *exploratory*, involving variations that are radical departures from the established norm (Boden 1990). Exploratory creation can proceed by *association*, where new relationships are established between two concepts in a kind of folding of *idea space*. In this way, previously disparate ideas are associated by identifying similarities between them, or transposing an idea from a familiar context to a new one. Finally, *inventive* creativity is the innovation of pattern in a foray into previously unexplored regions of idea space. Methods of implementing creative learning in ecocyborgs are speculative at this point, but a certain amount of consciousness would certainly increase the effectiveness of some associated activities, such as evaluating new phenomena or activities that directly involve the ecocyborg.

4.6 Consciousness

Although there is no universally accepted definition, consciousness is generally conceded to involve the ability to observe and reason about oneself. Based on this, a proposed working definition of *consciousness* is the maintenance by an entity of a self-referential model; i.e. a model that includes some representation of the entity itself, thus enabling it to reason about itself in relationship to its environment (Chalmers 1990; Lacroix and Kok 1991). The abilities that are necessary for consciousness in an ecocyborg are shared among all the mental faculties. Since consciousness is based on the creation of models, it requires, for instance, the perception of phenomena, the identification of patterns in the resultant information, and the creation of formal constructs that are similarly patterned.

The degree of consciousness of an entity can be measured on a continuous scale, as opposed to being regarded as a discrete, binary attribute. Human mentation, for instance, is sometimes deliberate, explicit, and transparent, but more often it is not directly observable by the reasoner himself. The human reasoner is therefore unable to generate a complete self-model, and is thereby less conscious than he might otherwise be. Although humans and many animals display various degrees of consciousness, natural ecosystems are only very slightly conscious by comparison, since they appear to lack the required abilities, virtual machinery, and corresponding physical substrates. It might be possible to make ecosystems more conscious, however, by cyborging them with technological control networks.

Once a self-referential model has been generated it can be used in prediction, reflection, and imagination. *Prediction* is mentation about how real events might unfold in the future; *reflection* concerns how they developed in the past; and *imagination* deals with hypothetical alternatives to actual situations. Variations on this basic theme allow for more sophisticated mentation. The recursion of consciousness, for instance, involves the creation of models representing the entity in enough detail so that the existence of the self-referential model is also denoted. Accordingly, a model that provides an ecocyborg with a representation of itself, but from which any representation of consciousness is excluded, endows the ecocyborg with *primary consciousness* (Lacroix and Kok 1991). An ecocyborg possesses *secondary consciousness* if the model does take itself into account, and so on for higher degrees of recursion. Ecocyborgs might also be engineered so as to be able to simultaneously instantiate a number of self-referential models, and so consciously reason in parallel about various problems and possible solutions. An ecocyborg that is able to reason consciously is likely to be more effective in its response to external phenomena than one that cannot do so. It would have a superior capacity to regulate its own internal state and to formulate appropriate external responses. This would increase its autonomy by making it more effective in the intentful pursuit of its own goals.

4.7 Autonomy

Autonomy is the independence of comportment that emerges when a sufficiently conscious mind can be described as possessing, to some degree, several defining characteristics (Kok et al. 1995; Bourgine and Varela 1992). The first is *automation*: the capacity to operate without outside intervention. Although necessary, this alone is insufficient for significant autonomy, since even a clock, for example, is capable of indefinite operation without outside involvement. The second required characteristic is *volition*, or choice in action or thought. A highly automatic, volitive mind can respond to its environment in a flexible manner by defining its own goals and then formulating and executing strategies for attaining them. Advanced greenhouse control systems are being developed, for instance, that are capable of limited volition in fulfilling their operating requirements (Lacroix 1994). Finally, in order to be significantly autonomous an entity must be *intentful*, and actually exercise its volition. Since the intentful pursuit of goals is involved, one could say that increased autonomy is equivalent to a greater degree of deliberate self-control (Conant and Ashby 1970). In general, these goals minimally include the survival of the biosystem. In the case of engineered biosystems such as ecocyborgs, they could also include other design objectives.

Like intelligence and consciousness, autonomy should be measured on a continuous scale. Moreover, although autonomy is dependent on mind and consciousness, their presence to any extent is not in itself sufficient to ensure significant autonomy. Even a highly intelligent and conscious ecocyborg, for instance, could be extremely curtailed in its autonomy if it were engineered to pursue a very specific set of objectives, explicitly defining the necessary subgoals, and putting in place a rigid set of rules that governed its allowable attainment strategies. In contrast, an ecocyborg would be a great deal more autonomous if it were bound only by broad, long-term objectives and a loose set of guidelines. In calibrating such a continuous scale for autonomy, one might think it appropriate to use a theoretical maximum as a standard. This leads to a paradox, however, since complete independence in an entity requires a structure that is free of any implicit design objectives or behavioral biases that might influence the definition or pursuit of goals. The actual behavior of such an entity, moreover, would have to conform exactly to its intent, and not be influenced in any way by the environment. In the limit, therefore, absolute autonomy would require that the entity be responsible for creating itself as well as its external environment, and an absolutely autonomous system would therefore have to be absolutely creative. Since humans are incapable of imagining what such an entity might be like, it is difficult to use it as a calibration standard. The average of some human population could be used instead, as is often done for the calibration of scales of intelligence.

Although absolute creativity is an unattainable goal, any entity with some degree of autonomy must be creative enough to formulate at least a few of its own goals and behavioral guidelines. A significant degree of autonomy is desirable in any ecocyborg that is engineered to achieve particular goals in an unpredictable environment. An uncreative ecocyborg would be dependent on preformulated action plans that might not be suited to new situations, whereas a more creative one would be capable of adapting to unforeseen situations by restructuring its goal tree and implementing new strategies in order to achieve its overall objectives. An automated greenhouse, for example, could vary the parameter values of its regulatory models and simulations in order to optimize them for the current situation. More radical creative measures could be implemented in ecocyborgs that were faced with more challenging environments, but in order for them to be useful to their designers their autonomy should be shaped so that they will not override their general design objectives.

4.8 Conclusions

Computers presently serve as the physical substrate for sophisticated virtual machinery that endows ecocyborgs with computational abilities that are superior to humans in some narrow domains. Such artificial constructs are, however, still vastly inferior to human minds in most computational tasks, and are completely incapable of performing others. As a result, the autonomy of existing ecocyborgs is very rudimentary, and they can operate without human supervision only under routine conditions. Some automated greenhouses can employ predictive control techniques to adapt to bounded fluctuations in feedstock quality or ambient temperature, for instance, but they cannot deal with large, unforeseen departures from normal operating parameters (e.g., Lacroix 1994; Lacroix et al. 1996; Linker et al. 1998). In many circumstances it would be desirable to employ highly autonomous ecocyborgs that are capable of reasoning about themselves in the context of their environment, setting their own goals, devising strategies for their attainment, and executing them, all without human supervision. It is postulated that a high degree of autonomy is required of any unsupervised ecocyborg that must persist in an unpredictable environment.

The coherent lexicon and philosophy presented here facilitate the characterization and engineering of significantly autonomous systems, such as ecocyborgs. The creation of these will serve some practical purposes, but will also have an impact well beyond the utilitarian sphere. Highly autonomous ecocyborgs could be employed, for instance, to mediate the increasing human impact on extant natural ecosystems, and thus have a profound impact on human society. Entirely artificial ecocyborgs could also be created to serve a variety of other purposes, such as the production of food, fiber, and other biological products. Large, self-sufficient ecocyborgs could even provide a base for habitation and industrial expansion in space. Once proven technology has been developed for the construction of such entities, it may be possible to create them in great numbers, and perhaps even to make them capable of replicating themselves. In sufficient numbers, they might develop their own societies, collective structures that might evolve as computational systems in their own right, complete with economies, philosophies, and theologies. These societies might also be subject to engineering practices, in which case researchers can look ahead to shaping new structures not only at the level of individual ecocyborgs, but also at higher levels of conglomeration. The lexicon and philosophy provided here provide a language and framework with which such endeavors can be envisioned, planned, and executed.

4.9 Acknowledgments

The authors gratefully acknowledge the support of the Natural Sciences and Engineering Research Council of Canada, the J.W. McConnell Foundation, and McGill University. Appreciation is expressed to M.C. Stacey and P.C. Clark for reviewing drafts of this article.

4.10 References

Boden, M.A. 1990. *The creative mind: myths and mechanisms*. New York, NY: BasicBooks.

Bourgine, P. and F.J. Varela. 1992. Towards a practice of autonomous systems. In *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life,* 1st ed., ed. F.J. Varela and P. Bourgine, xi -xvii. Cambridge, MA: MIT Press.

Chalmers, D.J. 1990. Consciousness and cognition. *AI CD-ROM*, Revision 2. Irving, TX: Network Cybernetics Corporation.

Conant, R.C. and W.R. Ashby. 1970. Every good regulator of a system must be a model of that system. *International Journal of Systems Science* 1:89-97.

Dyson, G. 1997. *Darwin Among the Machines*. Redwood City, CA: Addison-Wesley.

Engelberg, J. and L.L. Boyarsky. 1979. The noncybernetic nature of ecosystems. *The American Naturalist* 114:317-324.

Gardner, H. 1993. *Frames of Mind*, 2nd ed. New York, NY: BasicBooks.

Goldman, A.I. 1986. *Epistemology and cognition*. Cambridge, MA: Harvard University Press.

Greenfield, S.A. 1995. *Journey to the centers of the mind*. New York, NY: Freeman.

Kok, R., L. Parrott, O.G. Clark, R. Molenaar and R. Lacroix. 1995. The engineering of autonomous biosystems: the EcoCyborg Project. CSAE Paper No. 95-602. Saskatoon, SK: CSAE.

Kok, R. and R. Lacroix. 1993. An analytical framework for the design of autonomous, enclosed agroecosystems. *Agricultural Systems* 43:235-260.

Lacroix, R., R. Kok and O.G. Clark. 1996. Use of a mulitasking operating systems as a setting for the simulation of an enclosed agro-ecosystem under cognitive control. *Canadian Agricultural Engineering* 38:129-138.

Lacroix, R. 1994. A framework for the design of simulation-based greenhouse control. Unpublished Ph.D. thesis. Department of Agricultural Engineering, McGill University, Montreal, QC.

Lacroix, R. and R. Kok. 1991. A cognitive controller with recourse to simulation. Unpublished paper presented at the Automated Agriculture for the 21st Century ASAE Symposium. Chicago, IL, December 16-17.

Linker, R., I. Seginer and P.O. Gutman. 1998. Optimal CO_2 control in a greenhouse modeled with neural networks. *Computers and Electronics in Agriculture* 19:289-310.

McNaughton, S.J. and M.B. Coughenour. 1981. The cybernetic nature of ecosystems. *The American Naturalist* 117:985-990.

Minsky, M.L. 1985. *The Society of Mind.* New York, NY: Simon and Schuster.

Parrott, L., O.G. Clark, R. Kok and R. Molenaar. 1996. The EcoCyborg Project: new tools for the biosystems engineer. NABEC Paper No. 9655. St. Joseph, MI: ASAE.

Patten, B.C. and E.P. Odum. 1981. The cybernetic nature of ecosystems. *The American Naturalist* 118:86-895.

Pinker, S. 1994. Mind Design. In *The Language Instinct*, 1st ed., 404-430. New York, NY: William Morrow and Company.

Pinker, S. 1997. *How the mind works.* New York, NY: W.W. Norton.

Sagan, C. 1994. *Pale blue dot: a vision of the human future in space.* New York, NY: Random House.

Vigilant, L., M. Stoneking, H. Harpending, K. Hawkes and A.C. Wilson. 1991. African populations and the evolution of human mitochondrial DNA. *Science* 253:1503-1507.

CONNECTING TEXT

Chapter 5, **The EcoCyborg Project**, was authored by L. Parrott, O.G. Clark and R. Kok. At the time this thesis was prepared, the text of this chapter was being readied for submission as an article to the journal *Canadian Agricultural Engineering*.

Chapter 5 is the confluence of the themes of *biosystem* and *autonomy*, which were introduced in Chapters 3 and 4, respectively. The EcoCyborg Project, referred to in the preceding chapters, is here discussed in more detail. The underlying philosophy of the project and the engineering approach that has been adopted are described, as is the type of hypothetical system that is the case study in the project. Such systems are called *ecocyborgs*, and are envisioned as combinations of an ecosystem and a technological control network. The autonomy of this kind of hybrid system is greater than that of a natural ecosystem alone. In this project, configurable computational models of the ecocyborgs are being developed, and implemented in simulations of their comportment.

CHAPTER 5. THE ECOCYBORG PROJECT

Abstract

The EcoCyborg Project, described in this article, is a research program for which the long-term objective is to develop a general theory of biosystems engineering, with emphasis on substantial system autonomy as a design criterion. Within this context, the short-term goal is to create tools for the modeling, simulation, and characterization of a particular type of biosystem, called an *ecocyborg*. Such systems consist of a large number of biological and technological components that are integrated (*cyborged*) at an organizational scale similar to that of an ecosystem, with some of the technological components fulfilling system control functions. The article is divided into four main parts, the first of which is a discussion of various philosophical issues related to the project. To start, a *biosystem* is defined as an entity that is substantially alive in the sense that it is *autopoietic*, or self-producing. Next, a conceptual framework is elaborated that comprises pairs of complementary descriptors, with each pair corresponding to one axis of a hypercube that can be used to reference possible kinds of systems. Following this, a philosophy of engineering is presented as complementary to that of science, where the former is prescriptive in nature and the latter is explanatory. The aforementioned philosophy is then cast in terms of biosystems engineering. The second part of the paper is a description of the objectives and engineering approach adopted for the EcoCyborg Project. Cyborging is discussed as a means of creating substantially autonomous biosystems, with computer modeling and simulation being the method of study currently used. This work also requires the development of characterization methods, which are essential in this context for: self-observation by the ecocyborgs themselves (if these are to be at all autonomous), explanatory description by scientific investigators, and prescriptive description by engineers. The third part of the paper is a description of the type of system currently being studied in the project. This particular type of hypothetical ecocyborg is described as an orbital space station comprising an enclosure, an ecosystem, and an artificially intelligent control system, all influenced by forcing functions (e.g., solar radiation). The final part of the paper contains a description of the computational models

created to represent these various aspects, and of the way these models are implemented in simulation.

5.1 Introduction

This article is a description of the philosophy and methodology of the EcoCyborg Project, which has been established to develop a general theory for the engineering of biosystems. Accordingly, the overall project objective is to learn how to create and modify biosystems for particular purposes. In this, substantial system autonomy is a design criterion of special interest. This long-term objective is currently being pursued through the development of modeling, simulation, and characterization tools for the study of a specific type of biosystem, called *ecocyborg*. These are biosystems of the ecosystem scale that are composed of large sets of both biological and technological components which function in an integrated manner. The term ecocyborg was originally derived from the concept of combining (*cyborging*) an ecosystem with technological systems. Many facilities in existence today, such as greenhouses, can be considered as primitive ecocyborgs. The thrust of the EcoCyborg Project is, however, to deal with much more complex, sophisticated, and autonomous systems. Short-term applications of this work include the remediation of natural ecosystems, the enhancement of agricultural production in an environmentally sustainable manner, and the construction of space habitats for humans. This may lead in the future to the creation of entirely new types of biosystems for particular purposes and, as a possible far application, the construction of biosystems capable of assimilating and generating knowledge beyond human comprehension.

This article is divided into four major parts, the first of which is an overview of the philosophy underlying biosystems engineering. This begins with a definition of biosystem as a type of highly complex, adaptive system that is alive as a whole to a substantial degree. *Alive* is considered here to be equivalent to *autopoietic*, a quality of comportment whereby the interactions of the system components combine as an overall network that is self-producing, or homeostatic for itself (Clark and Kok 1999a, presented here in Chapter 3; Maturana and Varela 1980). Following this definition, it is briefly described how biosystems can be characterized with pairs of complementary descriptor variables, which are analogous to orthogonal axes that form a hypercubic space of possible modes of existence. Next, engineering is discussed as being based on a philosophy oriented toward the creation of entities, such as biosystems, in order to fulfill particular design objectives. This philosophy is considered as complementary to, but distinct from, that of science, which is more oriented toward the explanation of existing phenomena. Finally, the themes of biosystems and engineering are brought together in a discussion of biosystems engineering as a unique discipline.

The second part of the article is a review of the objectives and engineering approach for the EcoCyborg Project. There is an exploration of the design objective of substantial system autonomy. This includes an explanation of associated terms, such as *mind*, *virtual machinery*, *mental abilities*, and *consciousness* (Clark et al. 1999, presented here in Chapter 4). Next, there is a description of the approach, called *cyborging*, adopted in this project for the engineering of biosystems. Cyborging is the creation of aggregate entities, such as ecocyborgs, comprising both biological and technological components, so that they fulfill prescribed objectives like substantial autonomy.

The third part of this article is an account of ongoing research in the EcoCyborg Project. The particular type of ecocyborg studied in this project is discussed, as are the modeling, simulation, and characterization tools that are being developed for this work. Ecocyborgs of the type that are the current focus of the project are envisioned as materially closed space stations in orbit about a Sun-like star. Such an entity comprises a community of biological organisms, similar to that found in a temperate woodland on Earth, together with their abiotic surroundings. It also includes a network of technological components intended to guide its dynamics and to endow the overall system with substantial autonomy. As well, it is influenced by some factors that cannot be controlled by the guidance network. These factors are so-called *forcing functions*, and include rain, solar radiation, and ambient temperature.

The physical construction of such ecocyborgs is clearly impossible in the short term and so, as is usual in such cases, they are being studied by means of a modeling and simulation approach. Accordingly, a number of virtual tools are being developed for use in the EcoCyborg Project, and these are described in the final part of this article. First, an object-based computer model is being written that can be configured to represent a given biotic community, its abiotic surroundings, the enclosure that contains these, and any guidance components that are internal to that ecosystem (Parrott 1995). Configurable models are also being developed to represent the forcing functions. As well, a control system is being created that can be configured to emulate the more sophisticated components of a given ecocyborg's guidance network. All of this software can be simultaneously implemented in a dynamic simulation, so as to emulate the comportment of the particular hypothetical ecocyborg. Finally, characterization methods are being identified, which can also be implemented as computer software, for use in studying the ecocyborgs and the computational models thereof. Such methods are required for three purposes: (1) to endow ecocyborgs with the capacity to observe and control themselves, so that they can be substantially autonomous, (2) for the explanatory description of the ecocyborgs by external, scientific observers and (3) for the effective specification by engineers of systems that they propose to create.

5.2 Philosophy of biosystems engineering

5.2.1 Biosystems

Biosystems are, as mentioned, systems that are alive as a whole to a substantial degree. To state that they are systems implies that they are of an aggregate nature, comprising a number of components that interact relatively strongly with one another so as to form a unitary whole (Clark and Kok 1999a, presented here in Chapter 3). As with any system, the discrimination of a particular biosystem as a discrete entity is somewhat arbitrary, often being accomplished by defining a boundary that encompasses a particular set of components. Accordingly, such a boundary is chosen so that external phenomena are largely uninfluenced by those which are internal. Nevertheless, biosystems need not be completely isolated. Thus, although a system boundary is usually chosen so as to transect lines of relatively weak mutual influence, there may be some exchange of mass, energy, momentum, or information across it. As discussed below, such an exchange is in fact characteristic, and necessary for the persistence, of most biosystems. Moreover, some external influences, such as forcing functions, might act unilaterally upon the system.

Biosystems are a subset of the larger group of substantially complex systems. Hence, a working definition of biosystem can be obtained by appending the adjective "living" to the definition of a complex system as a "network of interacting objects, agents, elements, or processes that exhibit a dynamic, aggregate behavior" (Bonabeau and Theraulaz 1994, p. 305). Biosystems may be of various physical sizes and cover the entire gamut of organizational scales. They are composed of large numbers of sparsely interconnected component entities. Every component, set of components, and even the system as a whole have both virtual and physical aspects, and either aspect may predominate in particular circumstances. Thus, they may be considered either as information-oriented constructs (that must always reside on a physical substrate), or as predominantly physical entities (that always have a virtual aspect). A given system might comprise components which are primarily physical or virtual in nature.

Another feature of biosystems and other complex systems is substantially *emergent* comportment, which arises from the interrelationships between their many internal components, as well as with their surroundings. Although this quality of comportment arises from local processes, it is not understandable without taking into consideration interactions at a variety of spatio-temporal scales, and is evident only when the system is observed as a whole. Emergent comportment can be of various types including, for example, chaotic modes in which small-scale, local dynamics combine in a deterministic but inherently unpredictable way to influence system-level features.

The comportment of biosystems and other kinds of significantly complex systems is usually not only highly emergent, but substantially adaptive as well. *Adaptive* systems react to their external environment so as to ensure that particular features of themselves either are maintained or change in a manner that is at least somewhat independent of external forcing functions. If a system is adaptive to a given degree, then it can also be considered as autonomous to some extent since it is, in effect, actively pursuing internal goals, some of which might have been originated by the system itself (Clark et al. 1999, presented here in Chapter 4). Autonomy, together with some related concepts and preconditions, is discussed later in this paper.

The defining characteristic that is particular to biosystems is life, which can be considered as equivalent to *autopoiesis* (Clark and Kok 1999a, presented here in Chapter 3; Maturana and Varela 1980). This particular kind of comportment is both substantially emergent and adaptive. Autopoiesis is the interaction of the components of a system so that their combined effect is the continual production of the components themselves and the maintenance of the overall system structure. This dynamical mode occurs to some degree in a variety of kinds of systems. It is usually evidenced by the simultaneous import of order from, and export of disorder (characterized, for example, as thermodynamic entropy) to, the surrounding environment (Boltzmann, as quoted in Broda 1983, pp.79-80; Schrödinger 1955; Brillouin 1951). To persist over the long term, a biosystem must, therefore, be open to a medium for this exchange, such as a flow of mass or energy (or else confine increasing amounts of entropy within some part of itself). In entities like computer-based systems, which have primarily virtual dynamics, this exchange might not be as evident as in systems that display comportment of a predominantly physical nature.

The definition of *biosystem* given here is expansive enough to accommodate a variety of different system types, physical sizes, and degrees of organizational intricacy. Examples include individual cells, multicellular organisms, self-directing factories, space stations, human societies, and artificial minds. Some of these systems might not even contain any components that are traditionally thought of as biological, whereas others might include components that can be considered as biosystems in their own right. In fact, large-scale biosystems are usually somewhat self-similar between scales in that they contain a variety of the latter kind of components, organized in a hierarchical manner. A natural ecosystem, for example, contains many types of organisms, which are themselves agglomerates of large numbers of different types of cells. Similarly, a human society consists of many progressively smaller groups such as nations, cities, communities and families. In the face of such diversity, some means must be provided for the effective observation, description, and specification of biosystems. A conceptual framework has been developed for this purpose, and is described below.

5.2.2 The hypercube of existence space

The conceptual framework developed for the EcoCyborg Project is based on a set of characteristics that are important from a biosystems engineering perspective. These can be associated with pairs of complementary adjectives. Each of these pairs can, in turn, be thought of as corresponding to one of a set of mutually orthogonal axes that defines a *hypercube of existence space*. A particular set of five such adjectival pairs is described below; different or more expansive sets might also be assembled for use in other circumstances. Some other pairs that might be useful, for example, correspond to characteristics such as the closure of a system (*closed/open*) or the time of its existence (*past/future*). Thus, it is possible to define an existence space of arbitrary dimensionality that is applicable to the characterization of one or a set of biosystems.

Given the five axes that form the basis for the hypercube chosen here, a particular biosystem can be described with a coordinate vector of five corresponding elements. An underlying premise of the philosophy used in this project is that a system characteristic is appropriately described as a continuous variable (as opposed to a discrete one that only has a distinct and finite number of possible states). The value of each element can, therefore, vary in a continuous manner, indicating where the point representing the biosystem is located in the existence space with respect to the associated axis. The chosen descriptor pairs are explained next.

The Real/Imaginary Axis. A purely imaginary system exists only as a hypothesis or idea in the mind of a cognitive entity. In contrast, a very real system has an objective existence. For example, the enchanted Old Forest in J.R.R. Tolkien's (1966) classic fantasy tale *The Hobbit* is a purely imaginary (albeit physical) biosystem. Contrariwise, an ant colony model implemented in a simulation environment like SimAnt (Maxis, Orinda, CA) is a real, yet virtual, biosystem. Similarly, ecocyborgs of the kind discussed here are imaginary, whereas the computational models being developed to represent them are real systems.

The Natural/Artificial Axis. The degree to which a biosystem, such as an ecocyborg, is natural or artificial is determined by the amount that humans have influenced its composition and structure. A completely natural biosystem is one that has been formed independently of human direction. A completely artificial system, in contrast, is one that has been designed and constructed entirely by humans, and may be composed of natural or manufactured components or some combination thereof. Agricultural systems, ski slopes, and water reservoirs are examples of somewhat artificial, primitive ecocyborgs that happen to contain many natural components, but that also generally comprise a great number of technological ones which are largely of human origin. Few large-scale biosystems remain that qualify as completely natural. For example, some forests in Canada's more isolated national parks are still essentially natural ecosystems, but even most of these have been modified through modern industrial and recreational activities, as well as by indigenous peoples over thousands of years.

The Physical/Virtual Axis. A predominantly physical biosystem is one in which the material or energetic aspects are considered to be of primary relevance, whereas in a virtual one the informational aspects are considered as the most important. Most humans are likely to think of large-scale biosystems, like forests, coral reefs, or the hypothetical space stations being studied in the EcoCyborg Project, as primarily physical entities. There are, however, an increasing number of predominantly virtual systems that could be considered as biosystems by virtue of their structure and comportment. Examples include the computational models described in this article and artificial life software created by other researchers (Taylor and Jefferson 1994; Ray 1994). On a large scale, the entire Internet can perhaps be considered as a biosystem (Dyson 1997).

The Living/Non-living Axis. As with other system characteristics, the vitality of a biosystem can also be measured according to a continuous scale that accommodates any value within a given range. Hence, rather than assign rules for discrete classification in this regard (which ultimately leads to difficulty in classifying entities such as biological viruses), a system can be considered as alive to a relative extent. There are no convenient, explicit measures of life (i.e., autopoiesis) as such, but some commonly accepted (albeit insufficient) indicators include: the maintenance of an internally ordered state through the export of entropy; adaptation to environmental forcing functions; and growth and reproduction. For example, organisms that are considered to belong to one of the traditional kingdoms of living creatures would, according to these criteria, be assigned an accordingly high value on the life axis. Biological viruses, however, would be given an intermediate value, since they have no internal metabolism by which entropy is exported and they depend entirely on host cells for reproduction. Similarly, most natural ecosystems would also be assigned an intermediate value, since they generally lack an overall mechanism for reproduction and, relative to an organism, they do not comprise a very cohesive autopoietic network. Also, their ability to adapt is usually relatively limited, although some particularly robust ecosystems might be able to persist under a variety of changing conditions. Such a scale could be used to evaluate even non-carbon-based life forms of all types.

The Guided/Unguided Axis. The term *guidance* refers here to the intentional manipulation of a system so as to make it behave in a particular manner. Thus, a system that is subject to a great deal of guidance is one whose autonomy is limited either by inputs or internal components (i.e., control mechanisms), which are intended to modify the system's behavior so that it fulfills requirements that are not of its own devising. Conversely, a completely unguided system is one in which comportment proceeds entirely without intentional control. Guidance is often of benefit (to the controlling agent) in production systems, for example, since it can greatly enhance their effectiveness and efficiency.

Overall, a hypercube defined with appropriate descriptor pairs can be used to succinctly characterize the nature of a particular biosystem. This approach can also be used to characterize any changes that might occur in the nature of the system with respect to time, or some other independent variable. For example, a biosystem that is originally extremely natural but which is subsequently managed by humans often develops an increasingly artificial character. Conversely, a largely artificial biosystem, such as a reforested slope or a backyard garden, that is left unguided will often become similar to a natural system. Thus, as a system changes, its position along the axes corresponding to any of the paired descriptors may also change. In engineering, existing systems may be intentionally modified to bring about such changes, or entirely new systems may be created, in order to achieve a particular purpose. The philosophy underlying such activity is discussed in the section below.

5.2.3 Engineering and science

As mentioned above, the EcoCyborg Project is mainly an engineering endeavor. The general public often confuses engineering with science, as do even many engineers and scientists. The two pursuits involve, however, quite different objectives and methods. On the one hand, science is oriented toward the observation and description of existing phenomena for the purpose of understanding and explanation. Some general scientific activities include, for example, taxonomy, experimentation, and analysis. Engineering, on the other hand, is practically the antithesis of this. It involves the design and construction of new systems and objects according to predefined specifications, as well as their subsequent control, maintenance, repair, and upgrading. Scientific knowledge is often applied in these tasks, but only if the explanatory scientific descriptions of "how things work" can be inverted to establish rules for the creation of things which fulfill preconceived goals.

The difference between science and engineering can be illustrated by considering an observed system as a "black box" containing a number of components that interact in a particular manner. The system's *composition* (number and kinds of components) and *structure* (inter-relationships between them) are collectively referred to as its *constitution*. The *state* of the system at a given time is the value of any changeable attributes of the components at that moment. The state of the system might change over time, and the manner in which this occurs is referred to as the system's *comportment*. Comportment may occur in response to *inputs* from the external environment (e.g., forcing functions) and, conversely, the external surroundings may be affected by the system. The latter influences are the system's *outputs*, and include any outward transfers of mass, energy, momentum, or information.

Engineering activities such as design, construction, etc., are usually intended to result in a system that fulfills a particular set of objectives. Hence, in such a goal-oriented exercise, the system itself is really only a means to an end. The effectiveness of such an endeavor depends largely on the engineer's understanding of the kind of system that is to be created. For this, general knowledge is first required of how the constitutions of similar systems give rise to their comportment under the influence of forcing functions. This understanding is equivalent to having an appropriate explanatory model of such systems, the development which lies primarily within the domain of experimental and analytical science. Second, an inverse model must also be formulated, so that a system can be specified that will give rise to the desired comportment in the face of the expected forcing functions. Hence, the design procedure is facilitated if an appropriate model exists of the type of system that is desired. It is to be noted that even unsuccessful engineering efforts can, however, be valuable because the lessons learned from such attempts often lead to the improvement of engineering procedures and to the refinement of models of poorly understood systems.

In the design phase of engineering, only a system's composition, structure and initial state are specified. If the design is effective, the system will then respond to external forcing functions in a manner that corresponds to the desired comportment. However, in some cases the exact nature of the forcing functions might be unknown, and the engineer's knowledge of them may be limited to only a range of probable values. In such a case, the system might be overdesigned by a factor of safety chosen to reflect a particular degree of uncertainty. Nevertheless, even the most robust design cannot accommodate all eventualities, and so engineering also involves several subsequent phases.

The phases of engineering which follow design and construction include control, maintenance, repair and upgrading. *Control* is the explicit guidance of a system so as to maintain its comportment within an allowable envelope (i.e., to achieve a set of goals). This is done by compensating for forcing functions or unexpected internal dynamics to which the system could not inherently respond in the absence of guidance. Control includes regulation, operation, and management, each of which is relevant over respectively longer time scales (Kok and Lacroix 1993). In the past, these activities have depended solely on humans, but they can now be increasingly achieved with various types technological devices. This aspect of engineering is discussed in more detail later in this article.

Even if a system includes a very advanced control network, it cannot be made infallible, and so engineering extends even beyond the control phase. Some *maintenance* is generally required, for instance, in which system components are regularly replaced and the system structure readjusted in order to ensure consistent generation of the desired comportment. *Repair* involves similar, but unscheduled, activities in response to damage from unforeseen component failure or external influences. Finally, *upgrading* is the planned alteration of a system in an effort to improve its performance or adapt it to changing conditions. The application to living systems of any of these engineering activities, from design to upgrading, constitutes the discipline of *biosystems engineering*.

5.2.4 Biosystems engineering

The study of biosystems is challenging for a number of reasons. For instance, biosystems differ in their internal structure from systems that are dealt with in more traditional science and engineering disciplines. They are highly complex, adaptive, living systems that comprise very large numbers of components, which interact over multiple temporal and spatial scales. Consequently, the autopoietic comportment of biosystems is highly non-linear and difficult to describe with traditional analytical methods. Nevertheless, forward, explanatory relationships are the subject of scientific inquiry in several fields of study, such as artificial life, complex systems science, and the cognitive sciences (Bourgine and Varela 1992). Experimentation has shown that the mechanisms that effectuate the emergence of autopoiesis in biosystems appear to be associated with certain structural qualities. For instance, cellular automata and other systems with similar networked structures tend to display emergent behavior if their degree of connectivity falls within a certain range (Kauffman 1995; Green 1993).

Generally, research of the kind mentioned above is undertaken from a scientific perspective, not an engineering one. In most of the experimental work, for instance, a system is initialized with an (often random) state and the various stages of morphogenesis are observed as it develops. Although studies of this nature yield information about the forward problem of how dynamics depend upon composition, structure, and state, they generally do not address the reverse problem, crucially important to engineering, of how to determine what kind of constitution will result in a given comportment. Related work that does have an engineering orientation, such as the development of advanced life support systems, has been very narrowly focused on technical issues. Overall, little work has been done on the development of a more general approach to the engineering of highly complex, adaptive systems, let alone on how such systems might be rendered substantially alive or autonomous. The EcoCyborg Project is, therefore, novel in this respect.

5.3 Project objectives and engineering approach

To reiterate, the EcoCyborg Project is oriented toward the long-term objective of developing a general theory for the engineering of biosystems. The class of all biosystems is too broad to be an immediate focus for the project and so a particular type of ecocyborg has been chosen for study over the short term. The reasons underlying this choice, and a more detailed description of the particular type of ecocyborg, are discussed in later sections. Finally, in this work, there is an emphasis on substantial autonomy as a design objective, an aspect of the project that is somewhat unique in the context of large, aggregate biosystems such as ecocyborgs.

The characteristic of autonomy is of interest here because it is deemed to be beneficial in any system that must persist, without external guidance, in variable surroundings. *Autonomy* is independence in the establishment and pursuit of one's own goals (Clark et al. 1999, presented here in Chapter 4). The overall extent to which a system may be autonomous depends on the degree to which the system is capable of using a model of itself in its computational activities. The formulation of realistic goals and strategies, for example, will heavily depend upon the use of such a model. The use of a self-model in this way is the working definition of *consciousness* adopted in this project. These kinds of activities can, of course, be considered as computational in their essence. Here, the apparatuses or contrivances (*virtual machines*) that give rise to computation in a particular system are collectively referred to as the *mind* of that system.

Hence, in order to create a system that is autonomous to some degree, it must be engineered to possess an appropriate mind. The biotic portion of the kind of system currently being studied in the EcoCyborg Project is similar to a natural ecosystem, which does not inherently possess the type of mind necessary for substantial autonomy. Cyborging has been adopted here as the approach to endowing these large-scale, aggregate biosystems with the required virtual machines. These are envisioned as residing on the technological components of the control system, which are physical devices such as digital computers, sensors, and actuators.

Such large, physical systems are, for reasons of convenience and practicality, inappropriate study subjects at this stage, and so computational systems are currently being created for this investigation. These include computational models and simulation tools for representing the static and dynamic aspects of the hypothetical ecocyborgs. Such virtual tools enable design and experimentation while obviating the need for expensive physical instantiations of the ecocyborgs being studied. Work can also be conveniently performed in the virtual realm that would, in fact, be impractical, immoral, or even impossible in a primarily physical context. Characterization methods, also implementable as computer software, are also being developed. These are necessary in a number of capacities, such as the formulation of scientific (explanatory) and engineering (prescriptive) descriptions of the ecocyborgs and the computational models. Many of the characterization methods apply to both system types because, in order to be useful, the computational models must evidently have features that are similar to essential aspects of the ecocyborgs.

There is another very important capacity in which tools for characterization, modeling, and simulation are necessary. In order to have substantial autonomy, ecocyborgs must, as indicated, be significantly conscious, and this requires that they have at their disposal a model of themselves. Although an existing model might be supplied, ecocyborgs with more sophisticated minds would generate and maintain their own self-descriptions. The formulation of a self-model is, of course, a characterization exercise; thus, substantially autonomous ecocyborgs (and the computational systems created to represent them) must possess virtual machinery that is appropriate for this. Such systems might use the models in simulations to predict future events, examine hypothetical situations (e.g., comparison of alternative control strategies), and reflect on past scenarios. They might even be able to implement a number of models in order to simultaneously consider different problems, or analyze the same problem using various approaches.

5.4 The ecocyborg type

The present case study for the project is a purely hypothetical type of ecocyborg that consists of a relatively small, materially closed ecosystem, an extensive network of control mechanisms, and the enclosure that contains them. It is envisioned as a space station, a setting chosen partly because of its interest to the members of the research team. By virtue of its control system, this type of ecocyborg is highly conscious and autonomous. In order to learn more about the engineering of such entities, the relationship between the constitution, initial state, and comportment of this specific kind of system, as affected by forcing functions, is being investigated in detail. In our treatment, this type of system is seen as being composed of three main parts (i.e., the enclosure, ecosystem, and control system) which are driven by the forcing functions (Figure 5.1). A further description of each of these is given in the following sections.

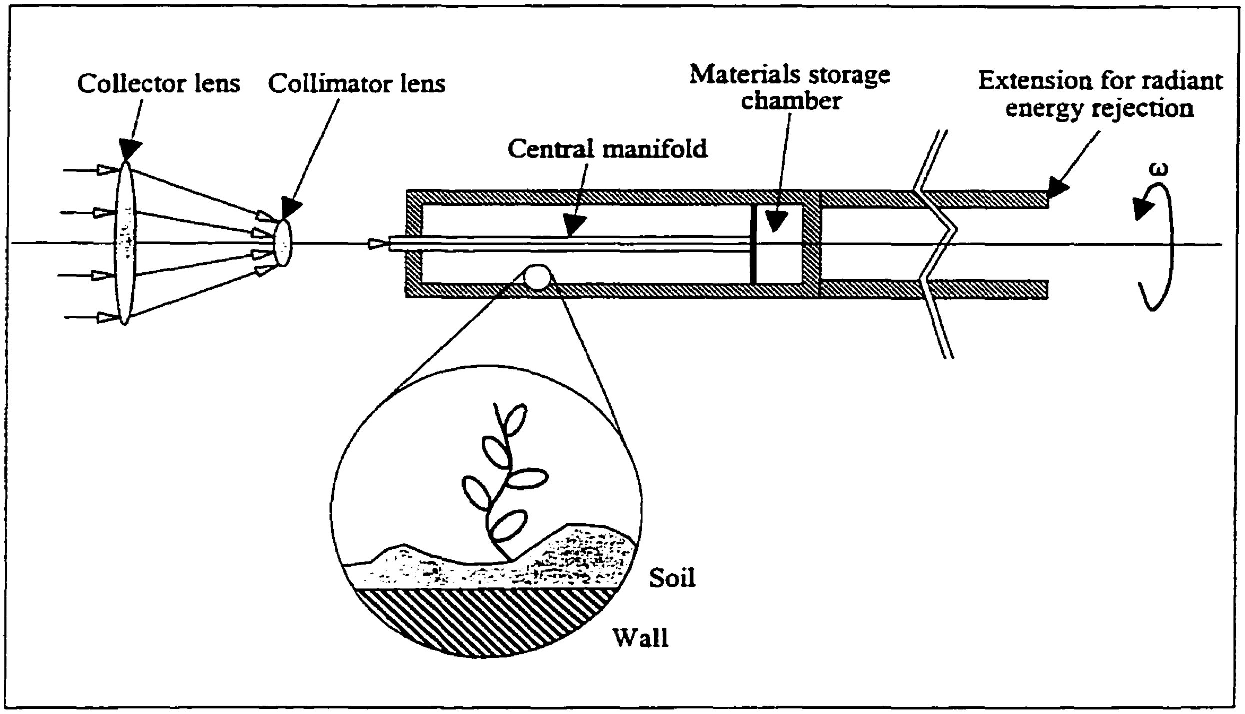

Figure 5.1 A schematic showing an instance of the type of biosystem being studied in the EcoCyborg Project.

5.4.1 Enclosure

The enclosure is a cylindrical shell with a total length of approximately 1000 m and a diameter of about 160 m. A central core, or manifold, concentric with the outer shell but of a much smaller diameter (about 25 m), contains service subsystems for functions such as energy management and atmospheric conditioning. Spokes radiate from the manifold to the outer shell, serving both as structural members and as conduits for the transport of materials. The cylinder's main chamber contains the ecosystem and is about 500 m long. A materials storage chamber extends from one end of this, and is divided into four compartments for mass storage. Beyond this, the outer walls extend further to provide additional surface area for heat rejection. The entire station spins about its longitudinal axis so that the resulting radial acceleration of an object at the outer shell is equivalent to that caused by Earth-normal gravity.

The enclosure of the type of ecocyborg envisioned here serves as both a physical and conceptual system boundary. Thus, such a system is energy-open but mass-closed, with high-grade radiation being converted to low-grade thermal energy. The source of this energy is a Sun-like star around which the ecocyborg orbits at a distance of about 47 million km. At this distance, the stellar radiation is assumed to have an intensity of about 10 kW/m2. The enclosure is oriented so that its longitudinal axis is always radially oriented toward the star. Perpendicular to this axis, between the station and the star, is a collector lens that entirely shades the ecocyborg from direct exposure to radiation. The concentrated energy from the collector lens is directed through a collimator lens and receiver into the central manifold. Assuming a collection efficiency of 90%, this makes available a continuous power supply of 7.0 GW for use within the ecocyborg. The incoming radiant energy can be redirected to illuminate the interior, or can be used to generate electricity to run various internal subsystems, etc. Energy is rejected from the ecocyborg by radiant cooling through the outer shell of the cylinder which, as previously mentioned, is extended in order to provide a surface area adequate for this purpose (see Figure 5.1).

5.4.2 Ecosystem

The ecosystem, most of which is located on the interior walls of the cylindrical enclosure, is materially closed, but open to energy. It is not production oriented, nor is it intended specifically to support human beings in space. Rather, it is envisioned as being similar in constitution to a temperate open parkland or forest edge terrestrial ecosystem. It is considered to be composed of three main parts: the biological component, encompassment, and materials storage realms. The biological component realm contains all of the living organisms in the ecosystem, including species of plants and animals representative of all trophic levels. Thus, although artificially constructed, its constitution is fairly similar to a natural ecosystem (Pimm 1991; Patten 1959). The encompassment is the abiotic environment in which the biological components exist, consisting of an atmosphere and a terrain (soil and water). Its terrain contains soils typical of a parkland with a rolling topology that drains into a small pond. The atmosphere is composed of nitrogen, oxygen, and carbon dioxide in proportions similar to terrestrial air and is maintained at standard atmospheric pressure. Also, some types of control mechanisms are considered as part of the encompassment. Since the material cycles of such a small system will operate on a much shorter time scale than on Earth (Nelson et al. 1992), and since the ecosystem itself is not of sufficient size to include large material buffers (e.g., like oceans on Earth), a number of essential compounds are contained in the compartments of the aforementioned materials storage chamber. The stored compounds may be accessed as required for control purposes and are considered as part of the ecosystem.

5.4.3 Control system

The type of ecocyborg being studied is envisioned as being substantially conscious and autonomous. This is achieved via a sophisticated control system that endows it with the required mental abilities. In general, any object or mechanism that has been deliberately added with the intent that it influence the comportment of the ecocyborg is considered to be part of the control system. In this case, the control system is a semihierarchical network of mechanisms, which is highly integrated with the ecosystem. Such mechanisms may be included within the ecosystem, in which case they are *intrinsic*. Humidity sensors, irrigation sprinklers, and gas diffusers are examples of such intrinsic control mechanisms. Depending on the degree of integration it can, however, be difficult to clearly distinguish between intrinsic mechanisms and the aspect(s) of the system that they guide. For instance, passive temperature regulation might be achieved by means of installing a massive object, but in this case it is only the engineer's intent that differentiates the object from other system components. In some cases intrinsic mechanisms may also play multiple roles, and might therefore be considered as having membership both in the encompassment of the ecosystem and in the control system; a large body of water could simultaneously control humidity, for example, but also serve as a home for many aquatic species. Other control mechanisms are *extrinsic* to the ecosystem.

The control mechanisms of the ecocyborg are semidifferentiated with respect to function and include, for example, memory, reason, and learning mechanisms, as well as a number of *perceptors* (sensory devices) and *effectors* (actuator devices), which interface with the ecosystem. The sophistication of these various devices can range from very simple to extremely complicated; they may be as elementary as a timer that turns a switch on or off, or as intricate as part of a biological nervous system. They can also, therefore, be classified on this basis, and grouped according to four implementation levels: *cognitive*, *Pavlovian*, *instinctive* and *basal* (Kok and Lacroix 1993). The set of each of these four different kinds of mechanisms can be referred to collectively as a *controller* of the same type as the associated implementation level. The basal and instinctive controllers are considered here to be intrinsic to the ecosystem, whereas the Pavlovian and cognitive controllers are considered to be extrinsic (although intrinsic to the ecocyborg as a whole) and are envisioned as being implemented in the form of software resident on digital computers. Different goal types, priorities, and activity classes require various kinds of computational resources, and so these are generally concomitant with the implementation levels (Kok and Lacroix 1993). The overall result of this semihierarchical, semidifferentiated approach is a completely integrated network that maintains the ecosystem (and, therefore, the ecocyborg) in a viable state by ensuring that its comportment remains within a chosen envelope. The different controllers play specific roles in this.

The basal and instinctive controllers are capable only of simple and inflexible responses. Pavlovian control mechanisms carry out routine activities without analyzing the possible consequences, although they may possess some retrospective learning capacity (Kok and Lacroix 1993). This may involve either physiological or situational control related to system regulation and operation. The former refers to the direct manipulation of an effector in response to the values of a small set of input variables and parameters. The type and magnitude of the response is related to the input set by relatively simple transfer functions (e.g., on/off, proportional, PI, PID, and cascade). Situational control implies the consideration of pattern in a larger collection of event data, or in a time series of one or more variable values. The cognitive controller performs other activities, related to operation and management, wherein its mechanisms formulate long-term strategies as well as tactics by which these might be implemented (Kok and Lacroix 1993). This may involve much more sophisticated computation than is required by the Pavlovian controller. Thus, most substantially conscious reasoning is performed by the cognitive controller. In this, self-models of the ecocyborg are implemented in simulations about past (reflective reasoning), future (predictive reasoning) and hypothetical (imaginative reasoning) scenarios. In many cases, cognitive mechanisms implement their strategies in an indirect manner by defining set point values for Pavlovian mechanisms.

5.4.4 Forcing functions

An ecocyborg of the type studied here is driven by three climate-related forcing functions which affect the comportment of the ecosystem: radiation intensity, rainfall, and temperature. These are not under the control of the ecocyborg although they are implemented by various technological subsystems that belong to it and may involve the manipulation of material within it. For instance, for rain to occur in the ecosystem, water must be moved from a storage compartment to the main chamber. In the engineering of different ecocyborgs, the forcing functions might be specified in different ways, but in the type of system discussed here they are assumed to vary temporally with patterns similar to those observed for some kinds of terrestrial ecosystems on Earth. Moreover, in order to further ensure that they remain similar to earthly weather patterns (e.g., in their long-term unpredictability) they cannot be directly influenced by the ecocyborg's control system. The conditions that arise as a result of these imposed, climatic forcing functions affect the values of other variables (e.g., relative humidity, vapor pressure, total pressure, soil water table level and soil available water) and, hence, the overall comportment of the ecocyborg. Other, truly external phenomena, that might be taken into account include: meteor impacts, solar flares, wear and failure of subsystems, and intervention by other entities (perhaps humans or other ecocyborgs). For the sake of simplicity, however, none of these are considered here.

5.5 Ecocyborg model

A modular style has been adopted in the model of the type of ecocyborg described. Thus, the overall computational model comprises a number of smaller, configurable models of the ecosystem, the control system and the external forcing functions, each of which is written separately and later implemented together with the others in a simulation. The enclosure is not explicitly modeled, but instead is implicit in the overall parameters, boundary conditions, and functionality of the other models. The different approaches that have been selected for the representation of each aspect of the ecocyborg are discussed below.

5.5.1 Ecosystem model

The ecosystem is modeled with a fairly high-resolution, object-based approach, in which all of its parts are represented by collections of distinct objects. The ecosystem model is a representation of the three realms of the ecosystem, and can be implemented in simulation so as to portray the processes that occur within and between these over time. It is completely configurable, so that different ecosystem constitutions and initial conditions (species lists, initial population sizes, mass allocation, topography, etc.) can be specified.

The representation of the biological component realm of the ecosystem adheres loosely to an individual-based, object-oriented paradigm, in which each organism (or small lump of organisms) is an instance of a species class (Hogeweg and Hesper 1990). Instances are modeled as distinct objects with attributes and methods that describe their states and possible behaviors, respectively. Each plant and animal population, therefore, is represented as a collection of instances of a corresponding species. In all, there may be up to 100,000 instances, representing as many as 1000 different species. Although the structure of the model can theoretically accommodate an unlimited number of species, 1000 has been set as an upper bound for this project, due to computer resource constraints. In order to limit the number of species that needed to be explicitly modeled, neither viruses nor other kinds of micro-organisms are explicitly represented. Instead, their activities are taken into account as an aspect of the environment in which they reside. For example, the soil is assumed to have an innate ability to decompose organic matter and fix nitrogen.

The encompassment and materials storage realms are also modeled as collections of discrete objects. The encompassment consists of a variegated terrain, made up of soil and water, and the ecosystem's atmosphere. The terrain is modeled as a grid of spatially explicit cells, covering a rectangular area equivalent to that of the unrolled cylindrical shell of the space station. Each cell has a number of properties that define its state (surface elevation, mass of decomposing plant material, etc.). Terrain processes that are modeled include subsurface water flow, nitrogen fixation, and decomposition. In contrast to the terrain, the atmosphere is modeled as a single object, and it is therefore assumed to be uniform throughout the ecosystem. The materials storage realm consists of four masses of reserve compounds in the solid state and is modeled with four corresponding objects,

The ecosystem model is, therefore, based upon interactions between discrete objects, both biological and non-biological. Although the model is not intended to be an exact representation of any particular type of ecosystem, the modeled biological components do, in general, mimic the traits and activities of terrestrial organisms, and the components in the other two realms allow for a reasonable representation of the biogeochemical processes that would occur in a similar physical system. Although each object, when considered independently, may not necessarily be an accurate representation of its physical counterpart, the combined interactions of all the objects does lead, when the model is implemented in simulation, to system-level comportment that exhibits at least some of the features that are common to all large ecological assemblages. Further details regarding the current implementation and architecture of the ecosystem model are given in Parrott and Kok (1999).

5.5.2 Control system model

Basal and instinctive mechanisms are considered as intrinsic to the ecosystem and are, as mentioned, represented accordingly in the ecosystem model. The remaining (Pavlovian and cognitive) control mechanisms are envisioned as extrinsic to the ecosystem, and are therefore modeled independently from the ecosystem. Since these would likely be implemented in an ecocyborg as computer software, they might well be very similar to the computational models now being created to represent them. Thus, the extrinsic control components of the hypothetical ecocyborgs and the configurable model thereof can be engineered and characterized in essentially the same way.

Models of the Pavlovian and cognitive controllers have been developed (Molenaar 1998; Molenaar and Kok 1995) based on previous work by Lacroix et al. (1996), and a more sophisticated cognitive controller is currently under development. The Pavlovian controller that Molenaar developed (Molenaar and Kok 1995) consists of a main body of computer software that, during simulation, coordinates the activity of a network of many other constituent mechanisms. Similar to the ecosystem model, the overall controller is completely configurable, allowing the number, type, connectivity, and other attributes of the constituent mechanisms to be specified.

A prototypal cognitive controller model was developed by Molenaar and Kok (1995) for testing in conjunction with the overall, composite ecocyborg model. This preliminary controller was based on an expert system that used rule sets to make decisions. In the future, cognitive control will be implemented according to the mental architecture proposed by Clark and Kok (1999b, presented here in Chapter 6), in a which a distributed, semihierarchical network of semidifferentiated control mechanisms provides the necessary infrastructure to support mental faculties such as perception, expression, memory, reason, and learning.

5.5.3 Forcing function model

The forcing functions (radiation, temperature and rainfall) that drive the ecosystem are also independently modeled. This model can be configured to emulate any given climate type, specified with a set of parameter values. The climates modeled here are, as mentioned, similar to weather observed on Earth, in accordance with the environmental requirements of the biological components in the ecocyborg. Temperature values are generated with a Fourier transform technique described in detail by Parrott et al. (1996b). Radiation and rainfall models have been developed based on a similar approach. These three variables are currently modeled independently, without consideration for the correlation between them in terrestrial climates (e.g., radiation intensity and rainfall rate).

5.5.4 Simulation of ecocyborg comportment

As described above, independent, configurable models have been created to represent the ecosystem, Pavlovian controller, cognitive controller and forcing functions of an ecocyborg. The overall comportment of such a system can then be portrayed via the parallel implementation of these models in simulation. This has been an effective approach, since it has enabled the independent development and testing of the different models. Simulations based on the ecosystem and forcing function models are, for example, currently being implemented in "open loop mode" (i.e., without control) as a means of assessing the performance of particular ecosystems in the absence of extrinsic control (Parrott and Kok 1999).

The selection of hardware and operating systems, and the development of a simulation platform for the effective implementation of the models are ongoing. Simulations have been executed on PC machines running IBM OS/2 Warp and Microsoft Windows 95/98, and on Apple Macintosh machines running System 8.5. A simulation workbench was developed by Lacroix et al. (1996) and further elaborated by Molenaar et al. (1995). This software makes use of the multitasking characteristics of OS/2 to allocate CPU time and memory space to the various models during the course of a simulation. Data is shared between the various models via the use of predefined shared memory segments. This workbench was used by Molenaar, for example, to test and develop his Pavlovian controller model in conjunction with somewhat rudimentary ecosystem, forcing function and cognitive controller models. Although this approach has been quite effective, it is anticipated that future simulations will require different hardware than what has been utilized to date, such as parallel-processing computers. For this reason, other methods of implementing the models in simulation are now being explored.

5.6 Conclusions

The EcoCyborg Project is about the creation of substantially autonomous biosystems, using a particular type of cyborged ecosystem as a case study. In contrast to recent scientific efforts in the field, the EcoCyborg Project is an engineering exercise, the objective of which is to learn how to specify the initial structure and composition of a biosystem such that its comportment will meet certain pre-defined criteria. In the case of biosystems of the ecosystem scale, examples of such criteria could include production quotas, long-term survival expectations, or minimal acceptable levels of species diversity. The future course of the EcoCyborg Project will be influenced by the progress made towards achieving the current, short-term research objectives. The specific objective of developing modeling and simulation tools with which to explore the engineering of ecocyborgs (and biosystems in general) will continue to be pursued through the ongoing improvement and refinement the computational models described here. As well, research will continue into effective methods for the characterization and comparison of various aspects of both these computational models and the modeled ecocyborgs (Clark and Kok 1999c, presented here in Chapter 7). Additional work that might be carried out includes the construction of a lab-bench physical ecocyborg and the creation of more sophisticated control systems that include advanced cognitive controllers. Such control systems will be integrated into ecocyborgs (of both predominantly physical and virtual construction) as a means of exploring a key premise of the research project: that the augmentation of a biosystem's mind in an appropriate manner is one means of increasing its viability (or autopoietic potential).

Due to the highly complex nature of biosystems, they are difficult to describe with current analytical approaches. Thus, a large part of the EcoCyborg Project is oriented towards the development of a philosophical framework, lexicon, and methodology for the generation and communication of knowledge about biosystems and how they may be engineered. These conceptual tools will provide engineers with a means of developing design principles specific to substantially autopoietic systems. Hence, they will be able to more effectively treat biosystems engineering projects. It is the intent that the work being pursued in the EcoCyborg Project will lead to the establishment of such general design principles, and that these will be applicable not just to ecocyborgs, but also to the wider class of biosystems. This knowledge will facilitate a broad range of applications, including the remediation of damaged terrestrial ecosystems, the development of more efficient production-oriented biosystems, the design of completely novel living systems.

5.7 References

Bonabeau, E., and G. Theraulaz. 1994. Why do we need artificial life? *Artificial Life* 1(3):303-325.

Bourgine, P., and F. Varela. 1992. Introduction: Towards a practice of autonomous systems. In *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, ed. F. J. Varela and P. Bourgine, 143-152. Boston, MA: MIT Press.

Brillouin, L. 1951. Physical entropy and information. II. *Journal of Applied Physics* 22:338-343.

Broda, E. 1983. *Ludwig Boltzmann: Man, Physicist, Philosopher*. Woodbridge, CT: Ox Bow Press.

Clark, O.G., and R. Kok. 1999a. Characterizing biosystems as autopoietic entities. *Oikos*, (submitted for publication).

Clark, O.G., and R. Kok. 1999b. Mental architecture of a cyborged biosystem. *Engineering Applications of Artificial Intelligence*, (submitted for publication).

Clark, O.G., and R. Kok. 1999c. Characterization of biosystems. *Biometrika*, (submitted for publication).

Clark, O.G., R. Kok and R. Lacroix. 1999. Mind and autonomy in engineered biosystems. *Engineering Applications of Artificial Intelligence* 12(3):389-399.

Clark, O.G., R. Kok and P. Champigny. 1997a. Generation of a virtual terrain. *Environmental Software* 12(2/3):143-149.

Clark, O.G., R. Kok, R. Molenaar and L. Parrott. 1997b. The implementation of mind in engineered biosystems. In *Proceedings of the 1997 Conference of the Canadian Society for Civil Engineering, Sherbrooke, Québec, May 27-30, 1997, Volume A: Agricultural, Food, and Biosystems Engineering*, 582-589. Montreal, QC: CSCE.

Clark, O.G., R. Kok and R. Lacroix. 1996. Engineering intelligent biosystems. NABEC Paper No. 9662. St. Joseph, MI: ASAE.

Dyson, G. 1997. *Darwin Among the Machines*. Redwood City, CA: Addison-Wesley.

Green, D.G. 1993. Emergent behavior in biological systems. In *Complex Systems - From Biology to Computation*, ed. D. G. Green and T. J. Bossomaier, 24-35. Amsterdam, the Netherlands: IOS Press.

Hogeweg, P., and B. Hesper. 1990. Individual-oriented modeling in ecology. *Mathematical and Computational Modeling* 13(6):83-90.

Holland, J. H. 1995. *Hidden order: How adaptation builds complexity*. Reading, MA: Addison-Wesley.

Kaneko, K. 1994. Chaos as a source of complexity and diversity in evolution. *Artificial Life* 1(1/2):163-177.

Kauffman, S. 1995. *At home in the universe: the search for the laws of self-organization and complexity*. New York, NY: Oxford University Press.

Kok, R., and R. Lacroix. 1993. An analytical framework for the design of autonomous, enclosed agro-ecosystems. *Agricultural Systems* 43:235-260.

Kok, R., L. Parrott, O.G. Clark, R. Molenaar and R. Lacroix. 1995. The engineering of autonomous biosystems: The EcoCyborg Project. CSAE Paper No. 95-602. Saskatoon, SK: CSAE.

Lacroix, R., R. Kok and O.G. Clark. 1996. Use of a multitasking operating system as a setting for the simulation of an enclosed agro-ecosystem under cognitive control. *Canadian Agricultural Engineering* 38(2):129-137.

Maturana, H.R., and F.J. Varela. 1980. *Autopoiesis and Cognition: the Realization of the Living*. Dordrecht, the Netherlands: D. Reidel Publishing Company.

Molenaar, R. 1998. Design and implementation of biosystem control for biosystem simulation. Unpublished Ph.D. thesis. Department of Agricultural and Biosystems Engineering, McGill University, Montreal, QC.

Molenaar, R., R. Kok and R. Lacroix. 1995. Structure of an EcoCyborg simulation under OS/2. CSAE Paper No. 95-604. Saskatoon, SK: CSAE.

Molenaar, R., and R. Kok. 1995. Development of a Pavlovian controller for the guidance of autonomous biosystems. NABEC Paper No. 9523. St. Joseph, MI: ASAE.

Nelson, M, P. Hawes and P. Augustine. 1992. Life systems for a lunar base. In *The Second Conference of Lunar Bases and Space Activities of the 21st Century*, vol. 2, NASA Publication 3166, 513-518. Houston, TX: NASA.

Parrott, L. 1995. The EcoCyborg Project: A model of an artificial ecosystem. Unpublished M.Sc. thesis. McGill University, Montreal, QC.

Parrott, L., and R. Kok. 1999. Baseline performance of an ecosystem model and simulation for the EcoCyborg Project. ASAE Paper No. 995041. St. Joseph, MI: ASAE.

Parrott, L., R. Kok and R. Lacroix. 1996a. Daily average temperatures: Modeling and generation using a Fourier transform approach. *Transactions of the ASAE* 39(5):1911-1922.

Parrott, L., R. Kok, O. G. Clark and R. Molenaar. 1996b. Crafting new life engines. *Resource* 3(3):8-10.

Patten, B.C. 1959. An introduction to the cybernetics of the ecosystem: the trophic-dynamic aspect. *Ecology* 40(2):221-231.

Pimm, S.L., J.H. Lawton and J.E. Cohen. 1991. Food web patterns and their consequences. *Nature* 350:669-674.

Ray, T. 1994. An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artificial Life* 1(1/2):179-209.

Schrödinger, E. 1955. *What is Life? The Physical Aspect of the Living Cell.* Cambridge, MA: Cambridge University Press.

Shannon, C. 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27(July-October):379-423, 623-656.

Tolkien, J.R.R. 1966. *The Hobbit: Or There and Back Again.* New York, NY: Houghton-Mifflin.

Taylor, C., and D. Jefferson. 1994. Artificial life as a tool for biological inquiry. *Artificial Life* 1(1/2):1-13.

CONNECTING TEXT

Chapter 6, **Mental architecture of cyborged biosystems**, was authored by O.G. Clark and R. Kok. At the time this thesis was submitted, this chapter had been sent for review to the editors of the journal *Engineering Applications of Artificial Intelligence*.

This chapter is an extension of Chapter 5, where the conceptual streams of *biosystem* and *autonomy* were brought together in the guise of the EcoCyborg Project. Here, an explicit description is given of how the technological control system of a cyborged ecosystem might be structured so as to give rise to the mental abilities necessary for substantial autonomy. Some previous research has been done in this respect by Kok, Desmarais, Gauthier, Lacroix, and Molenaar (see the literature reviews in Chapters 2 and 6), but in this chapter an architecture is explicitly described so as to set the direction for future work in the EcoCyborg Project.

CHAPTER 6. MENTAL ARCHITECTURE OF CYBORGED BIOSYSTEMS

Abstract

This paper is a discussion of how biosystems can be combined (cyborged) with control systems to create new, substantially autonomous biosystems. Here, a biosystem is considered as an integrated group of component entities that is collectively alive to some degree, and a control system as a set of components meant to influence its comportment. The cyborging of ecosystems for substantial autonomy is examined as a case study within this context. In order for a cyborged ecosystem to be substantially autonomous, its collective computational apparatus (i.e., mind) must enable it to reason using a self-model, so that it might formulate and actively pursue its own goals. A mental architecture is described that would accommodate this, comprising an object-oriented knowledge base and cybernetic mechanisms that process the information contained therein. This architecture is characterized in terms of functionally semidifferentiated, intermediate-scale components arranged in a semihierarchical control organization. A detailed description is given of how a mind based on such an architecture could give rise to the required abilities of perception, expression, memory, reason, and learning. Some progress has already been made toward implementing such a cyborged ecosystem in a completely virtual setting, and this paper outlines how this work might be furthered.

6.1 Introduction

This paper is a discussion of how an *original biosystem* and a *control system* can be combined (*cyborged*) to create a new, substantially autonomous biosystem. Here, the term *biosystem* refers to an integrated group of component entities that is collectively alive to some degree (Clark and Kok 1999), and a *control system* is a set of components that are intentionally added to a system in order to influence its comportment in a particular, predetermined manner. The context of the paper is the EcoCyborg Project, in which the long-term goal is the development of a general theory for engineering biosystems so that they meet specific design objectives, such as substantial system autonomy (Molenaar 1998). One approach that can be adopted for the engineering of such entities is the cyborging of an existing biosystem with a control system.

Accordingly, the short-term goal of the EcoCyborg Project is to investigate how to cyborg a particular kind of original biosystem, an ecosystem, so as to yield a substantially autonomous *ecocyborg*. Much of the reasoning presented in this paper has, therefore, been developed in the context of this short-term goal and is illustrated with reference to the engineering of ecocyborgs. The concepts are nevertheless intended to be relevant to the engineering of a wide variety of cyborged biosystems.

The achievement of substantial autonomy in a cyborged biosystem requires that it transform input signals into output signals (i.e., compute) so that the overall effect is the formulation and active pursuit of its own goals (Clark et al. 1999). Computation is considered here to arise from virtual machines, which are apparatus or contrivances, such as algorithms, that perform informational functions (Clark et al. 1999). Collectively, all of the virtual machinery encompassed by a biosystem is interpreted as mind, the potential activities that might arise from this as mental abilities, and the computation that actually takes place as mentation. Thus, when an original biosystem is combined with a control system, the result is a cyborged biosystem whose mind consists of the sum total of the virtual machinery resident on both parts, and that is oriented toward the fulfillment of certain design objectives, such as substantial autonomy.

The first part of the paper begins with a more detailed review of the lexicon used, which is set in the paradigm of cybernetics. There is then some exploration of how the characterization and comparison of biosystems can be impacted by the resolution at which they are observed. As well, there is a discussion of how characterization and comparison can be affected by the breadth at which the lexicon is interpreted, whereby the meanings of the terms are taken to be more or less inclusive. Finally, all of these ideas are used in the description and justification of the engineering approach that has been adopted in the EcoCyborg Project.

The subsequent part of the paper is a discussion of *mental architecture*, or the general plan after which a mind might be patterned. It is surmised that since humans are the most autonomous entities that are readily available for study, the minds of cyborged biosystems will be engineered to emulate some aspects of the human mind. Mental architecture is therefore discussed with reference to humans and computational models of the human mind. The three principal aspects of mental architecture that are dealt with are *distribution*, *differentiation*, and *control organization*. The first aspect refers to the number of components that a mind comprises, the second to the degree of specialization of these components, and the third to the flow of causality among them.

A specific architecture is then proposed for the minds of cyborged biosystems, and described in terms of its applicability to ecocyborgs. The overall character of the proposed mental architecture is described with reference to the three aspects introduced previously. It is further proposed that this architecture be based on the interaction of two broad types of virtual machines: *knowledge objects* that form a shared, object-oriented knowledge base; and *cybernetic mechanisms* that transduce signals into information and vice versa, and also process the information contained in the knowledge base. In the proposed ecocyborgs, all the components of mind will reside entirely on the control system, although this need not be so for all types of cyborged biosystems. It is so for ecocyborgs because natural ecosystems do not possess anything comparable to the virtual machinery of the human central nervous system.

After this, a detailed description is presented of how an ecocyborg could be engineered to have a mind, based on the proposed architecture, that gives rise to mental abilities similar to those of humans. The abilities are grouped into sets called *faculties*, of which five are required in order to successfully emulate a human: perception, expression, memory, reason and learning. Each of these faculties is examined in turn, followed by a final discussion of how an ecocyborg with such a mind might fulfill the design goal of significant autonomy.

6.2 Background

The engineering of biosystems has applications ranging from the improved management of natural ecosystems to the creation of highly autonomous, task-oriented, artificial biosystems of all kinds. For such work, a body of engineering theory is required, the development of which is the long-term goal of the EcoCyborg Project. This goal is being pursued in two main phases. The first phase is an investigation of the relationship between biosystem *constitution* (which encompasses *composition*, or the number and kinds of components, and *structure*, or the interrelationships between the components), *initial state* (the starting values of the attributes of the components), and the resultant *comportment* (how the state changes with time) under external influences, such as forcing functions. This phase is essentially a scientific effort, since it addresses the "forward problem" of formulating general, conceptual models, as well as detailed, computational models on which simulations can be based. Accordingly, in the early stages of the project, the focus has been on the creation of methods and tools for the characterization of biosystems, as well as for the computer-based modeling and simulation of the latter (Clark and Kok 1999). As the methods and tools become more refined, they are providing the knowledge and insight necessary for the second phase of the project, an engineering effort that addresses the "reverse problem" of how biosystems might be designed, constructed, maintained, repaired, or upgraded in order to fulfill certain objectives. Some of the questions that are being addressed in this phase include: (a) what ranges of composition, structure, and initial states will lead, under a given set of forcing functions, to a specific, desired kind of comportment, and (b) which particular possibilities in these ranges are the optimum ones, as evaluated according to some given criteria? Since the desired comportment that is addressed in this article is mentation that gives rise to substantial autonomy, the engineering of the biosystem's mind is implied in these objectives.

An essential prerequisite for the effective engineering of mind is a paradigm within which it can be characterized, that is to say, within which adequate descriptions can be generated. The stance adopted here corresponds well with a paradigm called *cybernetics*, promulgated by one school of thought in cognitive science. In this, mind is considered to be a network of control and communication mechanisms (von Neumann 1963). Although it is not universally agreed that the cybernetic paradigm is actually suitable for the description of all aspects of mind (Penrose 1989), the discussion in this paper is based on the assumption that it does accommodate the description of those aspects of mind that give rise to substantial autonomy both in original biosystems, such as humans, as well as in cyborged biosystems, such as ecocyborgs. The virtual machinery of which mind is composed therefore can be taken to consist of cybernetic mechanisms and knowledge objects.

Once a suitable paradigm has been adopted for the characterization of mind, a lexicon is required for the description of how mind gives rise to mentation in biosystems and, ultimately, fulfills the long-term goal of substantial autonomy. Accordingly, the development of such a lexicon has been one stream of the EcoCyborg Project (Clark and Kok 1999; Clark et al. 1999). In this lexicon, every biosystem is considered to possess a mind that comprises all of its virtual machinery. The sophistication of the mental abilities that arise from these is *intelligence*. Thus, mind gives rise to an epistemic space that is spanned by these mental abilities and within which actual mentation takes place. A biosystem of appropriate intelligence is able to mentate using a model of itself in the context of its environment. This aspect of mentation is defined as *consciousness* and, in turn, is the basis for *autonomy*. A substantially autonomous biosystem is *automatic* (acts without external guidance), *volitive* (capable of formulating its own goals and strategies for achieving them), and *intentful* (actively pursues those goals and strategies) (Clark et al. 1999). All of these characteristics are considered here to be measurable on continuous scales, rather than as discrete, binary properties that are either absent or present. Accordingly, a system can be conscious, autonomous, etc., to any given degree.

The paradigm and lexicon outlined above can be used to characterize and compare the minds of biosystems. Because the lexicon defines phenomena of a primarily virtual nature, it can be used to formulate descriptions and make meaningful comparisons of biosystems that are physically very different. This is to say that, although the virtual machinery of which any mind is composed (and, equivalently, the information content of input or output signals) must always reside on a physical substrate, the nature of that substrate is not actually of theoretical importance in the context of cybernetics because it does not determine the essential nature of the machinery (or, equivalently, of the information inherent in the signals). Thus, biosystems that are physically different from one another can possess minds that are qualitatively quite similar, and vice versa. This is illustrated below with reference to ecosystems and humans, where the former are the original biosystems being used in the creation of ecocyborgs, and the latter serve in this paper as the archetypes of substantially autonomous biosystems.

Entities such as ecosystems and humans can be characterized and compared at various scales and the choice of scale can strongly impact such analysis. For instance, at extremely fine scales the meaning of even the most fundamental terms defined above, such as *computation* or *mentation*, can be unclear, so that the given lexicon becomes less useful. At slightly coarser scales, distinct virtual machines might be identifiable and their activities more definitely interpretable as mentation, but other concepts might not yet be applicable. For example, there might be no way to establish any correspondence between a particular virtual machine and another phenomenon, in which case no symbolic representation could be inferred and syntactic operations of symbol manipulation would not be relevant. In this case, mentation at scales finer than a certain threshold is therefore referred to as being *subsymbolic*.

The distinction between symbolic and subsymbolic mentation is not always associated with scale, but in fact depends on the establishment (or not) of correspondence between some (usually variable) aspect of the physical substrate with another phenomenon. Although this is often determined by the scale at which an observer can resolve the features of a system, this is not necessarily so. There is, for instance, mentation inherent in the interactions between the components of a large-scale ecosystem, such as the organisms that it includes (Patten and Odum 1981; McNaughton and Coughenour 1981). This is not, however, normally considered to be symbolic simply because present-day, scientific observers do not attribute meaning to them. Human mentation is similarly founded on the interaction of myriad components. This interaction involves not only the activity of the central nervous system, but also of the endocrine and digestive systems, the musculature, and so on, to include the functioning of the entire body, much of which is subsymbolic. Natural biosystems such as ecosystems and humans might also, as a result of their having been shaped by selective evolutionary forces, encompass intermediate-scale components that are identifiable as distinct virtual machines (Pinker 1997; McNaughton and Coughenour 1981; Patten and Odum 1981). Finally, at coarse scales, both types of biosystems transform input signals into output signals and so, using the given lexicon, instances of either type can be considered as virtual machines at the system level. Hence, the minds of humans and natural ecosystems are not entirely dissimilar from one another, because both types of biosystems comprise virtual machinery at a range of scales.

The mental architecture proposed later for ecocyborgs is based on an intermediate-scale characterization of mind. For engineering purposes, characterization at such a scale is more appropriate than at coarser scales because it allows the internal composition and structure of mind to be resolved and, therefore, to be modified. It also allows the explicit characterization of the kind of virtual machinery that is, as described below, apparently unique to highly autonomous entities like humans. Characterization at intermediate scales is also more advantageous than at very fine scales since the latter can, at best, yield only an implicit description of the virtual machinery of interest whereas, for the purposes of engineering, an explicit description is required. Hence, although different biosystems can be characterized and compared at a variety of scales, it is not necessarily useful do so in a given context.

The characterization and comparison of biosystems is affected not only by the choice of scale, but also by the breadth with which the given lexicon is interpreted. The terms that the lexicon comprises can be defined more or less inclusively, to suit the context in which they are employed. When dealing with the general characteristics of an expansive class of biosystems, for instance, it might be convenient to use the very broad interpretation already introduced above, whereby *mind* is given an extremely inclusive definition that encompasses all of the virtual machinery of a biosystem. Following the example introduced in the previous paragraphs, the intermediate-scale virtual machinery of ecosystems and humans might then be compared according to this interpretation. Since, in this case, mind is taken to include all virtual components without discrimination on the basis of type, the minds of the two kinds of biosystems are not found to be dissimilar. In a more circumscribed context, however, a narrower interpretation of the lexicon might be used, whereby *mind* could be taken to denote only a particular subset of virtual machines. For instance, the term *mind* is often taken to denote only virtual machines that are similar to those embedded in the human central nervous system and, accordingly, only the activity arising from such machinery is considered to be mentation. (The particular kind of mentation that occurs at intermediate scale in the central nervous system of a human is referred to as *thought*.) When the lexicon is interpreted in this manner, a comparison of humans and ecosystems yields a quite different result from that obtained before. Whereas certain kinds of mentation, such as symbolic processing, are prominent in human thought, ecosystems do not seem to mentate in this way at all. Evidently, ecosystems lack the required intermediate-scale virtual machinery to support this kind of mentation.

The central nervous systems of humans (and apparently of all other highly autonomous natural biosystems, such as other mammals), supports intermediate-scale virtual machinery that is specialized for certain kinds of abstract mentation. This seems to indicate that substantial autonomy might actually depend on their presence. It is therefore speculated that biosystems, such as natural ecosystems, that do not seem to be highly autonomous at the system level are in fact less autonomous precisely because they lack such virtual machines. Hence, the engineering approach adopted in the EcoCyborg Project is to endow them with virtual machinery of this type. Accordingly, a correspondingly narrow interpretation of the given lexicon is used later in this paper in the description of the mental architecture proposed for substantially autonomous ecocyborgs.

The engineering approach described here involves the alteration of existing virtual machines, or the addition of new ones. On the one hand, as mentioned, natural ecosystems probably already possess intermediate-scale virtual machines, but not of the kind that give rise to the desired degree of autonomy at the system scale, nor do they necessarily fulfill any other design objectives. Biosystems that are largely or entirely artificial, on the other hand, result from human design and construction, and have not been shaped by evolutionary forces at all. Although their physical substance does support fine-scale virtual machinery by virtue of its ability to transform inputs signals into outputs, it is highly improbable that they would possess any machines of intermediate or coarse scales unless these were intentionally included. For both natural and artificial biosystems, therefore, a deliberate approach to the engineering of their virtual machinery must be adopted in order to achieve the goal of substantial autonomy.

Various approaches might be taken to engineering the virtual machinery of biosystems. One approach is to employ passive methods, whereby systems are created whose constitution can adapt according to circumstance, and that can thereby modify their own existing virtual machines, or acquire new ones as they "learn through experience". This can be an effective approach in situations, such as genetic programming, where the system to be engineered is inexpensive to modify, and is not of great initial intrinsic value. In other circumstances, the limited use of this approach is actually an essential aspect of a system's mentation. This is the case, for instance, if a system is to be adaptable or able to learn. However, for large, physical systems it can be very slow and expensive. Its exclusive application would be especially inappropriate, for example, in the engineering of the physical aspect of a valued natural ecosystem that was being modified in hope of its preservation.

An alternative approach is to employ active engineering methods. Although these could be used at any of the various scales at which mind can be characterized, in terms of the objectives of the EcoCyborg Project there are, as discussed above, disadvantages to working at very fine or very coarse scales. The approach adopted is, therefore, to actively engineer the mind of biosystems (such as ecosystems) at an intermediate scale. The intent is to create significantly autonomous entities by incorporating the required virtual machinery into cyborged biosystems in the form of control systems. The overall architecture of this virtual machinery is explained later in terms of some particular design considerations, described in the next section with reference to the human mind. In accordance with the adopted engineering approach, a narrow interpretation of the lexicon is used in the remainder of the paper, whereby *mind* is taken to refer only to intermediate-scale virtual machinery similar to that resident on the human central nervous system.

6.3 Architectures of mind

Because human beings are the most autonomous entities that are readily available for study, notions about their minds greatly influence efforts to create other kinds of substantially autonomous entities. Accordingly, a productive interaction exists between the fields of human psychology, neurology, and artificial intelligence (AI). The engineering of cyborged biosystems will likely follow in this pattern and, hence, their minds will probably be at least somewhat similar to those of humans. Despite the likely similarities, however, those minds will also be shaped by technological availability, practical expediency, and the substrate in which they are embedded, so that they will also differ considerably from the human archetype (Dyson 1997).

In order to set the context for proposing a mental architecture for cyborged biosystems (particularly for ecocyborgs), the general plan of a mind, i.e. *mental architecture*, is discussed in this section. The three aspects that are dealt with are distribution, differentiation, and control organization. Each of these is explained with reference to the human mind, as well as to AI models thereof. As mentioned, an intermediate scale is emphasized in this section because the characterization of mind at this scale is most appropriate for the engineering approach adopted in the EcoCyborg Project. The use of a particular scale in this way affects how mind is perceived, and so this impact is also considered.

6.3.1 Distribution

One aspect of mental architecture is *distribution*, or the number of components that a mind comprises. The components that are referred to here are virtual machines (cybernetic mechanisms and knowledge objects). Although these might sometimes correspond closely to elements of the physical substrate on which they reside, this need not always be the case; a single physical element might host many distinct virtual machines, or a single virtual machine could be spread across several physical elements.

A mind can be, at one extreme, an undivided entity or, at the other, an aggregate of very many components. Of course, the number of components that can be perceived in a particular mind is limited by the resolution at which that mind is observed. Thus, at a coarse resolution of observation it is impossible to distinguish fine-scale components, so that even if a mind is in reality comprised of very many of these, it will nevertheless be perceived and characterized as a single, undivided entity. With greater resolution of observation, however, components of finer scale can be distinguished. Under these circumstances, if the number of distinguishable components approaches the number that could potentially be discriminated, then the mind being examined has a *highly distributed* architecture at that scale. The term *semidistributed* then refers to a mind that is characterized as being composed of numerous, distinct virtual machines, but considerably fewer than the maximum number that is potentially discernable. As well, these might be of a coarser scale than the finest that could be distinguished. Finally, a *unitary* mental architecture is one in which a mind is characterized as being a single, cohesive unit.

The human mind has, in the past, generally been characterized at an extremely coarse scale. This is because methods and technologies for distinguishing fine-scale physical elements and the virtual machines that they support have not been available until relatively recently. The human mind has therefore traditionally been considered to be a unitary entity. This perception has probably been reinforced by the image that most people have of themselves as possessing a single consciousness that can be effectively focused on only one task at a time. This description is not entirely false since, in most cases, a human mind does function as a relatively cohesive unit.

Unitary mental architecture, in which one agent performs all the computational tasks, was the basis of many early AI applications. Even today most commercially available expert systems, for instance, still consist of a central procedural algorithm (an inference engine) that responds to inputs by serially firing individual rules from a rule base. This approach continues to be sustained by the available technology; to date, the design of the vast majority of digital computers has been based on the serial architecture first implemented by Von Neumann, following Turing's theory of computing automata, in which all computation is performed by a single, central processing unit (Dyson 1997; von Neumann 1963). Once such a bias is established, it tends to be self-perpetuating because it is more practical for designers to opt for a familiar architecture rather than experiment with new ones.

When the human nervous system is observed at increasingly fine resolution it becomes evident that nerve tissue is composed of vast numbers of individual cells. Thus, as the resolution of observation increases, so does the number of physical elements that can be descried. Even small areas of tissue appear to be substantially functionally independent from one another, as evidenced by, for example, the study of surgery patients and accident victims (Sacks 1998). This notion is being further corroborated as improved technology enables the observation of metabolic and electromagnetic activity in greater detail inside the tissues. The resulting evidence seems to indicate that there is also some merit in a description of mental activity and virtual machinery in humans as being highly distributed, corresponding directly to the physical distribution of the nervous system.

The characterization of human mental architecture as being highly distributed has inspired the *connectionist* approach to AI. This was originally conceived when digital computers were first being constructed (McCulloch and Pitts 1965) and has subsequently led to the development of a variety of models of the human mind. These *artificial neural networks* (ANNs) comprise large collections of relatively simple components, which are interconnected so that the output signal of each is received as an input by one or many others. The computational character of an ANN is inherent in the strength and configuration of the connections between the components, and the way in which each component transforms the signals that it receives. As a whole, therefore, an ANN can perform quite sophisticated tasks.

As understanding of the human mind has improved it has become evident that although its architecture is ultimately highly distributed at fine scales, it is, nevertheless, the activities of distinct, virtual machines of intermediate scale that actually are of paramount importance in human thought (Pinker 1997). These virtual machines are each much more sophisticated than those that are resident on individual neurons, but are still considerably simpler than the overall mind. They perform the high-level mentation, such as language processing, that is commonly deemed to be characteristic of human thought. If it is maintained, as it is here, that the machinery at intermediate scales is the most crucial to substantial autonomy, then characterization of the human mind as having a semidistributed architecture is in fact the most appropriate.

The acknowledgement of the significance of mental phenomena at intermediate scales is reflected by the fact that many semidistributed AI models of the human mind have been created, and that successful emulation of the human mind has been based on these. Work in this area was significantly influenced by experiments with communities of relatively sophisticated machines called Perceptrons (Minsky 1985, p.330; Minsky and Papert 1969). These can learn to evaluate evidence for the presence of particular patterns in data. Individually, they are not very effective, but in community some can discern relationships between the simple patterns recognized by others. Abstraction of metapattern from pattern in this way enables the community to perform more complicated computational tasks than can be performed by any individual agent. Another significant contribution to the intermediate-scale modeling of the human mind was the development of *blackboarding*, first implemented in the Hearsay II program, wherein many semi-independent agents share a single, distributed information storage construct (Lesser et al. 1975). More recently, Hofstadter (1995) created programs such as Copy Cat and Table Top, which consist of aggregates of information-processing agents that transform input signals into outputs. They operate in abstract microdomains, such as spaces of alphanumeric character strings, and identify analogies by recognizing relationships between informational constructs. Although rudimentary in comparison to human thought, this activity depends on the implementation of many of the mental abilities that underlie autonomy in humans.

As mentioned above, the mental architecture of cyborged biosystems will probably initially be patterned after that of humans. In the case of an ecocyborg, its mind will consist almost exclusively of a control system that, in the short term, will likely reside on physical elements such as digital computers. Thus, when characterized at very fine scales, the mental architecture of ecocyborgs will also be highly distributed, being composed, for instance, of myriad binary switches resident on the transistors of the computers. As pointed out it is, however, currently infeasible to engineer mind at a very fine scale without consideration of components of intermediate scale. On the other hand, however, it is maintained here that similar intermediate-scale virtual machines can be engineered without regard to the exact nature of the fine-scale components of which they are composed. Since it appears, moreover, that virtual machinery of intermediate scale is that which is of paramount importance in the human mind, this is likely the scale that will be emphasized in the characterization and engineering of artificial minds in the near future. The proposed mental architecture is, therefore, of a semidistributed nature, i.e., the mind will consist of numerous components. These could all be functionally identical to one another, or they might each give rise to slightly or even radically disparate abilities. This aspect of mental architecture is examined in the next section.

6.3.2 Differentiation

Differentiation, in the context of mind, is the degree to which virtual machines are specialized so as to give rise to particular mental abilities. Like distribution, the utility of this attribute depends somewhat on the resolution that is used in the observation of mind. Thus, when a very coarse resolution of observation is used, characterization is only possible at extremely coarse scales. The concept of differentiation is therefore not really applicable in this case because individual mental components cannot be distinguished. At fine scales of characterization, on the other hand, the concept of differentiation is relevant, and virtual machines might be differentiated to any extent. Nevertheless, as pointed out, characterization of mind at very fine scales alone is not appropriate in the context of the EcoCyborg Project. Differentiation is, however, both applicable and appropriate at the intermediate scale that is of interest here. When thus characterized, the mind of a biosystem might be revealed, at one extreme, to be very highly differentiated, so that each mental ability arises from a separate virtual machine. At the other extreme, the mind might be completely undifferentiated, so that each virtual machine gives rise to an identical set of mental abilities. In general, highly autonomous natural biosystems appear to have a semidifferentiated mental architecture that falls somewhere between these two extremes.

The human mind seems, for instance, to be semidifferentiated over a range of scales, as evidenced by the functionality of the various parts of the central nervous system, on which it resides. At a very coarse scale, the two hemispheres of the brain are somewhat specialized for particular mental abilities, although there is a good deal of overlap between them. At a very fine scale, there exist a number of different kinds of neurons, and even those of the same type vary physiologically, which seems to indicate that there are functional differences between them (Penrose 1989). There are also indications that the human mind is semidifferentiated at an intermediate scale. When localized regions of the brain are damaged, for example, the impairment of very specific mental abilities often results (Sacks 1998). Furthermore, imaging techniques such as positron emission tomography (PET) scanning show that activity in the various regions of the brain (and, presumably, the activity of the virtual machines that these support) is often correlated with very particular kinds of thought.

Semidifferentiated mental architecture can balance the advantages as well as the disadvantages of both undifferentiated and highly differentiated arrangements. An undifferentiated architecture can be costly to maintain, especially if resources are limited. For instance, if some mental abilities are needed at the local scale only infrequently, and if they arise from many undifferentiated virtual machines, then these machines will be underutilized much of the time. Some degree of differentiation is, therefore, advantageous in such a situation because specialized virtual machines would be more fully utilized, and fewer would be required if the corresponding tasks were performed centrally. If, on the other hand, some tasks are locally very common, then their execution by relatively undifferentiated, local machines requires less time than is necessary to communicate them to a single, specialized machine, await their completion, and retrieve the results. Moreover, the completion of such tasks is not delayed until a specialized machine becomes available. Any other mentation that depends on their prior execution can, therefore, also proceed more expediently. Hence, extreme differentiation is not necessarily beneficial in all circumstances.

For these reasons, therefore, most artificial minds are likely to be based on semidifferentiated architectures, and this is the approach taken here. The primary differentiation among the virtual machines is between knowledge objects and cybernetic mechanisms. The first type of machinery stores information in a relatively passive manner and forms the object-oriented knowledge base (OOKB) of the ecocyborg's mind (Gauthier and Guay 1998; Zeigler 1990). Machines of the second type actively transform information, and interact with one another via the OOKB. They can be further differentiated according to activity class, goal type, priority, and implementation level, as discussed below in more detail (Kok and Lacroix 1993). This differentiation is illustrated in Figure 6.1. The organization of the interactions of the machinery is the subject of the next section.

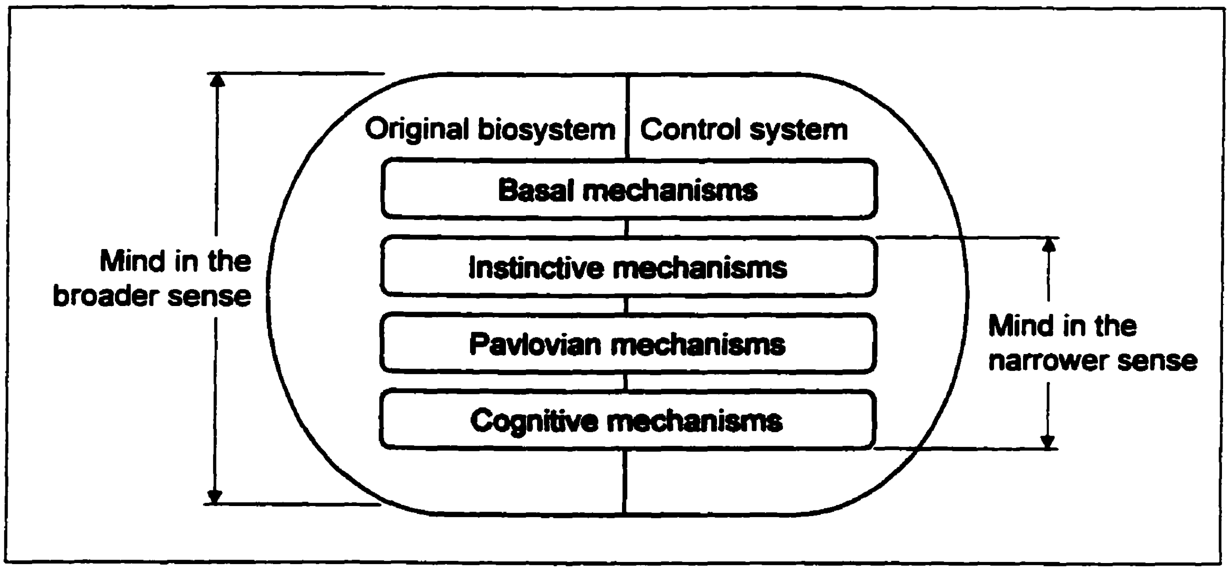

Figure 6.1. Virtual machinery of a cyborged biosystem shown differentiated according to implementation level.

6.3.3 Control organization

As discussed, a sophisticated mind is likely to comprise numerous intermediate-scale, semidifferentiated, virtual machines. The effective engineering of such a mind requires consideration of control organization, or the flow of causality among these machines. One extreme type of organization is centralized control, in which all mentation is orchestrated by a single agent. A variation of this is a linear hierarchy similar to that of, for example, a business or military unit in which a series of middle managers resolves high-level strategies into low-level tactics. In such an arrangement, ultimate authority still resides with the agent at the top of the hierarchy, but some decision-making power is delegated to subservient machines.

A number of AI programs have been based on architectures with centralized or linearly hierarchical control organizations. In Winograd's SHRDLU program, for example, directives originating from an external (human) operator are parsed and passed on to subroutines for further parsing until the level of physical expression is reached (Waldrop 1987). Supervisory digital control schemes are also based on this type of organization, wherein the overall objectives are presented to a central machine that determines optimal setpoints for subsidiary controllers which, in turn, issue directives to final control elements.

The extreme opposite of linearly hierarchical or centralized control is a flat organization in which many agents share power equally. Although it is often assumed that order in a system must stem from a central or hierarchical control organization, a flat control structure does not at all imply a lack of order. On the contrary, systems with the latter kind of architecture frequently display highly ordered comportment. Resnick (1994) has pointed out that pattern in ant colonies and traffic jams, for example, often emerges as the synergistic product of interactions among many relatively independent components, none of which has a disproportionate amount of influence on the comportment of the overall system. When the human mind is characterized at a very fine scale it becomes evident that it too has an approximately flat control organization at this scale, since no particular neuron has significantly greater influence on the overall system than does any other.

A number of authors have devised models of the human mind that are composed of collections of agents arranged in an approximately flat control organization. Minsky (1985), for example, has described an egalitarian "community of mind" in which the component agents are not subject to any hierarchical control at all; Jackson (1987) has written about how these agents might coexist either passively, cooperatively, or competitively.

At the intermediate scale that is of interest here neither centralized, linearly hierarchical, nor flat control organization is an entirely satisfactory description of the human mind. It appears that control at this scale is partly, but not completely, concentrated in specific components, and a semihierarchical description of the mind's control organization is therefore probably the most accurate. Moreover, the flow of causality in the human mind does not seem to be entirely rigid, but to shift with changing circumstances. For instance, when there is time to contemplate a novel problem, slower agents that perform conscious reasoning might maintain authority, while others enact whatever strategies are conceived. When the mind is confronted with a routine or extremely time-critical task, however, very fast virtual machines that act without conscious mentation might assume control, as in reflex responses where nerve impulses do not even reach the brain before action is taken.

Models of the human mind have also been devised that are based on this kind of flexible, semihierarchical control organization, as exemplified by the computer programs Copycat and Tabletop (Hofstadter 1995). In these programs there is a network of generative cybernetic mechanisms, each of which creates local mechanisms that perform certain information-processing tasks. Depending on the relative frequency with which the local mechanisms successfully execute their tasks, the corresponding generative mechanisms produce local ones at a faster or slower rate. This interaction of the generative and local mechanisms thereby causes the overall system to mentate in a context-sensitive manner.

The control organization proposed for the minds of ecocyborgs is also a semihierarchical one, and is somewhat similar to that implemented in Copycat and Tabletop. Hence, the selective replication of virtual machines will be key to making the mind context-sensitive. It is also envisioned that control will be shared among the cybernetic mechanisms of an ecocyborg's mind in a way similar to the flexible organization of the human mind. Thus, the flow of causality, too, will shift depending on the context. These ideas are described below in greater detail.

6.4 Mental architecture of an ecocyborg

The long-term goal in the EcoCyborg Project is, as mentioned, to develop a general theory for engineering biosystems so that they meet specific design objectives such as substantial autonomy. The approach that is being followed in the short term is to combine existing biosystems and technological components so that the resulting entities have minds (in the narrow sense) that are reminiscent of those of humans. In the general case of a cyborged biosystem, the mental abilities required to emulate human thought might arise from both the original biosystem and the control system. Many kinds of biosystems, such as animals, have sophisticated nervous systems that host virtual machinery which would contribute significantly to the mind of a cyborged entity. Natural ecosystems, however, generally encompass little or no virtual machinery that might give rise to these kinds of abilities. In the mental architecture that is envisioned, therefore, an ecocyborg's mind (in the narrow sense) resides almost exclusively on the added technological components.

As discussed above, the added components are envisioned as hosting virtual machines of intermediate scale that engage one another in a flexible, semihierarchical control organization (Lacroix and Kok 1999; Molenaar 1998; Kok and Lacroix 1993; Gauthier and Kok 1989; Kok and Desmarais 1988). Moreover, that machinery is also semidifferentiated, divisible first into passive knowledge objects and active cybernetic mechanisms. The first set of machinery encodes a shared pool of information and serves to mediate communication among members of the second set which, as described further below, is roughly subdivided into the instinctive, Pavlovian, and cognitive controllers. The mental abilities to which all of this virtual machinery gives rise can likewise be subdivided into (possibly overlapping) sets called *mental faculties*, not unlike the human ones described by Gardner (1993). For an ecocyborg to successfully emulate the autonomy of humans to a reasonable degree, it should be engineered to possess at least five mental faculties (Figure 6.2): *perception*, *memory*, *reason*, *learning*, and *expression*

(Clark et al. 1999). Some progress has already been made toward implementing, in a completely virtual setting, an ecocyborg with such faculties (Molenaar 1998). The remaining sections are a description of how this work might be furthered.

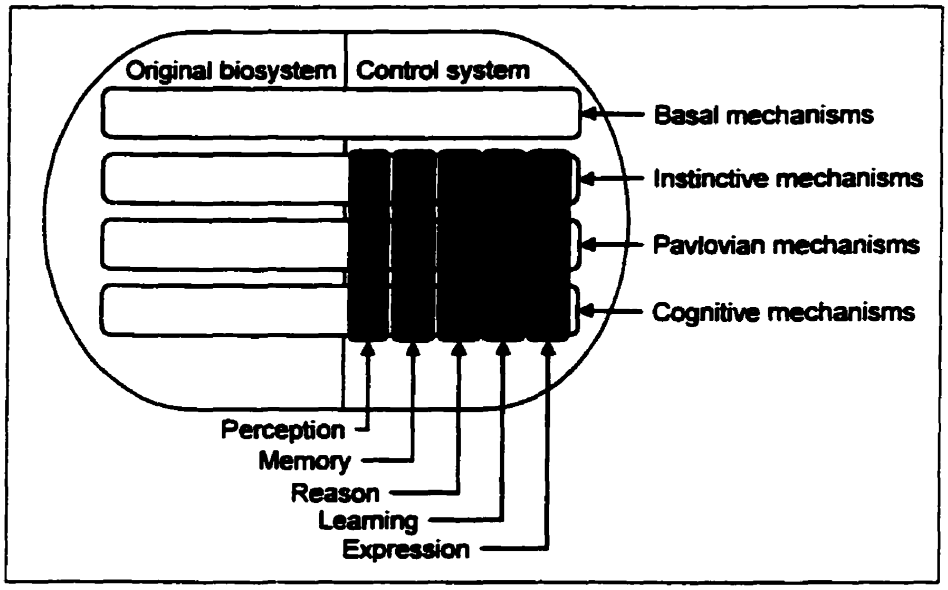

Figure 6.2 Mental faculties of an ecocyborg, shown as arising from the installed control system.

6.4.1 Perception

The faculty of *perception* includes all mental abilities involved in the transduction of input signals to information in the mind. Firstly, perception in an ecocyborg is influenced by the nature of the input signals that impinge upon it, as determined by their source. Either the physical or virtual aspects of these signals might be of importance in a given circumstance. Furthermore, if the ecocyborg is to be substantially conscious and autonomous, then these signals must originate not only from its surroundings, but also from within the ecocyborg itself. Secondly, perception is largely determined by how signals are transduced. This depends on the *perceptual intelligence* of the ecocyborg, which refers to the quality and sophistication of all of its perceptual abilities. The aspects of intelligence that are considered here are perceptual breadth and depth, due both to the cybernetic mechanisms that actually generate information from the input signals (*perceptors*), as well as to any adjunct mechanisms that subsequently process this information. The way an ecocyborg can transduce signals in a given setting partly shapes the kind of comportment that it is able to display. Furthermore, coming full circle, the ecocyborg's comportment can also influence the signals that reach it.

As stated, the information that an ecocyborg can acquire depends on the input signals that impinge upon it. These signals have both virtual and physical aspects, just as do the components of the ecocyborg. Accordingly, the virtual aspect of a system component is the virtual machine, or information-processing apparatus, that is resident on the physical substrate. Equivalently, the virtual aspect of a signal is the information inherent in the variable state of some physical quantity. In essence, neither the virtual machine of a system component nor the information content of a signal depend on the nature of their respective physical substrates. In the latter case, for instance, the same information could be encoded as changes in either the amplitude or the frequency of an electromagnetic field, fluctuations in air pressure, or as differences in the shapes and positions of ink marks on paper. Regardless of which aspect (or aspects) of a signal is taken into account, the end result of all perception is the transduction of signals into information within the mind of the ecocyborg.

Either or both of the virtual and physical aspects of the signals impinging on an ecocyborg can be of significance, their relative importance depending largely on the context. For example, only the virtual aspect of the radio signal carrying daily auction prices is likely to be useful to a semi-autonomous, market-oriented, agricultural production ecocyborg. Conversely, it is the physical aspect of the solar radiation impinging on such a unit that is more likely to be important, because it supplies energy for photosynthesis and has a direct impact on the temperature inside. Some signals, like those already mentioned, impinge on the perceptors from outside the ecocyborg, but there are also signals that originate internally, deriving from the ecosystem. In this latter case, the physical aspect is likely to be the more important.

The kinds of signals that an ecocyborg is able to transduce depend in part upon the breadth of its perceptual abilities. Clearly, an ecocyborg requires enough perceptors of sufficient diversity to transduce both the physical and virtual aspects of a variety of signals that originate internally as well as externally. Human perception is based on arrays of many millions of sensory neurons, and physically extensive ecocyborgs in particular will probably possess similarly large numbers of perceptors. These will be qualitatively quite diverse. Qualitative diversity here refers not only to the heterogeneity of signals that can be transduced, but also to the regions of the signals' spectra and the ranges of intensity that can be dealt with. For instance, both the vision and hearing perceptors of humans are able to transduce only narrow regions of the available spectra, and can accommodate only a fairly small range of intensities. One advantage of perceptual diversity is that it allows for the generation of a variety of information streams about the same phenomenon. The human sensory system, for instance, is able to generate information about a single phenomenon through gray-scale vision, color vision, smell, taste, hearing, touch, etc. Ecocyborgs are foreseen to likewise possess very diverse perceptors for the simultaneous sensing of various aspects of one phenomenon, (e.g., temperature, images, concentrations of atmospheric gases, soil nutrients and moisture content, populations of organisms, etc.). Generally, the wider the regions of the spectra and ranges of intensity that can be transduced, the more perceptually intelligent the ecocyborg and the more information available to it. Minimally, in designing an ecocyborg it will be essential to endow it with a breadth of perception adequate to meet the requirements for the desired degree of autonomy.

The structure of the overall assemblage of perceptors also impacts an ecocyborg's perceptual intelligence. At one extreme, for instance, a perceptor array comprising many elements can be physically very dispersed, like the temperature and tactile sensors of human skin. At the other extreme, perceptors can be concentrated into very localized assemblies, as are the hair cells in the cochlea of the ear. An ecocyborg will likely have similarly configured perceptor arrays with, for instance, many temperature sensors distributed throughout the original ecosystem as well as, perhaps, panels of photodetectors concentrated inside digital cameras. Aside from the physical configuration of an array, the manner in which perceptors interact is also extremely important. Signal processing methods can be implemented in an ecocyborg's mind by interconnecting perceptors so as to essentially create coarser-scale virtual machines. Just as significant patterns like hands and faces are detected by the human visual system at very early stages of perception (Pinker 1994), so a very large agricultural ecocyborg might be made capable of recognizing hail clouds in order that they could be promptly seeded in order to circumvent crop damage. Instead of being entirely rigid, moreover, the structure of a perceptual array can also be made somewhat flexible, allowing for the redistribution of resources so as, for example, to focus on a particular phenomenon that is important at a given moment.

Perceptual depth corresponds to the precision and accuracy of the information generated by the perceptors. In a numerical context, precision is the number of digits that a measurement yields. More generally, it is equivalent to the resolution of observation, whether in a temporal or a spatial sense. Accuracy, on the other hand, is the degree to which a measurement actually corresponds to reality. In other words, it is the reliability of the information that is generated by the perceptors. There are two main approaches to achieving a particular perceptual depth: the first is to use many perceptors (or a high sampling rate) to generate a large number of lower-quality data, and then to condition this data; the second is to use only a few perceptors (or a lower sampling rate) to generate fewer, higher-quality data that do not require as much conditioning. If the signal is highly distributed, then the use of many perceptors would seem to be the more fitting. This might apply, for instance, to signals such as temperature and soil moisture in the ecosystem. If the signal is very focused, however, it may be more appropriate to have only a few, robust perceptors of high precision and accuracy. This might be true, for example, of sound in the ecosystem. Thus, only several perceptors might be required to listen to the noise of birds and insects, so as to monitor their health. In choosing between these approaches, the incremental cost of improving the quality of a perceptor must be compared to that of creating, installing, and managing more, lower quality perceptors and conditioning the information that would be available from these.

In the mental architecture proposed here, adjunct mechanisms are permanently associated with particular assemblies of perceptors and perform the kind of data conditioning discussed above. Such conditioning might involve, for example, the analysis of information generated by many perceptors in parallel, allowing the ecocyborg to obtain better quality information from a given signal than would be possible if only a single perceptor of that same quality were available. Adjunct mechanisms are also used to prepare information for storage, and for subsequent processing by mechanisms associated with other faculties. These abilities overlap somewhat with those in the faculties of reason and memory but, because they have a relatively permanent and inflexible association with particular perceptors, they are considered to have a high degree of membership in the faculty of perception. Evidently, they can very strongly affect the perceptual intelligence of an ecocyborg.

Appropriate perceptual intelligence enables an ecocyborg to display certain comportment in a given setting, corresponding, for instance, to the pursuit of a set of design goals. As an example, the aforementioned agricultural ecocyborgs are very likely to be oriented toward the maximization of profit, but the pursuit of such goals is possible only if the relevant mental abilities are present. One design goal that is of particular interest here is substantial autonomy, which depends on a significant degree of consciousness. Minimally, therefore, the ecocyborg must be able to use a model of itself in mentation. Thus, a highly autonomous ecocyborg is envisioned as possessing perception that is broad enough to transduce signals of both internal and external origin, so that it can create and maintain the required model. Moreover, the greater the diversity of perceptors that the ecocyborg possesses, the more extensive the mental model that it can manage, and the richer the information that is available for effective prediction of a given phenomenon. Hence, the potential autonomy of the ecocyborg is correspondingly greater. A similar argument holds for depth of perception.

Finally, the comportment of an ecocyborg can influence the kinds of signals that impinge upon it. For instance, a cyborged biosystem that is mobile must deal with signals that change continuously as a result of the relative motion between it and its surroundings. If it were substantially autonomous, then it might decide to move to surroundings more conducive to the achievement of its goals, or even directly modify its surroundings to its own ends. The ecocyborgs envisioned here are, however, largely "immobile robots" (*immobots*) like the agricultural production ecocyborgs mentioned previously, and are more likely to have immediate surroundings that are relatively predictable over the short term. Such entities are, therefore, largely oriented toward the management of their own internal state, and not toward the modification of their surroundings (Williams and Nayak 1996). An important part of managing the internal state of an ecocyborg is to incorporate into its OOKB the information that is generated by its perceptors. The abilities required for this and related tasks are part of the faculty of memory.

6.4.2 Memory

The mental faculty of *memory* comprises the abilities that enable an entity to organize and store, retrieve, and delete information. Evidently, for an ecocyborg, this faculty should be of an intelligence appropriate to the fulfillment of the design objectives. In the mental architecture proposed here, memory is implemented as a set of cybernetic mechanisms whose purview is the OOKB. The faculty of memory includes abilities involved in organizing new information and incorporating it into the OOKB, and in this it overlaps somewhat with the faculty of perception. As well, memory comprises the ability to reorganize existing information so that it can be more effectively stored, or be utilized better by particular mechanisms. In a mind of limited storage or management capacity, this could involve strategies such as compression and decompression, partial or complete de-indexing, or even the erasure of information when it becomes obsolete or redundant.

As part of the proposed mental architecture, the OOKB serves as a pool of information that is accessible to all the mechanisms of mind, in accordance with the blackboard concept of memory (Lesser et al. 1975). As well, it provides a medium for communication among the cybernetic mechanisms. It is based on *knowledge objects* (Gauthier and Guay 1998; Gauthier and Néel 1996; Gauthier and Kok 1989), the most fundamental of which are qualitative or quantitative (possibly numerical) data, but which can also be larger items like images, procedural models, and other composite objects. Any object can have attributes (like 'weight'), which are also knowledge objects in themselves. These attributes can have values, which might be either qualitative (e.g., 'heavy') or quantitative (e.g., '48'), and they may also have attributes themselves. A particular type of attribute is linkage with another type of object and the value of that link is equivalent to its strength or quality. As well, any object can have, as an attribute, the specification that it is associated with (i.e., is an *instance* of) another type of more abstract object called a *class*. If this association can have a variable value then all of the instances of the corresponding class are a *fuzzy set* (Kosko 1993). Instances of a particular class share a number of common *class attributes*. The class itself, as an object, can also have attributes, including links to other objects and the status of being an instance of another class. Thus, an object might be a single datum or larger item, a class, or an attribute (either quantitative or qualitative, and including links, as well as the condition of being an instance of a class). Very intricate composite objects can result from the recursive association of smaller objects in these ways. Such informational constructs can represent knowledge about the environment, the original ecosystem, and even the mind of the ecocyborg itself, if the latter is significantly conscious.

The organizing and storing of new information, which is the first activity mentioned above, can be considered as synonymous with indexing the knowledge in the OOKB. This is based on the association of similar patterns by linking or classifying knowledge objects. Objects can be linked routinely by adjunct mechanisms to organize, for instance, the subsequent data generated by certain perceptors as a time series. Alternatively, the data can be made instances of a particular class. Such adjunct mechanisms might also continuously search through the OOKB for analogies, similar to the operation of the Copycat and Tabletop programs (Hofstadter 1995). They could then further organize the OOKB by making explicit any discovered analogies, in the form of links between extant knowledge objects or by associating these with a particular class.

The retrieval of information from the OOKB is also accomplished by associating knowledge objects with one another. In this case, however, one of the objects involved is a template corresponding to one particular datum or larger item, and representing an internal request from a cybernetic mechanism. Retrieval mechanisms attempt to match existing objects in the OOKB with the template. Effective indexing makes this task easier because the required information might, for example, already be linked to similar objects, or be an instance of a particular classes. If a retrieval mechanism is successful, it can then communicate the required information (or its location) to the mechanism that issued the request. If the stored information has been compressed or otherwise encoded, it might be necessary to first transform it so as to make it conform to a particular, usable format.

The final task mentioned above, the deletion of obsolete or redundant information, is necessary in any system whose capacity for storage or management becomes insufficient if that system is to continue to assimilate information. As the limits of memory are approached, it could be necessary to partially de-index, or to archive, information. In this case, some of the indexing attributes, such as links or class memberships, are removed. Thus, the data are indexed in less detail and are stored together as larger units. They can still be accessed if necessary, but only by searching through the larger storage units. As demand on storage and management abilities intensifies, data might be yet further de-indexed until no indexing information remains at all. In that case, all of the archived data would have to be searched to locate a specific item. A complementary strategy to de-indexing is to compress information by removing redundancy. This usually also results in slower processing. As the situation becomes more critical, more rigorous compression schemes could be used to remove not only redundant data, but also to remove detail from the information. Eventually, it might be necessary to completely erase some information. A strategy that might be used to retain data that are most useful would be to assign attributes indicating the time at which they were generated and when they were last accessed. They might then be removed if they exceeded a certain age and had not been recently utilized. Of course, such a strategy has a trade-off value, in that the inclusion of the extra attributes places an added demand on memory resources.

The way that the faculty of memory is implemented in the proposed architecture bears some resemblance to *knowledge base management systems* (KBMS) that have been devised for very large computer data bases (Dubitzky et al. 1996). In both cases, the arrangement must be flexible enough to catalogue diverse kinds of information, and to abstract new information from it. In a KBMS this is also accomplished by applying an object-oriented management strategy to knowledge representation, in an approach similar to the one described here. Thus, information is represented by knowledge objects that are associated with more abstract objects (*classes*) and these are, in turn, hierarchically ordered. The proposed OOKB differs, however, from the KBMSs in some significant ways. For instance, objects in a KBMS are capable of performing certain characteristic activities called *methods* (e.g., the encoding, comparison, or updating of information). These are often based on dedicated rule sets that are each associated with a particular class. In an ecocyborg's OOKB, such information-processing activities are not performed by the knowledge objects, but are instead implemented as mental abilities arising from independent cybernetic mechanisms. The knowledge objects and cybernetic mechanisms are therefore more differentiated with respect to passive and active roles than are the objects in a typical KBMS. This arrangement is more flexible and combines a KBMS's capacity for sophisticated organization of information with the potential for very rapid, highly parallel processing of that information.

6.4.3 Reason

With regard to overall system mentation, the faculty of *reason* enables a critical phase in the continuum of information-processing, linking the perception of input signals to the generation of outputs. In the proposed mental architecture, the ecocyborg's faculty of reason enables it to mentate about a wide variety of situations, ranging from the well-defined to the more abstract, and perhaps even to the nebulous. Like the other faculties, reason is a diverse set of mental abilities that arise from a flexible and expandable set of semidifferentiated cybernetic mechanisms. As part of their mentation, the mechanisms create and utilize a diversity of knowledge objects including, for example, strategies and tactics. These are subsequently parsed into directives to particular kinds of cybernetic mechanisms that generate output signals (*effectors*). Some reasoning mechanisms can also create such directives explicitly, resulting in the immediate generation of signals either in the surroundings of the ecocyborg, or in its ecosystem. Other kinds of objects that might be generated include databases, rulebases, models, goal trees, and judgments, all of which can have a far-reaching influence on the comportment, although this might not necessarily be direct nor immediate.

Cybernetic mechanisms can be created to correspond to different implementation levels, goal types, priorities, and activity classes (Kok and Lacroix 1993). Accordingly, highly optimized, relatively inflexible control mechanisms called *basal* devices can be created that react directly to environmental stimuli. These are used to deal with routine situations in an elementary way. *Instinctive* devices implement more complicated control sequences, but still act very much in a repeatable manner. Typically, they are used for prevention and assurance type goals. *Pavlovian* devices are able to adapt so as to produce increasingly optimal outputs (Molenaar 1998). Finally, *cognitive* mechanisms perform in-depth analysis, deal with unforeseen situations, create strategies for attaining long-term goals, etc. They make extensive use of symbolic processing in activities like modeling and simulation. When grouped according to implementation level, these mechanisms form sets that are referred to as the basal, instinctive, Pavlovian, and cognitive *controllers*, respectively.

In the narrow interpretation of the lexicon, it is the latter three controllers and the OOKB that form the ecocyborg's mind, whereas the basal controller is not considered part of it. The basal controller does, however, form part of the ecocyborg's environmental interface and might comprise very sophisticated machinery. It corresponds to, for example, human physical responses such as tanning. Although the abilities engendered by the instinctive and Pavlovian controllers are likely to have high membership in the faculty of reason, they are also likely to have high degrees of memberships in the faculties of perception, memory, and expression. They are optimized and rather inflexible, and therefore give rise to relatively consistent mentation. The cognitive controller, on the other hand, is a very openly structured community of mechanisms which are very likely to pertain exclusively to the faculty of reason. They give rise to abilities that are much more flexible than those arising from instinctive and Pavlovian mechanisms, and that are central to the autonomy of an ecocyborg.

As pointed out, the activity of cybernetic mechanisms involves the creation and manipulation of various types of objects in the OOKB. Tactical objects and directives to effectors will most often be produced by the instinctive and Pavlovian controllers. Thus, these will have a direct and immediate impact on the generation of output signals. The instinctive controller is oriented partially toward ensuring the system's long-term persistence and safeguarding its survival in emergency situations, whereas the Pavlovian controller is oriented more to system operation and the optimization thereof. The cognitive controller, on the other hand, is primarily involved in the production and manipulation of more abstract knowledge objects. As part of its abilities it can generate objects like goals and strategies, and then judge and rank these with respect to priorities and moral standards, which are themselves instantiated as objects. It is through this kind of mentation that self-referential models are generated and utilized, and so the activities of cognitive control mechanisms are crucial to consciousness and to substantial system autonomy.

All interaction between cybernetic mechanisms occurs via their manipulation of the OOKB. In this way, particular sets of mechanisms can be more or less rigidly associated with one another. Thus, relatively independent instinctive mechanisms could directly receive information transduced by certain perceptors, implement a discrete procedural algorithm, and feed directives to particular effectors. Pavlovian mechanisms could perform more flexible activities, such as searching out and linking items in the OOKB that bear a certain relationship to one another. Some cognitive mechanisms might create generally available resources that enhance the overall mind. For example, they might implement a rudimentary consciousness by examining a great variety of objects and mechanisms, and then model some aspects of these with an ANN (Kok and Lacroix 1993). Objects in the OOKB are therefore subject to iterative and serial transformations by a variety of mechanisms, making possible a great variety of sophisticated feedback, feed-forward, inhibitory, and excitatory control circuits.

Even though cybernetic mechanisms are semidifferentiated and may give rise to different mental abilities, they might still address the same issues. Thus, conflict between mechanisms in the mind is quite likely to occur. The differentiation of mechanisms according to goal type might help to circumvent this; for instance, particular goals might be addressed exclusively by a certain group of mechanisms. Even with this arrangement, however, it is still possible that several mechanisms will address the same goal simultaneously, resulting in a number of alternate strategies. Thus, some means of conflict resolution is required. The ability to make sophisticated judgements about conflict is most appropriately accommodated in the faculty of reason but, for the sake of expedience, some straightforward conflict resolution might also be implemented in the environmental interface as part of the faculty of expression.

6.4.4 Expression

Expression is the complement of the faculty of perception, being the set of abilities that enable an entity to generate output signals. This faculty includes the abilities engendered by effectors, and also encompasses the abilities of any adjunct mechanisms that are permanently associated with these and that facilitate their mentation. In the case of an ecocyborg, outputs can impinge either on the surroundings, or on the original ecosystem.

As with inputs, either the physical or virtual aspects of these signals can be important, depending on the context in which they are generated. Output signals are ultimately synonymous with the interaction of the control system with the external surroundings and the ecosystem. For the substantially autonomous ecocyborgs envisioned here, output signals will result to some extent in the realization of the ecocyborgs' goals. The expressive intelligence of an ecocyborg therefore has a direct effect on its autonomy.

The many effectors of a sophisticated ecocyborg are likely to be of a wide variety of different types, making possible the breadth and depth of expression required for substantial autonomy in a given context. An ecocyborg might generate a signal that impinges on its surroundings and physically modifies them as a result. A highly autonomous ecocyborg with such capabilities might modify its surroundings to suit its own purposes. Ecocyborgs as envisioned here, however, will not be overly preoccupied with external expression of this kind. Instead, they will mostly generate output signals that impinge upon and influence the state of their own, integral ecosystems. In the case of an enclosed agricultural ecocyborg, these outputs might involve the application of fertilizers, the control of lighting conditions, or the injection of carbon dioxide into the inside atmosphere.

Outputs are always generated through the manipulation of physical phenomena but, as with inputs, either the physical or virtual aspects of the resulting signal can be of importance. An immobot, as mentioned, is oriented primarily toward the maintenance of its own internal state, but it might also generate external output signals whose virtual aspects are of significance. This is to say that such an ecocyborg could transmit signals as a means of communication with other entities like, for example, a human production manager. Internal signals, on the other hand, are more likely to be manipulations of physical aspects of the ecosystem. Also, there could very well be internal messages transmitted between semiautonomous components of the control system. Since these components are, however, considered here to be part of the ecocyborg's mind, such messages are not interpreted as outputs.

Like perceptors, effectors can be arranged in a variety of ways. An ecocyborg with a physically extensive ecosystem will probably have many effectors configured variously as individual mechanisms, highly dispersed collectives, or integrated assemblies whose components work in concert to perform certain tasks. In an agricultural ecocyborg these could include, for example, a few individual maintenance robots, a large number of independent irrigation sprinklers, and a coordinated assembly of valves to regulate water flow in distribution pipes.

As well as effectors, the faculty of expression also includes adjunct mechanisms that facilitate the transduction of knowledge objects into signals. These could parse more general items from the OOKB, such as tactics, into an ordered sequence of directives that are then passed to effectors for execution. They can also resolve conflicts by, for instance, using fuzzy logic to combine multiple instructions into directives that represent the contributions of many different knowledge objects (Kosko 1993). As well, they are necessary for the coordination of sets of effectors that are integrated into cooperative assemblies. The performance of such tasks by adjunct mechanisms frees the rest of the mind from an otherwise unwieldy computational burden. Their abilities thereby greatly enhance the faculty of expression and enable an ecocyborg to respond more intelligently to its environment. The creation of signals in the surroundings or in its own ecosystem does not, however, ensure the substantial long-term autonomy of an ecocyborg in a changing environment. This requires that the ecocyborg be able to restructure its control system.

6.4.5 Learning

Learning is the ability of an entity to adaptively restructure its own mind. This faculty includes the capacity for long-term retention of acquired knowledge, thereby overlapping with memory. As well, it enables the optimization of routine behavior to various degrees, thus also overlapping with some aspects of reason. In scenarios where existing virtual machinery is rendered inadequate, a system that is sufficiently intelligent in terms of learning is able to modify or expand its own mind. This enables it to mentate in an original manner, or to apply existing methods in different contexts. Finally, an entity may possess recursive learning abilities, which are especially potent. In fact, human genius has been defined as the result of having learned better means to learn, integrate, and manage knowledge (Minsky 1985, p.80).

One basic aspect of learning is the subsumption of transient information into a more permanent form. This process is, in other words, the transferal of knowledge objects from short-term to long-term memory, and so the required abilities also have membership in that faculty. These objects might be simple data or larger items, such as images or rule sets. Before being given special status, candidate objects must first be identified and evaluated. As mentioned in the section on memory, transient knowledge objects in the OOKB might be given various attributes related to how long they have existed and how frequently they have been utilized, and learning mechanisms could base their decisions on the values of these. Thus, if an object is deemed sufficiently useful, it can be exempted from de-indexing or erasure and, perhaps, made more easily accessible for use in mentation.

A second aspect of learning, which primarily impacts cybernetic mechanisms rather than knowledge objects, is the optimization of routine behavior. First, this includes the tuning of virtual machines for the most effective comportment possible in a given set of circumstances, as exemplified by the training of neural networks for the execution of particular tasks. This might also involve explicit procedures as in, for instance, the tuning of the gain coefficient of a feedback controller. A second level of optimization involves the replacement of a control mechanism with one of another type, which is deemed to be more suitable in a given context (Molenaar 1998). Thirdly, entire Pavlovian control circuits could be redesigned so as to be more effective or appropriate, a task that is likely to be carried out by cognitive mechanisms.

A third aspect of this faculty is the capacity for creative learning. Creativity is the making of something new, and often implies an element of unpredictability. In the context of learning, it involves the modification of the mind. This might enable an ecocyborg to devise original methods of solving existing problems and, hopefully, to respond effectively to novel situations.

Creativity can be implemented in a number of ways, one of which is to generate novel virtual machinery patterned after external phenomena. This requires abilities that might also have some degree of membership in the faculty of perception. One way to achieve creativity might be to train a neural network to associate causal events with their results, and then prune the trained network so that only a relatively simple set of relations remained. These relations could then be tested with statistical tools, and retained if found to be effective in representing the observed phenomena. This set of relations might next be formulated as a hypothesis and tested according to formal experimental protocol, possibly resulting in a model of a higher level of abstraction than the original information. Other examples of generating virtual machinery based on external phenomena can be drawn from image processing techniques, in which perceived images are translated into representative idealizations composed of geometrical shapes or other mathematical constructs.

A second way of implementing creativity is not to generate new machines, but to alter the structure of the mind by associating existing machines with each other in new ways. This could be performed by specialized mechanisms that recognize similarities, or analogies, between different sets of virtual machines. Thus, upon recognizing similar patterns of relationships, a particular kind of mentation might be transposed from one set to the other. This is, in effect, the equivalent of processing familiar information in a different way. The ability to discover such analogies is central not only to creative learning, but also to the indexing activities of memory, as described previously, and to the creation of effective models.

A third way of implementing creative learning is to use an adaptive approach patterned after Darwinian natural selection. This approach requires that machines be replicated. As already described, the OOKB is continually modified by the activities of cybernetic mechanisms, and the same can also be true of the cybernetic mechanisms themselves. As well, mechanisms can be made capable of replicating themselves or, alternatively, they might be replicated by others. If the replication rate of a particular type of mechanism is related to the success with which it performs its task, then a positive feed-back loop is established that intensifies currently successful activities, and the relative populations of the different kinds of mechanisms will change in a context-sensitive manner.

The scheme as described above is adaptive but, in order for it to be truly creative, there must be some variation in the virtual machines that are generated. This can be achieved by a number of methods. First, new machines can be patterned after existing ones, with some superficial changes being generated in a systematic (e.g., pseudorandom) way. Second, a recombinant approach might be used in which new machines are given features drawn from several existing ones. Third, new machines could have entirely novel features somehow adapted from or inspired by other phenomena. With the latter approach, existing possibilities are not merely revised or extended, but entirely new ones are created. Readers are referred to Hofstadter (1985) and Boden (1990) for more thorough discussions of creativity in mind.

The creation of new machines, referred to above, will result in an irreversible increase in their population unless there is also some means of destroying them. The rate of destruction must be proportional to the total population, in order that the latter be kept in check even in times of high stimulation and rapid replication. The machines could be selected for destruction pseudorandomly, or in some other systematic way. If an impartial method is used, the population will eventually be composed primarily of machines similar to those that are currently most active, and which are therefore replicated the most rapidly. A more selective strategy, however, is to also take into account how well the machines fulfill a set of explicit criteria. The evolution of the mind's constitution is thereby influenced by the goals which underlie those criteria. This could be exploited in engineering a substantially autonomous ecocyborg, as further discussed below.

Finally, learning becomes especially sophisticated if its abilities can be brought to bear on the very mechanisms that engender them. In this way an entity can optimize its learning methods to suit the current circumstances and create learning abilities to cope with new challenges, as well as retain learning abilities over the long term if they have proven to be particularly effective. The introspective capacity to modify the virtual machinery that performs these very modifications requires that the entity be substantially conscious. Learning of this kind depends on the ability to create models, and then to use these in mentation so as to determine how they might be improved. In recursive learning, these models must represent at least some aspects of the entity's own mind, and this is a fundamental prerequisite for autonomy. All of the above methods are foreseen as being implemented to some degree in the architecture proposed for the ecocyborg.

6.5 Emergence of autonomy

One of the principal objectives in the EcoCyborg Project is, as discussed, to learn how to render a system substantially autonomous. A necessary but insufficient prerequisite for autonomy is consciousness. An ecocyborg with a mind that is based on the architecture described can be made conscious to any degree. Consciousness can, in turn, give rise to some degree of automation, volition, and intent, as outlined below. These are three fundamental aspects of autonomy and, therefore, a system that possesses each of them is also autonomous to some extent (Clark et al. 1999).

As pointed out, consciousness is an essential precursor of autonomy. The working definition of consciousness used here is the possession of a self-model, and its use in simulation-based mentation. Such a model, and the simulations that are based on it, can be more or less sophisticated (Clark et al. 1999). The model might be either explicit or implicit, with each approach having some advantages over the other. An implicit model, inherent in an ANN, for example, might afford greater computational speed than an explicit one, but the latter is often more flexible and, therefore, more convenient (Lacroix and Kok 1999; Shukla et al. 1996; Kok and Lacroix 1993). In more sophisticated conscious mentation a number of models of different types might be implemented simultaneously in various simulation streams, with different aspects of the ecocyborg's constitution, state, or comportment being emphasized in each of these. Moreover, the various models could be constantly updated by mechanisms similar to the tools currently being developed in the EcoCyborg Project for the objective characterization of biosystems (Clark and Kok 1999). The models could also be made recursive, so as to include more or less detailed representations of themselves, as well as of other aspects of the ecocyborg (Clark et al. 1999; Kok and Lacroix 1993).

Consciousness gives rise to the various aspects of autonomy, one of which is automation, or the potential to bring about a particular chain of causal activities, without external guidance. This is in agreement with the assertion, put forward by some proponents of the cybernetic paradigm, that any effective controller of a system must include a model of that system (Conant and Ashby 1970). Accordingly, on the one hand, a marginally effective controller might include only a very restricted model of the controlled system (which is often entirely implicit in the controller's constitution). On the other hand, if a controller is to be highly effective, it must utilize a much more complete and sophisticated model. If control is to be robust in the face of unusual perturbations, for example, a flexible model is required (which, as mentioned above, is likely to represent the controlled system in an explicit manner). Highly sophisticated models of this kind, and rapid simulations based upon them are, for example, at the heart of the modern digital controllers that make extensive automation possible. In this discussion, the controlled system is the entire ecocyborg, so the controller is itself part of the controlled system. The causal activities being discussed are, therefore, the response of the ecocyborg to input signals originating either from the surroundings, or from its own ecosystem. This includes mentation (in the narrow interpretation), the resulting comportment of the controlled ecosystem, and any external expression that the ecocyborg might exhibit. Overall, the better the model the ecocyborg has of itself, the more conscious it is, and the more effectively it can control its own comportment. If it is highly automatic then it is able to accommodate a wide variety of input signals and still maintain a particular mode of activity.

An ecocyborg can be made substantially automatic by engineering the virtual machinery of its mind so that it comprises a causal chain extending from perceptors to effectors, and ultimately to the controlled ecosystem. The more independent this causal chain is from external influences, the greater the degree of automation of the ecocyborg. Complete isolation is, however, undesirable because, minimally, sources of low-entropy energy are necessary for the persistence of dynamic comportment in any physical system (Clark and Kok 1999). Hence, there will be input signals that impinge upon the ecocyborg, and the control system must be able to transduce information from these and incorporate it into the OOKB, possibly updating models in the process. This information must also be screened for inputs that might cause the comportment of the ecosystem to deviate from its trajectory and, if any are detected, then effector directives must be generated that result in compensatory output signals. In a highly automatic ecocyborg, the control system must have internal access to the mechanisms required to perform all of these transformations, so that they might be performed without external guidance. In current control systems, such mechanisms generally are explicit procedural algorithms that are included by design but, as discussed in the section about learning, future ecocyborgs will likely be able to generate some of their own mechanisms. Finally, the effects of the ecocyborg's compensatory output signals (and possibly the signals themselves) will be detectable as inputs, enabling, for example, feedback control.

For a system to be substantially autonomous it is not sufficient that it be highly automatic; it must also be volitive to a large degree, meaning that it must be capable of formulating its own goals and strategies for achieving them. For instance, modern onboard automotive control systems provide today's cars with sophisticated self-models, so that the cars' comportment is quite automatic. Since cars lack the ability to formulate their own goals they are, however, autonomous to only a very small degree. An example of a much more autonomous entity is a human artist who pursues her own drives and desires to the point where she abandons established tradition and forges her own style. However, even the behavior of a biological entity that is autonomous to this degree is founded on fundamental system goals. Such goals are, in fact, necessary in any volitive entity to provide direction to its comportment, although they also ultimately restrict its volition. In biological entities system goals such as self-preservation and reproduction are genetically entrenched by natural selection, but in artificial entities they will need to be built in by design. The nature and flexibility of these goals will depend on the particular purpose behind the creation of the system in question; for an ecocyborg they will likely relate to self-preservation, as well as more specific design objectives such as general production requirements. These fundamental system goals serve as the basis for the generation of subsidiary goals although, especially in relatively volitive systems such as humans, the relation between these might not always be direct.

Volition can be implemented in the architecture proposed here, on the basis of fundamental system goals that are embedded in immutable virtual machinery. These goals are envisioned as knowledge objects, perhaps highly distributed ones, that represent, in a very general way, a particular constitution, state, or comportment of the ecocyborg. More immediate subgoals are generated by comparing such a knowledge object with the ecocyborg's current representation of the corresponding aspect of itself. Any differences are analyzed and subgoals are formulated to represent hypothetical conditions of the ecocyborg in which these differences are reduced in some way. The virtual machines used in the generation of subgoals might be simple algorithms or more sophisticated machinery, and they might be included in the ecocyborg's control system by design or generated by the ecocyborg's faculty of learning. They might include not only models of the surroundings and of the ecocyborg, but also empirical rule bases, ANNs, etc. It is likely that many subgoals will be generated, which must then be evaluated for practicality and, if necessary, reformulated or discarded accordingly. Since much of this mentation requires access to effective models and the ability to implement these in simulation, the volition of the ecocyborg is heavily dependent upon its degree of consciousness.

Finally, if an entity is to be substantially autonomous it must be highly intentful, actively pursuing the goals that it formulates. Thus, the degree of intentfulness of an ecocyborg depends upon the extent to which it is automatic, volitive, and therefore indirectly, conscious. Moreover, intentfulness requires that the ecocyborg consider the possible consequences of its activities by evaluating and prioritizing the courses of action that it generates. It can thereby establish which tactics and effector directives might contribute most effectively toward the achievement of its goals. Such mentation is also based on the modeling and simulation of the ecocyborg in the context of its environment, and so intentfulness is also dependent on consciousness in a direct way.

Intentfulness might be implemented with either explicit or implicit approaches, or through a combination of these. Strategies, tactics, and effector directives could be prioritized based on the results of simulations predicting the eventual outcome of their implementation, and less effective ones could then be eliminated. At the strategy level, these simulations could be based on procedural models, but this might be a rather cumbersome way to evaluate tactics and directives, and so alternative methods, such as modeling with an ANN, could be used instead. The prioritization or elimination of knowledge objects in this way is equivalent to biasing the criteria used in adaptive mentation to select machinery for replication or destruction, as described in the section about creative learning. It is conceivable that such criteria might somehow be formulated so as to implicitly reinforce particular goals, but the judgement of machines based on explicit simulation, at least initially, appears to be easier to implement. Machines could afterward be generated that utilized more implicit methods; for example, an ANN could be trained through observation of the initial procedural modeling and simulation.

Another means of implementing intentfulness in adaptive mentation is to use explicit models to compare one of the ecocyborg's goals with the corresponding aspect of the ecocyborg itself, and then to adjust the rate of mutation in the replication of machinery according to the closeness of the match. If the situation did not correspond closely to the desired goal, then the rate of change could be increased. This would encourage different patterns of mentation to arise, hopefully leading to a more desirable scenario. On the other hand, if the situation was found to correspond closely to the goal, then the mutation of new machinery could be suppressed. Overall, this would reinforce the drift toward successful mentation that is already present in the selective reproduction of more active machinery. In Hofstadter's (1995) models of cognition a similar approach is used, whereby the mutability of virtual machines is inversely proportional to how satisfactory the current configuration of the program is deemed to be.

Overall, the mentation described with respect to automation, volition, and intent can give rise to substantially autonomous comportment through a roughly cyclical series of interactions involving a number of key types of knowledge objects (Figure 6.3). These are the models that the ecocyborg maintains of its own constitution, state, and comportment; representations of particular goals and subgoals; and models of the ecocyborg's surroundings. The interactions between these objects are mediated by numerous cybernetic mechanisms of the types described previously, and could also involve a variety of other knowledge objects. The pattern of interaction between these could vary according to the circumstances, thus corresponding to the flexible network of causal interactions mentioned previously in the discussion about control organization. The overall pattern of interactions is described here with the fundamental system goals as the arbitrary starting point. Broad subgoals are generated by comparing current representations of the ecocyborg and its surroundings to the system goals and then, through simulation, determining a series of intermediate constitutions, states, or modes of comportment that lead from one to the other. The subgoals are evaluated for practicality and the relative effectiveness with which they might lead to the achievement of the system goals. They might be iteratively reformulated during this process, and are finally prioritized according to their predicted value to the system. Strategies are then formulated for attaining them, again using simulation based on models of the ecocyborg and the surroundings, but with greater depth and less breadth. These are parsed into tactics, or ordered sets of directives that are finally passed on to the appropriate effectors. The effectors generate output signals that impinge on either or both of the external surroundings and the ecosystem of the ecocyborg, hopefully bringing about changes that will move the ecocyborg further along the anticipated series of subgoals. In turn, signals arise from the ecosystem and the external surroundings and impinge on the ecocyborg's perceptors, causing them to generate new information. This is then used to update the ecocyborg's models, and might also prompt some immediate instinctive or Pavlovian response. The cycle then begins anew with another comparison of the current self-model of the ecocyborg and its goals.

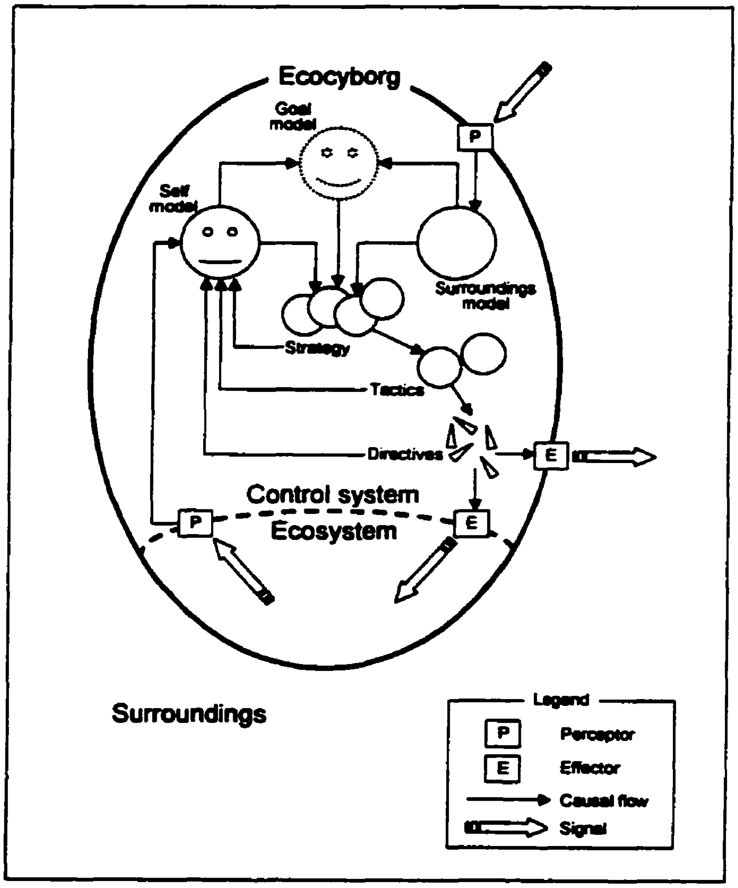

Figure 6.3 Schematic representation of mentation giving rise to substantial autonomy in an ecocyborg.

6.6 Discussion and conclusions

This article is a discussion of how an original biosystem, especially an ecosystem, can be engineered so as to be substantially autonomous by combining it with a technological control system. In many ways, this approach to the engineering of ecocyborgs is analogous to the manner in which some aspects of the natural world are thought to have come about. This is illustrated below with several examples. In the first one, the parallel is drawn between the cyborging of ecosystems with control components, and the manner in which the internal chemistry of modern eukaryotic cells may have resulted from the combination of separate protein and nucleic acid metabolisms that arose independently of one another. In a second example, the cyborging of ecosystems is compared to the endosymbiont theory of the origin of modern eukaryotic cells. Finally, there is also a brief discussion of the possible future relationship between ecocyborgs and humans.

The first example is based on the hypothesis of the double origin of life, due to Dyson (1988, p.92). In this, it is first proposed that various types of ancient protocells may have had quite different metabolisms. Specifically, he proposed that the metabolism of some may have been based only on proteins, whereas the metabolism of others may have been based on nucleic acids alone. The second part of the hypothesis concerns the possibility that these two types of protocells may have combined into a new type of cell in which the two kinds of metabolism symbiotically influenced one another. This new type of cell would have been better able to regulate itself and respond to changes in its environment, making it more viable than either of its precursors and leading to their disappearance. This idea prompts the analogy between protein metabolism and the dynamics of a wild ecosystem, and nucleic acid metabolism and the activity of a control system. Dyson himself, in fact, described protein metabolism by comparing it with an ecosystem, invoking Darwin's image of a wild community of plants and animals as a "tangled bank". Just as the combination of two cellular metabolic types may have resulted in a more effective overall metabolism, the addition of control components to an ecosystem might similarly transform the tangled bank of an unguided ecosystem into a highly autonomous ecocyborg.

A second example stems from the endosymbiont theory of the origin of modern eukaryotic cells (Margulis and Sagan 1986). According to this theory, cell organelles originally derived from free-living bacteria that infected larger hosts, and the vestiges of both kinds of ancient organisms now play a vital role in the modern eukaryotic cell. Centrioles and associated components like microtubules which, for instance, might be the vestigial remains of invading spirochetes, now manage the genetic apparatus of eukaryotic cells during mitosis. The addition of control components to an ecosystem might similarly result in a new type of entity with superior internal organization and a more effective structure.

The control system of an ecocyborg may ultimately include not only technological components, but biological ones as well. Humans are currently primarily responsible for the creation, replication and survival of cyborged biosystems, and in most cases human managers are key components of the control systems. These kinds of relationships will probably continue into the future. In fact, humans will likely be an integral part of many kinds of ecocyborgs, like those constructed for the colonization of space. Hence, even if such entities are very highly intelligent, conscious, and autonomous in themselves, it is likely that their fate will continue to be vitally intertwined with that of humans, who might be simply part of the ecosystem's biota, act as control components, or both.

6.7 Acknowledgments

The authors gratefully acknowledge the support of the Natural Sciences and Engineering Research Council (NSERC) of Canada, McGill University, and the Keith Gilmore Foundation.

6.8 References

Boden, M.A. 1990. *The Creative Mind: Myths and Mechanisms*. New York, NY: BasicBooks.

Clark, O.G., R. Kok and R. Lacroix. 1999. Mind and autonomy in engineered biosystems. *Engineering Applications of Artificial Intelligence* 12(3):389-399.

Clark, O.G., and R. Kok. 1999. Characterizing biosystems as autopoietic entities. *Oikos*, (submitted for publication).

Conant, R.C., and W.R. Ashby. 1970. Every good regulator of a system must be a model of that system. *International Journal of Systems Science* 1(2):89-97.

Dubitzky, W., D. Bell and J. Hughes. 1996. Generic, object-oriented case-knowledge representation scheme, and its integration into a wider information management scenario. *Expert Systems* 13(3):219-233.

Dyson, G. 1997. *Darwin Among the Machines*. Redwood City, CA: Addison-Wesley.

Dyson, F.J. 1988. *Infinite in All Directions*. New York, NY: Harper and Row.

Gardner, H. 1993. *Frames of Mind*, 2nd ed. New York, NY: BasicBooks.

Gauthier, L., and R. Guay. 1998. Using object-oriented database management technology in agricultural decision support software. *Canadian Agricultural Engineering* 40(3):219-226.

Gauthier, L., and T. Néel. 1996. SAGE: An object-oriented framework for the construction of farm decision support systems. *Computers and Electronics in Agriculture* 16:1-20.

Gauthier, L., and R. Kok. 1989. Integrated farm control software: I. Functional requirements and basic design criteria. *AI Applications* 3(1):27-37.

Hofstadter, D.R. 1995. Fluid concepts and creative analogies: computer models of the fundamental mechanisms of thought. New York, NY: BasicBooks.

Hofstadter, D.R. 1985. Variations on a theme as the crux of creativity. In *Metamagical Themas: Questing for the Essence of Mind and Pattern*. New York, NY: BasicBooks.

Jackson, J.V. 1987. Idea for a mind. *SIGGART Newsletter* 181(July):23-26.

Kok, R., and G. Desmarais. 1988. A hierarchical control system for an intelligent greenhouse. EurAgEng Paper No. 88-249. Silsoe, UK: Europoean Society of Agricultural Engineers.

Kok, R., and R. Lacroix. 1993. An analytical framework for the design of autonomous, enclosed agroecosystems. *Agricultural Systems* 43:235-260.

Kosko, B. 1993. *Fuzzy thinking: the new science of fuzzy logic*. New York, NY: Hyperion.

Lacroix, R., and R. Kok.1999. Simulation-based control of enclosed ecosystems – a case study: determination of greenhouse heating setpoints. *Canadian Agricultural Engineering* 41(3):175-184.

Lesser, V. R., R.D. Fennel, L.D. Erman and D.R. Reddy. 1975. Organization of the Hearsay-II Speech Understanding System. *IEEE Transactions on Acoustic Speech Signal Processing* (ASSP-23):11-23.

Margulis, L., and D. Sagan. 1986. *Microcosmos: Four Billion Years of Evolution from Our Microbial Ancestors*. New York, NY: Simon and Schuster.

McCulloch, W.S., and W.H. Pitts. 1965. A logical calculus of the ideas immanent in nervous activity. In *Embodiments of Mind*, ed. W.S. McCulloch. Cambridge, MA: MIT Press.

McNaughton, S.J., and M.B. Coughenour. 1981. The cybernetic nature of ecosystems. *The American Naturalist* 117:985-990.

Minsky, M., and S. Papert. 1969. *Perceptrons: An Introduction to Computational Geometry*. Cambridge, MA: MIT Press.

Minsky, M. 1985. *The Society of Mind*. New York, NY: Simon and Schuster.

Molenaar, R. 1998. Design and implementation of biosystem control for biosystem simulation. Unpublished Ph.D. thesis. Department of Agricultural and Biosystems Engineering, McGill University, Montreal, QC.

Patten, B.C., and E.P. Odum. 1981. The cybernetic nature of ecosystems. *The American Naturalist* 118:86-95.

Penrose, R. 1989. *The Emperor's New Mind*. Oxford, UK: Oxford University Press.

Pinker, S. 1997. *How the mind works*. New York, NY: W.W. Norton.

Pinker, S. 1994. *Mind Design*. In *The Language Instinct*, 404-430. New York, NY: William Morrow and Company.

Resnick, M. 1994. *Turtles, termites, and traffic jams*. Cambridge, MA: MIT Press

Sacks, O. 1998. *The Man Who Mistook His Wife for a Hat*. New York, NY: Touchstone Books.

Shukla, M.B., R. Kok, S.O. Prasher, G. Clark and R. Lacroix. 1996. Use of artificial neural networks in transient drainage design. *Transactions of the ASAE* 39(1):119-124.

von Neumann, J. 1963. The general and logical theory of automata. In *John von Neumann: Collected Works*, vol. 5, ed. A.H. Taub, 288-328. Oxford, UK: Pergamon Press.

Waldrop, M.M. 1987. *Man-Made Minds: The Promise of Artificial Intelligence*. Rexdale, ON: Wiley.

Williams, B.C., and P.P. Nayak. 1996. Immobile robots: AI in the New Millennium. *AI Magazine* 17(3):16-35.

Zeigler, B.P. 1990. *Object-oriented simulation with hierarchical, modular models: intelligent agents and endomorphic systems*. San Diego, CA: Academic Press.

CONNECTING TEXT

Chapter 7, **The characterization of biosystems**, was authored by O.G. Clark and R. Kok. At the time this thesis was submitted, this chapter had been sent for review to the editors of *Biometrika*.

This chapter is the last of the original material in the thesis, since the remaining ones are devoted to general conclusions, recommendations, and a summary of the originality of the work and contributions to knowledge. It contains a very general and theoretical analysis of the process of characterization. This is a fitting conclusion to the thesis because, overall, the work presented can be considered as being oriented toward the characterization of particular kinds of systems including, from general to specific, all biosystems (Chapter 3), substantially autonomous biosystems (Chapter 4), ecocyborgs (Chapter 5), and the mind of a substantially autonomous ecocyborg (Chapter 6). The concepts of *biosystem* and *ecocyborg* are, therefore, briefly reviewed in this chapter and the theoretical concepts are illustrated in this context. As part of this illustration, a number of characterization methods are presented that might be useful in further research in this field.

CHAPTER 7. THE CHARACTERIZATION OF BIOSYSTEMS

Abstract

This paper is about the characterization of *biosystems*, which are aggregate entities that are alive to some degree as a whole. Instances of the particular type of biosystems referred to here (i.e., *ecocyborgs*) comprise both biological and technological components, and have been engineered to be substantially autonomous. First, there is a general discussion of *characterization*, an epistemic process through which an observer transduces input signals into knowledge and might then express some part of this. Characterization is necessary for ecocyborgs that are substantially autonomous, since these must be able to characterize themselves, as well as for scientists and engineers so that they may generate explanatory and prescriptive descriptions, respectively. Characterization involves perception, discrimination, assimilation, conceptualization, and expression, and these are each discussed at some length. The knowledge resulting from characterization constitutes a *conceptual network* in the observer's mind. The constituent *knowledge objects* are of various degrees of *abstraction*, i.e., they may be somewhat removed from direct experience. Very abstract knowledge objects called *archetypal concepts* constitute a general schema (i.e., *paradigm*) for the organization of the observer's conceptual network. Other knowledge objects are *descriptions*, or abstract representations of some part of the conceptual network (as well as, perhaps, of external phenomena). Descriptions might be generated and possibly expressed as per formalized procedures called *measures*. These ideas are illustrated with particular reference to substantially autonomous ecocyborgs, which are the primary focus of the research project that is the context of this work. The research project and the systems being studied are briefly described. Next, a specific paradigm is presented that has proven to be useful in this research, and the last portion of the paper is a description of some measures that might be embedded in the paradigm. Relatively direct measures are described for quantifying the composition, structure, state, and comportment of ecocyborgs. More abstract procedures, such as measures of order and disorder, complexity, and emergence, are then discussed for the characterization of all of these various aspects in a more general sense.

7.1 Introduction

This article is about the characterization of systems that are alive to some degree. Members of this expansive class are called *biosystems*. They can range in scale from the molecular to the biospheric, and can be of natural, artificial, or combined origin (Clark and Kok 1999a, presented here in Chapter 3). Of particular interest here are biosystems that can be, or have been, engineered to be substantially autonomous, so that their comportment is relatively independent of external influences (Clark et al. 1999, presented here in Chapter 4; Clark and Kok 1999b, presented here in Chapter 6). The discussion is illustrated with specific reference to *ecocyborgs*. An ecocyborg is a composite entity that consists of an ecosystem combined with a control system, with the intent to make it substantially more autonomous than the ecosystem by itself. Such systems are the focus of the EcoCyborg Project, which forms the context of this research (Parrott et al. 1999, presented here in Chapter 5). The long-term objective for the project is the development of a general theory for the engineering of biosystems, with an emphasis on the design criterion of substantial autonomy. The short-term objective of the project is to develop computational models, simulations, and characterization tools for the study of ecocyborgs.

The second section of the paper is a general discussion of *characterization* as an epistemic process through which knowledge about a system is generated by an observer. There are three reasons why characterization is important in the context of substantially autonomous ecocyborgs, as well as for other kinds of biosystems. Firstly, in order for a system to have a high degree of autonomy it must be able to effectively guide its own comportment, and this requires that it be able to characterize itself (Clark and Kok 1999b, presented here in Chapter 6; Conant and Ashby 1970). Secondly, characterization is a necessary part of any engineering (prescriptive) activity, including the research encompassed by the EcoCyborg Project. Thirdly, it is central to any kind of scientific (descriptive) study.

Characterization can be construed as the way in which an observer transduces input signals into knowledge. In complement to this, the knowledge that is generated can also have an influence on the observer. The steps involved in the overall process of characterization include perception, discrimination, assimilation, conceptualization, and expression. *Perception* is the transduction of input signals into information, which may be segregated (*discriminated*) into coherent informational constructs called *knowledge objects*. These may then be associated with existing knowledge objects (*assimilated*) and thereby become integrated into the observer's *conceptual network*. The conceptual network as a whole consists of such experiential knowledge objects that are derived directly from perceived information, together with conceptual knowledge objects that are removed from perception to varying degrees, and which are therefore considered to be more *abstract*. The most abstract knowledge objects in a conceptual network are *archetypal concepts* that constitute a general schema (*paradigm*) for the organization of all of the knowledge contained therein. Some knowledge objects may consist of scripts for the processing of others. Such processing might result, for example, in the creation of a certain kind of knowledge object called a *description*, that is a more abstract representation of some part of the conceptual network. It may comprise, for instance, abstract information about a perceived phenomenon. Descriptions, or formalized fragments thereof, may be generated according to standardized procedures called *measures*.

Ultimately, a description or some other part of the conceptual network might be encoded in an output signal as an instance of *expression*. As pointed out above, the organization of a conceptual network may stem exclusively from the underlying paradigm, but this need not be so; the organizational flow can also move in the opposite direction. Thus, depending on the nature of the observer, knowledge objects of varying degrees of abstractness may be derived from less abstract ones through a process called *conceptualization*. In sophisticated observers, the entire conceptual network might, in fact, arise in this way.

These ideas are illustrated here with reference to the kinds of systems of particular interest in the context of the EcoCyborg Project. A more detailed description of these systems is presented in the third section of the paper. In the fourth section, a set of archetypal concepts is described. They constitute a specific paradigm which might be used in this context. Some related descriptive procedures (such as measures) are also explained. This is not an exhaustive catalog of all measures applicable to biosystems, nor even to just ecocyborgs. Instead, the intent is to illustrate the manner in which appropriate characterization tools might be chosen or formulated. The descriptive procedures that are mentioned form, in fact, only a small subset of all those that might be accommodated by even the single paradigm that is presented.

7.2 The epistemology of characterization

Epistemology is the study of the origin, nature, and limits of knowledge, as well as of methods for obtaining and generating it. *Characterization* is the process by which an observer generates knowledge about a phenomenon and it is, therefore, of an epistemic nature. The objectivity of the observer during the process of characterization is sometimes taken for granted. This may happen, for example, in the use of familiar procedures such as the standard methods of science. Consequently, the required internal activities of the observer during this process may be overlooked. On the other hand, the internal state of the observer is occasionally emphasized to the point where all experience is considered to be strictly limited to this domain (Maturana and Varela 1980). Such a viewpoint is regarded as being rather extreme and, in this paper, an observer is considered to be capable of interacting with external phenomena. It is, nevertheless, important to contemplate to what degree objectivity is possible in characterizing any phenomenon. The approach taken here is that, since there are multiple internal activities that must be performed by an observer during characterization, and since these activities can only be carried out in a limited manner, the descriptions that are formulated can never be complete, nor entirely unbiased. Therefore, observation can never be perfect, and no observer can be fully objective.

Perception, the foundation of characterization, is the interaction of two systems so that one is somehow affected by the other (Maturana and Varela 1980). One system is that which is being perceived, and the other is the observer, or that which is perceiving. Signals originating from the perceived system are said to be received as inputs by the observer, so that there is some kind of correspondence between them. The signal might affect the state of one of the perceptor's components, for example, or alter its structure or composition. The perceived system, on the other hand, need not be affected significantly by the act of perception. Often, only some part of the observer might be involved in perception and, in fact, many biological and technological systems have components (*perceptors*) that are specialized for this purpose. In the limit, a very simple observer might consist of nothing but a single perceptor. A piece of litmus paper can, for example, be interpreted in this way. Such a simple system has, however, a very limited capacity as an observer.

In reality, the two systems involved, as well as the signal relating them, must have both physical and virtual (informational) aspects, but it is the latter which is of primary importance in this discussion. From this perspective, therefore, the perceptor is somehow affected by the impinging signal so that it will contain information equivalent to some of that originally inherent in the signal. In a very simple observer system, such as the litmus paper by itself, mentioned earlier, this information remains in an implicit form. A more sophisticated observer can, however, *discriminate* the information, making it explicit by assigning meaning to the way in which the perceptor is affected. Some quality of the affected feature thereby becomes a symbol that explicitly corresponds to, or represents, some of the information transduced from the perceived signal. Thus, in the given example, the pH of the environment affects the color of the litmus paper, and the resulting shade might then be assigned meaning so that it becomes symbolic of the condition of the perceived phenomenon. The information that is finally represented by a particular symbol is an instance of a *knowledge object* (Clark and Kok 1999b, presented here in Chapter 6). Ideally, if perception and discrimination are effective, the information captured will be the same as the information inherent in the impinging signal and, therefore, in the phenomenon being perceived. Thus, discrimination of information is equivalent to discrimination of the phenomenon itself.

The process of discrimination can be applied iteratively so as to discern increasing detail, down to the limits of resolution (spatial, temporal, etc.) of the observer. A particular system might first be discriminated as being distinct from its surroundings. This is equivalent to the determination of the system's boundaries, or the establishment of other defining criteria (Clark and Kok 1999a, presented here in Chapter 3). It might then be further discriminated so as to distinguish between the system's internal components, for instance, or between the state of the system at different times.

The knowledge objects generated by an observer are examples of *virtual machines*, which are informational constructs that perform certain functions. In the case of a knowledge object, the function is the passive retention of information, but a virtual machine might also serve to compute, or somehow process information. Such active virtual machines are referred to as *cybernetic mechanisms*. All virtual machines must be resident on some physical substrate, although the same information can reside on many different substrates in any of a variety of different ways. For instance, the same body of information might be encoded as ink marks on paper, chisel marks in stone, or amplitude modulations in an electromagnetic field. As well, what information is associated with a system depends not only on the nature of that system, but also on the way it is observed. For example, any symbol can correspond to a particular knowledge, so long as that symbol is assigned the appropriate meaning. Thus, the same body of information might be encoded in different scripts, in different languages, or even with a single symbol. It is because of the multiformity of information that it can be transcribed from a perceived phenomenon, to a signal, to a perceptor, and to other virtual machinery in the mind of an observer. It is also because of this that equivalent cybernetic mechanisms, like summation algorithms, can reside on physically very different systems, such as electronic calculators, mechanical cash registers, and biological brains.

All of the virtual machinery together, including both knowledge objects and cybernetic mechanisms, constitutes the *mind* of an entity. Some systems, like the aforementioned, solitary piece of litmus paper, host such extremely simple virtual machinery that they are hardly describable as having minds at all, whereas others, like humans, host virtual machines of much greater sophistication (*intelligence*). Accordingly, knowledge objects may vary from being quite elementary to extremely complicated. Thus, they include the simplest kind of data that can be represented (e.g., in the context of digital computing, the particular state of a transistor represents a single bit of information), but they can also be large, composite constructs. Similarly, cybernetic mechanisms may range from extremely simple to very sophisticated devices. The informational activities that are actually performed by a mind are called *mentation*, and those to which it can potentially give rise are referred to as *mental abilities*. The latter include the various steps of characterization that are examined here. These terms were described in some detail by Clark et al. (1999, presented here in Chapter 4).

As described above, knowledge objects in the mind of an observer may originate, through perception and discrimination, from input signals that originate in the observer's surroundings. Input signals may, however, originate as well from the physical aspect of the observer system itself. These signals, too, may be perceived and discriminated to form knowledge objects that, in this case, are representative of some aspect of the observer. In both circumstances the resulting knowledge objects are directly experiential. Cybernetic mechanisms may also generate more abstract, conceptual knowledge objects (*concepts*) that do not arise directly from perception. The mind of a substantially autonomous ecocyborg, for instance, will therefore include experiential information about its surroundings and its own integral ecosystem, as well as conceptual knowledge of various degrees of abstractness that relates to that experiential information and to other aspects of the mind itself.

The various knowledge objects of all degrees of abstractness are associated with one another to form a branched and multi-referencing hierarchy, which is the observer's *conceptual network*. Depending on the nature of the mind in question, this network can be modified by cybernetic mechanisms in various ways and to varying degrees. A number of possible activities are referred to here and further below, and others were described by Clark and Kok (1999b, presented here in Chapter 6). One activity that cybernetic mechanisms may engage in, for example, is the *assimilation* of new, experiential-level knowledge objects into the conceptual network, so that the perceived phenomena to which they correspond become associated with concepts of various degrees of abstractness. As well, the cybernetic mechanisms may subject all of the knowledge objects to further processing, so as to modify the network's structure. In this way, for example, a number of knowledge objects might all become associated with a more abstract, representative concept which would then fill the role of a class object, so that the associated ones would be instances of that class.

Knowledge objects of greater abstractness may serve as templates for the structuring of less abstract, more experiential, knowledge. The entire conceptual network of an observer is seen, in fact, as being ultimately organized according to a set of very abstract *archetypal concepts*. Together, these archetypes constitute a *paradigm*, which is a general schema for structuring knowledge. One example of an archetypal concept is the equivalence relationship, which is an explicit representation of some similarity between two systems. Other examples of archetypal concepts are exclusion relationships, hierarchies, chronologies, and generative relationships (e.g., parent/child). A particular paradigm thereby accommodates the acquisition of certain kinds of knowledge, but at the same time limits what knowledge can be acquired. An observer might, in fact, be entirely incapable of assimilating information about a phenomenon in any manner that differs greatly from, or conflicts with, its current paradigm. A particular paradigm, the conceptual network that is built upon it, and the assimilated knowledge contained therein together constitute a model of reality in the mind of the observer. A sophisticated observer might, therefore, employ several alternative paradigms so as to be able to accommodate a number of different models of reality.

Whereas the existent conceptual network of an observer governs what kinds of knowledge can be acquired and how this might be done, it is also possible for conceptual knowledge objects to be derived from experiential knowledge. Thus, the cybernetic mechanisms may engage in a process called *conceptualization*. In this way, the organization of an observer's mind may arise, wholly or partly, from perceived information. This can happen in observers who are capable of substantially intelligent learning, examples being humans and (to a marginal extent) some of today's advanced computer-based learning systems (Michie 1999). When such an observer is confronted with new kinds of information, either as a result of perception or of internal processing, its mind may modify itself as necessary to incorporate that information. Observers who are less capable of learning may not be able to act in this way but may, nevertheless, still possess very sophisticated conceptual networks by design or inheritance, for instance. Hence, although they might be very intelligent in terms of perception or even reasoning, these faculties would not be very flexible (Clark and Kok 1999b, presented here in Chapter 6).

Thus, the mind of an observer may be influenced by its history, and the nature of the observer determines the extent to which this is so. The mind, however, is the virtual aspect of a system that also exists physically. Since virtual machinery is not entirely independent of the physical substrate on which it is resident, the properties of this substrate determine in part the kind of virtual machinery that the mind can comprise.

Thus, a particular substrate might be limited with respect to capacity or structural flexibility, for example, and make impossible the implementation of some kinds of knowledge objects or cybernetic mechanisms. Another factor that influences the observer's mind is the resolution at which the observer can resolve and manipulate (either consciously or not) the very substrate upon which its mind resides. Thus, the substrate limits the kind of paradigm that can be present, and therefore has a significant influence on the kinds of knowledge that can be acquired and generated.

Whereas the paradigms of human observers are usually implicit in the functioning of their minds, a paradigm can also be formalized and explicitly described with, for instance, a rule set. A number of paradigms have been very thoroughly formalized in this way including, for instance, those related to the various branches of modern science, and a number of religious world-views. The design of artificial minds is usually based on very explicit paradigms; in fact, essentially all information-processing devices available today can be characterized as having a very narrow and inflexible paradigm, although many of them are extremely intelligent and knowledgeable in their specific domains. Certainly, in terms of learning, such constructs are currently much less intelligent than humans. Accordingly, they are relatively incapable of conceptualization, especially in terms of generating archetypal concepts, and thereby modifying their paradigms. This is true, for instance, of the Object Oriented Knowledge Base (OOKB) and accompanying set of cybernetic mechanisms that has been proposed as the mental apparatus for substantially autonomous ecocyborgs (Clark and Kok 1999b, presented here in Chapter 6). In this case, the OOKB comprises all of the knowledge objects in the mind of such an entity, and is an example of an artificial construct based on a very particular paradigm, determined by the authors. It can be considered as a conceptual network that is somewhat flexible, but whose archetypal concepts are entirely fixed.

Once an observer has assimilated knowledge about a perceived phenomenon, its cybernetic mechanisms might further process that knowledge to create a set of representative objects called a *description*. Science and engineering often involve fragmentary descriptions created according to formalized procedures called *measures*. These can be made as objective as possible by specifying the way in which an observer is to perceive, discriminate, and assimilate knowledge about a phenomenon. They might, for instance, call for the use of an intermediate system called a *measurement device* in order to standardize perception and discrimination. A measurement device can be a very simple system consisting of only a single perceptor, such as the litmus paper mentioned previously, or it might be a much more complicated assembly. Assimilation can be made uniform by prescribing the use of particular archetypal concepts in structuring internal representations of the measured system, the measurement device, and the equivalence relationship between them. Certain equivalence relationships may also be specified between the state of the measurement device and the states of other carefully chosen constructs called *calibration standards* (possibly some aspect of a natural phenomenon such as a certain physical object).

As with other knowledge objects, descriptions (including measurements) can be of different orders of abstractness. Direct descriptions are most closely associated with perception, whereas more abstract ones are descriptions of other descriptions. A *mean*, for instance, is a somewhat abstract measure because it is based on more direct measures. Measures of different abstractness are suited to different purposes. Hence, the number and degree of abstractness of the measurements included in the description of a system depends on the goals of the observer. On one hand, for instance, the prescriptive specification of an ecocyborg for engineering purposes might involve relatively direct measurements such as the populations of different kinds of organisms in the ecosystem. On the other hand, for scientific descriptions used to compare ecosystems of very different kinds, these direct measures might serve only as a basis for calculating more abstract measurements, such as the homogeneity of the frequency spectrum of population change.

Abstract measures are, as mentioned, useful for characterizing and comparing disparate systems. In fact, entire bodies of theory have been developed around their use in this way. *Statistics*, for instance, is the study of large populations, with abstract measures that are based on sampling of individual members (Parker 1994). The construction of abstract measures for such purposes can sometimes be avoided by identifying phenomena that inherently summarize some of the local features of a system. It would be impossible, for example, to keep an accurate account of the metabolic activities of every individual organism in the biosphere but, because it is believed that they each influence the relative proportions of atmospheric gases, the ratios of these gases can be used as indicators of overall biotic activity. Lovelock, for instance, used this approach to characterize and compare the biospheres of Earth and Mars (Margulis and Lovelock 1974). He proposed that the atmospheric state of Earth, being farther from thermodynamic equilibrium than that of Mars, be interpreted as resulting from the cumulative effects of the biosphere, indicating that the Earth is the more biologically active of the two planets. The degree to which an atmosphere is removed from thermodynamic equilibrium was thereby employed as a measure in the characterization and comparison, at a highly abstract level, of the very different global biogeochemical aspects of the two planets.

Once a description of an observed phenomenon (or some other set of knowledge objects) has been formulated within the mind of an observer, it might then be communicated to another entity. Communication can only be successful, however, if a common conceptual basis exists that is shared by the two parties, and which adequately supports the knowledge that the observer wishes to convey. In other words, some part of the conceptual networks of the two entities must be based upon similar paradigms, the degree of similarity limiting the effectiveness of communication. The requirement for similar paradigms is often very evident when dealing with interactions between artificial entities, because of their relative inflexibility. This requirement also, however, holds true for humans. Whereas it is often taken for granted that they possess similar paradigms, and although this is likely to be true to some extent, especially if their experiences derive from a common culture, there can be substantial differences between individuals' paradigms. Overall, before knowledge can be transferred effectively, a mutual paradigm must somehow be specified beforehand, perhaps based upon an alternate one.

Thus, once a common paradigm has been established, and assuming that a channel of communication is available, it is possible for communication to take place. This involves the use of a *language* to encode some part of the knowledge of an entity onto an output signal. A language comprises a set of symbols and a grammar, and lies within a particular paradigm. Symbols, as mentioned, are features of a system or system component (often physical, e.g., uttered sounds or written figures) to which meaning is assigned. This is done by ascribing to the symbols some arbitrary correspondence to particular knowledge objects. A *grammar* is a set of rules for combining these symbols.

The set of all languages includes natural and formal languages: examples of natural languages are those generally used by people to communicate in their everyday lives; formal languages are analytic and synthetic conventions, such as those used to work with differential equations, cellular automata, etc. An *expression* is a string of symbols that is created in a language, and which represents some part of the knowledge of the observer. Since the initial choice of a symbol set is entirely arbitrary, an expression can take any convenient physical form, like vocalization or ink marks on paper. The virtual aspect of the expression is the information, or *message*, that is encoded in it, and this may contain one or more descriptions.

In order that a message might be generated that codifies a description (such as a measurement) in a relatively objective manner, measures may also include a procedure for expression. When a measurement device is used, for instance, the whole descriptive process might be abbreviated to the transcription or translation of some numeric symbols from the device to some other storage medium, like paper or a magnetic substrate. As part of the Western scientific paradigm, a large number of measures have been defined in great detail (e.g., Specter and Lancz 1986; Lyman 1982). These are meant to be used in scientific work, and are strictly codified methods in which certain kinds of measurement devices, calibration standards, equivalence relationships, procedures for expression, etc. are specified.

The terms used to refer to some measures in common parlance often have several kinds of meaning, which can lead to confusion: they may be used in reference to a procedure (as mentioned), a concept abstracted from such a procedure, or a measurement. Because humans tend to conceptualize to increasing degrees of abstraction, many procedural measures have become so familiar as to be transparent, i.e., often the associated concepts have been abstracted to the extent that they are no longer associated with a procedure at all. The same word, for example *length*, is therefore often used to refer to the procedure (or several equivalent procedures), the abstract concept which originally corresponded to that procedure, and to a numerical value. There are numerous other examples of familiar concepts that likely originated in this way, including height, weight, time, etc. Each of these probably had its origin in a comparative procedure, but an associated concept was formed that gradually became so reified that now it is considered as a feature or property that is commonly considered to have an objective, independent existence.

Furthermore, the same term might be used in reference to many procedures, concepts, or measurements, and this can also be a source of confusion. For example, several methods bearing the same name might be used to generate measurements. These might be equivalent, in which case the resulting values will be of the same magnitude and will be referenced to the same standard. Often, however, the procedures are not, in fact, exactly equivalent, and are referenced to concepts that are only similar at a very high degree of abstraction. For instance, one might refer to the length of a day, the length of a road, the length of a computer program, or the length of a song. However, the various concepts that the word *length* refers to in these cases are only distantly related to one another, and the procedural measures that might be used to generate descriptions corresponding to each of them are very different. Finally, measurements that are generated by equivalent procedures and that are referenced to analogous concepts might still not be equal if, for instance, those procedures were employed at different resolutions.

Hence, it is apparent that a cogent discussion of characterization first of all requires that the flexibility of the lexicon be acknowledged, and that not only the system of interest, but also the observer, and the interaction between the observer and the observed system be taken into consideration. Perception and discrimination are constrained by the nature of the perceptors and their adjunct cybernetic mechanisms (Clark and Kok 1999b, presented here in Chapter 6). Discrimination is also influenced by the paradigm employed, with its gamut of archetypal concepts and descriptive methodologies (such as measures), as are assimilation and conceptualization. Expression is limited by the language or languages available. All of the above are affected by the nature of the physical substrate upon which the observer's mind is resident. This results in the knowledge about a phenomenon being composed and structured in a very particular way, so that some aspects are stressed and others are not represented at all (Gould 1980). This is so when direct measurements are obtained, but is especially true when more abstract measures are used to generate a description because some information is sacrificed with every step of the descriptive process. The ubiquity of such biases must be recognized and considered when choosing characterization methods, so that they will be appropriate for the circumstances in which they are used, e.g., by an autonomous biosystem in characterizing itself, by a scientist in formulating a protocol for observing such systems, or by an engineer in specifying the design of such a system. The next section sets the context for this deliberation by describing the systems of interest and the observers involved in the EcoCyborg Project.

7.3 A context for characterization: the EcoCyborg Project

The overall, long-term objective in the EcoCyborg Project is the development of a general theory of biosystems engineering. Biosystems, however, constitute a very broad class. They differ greatly from one another in organizational scale, composition, and origin, ranging from tiny to immense, from primarily physical to mainly virtual, and from wholly natural to completely artificial (Parrott et al. 1999, presented here in Chapter 5). In order to make the project more manageable, therefore, its scope has been limited, for the time being, to the investigation of a very particular kind of biosystem: ecocyborgs, ecosystems combined with technological control networks. Moreover, these are not being studied in a direct way but, instead, through a modeling and simulation approach. It is felt that this methodology is a much more flexible and effective approach to the initial development of a theoretical basis for biosystems engineering than would be the direct study of large-scale, real (and presumably, primarily physical) systems. Correspondingly, as stated in the introduction, the short-term objectives of the project are related to the development of various computational methods or tools to model, simulate, and characterize ecocyborgs. However, although these tools are being created with the investigation of ecocyborgs in mind, their wider application is also continually considered, corresponding to the long-term project objective. This is especially true of the characterization methods. Hence, in the context of the EcoCyborg Project, these methods are considered in terms of their utility with respect to the set of all biosystems, as well as to the set of all ecocyborgs and computational models thereof.

With respect to the first group mentioned, the set of all biosystems, each member has both physical and virtual aspects, and upon consideration of these it has become apparent that the common characteristics among them are primarily *virtual*, or informational. One feature that is common to all living things, for instance, is that they are self-organizing, which is to say that, as a result of their structure, their states tend to remain within a bounded subset of those that would otherwise be theoretically accessible (Clark and Kok 1999b, presented here in Chapter 6). Most importantly, however, living entities are distinguished by a particular dynamical mode called *autopoiesis* (Maturana and Varela 1980). In autopoiesis, the components of a system interact at the local scale to form an organizationally-closed network of relationships at the system scale, so that every kind of component is replaced or regenerated, either from other components or from inputs. From a physical perspective, the maintenance of autopoiesis depends on a continuous flow of some kind of ordered input that can be degraded and rejected to the environment as a less-ordered output stream. Thus, a biosystem's maintenance of a far-from-equilibrium internal state results in the generation of thermodynamic entropy (Prigogine 1980; Clark and Kok 1999a, presented here in Chapter 3). The overall effect is that local interactions combine to yield a homeostatic structure that is an example of the "flowing balance" of nature (Capra 1996). Since the comportment of a particular entity can be autopoietic to any degree, vitality is a variable quantity and, on this basis, one biosystem can be considered as more alive than another.

The definition of life as being a primarily virtual phenomenon provides a useful basis for the discrimination of the otherwise disparate set of all biosystems from other phenomena. Thus, when the informational aspects are emphasized, the comportment of biosystems is seen as involving the acquisition, storage, transmission, and processing of information. From this perspective, autopoiesis, the defining feature of life, is the continuous computation of self-induced internal adjustment so as to approximately nullify the effects of any environmental influences which might disrupt the autopoietic structure itself. Hence, a construct may be substantially autopoeitic regardless of whether its apparently most significant aspects are primarily virtual or predominantly physical. Accordingly, a computer-based construct may be alive to any degree and therefore qualify as a biosystem (Clark and Kok 1999a, presented here in Chapter 3). Although it might not be immediately evident, because even such a predominantly virtual system must reside on a physical substrate, any degree of autopoietic comportment on its part still eventuates the generation and export of entropy. The computational systems being developed in the EcoCyborg Project are predominantly virtual systems of just this type.

As pointed out above, in the short term, characterization in the EcoCyborg Project is oriented specifically toward ecocyborgs and the computational systems that are being created to represent them. If the latter are to be useful representations of ecocyborgs, they must share with them the features that are of interest in the investigation. They should, therefore, adequately portray the relationships between the constitution (composition and structure), state (especially the initial state), and comportment of such ecocyborgs, as well as the way in which these are influenced by external forcing functions. Since one of the principal characteristics of ecocyborgs (or of any other biosystem) is that they are substantially autopoietic, the computational systems representing them should also demonstrate such comportment to some degree. This is, as mentioned, intended to be the case. The characterization methods being developed will, therefore, be applicable to both ecocyborgs and the computational systems representing them, as well as, in the case of the more abstract measures, other kinds of biosystems.

Like some other computational models (e.g., Ray 1994), those of the EcoCyborg Project depict ecological systems. In the organizational hierarchy of natural, physical biosystems, *ecosystem* falls between *population* and *biome*, and is traditionally being considered to comprise a biotic community (the *biota*) together with its *abiotic surroundings*. As part of the project work, prototype computational models have already been developed that can represent a number of populations (typically twenty, corresponding to species) as collections of several thousand organisms, each of which has a number of attributes (Molenaar 1998). A more sophisticated model is currently under development that will represent ecosystems composed of up to one thousand species of plants and animals inhabiting several hectares of terrain; a temperate, open woodland ecosystem is first being modeled (Parrott 1995). In order to simplify the modeling task, the ecosystem is assumed to be sealed inside an enclosure (the *envelope*) with energy being supplied from and rejected to the surroundings. The ecosystem is subject to several weather forcing functions (ambient temperature, radiation, and rainfall) and to guidance from a control system. All together, the components form an (hypothetical) ecocyborg which, in this case is an orbital space station. The various models (i.e., of the ecosystem, the enclosure, the forcing functions, and the control components) are used in simulation experiments to study the comportment of the ecocyborg under the influence of the forcing functions. All the models, as well as the simulation platform, are generally configurable so that a wide variety of ecocyborg constitutions can be represented under a range of circumstances.

There are three distinct groups of observers who have an interest in characterizing the systems being studied in the EcoCyborg Project. The first group comprises ecocyborgs and the computational systems representing them. It is intended that the models eventually represent highly autonomous ecocyborgs to an acceptable degree and in order to do that, the models will need to be highly autonomous also. This means that both types of systems must be effective in operating without external guidance (i.e., be *automatic*), capable to some extent of developing their own goals and strategies for attaining them (i.e., be *volitive*), as well as active in executing these strategies (i.e., be *intentful*). The approach that is being followed with regard to the ecocyborgs is to create them as assemblages of both biological entities and control networks, the latter comprising perceptors, effectors, and various other cybernetic mechanisms. For the computational systems to be adequate models of the ecocyborgs, they should have features similar to these. Both types of systems must, therefore, include control systems hosting minds that enable them to formulate and monitor their own progress toward their own goals, and this requires that they possess the general kinds of characterization apparatus that has been described previously. The other observers with a need to characterize are those who are studying the ecocyborgs and their computational equivalents. These observers, first of all, have a descriptive, scientific, agenda in that they wish to explain the constitutions, states, and comportment of the two types of systems as these respond to forcing functions. They also have the prescriptive, engineering agenda of learning to create computational models and, eventually, full-fledged biosystems (such as ecocyborgs) that fulfill predetermined criteria. All of these observers require measures of various degrees of abstractness, for reasons described below, and examples of these are presented in the rest of the paper.

7.4 Characterization of ecocyborgs

This section contains a brief description of a paradigm, or set of archetypal concepts, that might be used as a foundation for the characterization of the types of systems mentioned above, i.e., biosystems in general and, specifically, ecocyborgs and computational models thereof. There is also a discussion of a number of specific descriptive methods (such as measures) that are deemed to be appropriate for the characterization of such systems. These epistemic tools might be used in characterization for scientific or engineering purposes, as well as in self-characterization by substantially autonomous entities.

As it must, of course, the paradigm adopted in this article includes archetypal concepts according to which knowledge about systems can be organized. Here, a system is generally considered to have a *constitution* comprising all of its features that are approximately invariant (Clark and Kok 1999a, presented here in Chapter 3); some qualifications to this are discussed below. The two aspects of constitution are the number and kinds of components included in the system (*composition*) and the relationships between these (*structure*). The values of any changeable attributes of the system constitute its *state*, and the way that these values change with time is its *comportment*. Finally, any significant influences on the system that are not defined as being part of it are considered to be *forcing functions*.

The archetypal concepts that comprise this paradigm might be interpreted in different ways, as determined by the needs and capabilities of various observers. A system might, for example, be considered to persist as the same entity even if aspects of its constitution varied somehow with time. Many primarily physical biosystems, for instance, depend on an exchange of material components as a medium for the rejection of entropy, and an interpretation of constitution as being strictly invariant would necessitate the discrimination of a new system every time that such an event occurred. Also, the distinction between features that are variable (state or comportment) or invariant (constitution) can be affected by the temporal and spatial resolutions at which a given system is characterized. As well, the discrimination of the individual features of a system depends on the observer. In the specific case of ecocyborgs, for instance, particular components might be considered to be biotic by one observer, but abiotic by another, or a medium that is considered as homogeneous by one observer might be characterized by another as a heterogeneous mixture of several distinct media.

The paradigm described above is compatible with numerous areas of knowledge and associated methods and procedures that are useful for characterizing ecocyborgs, other kinds of biosystems, and computational models of these. For instance, in the following discussion several branches of mathematics are mentioned in which are couched numerous measures useful for generating formal descriptive fragments (measurements) of the types of systems of interest in the EcoCyborg Project. Some of these are better suited for use in some contexts than in others, depending on the needs and capabilities of a particular observer and on the nature of the observed system. For example, they might vary, as mentioned, in their degree of abstractness, or in their resolution.

System-scale or multiscale features are important not only in the study of particular systems but some, like autopoiesis, are common to broad sets of systems. Such features are generally not straight-forward combinations of the local features of system components, and correspondingly abstract measures are therefore required by both external observers and the ecocyborgs (and their computational models). The utility of more abstract measures is illustrated, for example, by recent trends in taxonomy. Phenotypic characters have historically been used to discriminate between species of biological organisms, but the statistical analysis of genotypes is now emerging as an alternative approach that is more abstract and appears to be superior in some respects. This has actually resulted in a substantial shift in the paradigm of biological taxonomy, and the relationships between some species are now characterized quite differently than before.

More direct measures are also very useful in the EcoCyborg Project, in the detailed characterization of a particular system or set of similar systems. For example, they might be used in the engineering specification of a proposed computational ecocyborg model. Also, sets of number of similar systems could be compared with direct measures, as in a scientific study of different instantiations of a computational model. The repeated characterization of the same system could also be achieved in this way, as in the self-evaluation by a substantially autonomous ecocyborg of its own progress toward achieving its goals.

The remainder of the paper is, therefore, a discussion of both some direct measures and some abstract measures. Particular aspects of composition are dealt with first: the biotic components of ecocyborgs, the abiotic surroundings, the virtual machines, and the envelope that contains them all. Some measures for characterizing structure are then presented, followed by several that pertain to state and comportment. The same kinds of measures that are used to describe comportment might also be applied to the analysis of the forcing functions that influence the systems. Finally, three types of more abstract measures are introduced, which might be suitable for the characterization of broader sets of ecocyborgs and their computational representations; these are measures of order, complexity, and emergence

7.4.1 Measures of composition

Composition refers to the numbers of components of different kinds that a system comprises. Following a strict interpretation of *system* as presented here, the composition of a system is invariant, which is to say that, if composition changes then the identity of the observed system also changes. This restriction can be relaxed somewhat for the sake of convenience and practicality, especially in the characterization of materially open systems such as most kinds of biosystems (Clark and Kok 1999a, presented here in Chapter 3). Characterization of the composition of the hypothetical ecocyborgs (and their computational models) discussed in this project is more straightforward in this respect than for many other biosystems because they are envisioned as being contained in a materially closed envelope (i.e., a space station). As with other kinds of measures, quantifiers of composition may be of various degrees of abstractness, with more abstract one being based on those that are more direct. As well, measures can vary with respect to the spatial resolution of observation at which they are employed. Although composition is, as mentioned, generally considered to be static with respect to time, the temporal scale over which a particular feature is observed does have an impact both on its value and on whether or not this is perceived to be static. Thus, temporal resolution is also an important consideration.

The most direct measures of coarsest resolution that might be used to characterize the composition of ecocyborgs (and their equivalents) are those that quantify them as a unitary entities. For instance, the internal surface area and the enclosed volume of a particular ecocyborg's envelope have unvarying values, although they might be different for various ecocyborgs (and their computational models might have been configured differently). Some abiotic features (atmosphere, water, and soil) of the ecosystem segment of these ecocyborgs can be considered as unchanging if characterized at sufficiently coarse resolution, at appropriate time scales. For instance, atmospheric pressure is maintained constant, and the relative proportions of atmospheric constituents remain within specified ranges suitable for the organisms included in the ecocyborg. Measures of slightly finer resolution of observation might also be utilized similarly. For instance, the global masses of various elements and compounds and, perhaps, their distribution among various spheres of the ecocyborg (storage, biotic, abiotic) might be approximately constant over relatively short time spans.

Similar kinds of relatively coarse-resolution, direct measures of composition might also be used to characterize mind. In an ecocyborg, the latter is likely to be manifested at least partly as computer software and will certainly take this form in the computational models of these. It can, therefore, be characterized in both these cases by quantifying the overall length of the associated software. This kind of measure is related to *algorithmic information*, which is discussed later in the article (Chaitin 1977). Of course, the computational models consist entirely of software (resident on a physical machine) and, so, such measures might also be applied to the model as a whole.

Finer-resolution, direct measures of an ecocyborg's composition might be based on the discrimination of smaller-scale biotic, abiotic, and mental components. It must be remembered that it is the observer who imposes a taxonomy on the components by perceiving, discriminating, and assimilating the information about the system. Care must be taken, therefore, to use an appropriate scheme that is not based on meaningless or unfounded categorization. The temporal resolution of observation is also important in this context because the discrimination between components is based on features that are approximately invariant at the temporal resolution being used. Moreover, as the spatial or taxonomic resolution of observation is increased, the rate at which these features change is likely to increase as well. The temporal resolution may therefore have to be increased correspondingly because, if the values of these features can be seen to change, then they can no longer be considered as invariant aspects of the system's constitution.

Thus, finer-scale measures of biota can be defined to function at various resolutions according, for instance, to taxonomic classification, from kingdom down to individual organism. Evidently, in the case of a computational model, the resolution of the measure is limited by the detail of the model and so, for example, some plants that are modeled as clumps could not be observed at resolutions finer than this. Some abiotic features, such as soil, are immobile and they too can be thought of as comprising discrete, interacting components with fixed locations. The identification of these might be done on a grid of arbitrary resolution (limited, as before, in the case of the model by the detail of that model) and they could then serve as the basis for observation of various related hydrologic, meteorological, and chemical phenomena. In the context of mind, the abilities of an ecocyborg might also be discriminated at finer resolutions, along the lines of human mental faculties such as perception, expression, memory, reason, and learning, for example, or according to even more restricted categories (Clark et al. 1999, presented here in Chapter 4). The minds of the ecocyborgs and, similarly, of the computational models will, as mentioned, be manifested largely or completely as computer software, and so the amount of code associated with each faculty might serve as an approximate indicator of the intelligence, or sophistication of those faculties.

The kinds of measures that have been mentioned up to this point may also have utility beyond the generation of direct descriptions of composition. First, the measures of composition, if they are appropriately selected, may be related to the number of parameters that might be varied in specifying the system at the chosen resolution. Second, direct measures serve as building blocks for more abstract ones. Abstraction can result in more concise description, and can also reveal features that are not immediately evident from more direct measures. Information is, however, inevitably lost in abstraction, and conceptual artifacts can be introduced if due care is not taken. Like direct measures, therefore, abstract one must be judiciously chosen. Some examples of more abstract measures that might be useful in the context of ecocyborg composition include grouping methods, such as discriminant analysis, and various linear, superficial, and volumetric densities.

7.4.2 Measures of structure

Structure refers to the overall set of interrelationships between a system's components (i.e., how they interact), including the *number*, *sense*, *magnitude*, *form*, and *type* of these relationships. The structure of a system may be thought of as the rule set that dictates the way in which the features of a system change with time, perhaps in response to forcing functions. Equivalently, the interrelationships within a system can be considered as the means by which information is stored, transmitted, and modified. When thus couched in virtual terms, it becomes evident that, as mentioned, the computational condition of autopoiesis can arise from extremely diverse types of phenomena, concurrent with the physical diversity of possible biosystems. In ecosystems, for example, interactions are mediated through large-scale biogeochemical cycling, predator/prey activities, the transport of trace chemicals through various media (e.g., air, water, soil), etc. Although these are usually considered primarily in a physical way, they can also, however, be regarded in terms of their equivalent, virtual aspects. Conversely, for other types of systems, such as an ecocyborg's control network, internal interactions are usually considered primarily in a virtual way, but these can also be regarded in terms of their equivalent, physical aspects, i.e., the interactions between electrical circuits in semiconductor wafers. This is also the case for systems like computational models of ecocyborgs. Measures of structure are intended to capture some aspect of the set of interrelationships between a system's components and may reflect either or both physical and virtual aspects, whatever is considered to be of primary interest from the perspective of the observer. It is to be noted that, like composition, structure is essentially non-varying although, as with composition, this restriction can also sometimes be relaxed for reasons of convenience and practicality. Hence, the measures that are formulated should quantify relationships that are approximately invariant at the spatial and temporal resolutions which are used. Theses measures can, however, be applicable over diverse ranges of resolution and, moreover, be of various degrees of abstractness.

A variety of languages can be used to formulate measures useful for composing descriptions of structure. These languages include, for instance, those corresponding to various branches of mathematics and consisting of formal symbol sets and grammars. Some examples of the descriptive expressions that can be created with these are directed graphs, matrices, cellular automata, sets of differential equations, and state space vector fields (Green 1993). These various languages may be used to create descriptive constructs that are isomorphic with respect to one another, although any given method will yield a description in which some system features are emphasized more than others. Some of these methods might therefore be well suited to the study of certain types of systems but not others, or they might correspond to the goals of one observer but not to those of another. The expressions that are generated may comprise descriptive fragments that can be interpreted as measurements, and these might serve as the basis for other measures, which would hence be of greater abstractness. Even in the limited context of ecocyborgs and their models, innumerable measures of structure can be formulated, concomitant with the great diversity of available languages, contexts for observation, and possible kinds of systems.

Minimal, direct descriptions of structure refer to the *connectance* of a system, and indicate the existence of any relationships between the components without specification of their type, magnitude, or sense. The number of identifiable relationships varies, of course, with the observational approach that is adopted. For instance, higher-resolution observation might reveal more relationships of lesser magnitude; similarly, the resolution can be decreased by using a *slicing parameter* to specify the magnitude above which relationships are considered as significant (Gould 1980). *Connectivity*, a related but somewhat more abstract measure, is the ratio of the number of actual relationships between components as compared to the number of possible relationships. Since it is based on connectance, the value that will be obtained will also vary with the observational approach (e.g., the value chosen for the slicing parameter).

With most systems, the study of one aspect will yield information about others. This is true for highly complex systems, whose structure and comportment have been found to be often closely related (Green 1993). For instance, when connectivity of a system is low (*subcritical*), its comportment tends to be relatively stable or even static, whereas when it is high (*supercritical*) the comportment tends to be fluid and unstable. At intermediate levels of connectivity, abrupt, qualitative shifts between extreme dynamical modes often occur, similar to phase changes in a physical material. (The importance of this phenomenon in the context of biosystems is that, in such cases, spatial and temporal patterns tend to arise that can correspond to sophisticated computation. This is related to *complexity*, a more abstract kind of measure, which will be discussed later.) Thus, knowledge of the connectivity of a system can sometimes be used to evaluate what type of comportment it is likely to exhibit, as limited by the observational approach used for the connectivity measure.

A number of other structural measures can also be used as indicators of the likelihood that the systemic "phase changes", referred to above, will occur. When a comparison is made of systems that have similar structures but differing values of connectivity, it is observed that the range of intermediate connectivity (concurrent with the greatest frequency of systemic phase changes) also corresponds to the greatest variation (with respect to connectivity) in the size of the largest group of contiguously connected components (*patch*) (Green 1993). As well, it corresponds to the greatest rate of decrease (again, with respect to connectivity) in the number of separate patches, and to the greatest number of relationships in the shortest chain spanning the largest patch in a system. All of these measures may therefore be useful in predicting system behavior. Evidently, if corresponding, reverse, prescriptive rules could be found, these would be very useful in the engineering of systems that could then be designed to have particular features, such as the aforementioned sophisticated patterns associated with the systemic phase changes.

Descriptive methods exist in which not only the presence, but also the *sense*, or directionality, of the interactions between components is taken into account. *Loop analysis* is one such method, which is based on the identification of any closed causal loops among components (Levins 1975). Such loops are important because they can be negative feedback circuits that contribute to the stability of a system, or even to the organizational closure of a substantially autopoietic entity. In loop analysis, system components are represented by nodes of a graph, and their interactions by the edges between these. It may sometimes be sufficient to indicate only the sense of the interactions as being either positive or negative, and neglect their magnitudes (the interactions included in the analysis having first been determined with the use of a slicing parameter, as per above). Based on this, a measure can then be constructed (Equation 7.1) to yield a *feedback constant*, which is equivalent to *gain* in the engineering literature. For a system to be stable, it is necessary, but insufficient, that the feedback constant be negative. (In the classical sense, this means that the system tends toward a steady state, such as equilibrium.) A positive value, on the other hand, indicates instability. Such a measure might be used by an ecocyborg, for instance, to determine the stability of its ecosystem segment.

$$F = \sum_{i=1}^{n} (-1)^{m(i)+1} \cdot L_i \quad (7.1)$$

F	feedback constant (unitless integer)
i	number of nodes in loop (unitless integer)
n	maximum loop length in system (unitless integer)
m	number of nonintersecting loops having i edges (unitless integer)
L	product of all the edges of all loops of length i (-1 or +1)

Yet another method of representing structure is to use an extension of set theory called *Q-analysis* (Casti 1994; Gould 1980). In this method, each component of a system is represented by a *vertex*. A group of linked vertices is called a *simplex*, which represents components that share some kind of relationship. Associated with each simplex is a *dimension value* which, in this context, is defined as one less than the number of vertices than the simplex includes. If two simplices share one or more vertices then they are *q-near*, where *q* is one less than the number of vertices that they share (Casti 1994). When a set of simplices exists in which all adjacent pairs are at least *q-near*, then it is a *q-chain* or, more generally, a *q-network*. Any two simplices in such a network are therefore *q-connected*, even if they are not adjacent to one another. All the vertices and simplices, which together denote the overall structure of the system, are called a *simplicial complex*. Q-analysis might be used, for example, to represent a food web in the ecosystem segment of an ecocyborg. Each species preyed upon by coyotes might be represented by one vertex of a "coyote prey simplex". If coyotes preyed on six different species, then the coyote prey simplex would comprise six vertices and have a dimension value of five. If foxes and coyotes shared four prey species, then the fox prey simplex and the coyote prey simplex would be connected in three dimensions, making them *3-near*. The numbers of independent q-networks that exist in the simplicial complex at each dimension together constitute a measurement called a *structure vector*. In the above example, a complete structure vector would reflect the integration of the food web by revealing whether or not higher-dimensional q-networks decompose into separate networks at lower dimensions.

Greater detail about the integration of a particular simplex with the rest of the complex can be obtained with a measure called *eccentricity*. This is done by subtracting the dimension value of the largest shared face of a simplex from the overall dimension value for that simplex, and then dividing the difference by the smaller value. Of course, as for the other measures described above, all these methods depend on the resolution of observation used, which will determine how many components and relationships are identified, and the values of any slicing parameters, which will determine how many of these are considered as significant.

A simplicial complex can also be made to reflect the *magnitude* of the interactions within a system. A set of rules, called a *pattern* in this context, is used to assign values to each vertex (such as the frequency with which coyotes feed on the corresponding species), and an overall value for the simplex is then calculated from these (Casti 1994). These rules constituting a pattern might be linear formulae, or more complicated mathematical functions. The resulting simplicial complex and associated numerical values are a description of the structure of the system at a given time. Q-analysis can also accommodate mappings other than functions, such as one-to-many mappings, and even relationships more general than mappings, such as nondeterministic influences (Gould 1980). With a strict interpretation of *system constitution* (i.e., constant composition and structure), a simplicial complex provides an unchanging *static backcloth* for system comportment. Thus, for example, although the number of species in a food web and their interactions might be assumed to remain constant, their populations might vary, resulting in changing numerical values associated with the vertices and simplices. Q-analysis can also, however, be used to study scenarios in which *system constitution* is defined somewhat more loosely, so as to allow some change in its composition or structure, without considering this as a transformation of the initial system into an altogether different one (Gould 1980). For example, over the long term, some species might be extirpated and the geometry of the simplicial complex would change correspondingly, altering the possible patterns of traffic. The magnitudes of the interconnections in the web might also change with many factors appropriately considered part of system comportment, such as the age of organisms, species abundance, and climate (Polis and Strong 1996). The resulting changes in the values of vertices and simplices are called *traffic* on the complex. Thus, q-analysis is a method for the characterization of both system structure and comportment.

Matrix notation can also be used to denote structural connectance, as well as the sense and magnitude of relationships in a web. A very simple instance of this is based on the representation of the composition of a single-species system according to discrete developmental stages of the organism, so that the sizes of the resulting subpopulations can be described with an *age vector*. The appearance and disappearance coefficients for each of these subpopulations might then be determined from either theoretical or empirical information and listed in a combined *transition matrix*. From this information, the dynamics of the system can then be predicted by multiplying the age vector with the transition matrix, summing the resulting population changes and the age vector, and iterating these two steps (Logofet 1993). This classical mathematical form of representing a single, structured population is called a *Leslie matrix model*. If interactions between the population segments are at least approximately linear, then much of classical stability theory can be brought to bear on this approach (Logofet 1993). If this is the case, some expectation of the stability of the system can be gained by calculating the characteristic roots (*eigenvalues*) of the transition matrix (Logofet 1993). Levins (1975) has described how the values of these measures can sometimes be estimated for a system according to a matrix model of this type, even if some of the coefficients are unknown. Such predictive characterization tools would be useful to an ecocyborg for the purpose of controlling the species populations in its ecosystem segment, for instance.

Communities that consist of several distinct species can be similarly described with more general matrix methods, where a *species population vector* replaces the age vector of the Leslie matrix model, and an *interaction matrix* replaces the transition matrix. As in Q-analysis, forms of interaction other than approximately linear ones may be represented with such descriptions by including functional forms as the matrix elements. Moreover, a number of parallel matrices can be used, with each representing a different type of interaction. For instance, one matrix might represent the predator-prey interactions of a food web while another represents less direct influences (Molenaar 1998).

As the complication of the observed system increases, matrix descriptions quickly become intractable to methods such as stability analysis. It would be very difficult to evaluate the stability of an ecocyborg's ecosystem segment directly from a Leslie matrix description, for example, because it is likely that many of the populations would be interacting in complicated, nonlinear ways. Nevertheless, such representations may still be useful for comparative purposes, and can serve as models for short-term prediction in simulation-based control. The early prototype ecosystem models used in the EcoCyborg Project were based, for instance, upon such population vectors and numerical interaction matrices. The more sophisticated computational models currently being developed are essentially based on this approach as well, but highly nonlinear elements can be included in the matrices (Parrott 1995).

7.4.3 Measures of state and comportment

Whereas the focus in the previous section was on the characterization of a system's constitution, in this section it is on the characterization of its state and comportment. Here, constitution is considered to be constant with respect to time, and the *state* of a system comprises the values of its temporally variable features. Ideally, the features that are chosen for evaluation are entirely independent from one another, thereby maximizing the amount of information derived from observing them; such optimal features are called *state variables*. As with constitution, any of a variety of methods can be used in the characterization of state variables, depending on the observer's preferences with respect to features of interest, resolution, and degree of abstractness. Hence, it may be possible to devise different measures that yield equivalent sets of measurements corresponding to the state variables. Such a set of values is called a *state vector*, and is a parsimonious characterization of the (independent) features of interest to the observer, at the desired temporal and spatial resolutions. Because all equivalent state vectors completely describe a system's state as observed according to the aforementioned preferences, they all address the same number of degrees of freedom of the system. Direct measures must initially be used for this purpose, similar to those discussed for constitution, and more abstract ones based on these can then be applied.

Abstract measures are commonly formulated as comparisons of different values of more direct measures. Thus, a series of direct measurements, inherently ordered with respect to some independent variable, such as space or time, can be analyzed to reveal changes that occur in the dimension corresponding to that independent variable. It is to be noted that sets of data ordered with respect to space, time, or any other variable, are not necessarily of essentially different character, and can often be described with similar methods. Data that are ordered with respect to time, however, are often dealt in a particular manner. This special treatment is mostly due to the paradigm in which characterization takes place and, here, the way in which a system's state changes with respect to time is even accorded a special name (*comportment*). Some of the methods discussed below are suitable to the description both of state and of changes therein, such as comportment, whereas others are suited only to the description of change.

One convenient way to characterize change in an ordered set of values is to approximate these with an algebraic function whose argument is the independent variable with respect to which the change occurs. For example, such data can often be approximated to any arbitrary degree of accuracy with a linear polynomial, whose coefficient values might be calculated from the data itself with a statistical procedure such as least squares regression (Barnes 1988). It may be more appropriate in some cases to use nonlinear functions, perhaps involving more than a single independent variable, having coefficient values estimated with any of a number of techniques. Regardless of the form of the function, in all cases the coefficients can be regarded as descriptive measures that are more abstract than the data themselves. The descriptive accuracy of these functions can be estimated using measures of even greater abstractness (e.g., goodness-of-fit measures) (Barnes 1988). Algebraic functions might be used in this way to describe an ecocyborg's comportment, for instance. As with any descriptive method, however, this approach should be used with caution because it can result in the inadvertent filtering out of much information. Moreover, many kinds of phenomena, such as cyclicity, cannot be adequately represented in this way.

Periodicity (i.e., regular cyclicity) is a very useful concept for the description of some kinds of change, even if the latter is not, in fact, strictly periodic. Aspects of many natural systems fit this description and it is generally recognized that, as for the algebraic functional approach described above, only certain information about such systems can be captured with this method and, even then, only to some limited accuracy. Fourier transformation (often implemented as the Fast-Fourier Transform on digital computers and applied to evenly spaced data) is a method commonly used to characterize data series in this way (Churchill 1969). A time series can, for example, be described with this method to any desired accuracy with a set of sinusoidal terms. In practice, a limited number of coefficients are used that correspond to the main frequencies at which the series is cyclical. The comportment of many natural biosystems has inherent cycles such as circulatory and circadian rhythms, for example, and the population dynamics of some kinds of simple ecosystems can be roughly cyclical (Shimada and Tuda 1996; Leven et al. 1987). As before, the coefficients of a Fourier transform can be regarded as measures descriptive of the data series that are more abstract than the data themselves, and other related measures, like goodness of fit descriptors, are even more abstract. Measures of even greater abstractness are also commonly formulated such as, for instance, regression coefficients of algebraic functions fitted to time series of the Fourier coefficients.

Whereas biosystems often generate signals that may closely approximate periodicity with respect to some key frequency or frequencies, practically all of these are revealed to be, in fact, aperiodic when observed at different temporal resolutions of observation. For instance, over a sufficiently long time span or at finer resolution of observation, the key frequencies may be found to shift slowly, to be cyclical themselves, etc. In such circumstances methods for the analysis of aperiodic signals are useful. Wavelet transformation, for instance, is a popular method based on functions with *compact support*. These, unlike the sinusoidal functions of the Fourier transform, are bounded in time (Graps 1995; Strang 1994). Many families of such basis functions exist, including, for example, Haar wavelets (which are square waves), Mexican Hat wavelets (second-order derivatives of the Gaussian function that are reminiscent of truncated sinusoids), and Daubechies wavelets (these are complicated fractals). Each type of wavelet is best suited to specific purposes (Strang 1994). In biology, wavelet transformation has so far been used most extensively for the compression and analysis of electromagnetic signals like electrocardiograms (Unser and Aldroubi 1996). Its use in the description of larger-scale biosystems is, however, rapidly increasing, with some initial applications including soil mapping (McBratney 1998), characterization of forest canopy structure (Song et al. 1997), very large-scale radar mapping of tropical forests (Simard et al. 1997), and the prediction of river flow (Prochazka 1997), while many other potential applications are also being explored (e.g., Graps 1995). Generally, wavelet transforms are useful in the analysis of aperiodic signals because they capture frequency information that is localized with respect to time (or any other independent variable according to which the data are ordered). As well, in wavelet transformation, the resolution of observation with respect to both time and frequency is varied so as to localize more precisely the higher frequencies with respect to time and to better resolve lower frequencies with respect to frequency. This trade-off is convenient in the characterization of biological phenomena because, in these, low-frequency events usually persist for longer than do high-frequency ones. Wavelet transformation is well-suited to the study of biosystems because it facilitates signal compression and therefore enables the parsimonious description of the features of interest. As well, information is preserved about both aperiodic and periodic aspects of phenomena, over a range of scales. This is very appropriate when dealing with biosystems such as ecocyborgs because interplay between a variety of processes of different scales is a hallmark of all living things. In the end, as with the other approaches discussed, the wavelet coefficients can be considered as more abstract measures of the system's state or comportment than the direct measures used to generate the original data.

The characterization methods mentioned up to this point are suited to the analysis of one-dimensional series of data (i.e., ordered with respect to a single independent variable). However, the characterization of the state and comportment of large biosystems often involves data that are ordered in several spatial dimensions, as well as in time. In an ecocyborg, for example, features that can be described in this way include topology, rainfall, water table level, various areal densities (e.g., of nutrients, vegetation and animal biomass), and the location of mobile components (Parrott 1995). Geostatistics and related fields of mathematics offer a number of methods, many of which are quite similar to the regression techniques mentioned earlier, that are useful for the characterization of spatial patterns. As well, variations of both Fourier and wavelet transformation are available for the analysis of two-dimensional distributions. These methods are well-suited to computer-based implementation, and are often compatible with tools such as geographical information systems (GIS), which might be applicable to the characterization of both the state and comportment of a system.

The characterization of ecocyborgs and other biosystems often involves many more than two or three independent variables. In theory, regression-type modeling can be carried out for any number of independent variables, and Fourier and wavelet transformation techniques are well developed for up to three variables. In practice, however, none of these methods is very convenient in analyzing data of more than three dimensions, and so other methods must be used in such cases. One very general descriptive approach that does accommodate this is based on the use of *state space*. As described above, state variables are the independent, changeable attributes of a system and a state vector consists of a complete set of measurements corresponding to these, so that a particular vector completely and parsimoniously represents the instantaneous state of the system (N.B., comprising, however, only the features of interest as observed at the chosen resolution). Each state variable can be seen as corresponding to an axis, and the complete set of these frames an abstract topological space. Hence, given an ideal choice of state variables, the axes of the space are mutually orthogonal. In this space, a state vector corresponds to a single point and the comportment of the system is, therefore, depicted by the locus of this point with respect to time. This is a convenient manner of accommodating all possible data in a single construct that lends itself to the organized production of abstractions. The depiction of phenomena in this way can, in some cases, be a powerful visualization method. For instance, a judicious choice of variables and the filtering of corresponding data with, for example, slicing parameters, may result in the collapse of the most significant changes into three or fewer dimensions. This is most likely to be effective at higher degrees of abstractness. An entity like a termite nest might, for instance, exhibit cyclical comportment in a low-dimensional space if appropriately characterized with sufficiently abstract measures.

Although sometimes it may be possible to meaningfully represent significant system change in a state space of relatively few dimensions, this is not usually the case. The information must then be dealt with as a trajectory in a high-dimensional state space. Measures that are appropriate for this include *Lyapunov exponents*, which are used to quantify the average amplification rate of perturbations in trajectories or, equivalently, the growth of the error-to-signal ratio in predicting trajectories based on initial measurements (Peitgen et al. 1992). Measurements can be obtained either for a single trajectory in state space that attains similar values more than once (corresponding to the nearly repetitive comportment of a particular system) or for a number of trajectories that start at very similar initial conditions. In both cases the values reflect the rate at which neighboring trajectories (or segments thereof) diverge with respect to specific state variables. Thus, Lyapunov exponents are indicative of the stability of the observed comportment, which is really a reflection of *potential* comportment. In this way they are closely related to the eigenvalues of the transition matrix, previously mentioned with reference to structure (Logofet 1993). One exponent value can be estimated for each state variable and, as with the eigenvalues, stable comportment is indicated only if at least one of these values is negative. *Stable* does not imply here that the system is static, but that its state remains within a bounded subset of the overall space. For example, chaotic comportment, indicated by one or more negative Lyapunov exponents in conjunction with some that are positive, remains in such a bounded region (Ruelle 1989, pp. 54-56). Obviously, if a high-dimensional state space is used in the characterization of a system, a large number of exponents will result (and a large number of state vectors will be required to arrive at an accurate estimate of their values). The observer may, however, be able to construct a state space of lower dimensionality by choosing, as with the other approaches discussed previously, more abstract measures of the features of interest at the desired resolutions. Lyapunov exponents will be revisited later on in the article.

Thus, state space can serve as a setting for the analysis of historical data as well as for the study of what kinds of comportment a system might potentially display. One way to achieve the latter is with the vector field approach. In a deterministic system, every state leads to a particular subsequent state (in the absence of external forcing functions) and these tendencies can be described with a surface, or *vector field*, in the state space. A vector field is, in fact, equivalent to a description of the structure of the system in that it is descriptive of all potential change in the system (Green 1993). Hence, if sufficient historical data are available they can be used to map the vector field over the entire state space of a system, thereby allowing insight into the rules that govern its comportment. If not historically available, such data might be generated on demand, either by means of experimentation on the system itself, or by simulation with a model thereof. A highly conscious ecocyborg with a good self-model could, for instance, generate such a surface and simulate the comportment that would follow from each of many hypothetical initial conditions. Ideally, either historical or simulated data would be available for all the state variables but, even if this were not the case, the vector field might still be inferred given a sufficiently long, but incomplete, series of data. With a process called *delay-coordinate reconstruction*, for instance, approximate state vectors can be constructed from a number of lagged measurements of a few observed variables, and this may be sufficient to give an indication of the comportment in more dimensions than are directly observable (Muldoon et al. 1993). A trade-off of this method is that the time resolution of the resulting characterization will be reduced. In any case, the vector field itself can be used as a description of the system, or it may serve as a basis for even more abstract measures based on its various features.

Examples of some features that a vector field may possess are *attractors*. These are topological forms in state space that correspond to persistent dynamical modes, toward which the comportment of a given system might tend. Each attractor is surrounded by a region in state space called a *basin of attraction*, which represents the range of initial conditions that will lead to this comportment (Ruelle 1989). There are a number of different types of attractors, including point, periodic, chaotic (strange), and complex. The ordered states maintained by an autopoietic system correspond to a subset of the latter type (Clark and Kok 1999a, presented here in Chapter 3). Attractors are useful in a number of ways. From a scientific perspective they can be used to characterize the comportment of a system in a very abstract manner; from an engineering viewpoint, they might be used in the design specification of the kind of comportment that is desired in a system; and from a control perspective, they allow for strategic planning. With regard to the latter, for instance, a significantly autonomous ecocyborg might reason about itself in terms of attractors in the vector field of its ecosystem, perhaps avoiding certain basins of attraction in its state space so as to not be drawn into a particular mode of behavior. Later in this article there will be some discussion of even more abstract measures that can be formulated to described the properties attractors. For example, measures can be devised to quantify their dimensionality, a feature that is closely related to the concept of *order*.

As mentioned, it is most probable that very many measures of low or moderate abstractness would be required for the characterization of even a relatively simple biosystem or model thereof, such as the prototype ecosystem models developed in the EcoCyborg Project State. Together, these measures correspond to a state space of high dimensionality. A more manageable state space, in which features of the overall system comportment would be evident, must therefore be based on relatively abstract measures. Autopoiesis is one such large-scale feature that is of particular relevance in the study of ecocyborgs and other biosystems. Although no measures currently exist for the explicit evaluation of autopoiesis, its presence can be induced with measures of related features. The following sections are devoted to families of such measures.

7.4.4 Measures of order and disorder

Both measures of order and disorder can be employed in the characterization of the constitution, state, and comportment of biosystems. *Order* is the degree of correlation between comparable features of a system and, conversely, *disorder* is the lack of correlation, or the degree of difference, between these. When the order or disorder of a biosystem are measured, particular features of interest are compared. When the measurement pertains to the system's state, for instance, the features compared are the attributes of the components at a given instant (subject to the constraints discussed above, e.g., the chosen temporal resolution); order is the degree of correlation between the values of these, whereas disorder is the degree of difference between them. Similar measures of order can be defined for constitution and comportment in terms of correlation; for constitution, measures would be based on the number and kinds of components or their interrelationships and, for comportment, on temporal changes in the attribute values. Measures of order and disorder can, in some cases, be constructed as complements, but this need not always be so. This is to say that, although absolute disorder implies an absence of order (and vice versa), for intermediate ranges, measurements of order and disorder need not always sum to a constant value. As for the measures discussed previously, the values of order and disorder that are obtained are also affected by the approach used for their measurement. They will depend, for instance, on the choice of features and on the resolutions (both spatial and temporal) used in observation, since the latter affects the number of meaningfully different values of the features that are resolvable.

Order and disorder are frequently quantified in terms of *information*, which is usually measured in *bits*. It can be defined as the number of distinctions required for as complete a description as possible of a particular phenomenon at a given resolution. (For instance, if a component attribute might attain one of four observable values, then two distinctions, i.e., two bits of information, are required to completely determine which value it has actually attained.) Thus the chosen features of a particular system, observed at a given resolution, will be maximally uncorrelated when there is as much variation among them as possible; in this case the disorder of the system will be at its theoretical maximum and the maximum amount of information will be needed to describe it. If, however, the same system is observed at a finer resolution, more possible values would be distinguishable for the features of interest, and a description of the maximally disordered situation would constitute more information. Evidently, in any other situation the disorder of the system will be less than this, and a complete description will constitute less information. If, for a particular system, there were such a theoretical maximum value for disorder, a measure of disorder might be defined as the difference between the actual and theoretically maximal values of order, so that the two measures would be complements of one another.

Any number of different measures can be defined to quantify order and disorder, one of them being thermodynamic entropy. This is another example of a measure (or number of equivalent measures) that is associated with a reified but highly abstract concept. It was originally defined by Boltzmann in a statistical way to link the molecular theory of matter and the concept of unavailable thermodynamic work (Broda 1983, p. 81). It quantifies the degree of uniformity of the distribution of attribute values of the molecules in a closed, ideal gas system. Given a system that has particular global properties (i.e., is in a particular *macrostate*), an entropy measurement reflects the uniformity of the probability distribution associated with the range of possible overall molecular configurations (*microstates*) corresponding to that macrostate. The broader that this probability distribution (*ensemble*) is, the higher the entropy of the corresponding macrostate. Thus the probability that a system will be in a particular microstate (exactly which one being unknown) increases as the entropy of the macrostate decreases. Conversely, the probability that the system will actually be in a particular microstate decreases as the entropy of the macrostate increases. The maximum entropy value occurs when the probability distribution is completely uniform. In this case, because a theoretical maximum value exists for the measure, a complementary measure of order can be defined as the difference between the actual value and the theoretical maximum. Although thermodynamic entropy was originally defined in this sense for a closed, ideal gas system at equilibrium, this particular measure also has broader applications.

Measures of order and disorder are relevant to the characterization of biosystems because these are *dissipative*. This is to say that there is a tendency for such systems to become increasingly disordered with time. Thus, the values of a temporal series of measurements of the disorder of the features of interest (e.g., aspects of state and comportment, or of composition and structure, if these are defined loosely) will have a tendency to increase, in accordance with the Second Law of Thermodynamics. As mentioned, such an increase is physically manifested as a disordering or degradation of energy or of material components. If a living system is to persist, it must somehow counteract this tendency toward disorder and it generally does this by exporting its disorder to the surroundings, so that it can be considered to "produce" entropy. Since this occurs in time, entropy production can be considered as an aspect of the comportment of a biosystem. This idea was popularized by Schrödinger (1955) who expressed it in terms of a complementary quantity called *negative entropy*. Although it is not easy to obtain measurements of these thermodynamic quantities for large, distributed systems, empirical estimators of entropy have since been used in the life sciences. For instance, the higher the temperature of radiant energy the lower its entropy. Schneider and Kay (1994) used this relationship to study entropy produced by large-scale biosystems. They measured the difference between the temperatures of the radiant energy absorbed and rejected by different vegetative communities. Larger differences in temperature (and therefore in entropy) were found for more complicated communities (such as mature forests) than for less complicated communities (such as parking lots). This was interpreted as an indication that the former generated more entropy than did the latter, and this entropy was exported via the radiant energy. The characterization of biosystems in this way might be further extended by formulating more abstract measures based on the characterization of a time series of entropy measurements (Aoki 1995).

In an even broader sense, other measures of disorder (and order) similar to thermodynamic entropy (and negative entropy) can be defined in the context of numerous kinds of systems, with respect to a variety of different features, and used at any preferred resolution of observation. For instance, the *Shannon information measure* (Equation 7.2) was one of the first to be formulated in this way (Shannon 1948). It is related to the ensemble corresponding to the set of messages that might be sent during a particular act of communication and the probabilities associated with them being received. Thus each of the possible messages corresponds to a microstate and the Shannon information value is zero if it is certain that one particular message will be received, whereas it is maximized if there is a uniform probability of any of the possible messages being received.

$$H = -\sum_{i=1}^{n} p_i \cdot \log_2(p_i) \qquad \text{(7.2)}$$

H	Shannon information (bits)
n	total number of classes
i	index number of class
p	relative frequency of class

The particular information measure devised by Shannon has since been described in the context of ecology as being only one member of a family of *diversity measures* (Equation 7.3) (Baczkowski et al. 1997). An ecocyborg might utilize this kind of measure to characterize any of the numerous features of its ecosystem segment including, for example, aspects of state such as the heterogeneity of species populations.

$$H(\alpha,\beta)=\sum_{i=1}^{n}p_i^{\alpha}(-\log_2 p_i)^{\beta} \quad \text{...(7.3)}$$

H	diversity
α	parameter
β	parameter
i	index number of particular species
n	total number of species; index number of final species in list
p	relative species population

The Shannon information measure has also been defined as one of an even much wider spectrum of *Rényi information measures* (Equation 7.4) (Peitgen et al. 1992). When the definition of such a measure is recast in terms of the state space archetypal concept, an ensemble corresponds to a set of volume elements of state space and the associated probabilities with which the system state might fall within each of those elements. These probabilities are referred to as the *natural measures* of the volume elements.

$$H_q=\frac{1}{1-q}\log_2\sum_{i=1}^{n(\varepsilon)}P_i^q \quad \text{...(7.4)}$$

H	information (bits)
q	order of the information measure
i	index number of volume element
n	total number of volume elements in ensemble
ε	size of volume element
P	natural measure of a particular volume element

Rényi information measures also serve as the basis for more abstract measures of disorder called *Rényi dimension measures*. Dimension measures can be used to determine, for instance, how the comportment of a biosystem is constrained, providing an estimate of the number of degrees of freedom in which the dynamics of the system develop. This might reveal that an adequate description of the dynamics of a very complicated system might, in fact, require the use of only a relatively limited number of variables. A spectrum of *Rényi dimension measures* has been developed based on the ratio between the size of a volume element and its natural measure (Equation 7.5)

(Peitgen et al. 1992). Algorithms have been described for estimating the values of many of these measures from data. Some of the more commonly employed Rényi dimension measures include D_0 (Kruger 1996; Peitgen et al. 1992), D_1 (Peitgen et al. 1992), and D_2 (Ding et al. 1993), which are respectively named the *box-counting*, *information*, and *correlation dimensions*.

$$D_q = \lim_{\varepsilon \to 0} \frac{1}{q-1} \cdot \frac{\log_2 \sum_{i=1}^{n(\varepsilon)} P_i^q}{\log_2 \varepsilon} \qquad (7.5)$$

D	dimension
q	order of the dimension measure
ε	size of a volume element
i	index number of a specific volume element
n	total number of volume elements
P	natural measure of a specific volume element

Because the Rényi dimension measures are difficult to calculate for high-dimensional comportment such as that of most biosystems, various empirical estimators of dimension can also be used. One such estimator is based on the *Hausdorff dimension*, which in fact inspired Mandelbröt (1982) to generalize formal measures of dimension. This dimension measure is based on the rate with which the value of a measurement changes as it is evaluated at increasingly finer resolution (Peitgen et al. 1992; Mandelbröt 1982). The relation between the magnitude of the values and the resolution can be described in the form of a power law, the exponent of which is the Hausdorff dimension (Equation 7.6). (The formal definition of the Hausdorff dimension is somewhat more involved than this, and is given in Mandelbröt 1982). The Hausdorff dimension does not generally have an integer value except for platonic objects such as straight lines and planes and an object with a fractional Hausdorff dimension (which includes most natural objects) is called a *fractal* (Mandelbröt 1982).

$$y \propto x^D \qquad (7.6)$$

y	length
x	scale (related to the resolution of the measure) (m)
D	Hausdorff dimension

Another useful method of estimating the dimension value is to use a method based on the Lyapunov exponent. This estimator, called the *Lyapunov dimension*, can be found by arranging the Lyapunov exponents of a system in decreasing order of their signed magnitudes, from the largest positive value to the largest negative one, and assigning index numbers to the values according to that that order. A graphical method is to plot cumulative sum of the values with respect to their resulting index numbers. A convex curve results that will cross the ordinate axis only once (Figure 7.1), and the Kaplan-Yorke conjecture states that the abscissa of this point of intersection is approximately equal to the information dimension, D_I (Peitgen et al. 1992). A mathematical formula for calculating this value is given in Equation 7.7.

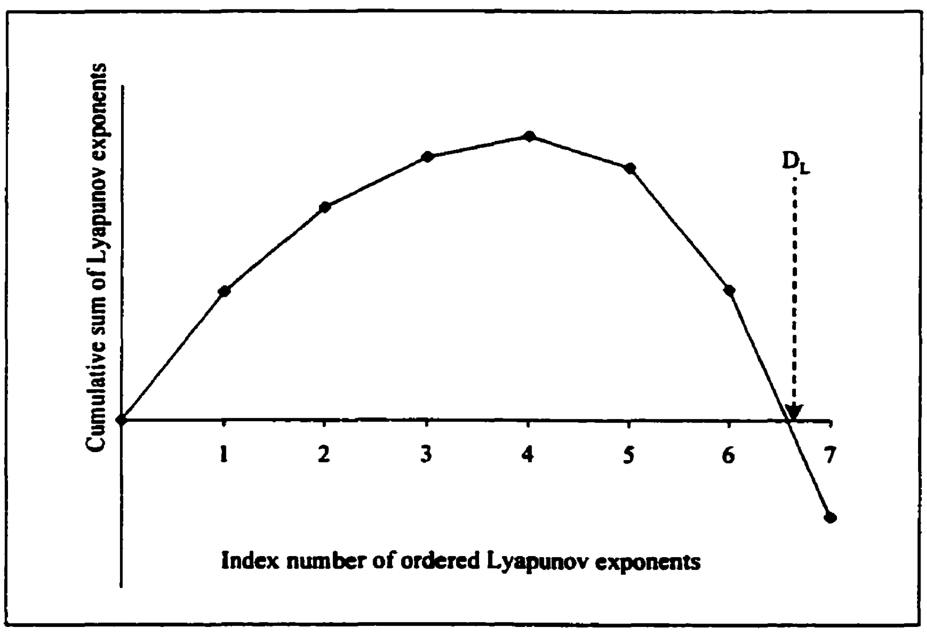

Figure 7.1 The curve formed by the cumulative magnitudes of the ordered Lyapunov exponents of a system. The abscissa of the point of intersection (D_L) equals the magnitude of the Lyapunov dimension.

$$D_L = n + \frac{1}{|\lambda_{n+1}|}\sum_{i=1}^{n}\lambda_i \quad \text{(7.7)}$$

D_L	Lyapunov dimension
i	index number of Lyapunov exponents, ordered by magnitude
n	index number where $\lambda_1 + \lambda_2 + \lambda_3 + ... + \lambda_n \geq 0$

Methods for estimating dimensionality must be used with caution. Often they involve statistical operations, such as least-squares regression, that were designed based on the assumption that the underlying process could be described with a linear or logarithmic function, with data affected by normally distributed fluctuations and measurement errors. These assumptions are not necessarily valid for the kind of characterization methods mentioned here. Even if the use of such descriptions is justified, they are usually valid only within a certain range of conditions (Peitgen et al. 1992). A more general drawback to all of the measures of order and disorder discussed in this section, such as entropy, and the families of information and dimension measures, is that they cannot be used to distinguish which part of variability is due to pattern and which is due, for instance, to external perturbations or error (Clark and Kok 1999a, presented here in Chapter 3). Pattern is equivalent to the relationship between features of the system, which might involve either change or invariance of these under different kinds of transformations (e.g., rotation, translation, scaling). The characterization of variation (or correlation) in this way is significant in the study of biosystems because autopoietic comportment, and the structure that gives rise to it, are patterns typical of all living things. The use of complexity measures is one approach to their characterization.

7.4.5 Measures of complexity

Complexity is, in its broadest sense, the difficulty of performing a given task (Li 1997). Here, complexity is more narrowly defined as the difficulty encountered by a particular observer in characterizing the patterns that exist in a system (or that are common to a class of systems). Measures of complexity can be used to distinguish, therefore, between any variations in a system that can be assimilated by the observer as instances of pattern, and those that cannot. Evidently, the natures of both the observed system and the observer influence the results of such a measurement. Any pattern that actually exists in an observed system will be characterizable only if the observer has an adequate conceptual network. On one hand, an observer with a conceptual network that is inappropriate for the characterization of that particular system would identify very little pattern among its features. On the other hand, an observer with a suitable conceptual network might perceive a great deal of pattern. This pattern may identify the system as a member of a particular class, such as the class of all biosystems or of all ecocyborgs. The remainder of the variability in the system would then be due to the (possibly random) features that uniquely distinguish it.

Given a conceptual network that enables the generation of very complete and accurate descriptions of the observed system, the observer has a maximal ability to predict the features of interest, with respect to some independent variable such as time or distance. Such an optimal description of the pattern present in a system will comprise a quantity of information that is characteristic of that system. This is because measures of complexity bear a relationship to information similar to that which measures of order and disorder share with this quantity. Hence, an optimal measurement of the variation associated with the pattern in a system must constitute a certain minimal number of distinctions. An optimal description does not, however, imply *complete* predictability, since some patterns are inherently unpredictable over the long term (or over large scales) and because some features are unpatterned (random).

Any number of measures could be devised to generate such measurements, based, for instance, on the composition, structure, state, or comportment of the observed system. According to Gunther et al. (1994), such measures should yield values that: *i*) are zero for a strictly ordered system, positive for intermediate values of order, and zero for total disorder; *ii*) do not increase for both of two independent systems as the consequence of the direct interaction of those systems; *iii*) do not increase with the simple enlargement of a system; and *iv*) give values dependent on the particular measure used to describe the system. The first of these recommendations is based on the assertion that entities which are dominated neither by ordering nor disordering influences are likely to more complex than those which are. For example, a system that is undergoing a phase transition, where one phase (e.g., solid) is a more ordered and the other is (e.g., gas) is less so, often displays sophisticated patterns (Langton 1990). This relationship is illustrated in Figure 7.2. The second recommendation is that complexity measures should reflect a principle of conservation, and the third is that they should be independent of the size of a system. The fourth recommendation underscores the importance of defining a measure that is appropriate to both the observed system and the observer, since different complexity measures are appropriate for different situations (Silvert 1995; Crutchfield 1994a, 1994b).

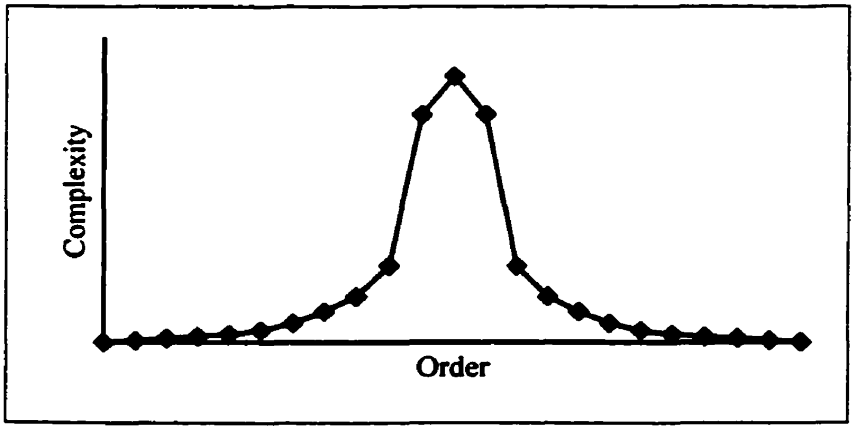

Figure 7.2 General relationship between measurements of order and complexity (in arbitrary units).

An empirical method of estimating complexity is based on the squared magnitude of the Fourier transform, the *power spectrum*. This indicates the way the power of a signal (represented as series of data) is distributed with respect to frequency, and is useful for distinguishing between different classes of comportment (Schroeder 1991). Strictly periodic comportment has a power spectrum in which all of the power is confined to one or a few sharply defined frequencies. Highly chaotic comportment, on the other hand, is even and continuous. In both scenarios, the power spectrum is relatively homogeneous. Highly complex comportment, however, falls between these two regimes. Its power spectrum is often very heterogeneous and can sometimes be approximated with a *power law distribution* (Equation 7.8) (Schroeder 1991). Bak and Chen (1991) consider the latter phenomenon to be evidence of a highly patterned state called *self-organized criticality*. It is speculated that the exponent of the power law distribution formula, called the *Mandelbröt-Weierstrass fractal exponent*, can be used in some instances as an estimator of the fractal dimension of the underlying system, although it must be used with caution in this capacity (Penn and Loew 1997). One drawback to the use of power series and the underlying Fourier transform for estimating complexity is that the basis functions are sinusoidal, and therefore do not have compact support, i.e., they are not localized with respect to the independent variable. One possible solution is to use a windowed Fourier transformation in which segments of the data series are analyzed independently, or a wavelet transformation. The spectra of the segments can then be compared to see if there are changes in the power distribution.

$$E(f) = k \cdot f^{\beta} \qquad \text{(7.8)}$$

E	squared value of Fourier transform
f	frequency
k	scaling constant
β	Mandelbröt-Weierstrass fractal exponent

A more formal complexity measure is called *effective complexity* (Gell-Mann and Lloyd 1996). This measure is the length of an optimal description of the identifiable regularities that define a system as belonging to a particular class. It is complemented by an information measure that quantifies the difficulty of describing the unique features of the system. The sum of the two measurements for a particular system is its *total information*, roughly analogous to the difficulty of describing its every detail. In order to be useful, this method must standardize the partitioning between the two measures. For instance, the two required measures might be calculated as *algorithmic information*, which is the length of the most parsimonious set of instructions that will cause a given universal Turing machine to generate a particular string of symbols (Chaitin 1977). In this idealized context, the information required to produce a particular output (the overall system description) is thus partitioned into two bodies. The first body of information corresponds to a *specialized* Turing machine (algorithm) which, regardless of input, will only generate a description of members of a general class of systems (such as all ecocyborgs). The second set of information is a set of input instructions for that Turing machine which causes it to generate a description of a certain unique system in the class (such as a particular ecocyborg). Gell-mann and Lloyd (1996) suggest that the combined length of the specialized Turing machine and the input instructions should approximately equal the length of a different set of instructions, this time for a *universal* Turing machine, which would cause that machine to generate the desired description (i.e. that of the specific ecocyborg). This measure can be made relatively objective by optimizing the first set of instructions, i.e. those intended for the specialized Turing machine. For this purpose, each particular system in the class might, for example, be identified by evaluating the relative likelihood of its occurrence, and then assigning the shortest codes to identifiers to systems that are most likely to occur (Anand and Orlóci 1996).

Anand and Orlóci (1996) have formulated a measure similar to effective complexity for use in characterizing plant communities. In evaluating this measure, the plant species list of a community *S* is first assigned a parsimonious coding. Next, the species are ranked in order of abundance and numbered in binary notation, so that the shortest identifying codes are assigned to those that occur most frequently. The Shannon information *H*(*S*) of the community is then calculated with Equation 7.2 (n = total number of species; p_i = relative frequency of species i). Next, the average code length *L*(*S*) is found for the species list (Equation 7.9).

$$L(S) = \sum_{i=1}^{n} (p_i - l_i) \qquad \textbf{(7.9)}$$

L	average code length (bits)
S	plant species list
n	number of species
p_i	relative frequency of species i
l_i	length of code for species i (bits)

H(*S*) and *L*(*S*) estimate total information and information describing a particular community configuration, respectively. The difference Δ(*S*) between these two measures is the information common to all configuration of the community, and this can be considered as the complexity of the community (Equation 7.10). This method can be adapted for estimating the relative complexity of any system, such as a computational model of an ecocyborg, that comprises a number of different kinds of components that occur with varying likelihood, i.e., for which a number of different configurations are possible.

$$\Delta(S) = L(S) - H(S) \qquad (7.10)$$

Δ	structural complexity
S	species list
L	average code length
H	entropy

7.4.6 Measures of emergence

Measures of *emergence* constitute additional methods for the abstract characterization of biosystems, including ecocyborgs. These quantify the influence that small-scale interactions between components have on large-scale system features, and thus reveal the effect of multiscale structural patterns on an entity. Autopoiesis is one example of a system-scale phenomenon that cannot be understood without consideration of the overall network of interrelationships among the components (Capra 1996). The results of complexity measures, as with the measures discussed previously, depend on both the objective properties of the observed system and on the mental capacities of the observer. Accordingly, they reflect the difference between an optimal description of the actual features of a system and a description generated by an observer with limited characterization abilities. A large value indicates a substantial difference between the complexity of the system in these two cases. There exists, therefore, relationship between measures of emergence and measures of complexity. This is not always a straightforward relationship, however; simple interactions at the local scale do sometimes belie great complication at the system scale, but the opposite case is also possible, whereby complication at the local scale underlies simplicity at the system scale.

As with order, disorder, and complexity, any number of measures can be devised to quantify emergence. Ideally, these would be based on a description of the composition, structure, state or comportment of a set of components in isolation from one another, and a corresponding description of an equivalent set that are engaged in mutual interaction as a system. Any differences between these two descriptions would necessarily be the result of the multiscale features of the system and could be quantified, for instance, in terms of information. Such a measure is quite abstract, since its value is calculated from more direct descriptions. Because emergence quantifies the interaction between multiscale features, any direct measures on which it is based must be meaningful at all of the various scales that are examined. The emergence of the comportment of a flock of birds, for example, could not be based on population, although this is a useful measure at the scale of the flock as a whole, it is meaningless with reference to a single bird. A more suitable direct measure is position, which is applicable at both the system and component scales. A value for emergence could, for example, be calculated from position data as follows. First, the path of each a set of isolated birds, perhaps a number of seagulls flying at different times down the same stretch of beach, could be recorded as a time series of position data. From this, a correlation value could be found for the path of each bird as compared with the average path of all of them, and a mean correlation coefficient calculated for the whole set of birds. Second, a similar method could be followed for a similar set of birds flying together as a flock. Third, the ratio of the two mean correlation coefficients would be a measurement of the emergence of the path of the flock. Thus if an observer in possession of a poor conceptual representation of the way the birds interact was to try and predict their paths as they flew together, the resulting description would likely be very different from an optimal one of the actual paths of the flock.

7.5 Summary and discussion

In this article the process of characterizing a biosystem is illustrated. The opening section of the article is an exploration of characterization as an epistemological process. There is a review of how characterization is founded on perception, or the interaction between two systems in which the state of the perceiving system is dependent on the state of that which is perceived. The dependence of characterization on the observer is emphasized: the observer acts as a perceptor; defines systems by discriminating between them and their surroundings; assimilates them by formulating mental representations that are integrated into a conceptual network; and describes them by codifying these conceptual representations using a language. Appropriate characterization therefore depends on the abilities of the observer, the paradigm on which the observer's conceptual network is based, and the archetypal concepts and measures that can be accommodated by that paradigm. The choice of characterization methods also depends, of course, on the intentions and preferences of the observer as well as on the systems that are being characterized.

The EcoCyborg Project is used as the context for illustrating the process of characterization. This project deals with a specific class of observed systems, comprising the somewhat autonomous, computational models of ecocyborgs used in the project, as well as the very inclusive set of all biosystems. The observers in this context include the ecocyborgs themselves, who must characterize themselves in order that they might be significantly autonomous, as well the scientists and engineers who are studying the ecocyborgs as examples of biosystems in general. A conceptual paradigm is described and examples of some archetypal concepts and associated measures are mentioned. Very direct, system-specific, measures are described for the characterization of systems like computational models and ecocyborgs, and more abstract measures for their characterization as members of the class of biosystems. Subsequent reporting about the EcoCyborg Project will employ some of these measures.

The abstract measures that are emphasized here quantify order, complexity, and emergence. It is interesting to speculate about the utility of these kinds of measures as they apply to ecocyborgs. Many different measures of order might be formulated, for instance, but what aspects of ecological or mental constitution should be reflected in such measures remains an open question. It would not be meaningful to characterize a system as being of intermediate order with respect to some set of arbitrarily chosen features. To be of use in science or engineering, measures of order should be based on features that are of fundamental importance in the study or creation of a particular class of systems. Understanding of these concepts is advancing; for instance, it has been proposed that systems which are of intermediate order tend to be of relatively high complexity. This relationship appears to be significant; systems that can retain essential pattern over the long term and at the same time adapt to environmental perturbations through short-term change are also highly complex, as quantified with measures such as those described here.

The study of ecocyborgs involves further issues relating to complexity. A highly autonomous ecocyborg is defined as one that is capable of guiding its own internal state toward self-determined goals in the face of environment variability. In order to do so it must perceive, discriminate, and create internal representations of perturbations, approximate their inverse functions, and use these to formulate and implement contravening responses. An ecocyborg might prepare contingency plans to be implemented should predicted perturbations occur, but if the perturbations are unexpected then the entire process must be performed quickly enough to prevent disruption of the ecocyborg's autopoietic potential. Finally, a highly intelligent ecocyborg should be capable of learning from experience so that it can increase its effectiveness. It has been proposed elsewhere that any effective controller must be, or must contain, a description of the system that it controls (Conant and Ashby 1970). An ecocyborg engaged in controlling itself must therefore generate models to represent both itself as well as potentially disruptive perturbations. It is interesting to contemplate how this requirement might impact the necessary complexity of the mind of the ecocyborg, and what implications the implicit recursion of self-awareness might have in this regard.

Finally, both the ecosystem segment and the control system of an ecocyborg can be considered as emergent entities. An essential structural feature of both is the mutual interaction of their many components. In the case of the ecosystem segment, the interaction between biotic and abiotic components results in overall comportment that is more or less homeostatic for autopoiesis. Ideally, this is reflected in the persistence of the ecosystem segment with respect to a given set of characteristics relating to biological diversity and activity. Similarly, the control system of the ecocyborg comprises many interacting computational components. Overall, it should constitute a responsive and adaptable control system that helps to guide the ecocyborg toward a desired goal state. The state and comportment of both these segments of the ecocyborg differ significantly from what one might expect of their individual components in isolation.

The ultimate goal of the EcoCyborg Project is to learn how to construct highly autonomous biosystems. As with any learning endeavor this project requires an appropriate conceptual paradigm, archetypal concepts, and measures for use in characterization. The utility of the methods that are chosen can be better assured if they are considered in context of the overall characterization process as described in this article. The paradigm outlined here seems to be appropriate to the project, and the archetypal concepts and measures suggested might well be of great utility. Appropriately constructed measures of order, complexity, and emergence, for instance, can be used to characterize certain states that are indicative of, or conducive to, greater autopoiesis and autonomy in biosystems. This article is not, however, an attempt to identify a complete set of measures for such an exercise. Appropriate measures will doubtlessly change along with the development of the biosystems employed (computational or otherwise) and the understanding and preferences of the investigators. The principal challenge of the EcoCyborg Project can in fact be described as the formulation of useful measures of autopoiesis and autonomy, and the development of an understanding of how such measures relate to one another and to the established scientific paradigm.

7.6 References

Anand, M., and L. Orlóci. 1996. Complexity in plant communities: the notion and quantification. *Journal of Theoretical Biology* 179:179-186.

Aoki, I. 1995. Entropy production in living systems: from organisms to ecosystems. *Thermochimica Acta* 250:359-370.

Baczkowski, A.J., D.N. Joanes and G.M. Shamia. 1997. Properties of a generalized diversity index. *Journal of Theoretical Biology* 188:207-213.

Bak, P., and K. Chen. 1991. Self-organized criticality. *Scientific American* January:46-53.

Barnes. 1988. *Statistical Analysis for Engineers and Scientists*. New York, NY: McGraw-Hill.

Broda, E. 1983. Ludwig Boltzmann: Man, Physicist, Philosopher. Woodbridge, CT: Ox Bow Press.

Capra, F. 1996. The Web of Life. New York, NY: Doubleday.

Casti, J.L. 1994. *Complexification: explaining a paradoxical world through the science of surprise*. New York, NY: HarperCollins.

Chaitin, G.J. 1977. Algorithmic information theory. *IBM Journal of Research and Development* 21:350-359.

Churchill, R.V. 1969. *Fourier series and boundary value problems*, 2nd ed. New York, NY: McGraw-Hill.

Clark, O.G., and R. Kok. 1999a. Characterizing biosystems as autopoietic entities. *Oikos*, (submitted for publication).

Clark, O.G., and R. Kok. 1999b. Mental architecture of a cyborged biosystem. *Engineering Applications of Artificial Intelligence*, (submitted for publication).

Clark, O.G., and R. Kok. 1998. Engineering of highly autonomous biosystems: review of the relevant literature. *International Journal of Intelligent Systems* 13(8):749-783.

Clark, O.G., R. Kok and R. Lacroix. 1999. Mind and autonomy in engineered biosystems. *Engineering Applications of Artificial Intelligence* 12(3):389-399.

Conant, R.C., and W.R. Ashby. 1970. Every good regulator of a system must be a model of that system. *International Journal of Systems Science* 1(2):89-97.

Crutchfield, J.P. 1994a. The calculi of emergence: computation, dynamics, and induction. *Physica D* 75:11-54.

Crutchfield, J.P. 1994b. Observing complexity and the complexity of observation. In *Inside Versus Outside*, ed. H. Atmanspacher and G.J. Dalenoort, 235-272. Berlin, Germany: Springer-Verlag.

Ding, M., C. Grebogi, E. Ott, T. Sauer and J.A. Yorke. 1993. Estimating correlation dimension from a chaotic time series: when does plateau onset occur? *Physica D* 69:404-424.

Gell-Mann, M., and S. Lloyd. 1996. Information measures, effective complexity, and total information. *Complexity* 2:44-52.

Gould, P. 1980. Q-analysis, or a language of structure: an introduction for social scientists, geographers and planners. *International Journal of Man-Machine Studies* 13:169-199.

Graps, A. 1995. An introduction to wavelets. *IEEE Computational Science and Engineering* 2(2):50-61.

Green, D.G. 1993. Emergent behaviour in biological systems. In *Complex Systems: From Biology to Computation*, ed. D.G. Green and T. Bossomaier, 24-35. Amsterdam, the Netherlands: IOS Press.

Gunther, R., B. Schapiro and P. Wagner. 1994. Complex systems, complexity measures, grammars and model-inferring. *Chaos, Solitons & Fractals* 4(5):635-651.

Kruger, A. 1996. Implementation of a fast box-counting algorithm. *Computer Physics Communications* 98:224-234.

Langton, C.G. 1990. Computation at the edge of chaos: phase transitions and emergent computation. *Physica D* 42:12-37.

Leven, R.W., B.P. Koch and G.S. Markman. 1987. Periodic, quasiperiodic and chaotic motion in a forced predator-prey ecosystem. In *Dynamical Systems and Environmental Models*, ed. H.G. Bothe, W. Ebeling, A.B. Kurzhanski and M. Peschel, 95-104. Berlin, Germany: Akademie-Verlag.

Levins, R. 1975. The search for the macroscopic in ecosystems. In *New Directions in the Analysis of Ecological Systems*, vol. 2, ed. G.S. Innis, 213-222. La Jolla, CA: The Society for Computer Simulation.

Li, W. 1997. The complexity of DNA. *Complexity* 3(2):33-37.

Logofet, D.O. 1993. *Graphs and Matrices: Stability Problems in Ecosystem Models*. Boca Raton, FL: CRC Press.

Lyman, W.J. 1982. *Handbook of chemical property estimation methods: environmental behavior of organic compounds*. New York, NY: McGraw-Hill.

Mandelbröt, B.B. 1983. *The Fractal Geometry of Nature*. New York, NY: Freeman.

Margulis, L., and J.E. Lovelock. 1974. Biological modulation of the Earth's atmosphere. *Icarus* 21:471-489.

Maturana, H.R., and F.J. Varela. 1980. *Autopoiesis: the Organization of the Living*. Dordrecht, the Netherlands: D. Reidel Publishing Company.

McBratney, A.B. 1998. Some considerations on methods for spatially aggregating and disaggregating soil information. *Nutrient Cycling in Agroecosystems* 50:51-62.

Michie, D. 1998. Learning concepts from data. *Expert Systems With Applications* 15:193-204.

Molenaar, R. 1998. Design and implementation of biosystem control for biosystem simulation. Unpublished Ph.D. thesis. Department of Agricultural and Biosystems Engineering, McGill University, Montreal, QC.

Muldoon, M.R., R.S. MacKay, J.P. Huke and D.S. Broomhead. 1993. Topology from time series. *Physica D* 65:1-16.

Parker, S.S. (ed.). 1994. *McGraw-Hill Dictionary of Scientific and Technical Terms*, 5th ed. New York, NY: McGraw-Hill.

Parrott, L., O.G. Clark and R. Kok. 1999. The EcoCyborg Project. *Canadian Agricultural Engineering*, (submitted for publication).

Parrott, L. 1995. The EcoCyborg Project: a model of an artificial ecosystem. Unpublished M.Sc. thesis. Department of Agricultural Engineering, McGill University, Montreal, QC.

Penn, A.I., and M.H. Loew. 1997. Estimating fractal dimension with fractal interpolation function models. *IEEE Transactions on Medical Imaging* 16(6):930-937.

Peitgen, H.-O., H. Jurgens and D. Saupe. 1992. *Chaos and Fractals: New Frontiers of Science*. New York, NY: Springer-Verlag.

Polis, G.A., and D.R. Strong. 1996. Food web complexity and community dynamics. *The American Naturalist* 147:813-846.

Prigogine, I. 1980. *From Being to Becoming: Time and Complexity in the Physical Sciences*. San Francisco, CA: Freeman.

Prochazka, A. 1997. Neural networks and seasonal time-series prediction. IEE Conference Publication. No. 440, 36-41. Stevenage, UK: IEE.

Ray, T.S. 1994. Evolution, complexity, entropy and artificial reality. *Physica D* 75:239-263.

Ruelle, D. 1989. *Chaotic Evolution and Strange Attractors*. Cambridge, MA: Cambridge University Press.

Schneider, E.D., and J.J. Kay. 1994. Life as a manifestation of the second law of thermodynamics. *Mathematical and Computer Modelling* 19:25-48.

Schroeder, M.R. 1991. *Fractals, chaos, power laws: minutes from an infinite paradise*. New York, NY: Freeman.

Schrödinger, E. 1955. *What is Life?* Cambridge, MA: Cambridge University Press.

Shannon, C.E. 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27(July-October):379-423, 623-656.

Shimada, M., and M. Tuda. 1996. Delayed density dependence and oscillatory population dynamics in overlapping-generation systems of a seed beetle Callosobruchus chinensis: matrix population model. *Oecologia* 105:116-125.

Silvert, W. 1995. Complexity. *Journal of Biological Systems* 4:585-591.

Simard, M., G. De Grandi, S. Saatchi, M. Leysen and K.P.B. Thomson. 1997. Processing and analysis techniques for continental scale radar maps of the tropical forest. *Proceedings of the International Geoscience and Remote Sensing Symposium* 4:1890-1892. Piscataway, NJ: IEEE.

Song, B., J. Chen and M. Rudnicki. 1997. The relationship between canopy structure and the pattern and process of the understory. *Bulletin of the Ecological Society of America* 78(4 Supplement):189.

Specter, S., and G.J. Lancz (eds.). 1986. *Clinical Virology Manual*. New York, NY: Elsevier Science.

Strang, G. 1994. Wavelets. *American Scientist* 82:250-255.

Unser, M., and A. Aldroubi. 1996. A review of wavelets in biomedical applications. *Proceedings of the IEEE* 84:626-638.

CHAPTER 8. GENERAL SUMMARY AND CONCLUSIONS

8.1 Philosophy of engineering

A comprehensive explanation of the philosophy that underlies the EcoCyborg Project is developed throughout the thesis (focused principally in Chapter 5). Although this project is primarily an engineering one, it also has a scientific aspect. These two paradigms are fundamentally different from one another. Science, on the one hand, is oriented toward the observation and understanding of phenomena that already exist and, accordingly, is descriptive or explanatory in nature. Engineering, on the other hand, involves creative activities such as design, construction, etc., that are usually conducted with the intent of obtaining particular objectives. A kind of mutualism often exists, however, between science and engineering whereby scientific inquiry may be driven by attempts to develop theories explaining heuristic engineering knowledge, and engineering theory is based partially on knowledge generated through scientific methods. The scientific aspect of the EcoCyborg Project involves the development of tools for use in studying the relationship between the composition, structure, initial state, and comportment of a biosystem, and any forcing functions that might impinge upon it. The understanding that is obtained through the use of these tools contributes toward the long term goal of the EcoCyborg Project, which is the formulation of a general theory for the engineering of biosystems. Thus, it is hoped that the tools and theory will eventually be improved to the extent that they may be used in the explicit characterization and creation of biosystems having particular features (such as substantial autonomy).

The tools being developed in the EcoCyborg Project include computer-based models and simulations, as well as suitable descriptive methods (such as measures), which are oriented toward the characterization of biosystems. An important aspect of this work is modeling, which is the identification of some equivalence between a phenomenon of interest (the modeled phenomenon), and another that can be more conveniently studied and manipulated (the model). Much of scientific knowledge comprises explanatory or descriptive models, which are usually simpler in most aspects than the phenomena that are being modeled. As a result of their relative simplicity, however, they clearly represent that which is of interest. On the other hand, in engineering, models often represent the features of some system that is intended to exist in the future. Some engineering models are predictive, and can thus be implemented in design simulations to determine the hypothetical comportment of the proposed system. This kind of design work is similar to a thought experiment, or *gedanken experiment*, but one in which the powers of the human mind are complemented by external devices. Other engineering models are prescriptive, exactly describing the chosen composition, structure, state, and (perhaps) comportment of a system as it is intended to be. The tools developed in the EcoCyborg Project so far are useful in a scientific (explanatory or descriptive) role, but are not yet suitable for engineering design (predictive or prescriptive) purposes.

The modeled phenomenon and the model that are involved in an equivalence relationship can be of very different natures. The nature of a phenomenon, such as a particular system, for example, are describable with mutually orthogonal pairs of variables (e.g., *real/imaginary*, *physical/virtual*, *artificial/natural*, *living/non-living*, *guided/unguided*, etc.), which can be considered as corresponding to the axes of a hypercube of possible modes of existence. In a scientific study, the intent is usually to model a real, often physical, system by creating another, often imaginary, one that is an analogous representation of it. For instance, an ecologist might formulate a hypothetical model of the food web in a real ecosystem that he or she has observed. Engineering models, on the other hand, are often real, physical or virtual systems, such as scale models or computer models, that represent imaginary systems which are to be built. The EcoCyborg, for instance, is an imaginary, physical system that is being modeled with a real, virtual system of computer programs.

8.2 Biosystems

Under the umbrella of the engineering philosophy that is elucidated in this thesis, two main lexical themes are also developed. One is related to the concept of *biosystem* and the other to *autonomy*. The first is pertinent because the overall EcoCyborg Project, which is the context of this thesis, concerns biosystems engineering. A clear definition of *biosystem* and the surrounding concepts has been lacking to this point, making it awkward to discuss the engineering of biosystems in a concise manner, and leading to some difficulty in clearly establishing the identity of this discipline. Hence, a biosystem is defined (Chapter 3) as a unitary assemblage of entities that is alive to some degree as a whole. This leads in turn to the question of what is meant by *alive*, an issue that is resolved by adopting *autopoiesis* as the defining characteristic of living entities. Autopoiesis is a homeostatic mode of comportment, originally described by Maturana and Varela (1980), in which the overall system is continually regenerated as a result of the interactions of its own components. This interpretation accommodates large-scale, combined, technological and biological systems such as ecocyborgs, as well as a wide range of other kinds of systems. This definition also determines the bounds of biosystems engineering as the discipline which treats the design, construction, operation, maintenance, repair, and upgrading of any system that is autopoietic to some degree.

The definition of biosystem (and the discussion of autopoiesis) can be set in a systems-theoretic context, which is suitable for engineering applications. Autopoiesis, being a mode of comportment, is therefore representable as an ensemble of trajectories in a state space defined by continuous measures. Such a state space might, for example, be constructed on the basis of measures such as those of order and disorder, complexity, and emergence. If the state space is appropriately defined, then these trajectories will tend toward a corresponding type of attractor. Points on these trajectories correspond to viable states of the biosystem that is being characterized and, accordingly, the strength and robustness of the attractor is related to the overall vitality of the system. Thus, autopoiesis is itself a variable quality that might occur to a greater degree in one system than in another.

8.3 Autonomy

The second main lexical theme that is developed relates to *autonomy*, which is the ability to formulate and pursue ones own goals in the absence of external guidance. This lexicon is clarified and extended principally in Chapter 4. From a human perspective, autonomy is a desirable system characteristic in any circumstance where external guidance is rendered impractical or undesirable, and the system of interest must fulfill objectives (such as persistence) in surroundings that are somewhat unpredictable. In the EcoCyborg

Project there is, therefore, a major emphasis on the engineering of biosystems that are substantially autonomous.

Autonomy has its basis in *mind*, which is the sum total of an entity's *virtual machinery*, or constructs and artifices, such as algorithms, that enable the storage and manipulation of information. Mind is one of three complementary interpretations of a computational entity that might be adopted from a cybernetic perspective. In the second interpretation such an entity is seen as giving rise to an epistemic space of *mental abilities*, or potential computational activities, which can be grouped into sets (*faculties*). In the third interpretation, *mentation* comprises those computational activities that are actually performed, of which human thought is a special instance. Particular qualities of mentation arise if the entity's faculties are of the appropriate type and degree of sophistication (*intelligence*).

One quality of mentation that is of particular interest is *consciousness*, defined as the use of a model of self in reasoning. To be substantially conscious, an entity must possess sophisticated faculties of *perception*, *memory*, *reason*, *expression*, and *learning*. The degree of consciousness determines, in turn, to what extent an entity can be autonomous. The three principal aspects of autonomy are: (1) *automation*, which is the capacity to persist in a particular mode of comportment without external guidance, (2) *volition*, which is the ability to formulate one's own goals and strategies for attaining them, and (3) *intentfulness*, which is the active implementation of those strategies. These qualities, too, are variable, so that they might be present in an entity to any degree. The clear definition of this lexicon facilitates the coherent discussion of the engineering of biosystems so that they may be substantially autonomous. The explicit application of these concepts in the context of the engineering of biosystems, and especially of large-scale ones such as ecocyborgs, is novel.

8.4 Cyborging as an approach to engineering biosystems

The biosystems currently under investigation in the EcoCyborg Project are *ecocyborgs*, hybrid systems that result from the combination (cyborging) of ecosystems with technological control components. These are large-scale entities comprising, or equivalent to, a community of biological organisms and their abiotic surroundings (the ecosystem part), together with the added components (the control part). This engineering approach can be applied not only to artificial, but also to natural ecosystems in order to make them correspond to specific design objectives.

One design objective that can be addressed through cyborging, and that is of particular interest in the EcoCyborg Project, is substantial system autonomy. Technological components that are installed in a biosystem with the intent of fulfilling this objective must host virtual machinery that is sufficient to endow the resulting cyborg with the appropriate type of mind. This is especially true in the case of ecosystems, since they are unlikely to possess any of the required virtual machinery. This virtual machinery must give rise to mental faculties, mentioned above, of sufficient sophistication so that the ecocyborg will be able to reason on the basis of a self-model, formulate its own goals, and actively pursue them. Since this machinery, in effect, guides the comportment of the resulting entity, it can be considered as a control network. The engineering of biosystems is not a new practice, but the cyborging of an ecosystem through the addition of a control network that explicitly enhances its autonomy is one novel aspect of the EcoCyborg Project.

8.5 Mental architecture for a substantially autonomous ecocyborg

If cyborging is to be used to create entities that are substantially autonomous, then the control network of the resulting system must be of a suitable design. The virtual machinery of such a mind is most appropriately characterized at an intermediate scale, for if the mind is resolved as a unitary entity, then its internal composition and structure cannot be engineered and, on the other hand, if it is resolved at a very fine scale then larger, integrated constructs cannot be created. Also, the most appropriate control organization of the virtual machines is a semihierarchical one, which ensures flexibility and robustness. As well, they should be functionally semidifferentiated for maximal effectiveness.

The five mental faculties essential for substantial consciousness and autonomy can be implemented with such an architecture by creating collections of virtual machines to perform the appropriate computational (mental) tasks. These machines would be of two principal types, the first (*knowledge objects*) forming an object-oriented knowledge base (OOKB) and the second (*cybernetic mechanisms*) processing the information contained therein. The five mental faculties required for substantial autonomy must arise from this virtual machinery. The mechanisms involved in perception (*perceptors*) transduce information from input signals into knowledge objects in the OOKB. Memory mechanisms index these objects for storage, retrieval, or use by other mechanisms, and otherwise maintain the OOKB. New knowledge objects may be created in this process, and obsolete ones destroyed. Reason arises from cybernetic mechanisms that process the information in the knowledge base into compound knowledge objects such as subgoals, strategies, tactics and directives for achieving the overall objectives of the ecocyborg. Expression involves cybernetic mechanisms (*effectors*) that translate these directives into output signals. Finally, learning mechanisms restructure the mind so that the ecocyborg can adapt to novel situations. A number of strategies that might be employed in each of these kinds of activities have been described in some detail.

8.6 Characterization

The engineering philosophy and lexicon described in this thesis are general tools for characterizing systems that are substantially autopoietic (i.e., biosystems), autonomous, or both. The specific mental architecture presented for ecocyborgs has been formulated by applying these tools, and is a prescriptive characterization of a system engineered with a particular approach (cyborging). Thus, all of the work presented in this dissertation is directly related to the characterization of biosystems, those specifically dealt with being ecocyborgs.

Characterization is an epistemic process that involves methods for obtaining, generating, and communicating knowledge. This process is important in the EcoCyborg Project because any substantially autonomous entity must be able to characterize itself, and also because scientific descriptions and engineering specifications are end results of characterization. The foundation of characterization is the interaction between an observer and an observed phenomenon. These two entities may be different from one another, as in the case of a scientist or engineer studying an ecocyborg, or they may be one and the same, as in the case of an ecocyborg observing itself. Several kinds of activity are required of the observer, including perception, discrimination, conceptualization, and expression. Because each of these can only be accomplished in a limited way, characterization can never be completely objective. The end result of the whole process may be the generation of a description, or the codification of some part of the knowledge acquired by the observer. This is accomplished with a language, which is framed in the context of a particular paradigm, as well as associated archetypal concepts and descriptive methods such as measures.

The characterization of substantially autonomous biosystems like ecocyborgs can be founded in a paradigm corresponding to the engineering philosophy set forth in this thesis. Thus, the nature of systems of interest can be described according to complementary pairs of continuously variable descriptors, placing them in the aforementioned hypercube of existence. The various aspects of these systems, such as their composition, structure, state, and comportment, can then be characterized according to archetypal concepts, such as equivalence relationships and trajectories in state space. Procedures such as measures, associated with these archetypes, can then be used to generate explicit descriptions of the various aspects of the systems. Any number of conceptual archetypes and descriptive procedures might be formulated, some of which are preferable to others in a given context. Thus, it is impractical to compile a comprehensive list of such characterization methods, even if it be only of those suitable for use in ecocyborg engineering. A number of different kinds of archetypes and measures are presented, however, that are convenient in this context.

CHAPTER 9. RECOMMENDATIONS

The work described in this thesis was conducted as part of the EcoCyborg Project, and so the following recommendations relate not only to the developments recounted here, but also to the short-term and long-term objectives of the project as a whole. To reiterate, the short-term objective of the EcoCyborg Project is to develop computational models, simulations, and characterization tools for use in the study of ecocyborgs. The long-term goal of the project is the development of a general theory of biosystems engineering, with an emphasis on substantial autonomy as a design goal. As part of this thesis, a theoretical explanation of biosystems as autopoietic entities has been given, and the surrounding concepts have been explained. As well, a lexicon has been developed for discussion about substantially autonomous systems. Also, the characterization process itself has been analyzed from a theoretical perspective. Finally, to help illustrate and clarify these various themes, a number of different measures have been described. Evidently, however, much remains to be done in the formulation and evaluation of tools for the characterization of biosystems, and following are some recommendations for this work.

1) **Evaluate the measures that have been identified to determine their relative usefulness in characterizing various kinds of biosystems.** This could be done by testing them on data from computer simulations, physical laboratory models, and large-scale natural and artificial systems. Ecocyborgs, both virtual and physical, are prime candidates as subjects in such investigations, since they are more convenient to work with compared to some other kinds of biosystems, such as organisms. Ecocyborgs are comparatively easy to instrument and modify (although large, physical ones can be rather expensive to construct and maintain). Also, there are few ethical or moral issues associated with them. This kind of work could be done immediately at an applied level.

2) **Formulate and evaluate other measures related to those that have already been described.** For instance, novel measures of order and disorder, complexity, and

emergence could be derived for use in the characterization of particular systems. Such measures might be found to be more appropriate than those presented here.

3) **Formulate novel families of measures corresponding to entirely different archetypal concepts.** For instance, methods of characterizing autopoietic comportment more directly would be extremely useful in the comparison of widely different types of. One intriguing possibility, for example, is that the level of autopoietic activity of biosystems might correspond to temporal patterns in their production of disorder (characterized as entropy or with similar measures). As well, means of characterizing the various aspects of autonomy (e.g., automation, volition, and intentfulness) would also be very beneficial. This kind of research is an intermediate stage in the translation of highly theoretical concepts into a more tangible form.

4) **Develop a general theoretical basis for the determination of the number and types of variables necessary for the effective, parsimonious characterization of highly complex systems.** For a system in which all the variables are correlated to some degree, the knowledge gained by observing each additional variable diminishes as more of these are taken into account. The creation of a model having complete equivalence to a particular aspect of a system is likely to be impractical or impossible. Even so, there is currently no objective way to determine which variables must be observed in order to create a model having a given degree of equivalence with the phenomenon of interest.

5) **Formulate a general theoretical basis for the understanding of the relationships between the composition, structure, initial state, and comportment of substantially autopoietic and autonomous systems, as they are affected by external forcing functions.** Partly as a result of the work described here, it is now possible to discuss these issues in a general but coherent way. The lack of detailed theoretical understanding of such systems, however, still impedes their effective engineering.

6) **Formulate a general theoretical basis for the determination of the kind of structure and initial conditions that will cause a biosystem to display a particular mode of comportment under the influence of a given set of forcing functions.** The development of such an understanding is the fundamental challenge confronting investigators who wish to create a general theory of biosystems engineering.

7) **Create ecocyborgs that have minds patterned after the mental architecture described in Chapter 6.** Control systems for both physical and virtual entities, such as the ecosystem models currently being developed, could be advanced by implementing such an architecture. This work must, however, proceed in step with the development of characterization tools since, if such minds are to be substantially autonomous, they will require methods of self-observation.

CHAPTER 10. ORIGINALITY AND CONTRIBUTIONS TO KNOWLEDGE

The work described in this thesis has resulted in the following original contributions to knowledge:

1) **The elucidation of a philosophy for the engineering of substantially autonomous biosystems.** A philosophy has been described that is suited to the engineering of complex, adaptive systems (such as biosystems in general, and ecocyborgs in particular). This involves characterization of phenomena in terms of orthogonal pairs of complementary descriptors (*real/imaginary*, *physical/virtual*, etc.) that define a hypercube of possible kinds of existence. As well, the natures of engineering and science are explained and contrasted as two fundamentally different ways of viewing phenomena thus described, where engineering has a prescriptive orientation and science an explanatory one.

2) **The development of a clear, coherent lexicon for use in the characterization of biosystems from an engineering perspective.** The term *biosystem* and related concepts have been explained, enabling the effective characterization of biosystems for the purposes of engineering them. A biosystem has been defined, within a systems-theoretical context, as any system that is alive to some degree. The sole criterion for being alive is *autopoietic* comportment, whereby the components of the system in question interact so as to continuously renew the system as a whole. Hence, autopoiesis corresponds to comportment that, if characterized with a state space based on appropriate measures, tends toward a particular kind of attractor. This work has also resulted, incidentally, in a clear (albeit rather broad) demarcation of the field of biosystems engineering.

3) **The development of a clear, coherent lexicon for use in the characterization of substantially autonomous systems from an engineering perspective.** The term *autonomy* and related concepts have been explained, so that systems might be effectively characterized in these terms for engineering purposes. This involved the

integration of concepts from the cognitive sciences (*autonomy*, *automation*, *volition*, *intentfulness*, *consciousness*, *mind*, etc.) and the development of working definitions for these in the context of biosystems engineering. Thus, a substantially *autonomous* system has been defined as one whose comportment is somewhat independent of external guidance in that it can, to at least some degree, operate persistently while formulating and actively pursuing its own goals.

4) **The description of a novel approach to the engineering of biosystems by combining technological and biological systems (*cyborging*).** It is described how this approach can be used to create new systems with particular, predetermined properties. Although this approach is generally applicable to all biosystems and a wide range of design goals, emphasis has been placed on the engineering of large-scale biosystems (such as ecocyborgs) with the explicit intent of creating substantially autonomous entities.

5) **The demarcation of a plan for the future development of substantially autonomous ecocyborgs** (and other large-scale cyborged biosystems). An architecture has been elaborated for the information storage and processing devices constituting the mind of a cyborged biosystem with substantial autonomy. It has been described in some detail how a mind patterned according to such an architecture can be constructed so as to give rise to the required abilities.

6) **The examination and illustration of the epistemic nature of characterization as applied to substantially autonomous biosystems.** The nature of characterization has been analyzed, with the role of a finite observer in this process being explicitly recognized. It has been explained how the inherent limitations of any observer constrain the objectivity that is possible in the process of characterization. In light of this explanation, the characterization process has been illustrated in the context of cyborged ecosystems, demonstrating how a suitable *paradigm*, *archetypal concepts*, and *measures* might arise in a system or might be chosen for use in a given context. This has resulted in the identification of a number of tools (e.g., measures of

composition, structure, state, and comportment) that are suitable for the characterization of such systems in an engineering context.

ADDENDA

This section includes the suggestions submitted by the external examiner (Dr. Stephen D. Murphy, Assistant Professor, Department of Environment and Resource Studies, University of Waterloo, ON) for improving the thesis. They are quoted verbatim from the Doctoral External Report which Dr. Murphy returned to the Faculty of Graduate Studies and Research, McGill University (dated 14 December, 1999). The only change that has been made in the list of suggestions has been to replace the original bullets with numbers, to facilitate reference to specific suggestions. Following the suggestions from Dr. Murphy are responses to those suggestions, written by O.G. Clark after the completion of his doctoral oral examination, which took place on 14 January, 2000.

1 External Examiner's suggestions for improving the thesis

1) I find that the literature used is rather populist in nature. Nothing wrong with citing references aimed at the public (or at least a segment of the public) but the thesis often omits some of the major peer-reviewed literature on subjects like complex systems. For example, James Kay (in my own department yet), Tim Allen, and C.S. (Buzz) Holling all have written extensively on the subject of complex systems. Additionally, the works of Bawden, Conway, Marten, Checkland, and Todd are relevant as well. This is, of course, the rationale for the first question I want to have addressed in Mr. Clark's verbal examination. Regardless of his answer, I would think that Mr. Clark needs to discuss/critique some of their papers to [*sic*] the literature reviews (especially chapters 2 and 3).

2) Admittedly, the question on Lovelock's Gaia hypothesis is leading and perhaps cynical but it is worth asking. I suggest Brockman, J. The third culture: beyond the scientific revolution. Simon and Schuster, New York. ISBN: 0-684-82344-6 would help start the answer.

3) In general, I think a deeper philosophical dialogue based on some of the work of Kuhn, Medawar, and Popper would help; I recognize that these are more scientific

philosophers but this is the context for my question on the difference between science and engineering. I suggest that Mr. Clark errs in his distinction between the two; certainly, little philosophical evidence or literature is offer [*sic*] to support such a key contention.

4) I also think Mr. Clark misses the main point of the reductionist-holist or individual-system debate. This is why I ask about what scale selection operates upon.

5) Similarly, the definition of living and consciousness in [*sic*] not well defended in the written thesis. My questions on these are meant to help here.

6) The sections on fuzzy sets and object-oriented modeling are highly relevant but in both cases, Mr. Clark missed a rather large amount of literature that is relevant to ecocyborg systems (apologies if this is redundant – perhaps just "ecocyborgs" will do). This is captured by my questions on fuzzy sets, object-oriented models and comparative question to Living Machines and Breathing Walls.

7) I find chapter 7 very interesting. However, I was a bit frustrated that I did not get a clear sense of exactly what even a model of an ecocyborg system would entail (at least not in detail beyond what is presented). My question on this attempts to probe into Mr. Clark's knowledge of this (i.e., I'm sure there's more he could tell me). This is considered in the context that I can find published work that tells me what equations I need to do spatially-explicit or fuzzy set modeling.

8) My questions on whether cultural conditions are relevant and ethics of ecocyborgs are meant to get Mr. Clark to consider an even more interdisciplinary approach because success or failure often hinges on whether humans accept the ethics presented or whether we account for human errors (actual mistakes or misuse via changing social mores).

9) As mentioned on page 1 of my comments, I would like to see the research question made more obvious; Mr. Clark does not need to be questioned on this.

2 Responses to the External Examiner's suggestions

Suggestion 1 relates to areas of literature that Dr. Murphy felt were missed or reviewed in insufficient depth. In the connecting text preceding Chapter 2 (Review of Relevant Literature) it is explicitly stated that "[D]ue to the multidisciplinary nature of this project, the bibliography was not intended to be comprehensive, but rather to present a general overview of literature associated with the relevant themes. The reference sections of the other chapters should, therefore, be consulted for more recent and specific citations relating to the corresponding topics." This statement is echoed in the introduction of Chapter 2, where it is stated that "[T]here are extensive bodies of literature directly and indirectly associated with each of the themes mentioned above, and to attempt a comprehensive review of all of them would exceed the bounds of this article. Therefore, most of the references that are presented are overviews, works of a general philosophical nature, or representative samples of the current state of knowledge in the relevant fields. Also cited are works that have had a particularly significant influence on the evolution of the EcoCyborg Project."

With reference to the specific authors mentioned by Dr. Murphy, the work of Kay and his associates is probably the most relevant to this thesis. Indeed, Schneider and Kay (1995) are cited in Chapter 7 with reference to their use of ideas from thermodynamics in the characterization of ecosystems. There is some forthcoming work (Kay and Regier 2000; Boyle et al. 1999) that is highly applicable to the topics discussed in Chapters 3 and 7. The unpublished thesis by Kay (1984) is of special relevance, being an in-depth discussion of the characterization of ecosystems from the perspective of thermodynamics.

Other, published works by Kay and his associates (Kay et al. 1999; Schneider and Kay 1995; Schneider and Kay 1994a; Schneider and Kay 1994b; Kay and Schneider 1994; Schneider and Kay 1993; Kay and Schneider 1992; Kay 1991) are also somewhat relevant to the discussion, but are generally presented in the context of the public management of ecosystems. Their content is not directly related to the characterization of biosystems as being autopoietic, nor are any formal measures (of, e.g., order or disorder, complexity, or emergence) offered that are substantially different from those already discussed in Chapter 7. The same is generally true of the other authors mentioned by Dr. Murphy. There are some relevant publications by Allen about ecological complexity as it relates to the hierarchical organization of biosystems (Ahl and Allen 1996; Allen 1987; Allen and Starr 1982). Holling (1999, 1987) has discussed complexity in a general way, but the bulk of his work is also oriented toward the public management of ecosystems (Holling 1995, 1994, 1993, 1986). Work by Bawden (1992, 1991, 1990; Bawden et al. 1984) treats the perception of, management of, and education about agricultural systems and, again, is oriented toward issues of public policy. It is possible that Dr. Murphy's mention of Conway is a reference to J.H. Conway, who first proposed the famous cellular automata popularly referred to as the "Game of Life". This seminal idea in the field of artificial life was explored extensively by others, including Wolfram (1984). In light of the other authors mentioned, however, it is more likely that the intended reference is to G.R. Conway (1990, 1985), whose work is in policy development for sustainable agricultural. No reference was found to Marten in a brief survey of the recent literature about complexity. Checkland (1998, 1994, 1992, 1990, 1981) has written primarily in the area of operational systems management. Dr. Murphy's reference to Todd may be to J. Todd, discussed further below in the context of Living Machines™ (Living Technologies Inc., Burlington, VT). Another possible reference is to the work of M.J. Todd (Khachiyan and Todd 1993) in the area of algorithmic complexity. This is a concept that is related to measures of order and disorder, as mentioned in Chapters 3 and 7, albeit quite narrowly focused on problems of computer science.

Suggestion 2 relates to the Gaia hypothesis of Lovelock, and (as specified by Dr. Murphy in his suggested questions for the verbal examination) how Lovelock's original ideas might have been presented in the popular literature so as to pander to popular moral opinions about cooperation. Reference was made in Chapters 2, 3, and 7 to the article by Margulis and Lovelock (1974), which was apparently the first about the Gaia theory to receive widespread attention. Previous journal articles (Lovelock 1972; Lovelock and Margulis 1973) were not referenced in this thesis, nor were the popular works about the theory (Lovelock 1979). The manner in which the Gaia hypothesis has been expounded in the popular literature is not a theme that is relevant to this thesis. Rather, the central concept of the Lovelock's hypothesis is presented, in which it is suggested that the biosphere of the Earth is an example of a very large-scale system that might be considered alive (i.e., self-perpetuating or autopoietic, to use the terminology adopted here) to some degree.

With respect to the Suggestion 3, one assumes that the "key contention" to which Dr. Murphy refers is that science is an explanatory endeavor, whereas engineering is a prescriptive one (as discussed in Chapter 5 and elsewhere). It is not clear that this assertion is contentious, nor in need of defence. Popper, often credited with providing science with a firm epistemic foundation in the concept of *falsifiability* (Popper 1961; Thornton 1997), considered science to be the pursuit of theoretical, predictive knowledge. In his view, new scientific theories should be evaluated and compared with other theories, in part, on the basis of their predictive power. On the other hand, even Popper made use of the term *engineering* in the sense of a prescriptive endeavor. For instance, he discussed "social engineering" (Popper 1959; Thornton 1997), by which he referred to the attempt by the members of a society to fashion that society so that specific objectives would be fulfilled. In the engineering literature, one need only look in any introductory textbook for a corroborating definition of engineering. Andrews and Kemper (1999), for instance, include a definition of an engineer as "a person who uses science, mathematics and technology, in a creative way, to satisfy human needs," and, later, as someone "usually concerned with creating devices, systems, and structures for human use." It is acknowledged in this thesis (Chapters 1, 5, and 8) that the same investigative *methodology* is often employed in science and engineering to generate knowledge and solve problems. The *intent* that motivates such activities, however, differs fundamentally in that, in science, it is to explain (or, by extension, to predict), while in engineering it is to specify (and by extension, to create) a system that fulfills a prescribed purpose.

Dr. Murphy, in Suggestion 4, does not specify what he believes "the main point of the reductionist-holist or individual-system" debate to be; one might conjecture by the list of suggested verbal exam questions to which he refers and his mention of the "scale selection operates upon", that he disagrees with the idea that a system might be considered as alive at a scale of resolution coarser than (or different from) that of an organism. In the verbal exam questions he states the opinion that "the criteria [Mr. Clark] give[s] to 'living' seems rather contrived and dependent on the idea of autonomy or some form of self-regulation." Further, he asks, "Is not self-replication via heritable characters the key to life?" and "If the above is true, does this not preclude [eco]systems from being alive, in any sense of the word (or conscious, again sensu lato)." In fact, the ideas to which Dr. Murphy has expressed his disagreement are exactly those stated and defended in Chapters 2 and 3 of this thesis. The *sole* criterion for life proposed here, as discussed extensively in Chapter 3, is autopoietic comportment, which is indeed a form of self-organization. Heritable characters are not considered here as the key to life, nor are they strictly defensible as such even in other contexts. Heritable characters (with variation) certainly are one of two necessary requirements for a (neodarwinistic) evolutionary process to occur (the other being the preferential selection of more-fit systems) and neodarwinistic evolution is one possible avenue by which life might have originated. The *origin* of life, however, is distinct from the *nature* of life, and the former is not discussed here. Nevertheless, if one were to abandon for the sake of argument the definition of life as autopoiesis, adhered to in this thesis, and accept that self-replication was indeed the "key to life", one could argue that there are ways in which ecosystems (assuming that Dr. Murphy is writing of ecosystems) do indeed replicate themselves. This might occur, for instance, through their expansion and subsequent fragmentation, as well as through the migration of species cohorts from an original ecosystem to a new location, resulting in the process of succession and the establishment of a new climax community in that location, as hypothesized by Clements (1916) and others.

Finally, Dr. Murphy asks, in the list of verbal exam questions, "If emergent properties define life, does this mean that a crystal (e.g. of snow) or a vortex (e.g., a tornado) are alive? Logically, your definitions in chapter 4 and 7 would seem to indicate they are; I beg to differ but would be interested in your defense of this." Indeed, in the spirit of Prigogine (1980), a vortex is considered to maintain itself in a far-from-equilibrium state and is therefore, in a vague sense, somewhat "self-producing" as described by Maturana and Varela (1980). A crystal is not generally considered to be a dynamic system, but the process of crystalline growth might be construed to be marginally autopoietic. The crux of this interpretation is to treat autopoiesis (or life) as a *variable* characteristic that a given system might possess to *any degree*, even a very marginal one.

With respect to Suggestion 5, the reader is referred to Chapters 3 and 4, where the definitions of living and consciousness are defended at length.

Suggestion 6 makes reference to fuzzy sets and object-oriented modeling. Fuzzy logic (Kosko 1993) was referred to in Chapters 3 and 6 to illustrate the idea of characterization with continuous variables, but it was not the intent to express any new ideas with respect to this well-known conceptual approach. Ideas from object-oriented programming (Dubitzky et al. 1996; Gauthier and Guay 1998; Gauthier and Néel 1996; Zeigler 1990) were used in the development of the mental architecture proposed in Chapter 6, but again there was no attempt to extend this well-known methodology, nor to employ it as a modeling tool (i.e., to create a predictive representation of a particular system).

"Living Machine" is a trademarked name (Living Technologies Inc., Burlington, VT) that refers a type of commercial biosystem based on the ideas of J. Todd, which are described in a series of popular books (e.g., Todd 1980, 1994). Such biosystems indeed fit admirably into the definition of *ecocyborg*, since they are highly artificial ecosystems that include technological components. These kinds of ecocyborgs have been developed for purposes such as the treatment of industrial effluents and municipal sewage. "Breathing Wall" refers, again, to a type of commercially available system (Genetron Systems, Downsview, ON), which could be considered as ecocyborgs. In such a system, air (e.g., in an office building) is circulated through panels of growth substrates that support a self-sustaining community of plants (e.g., mosses and ferns). These plants are intended to remove volatile organic compounds from the air and thus improve its quality.

With respect to Suggestion 7, the objectives stated in Chapter 7 did not include the description of a model of an ecocyborg; the intent instead was to discuss and illustrate characterization as an epistemic process. Neither was there any effort made to address cultural or ethical questions in this thesis (Suggestion 8), although, as Dr. Murphy undoubtedly intended when mentioning these topics, they became the subjects of much lively discussion during the oral examination.

Lastly, in Suggestion 9, Dr. Murphy asks that a clearer research question be stated. Granted, neither this thesis nor the body of work that it describes is patterned after the traditional (Popperian) scientific research format of stating a falsifiable hypothesis, performing experiments, and drawing conclusions based on the results of the latter. It must be remembered, however, that this is not a scientific thesis, but an engineering one. Moreover, although presented in a format that might be unfamiliar, the thesis meets all of the objectives as laid out in the 1998 revision of the *Guidelines for Thesis Preparation* (Faculty of Graduate Studies and Research, McGill University). Elements of the thesis that are considered to constitute original scholarship and an advancement of knowledge in the domains in which the research was conducted have been clearly indicated (Chapter 10, Originality and contributions to knowledge). Also included in the thesis are: a detailed table of contents; a brief abstract in both English and French; an introduction that clearly states the rationale and objectives of the study (Chapter 1, especially Section 1.5, Objectives); a review of the literature (Chapter 2 and the introductory sections of Chapters 3 through 7); a final conclusion and summary (Chapters 8, 9, and 10); and a thorough bibliography (the reference sections of Chapters 2 through 7). Overall, the research objectives are indeed clearly stated, (although not in the form of questions), an organized and coherent body of writing is then presented in which those objectives are addressed, and concluding sections give suggestions for further work and clearly summarize how the stated objectives were met.

3 References

Ahl, V. and T.F.H. Allen. 1996. *Hierarchy Theory: A Vision, Vocabulary and Epistemology*. New York, NY: University of Columbia Press.

Allen, T.F.H. 1987. Hierarchical complexity in ecology: a non-euclidean conception of the data space. *Vegetation* 69:17-25.

Allen, T.F.H. and T.B. Starr. 1982. *Hierarchy: Perspectives for Ecological Complexity*. Chicago, IL: University of Chicago Press.

Andrews, G.C., and J.D. Kemper. 1999. *Canadian Professional Engineering Practice and Ethics*. Toronto, ON: Harcourt Brace.

Bawden, R.J. 1992. Systems approaches to agricultural development: the Hawkesbury experience. *Agricultural Systems* 40:153-176.

Bawden, R.J. 1991. Systems thinking and practice in agriculture. *Journal of Dairy Science* 74: 2362-2373.

Bawden, R.J. 1990. Of agricultural systems and systems agriculture: systems methodologies in agricultural education. In *Systems theory applied to agriculture and the food chain*, ed. J.G.W. Jones and P.R. Street, pp. 305-323. London, UK: Elsevier.

Bawden, R., R.D. Macadam, R.G. Packham and I. Valentine. 1984. Systems thinking and practice in the education of agriculturalists. *Agricultural Systems* 13:205-225.

Boyle, M., J.J. Kay and B. Pond. 1999. Ecosystems as self-organizing holarchic open systems: narratives and the second law of thermodynamics. Available online at: <http://www.fes.uwaterloo.ca/u/jjkay/pubs/CRC/> [Accessed 20 January, 2000].

Checkland, P., and S. Holwell. 1998. *Information, Systems and Information Systems*. New York, NY: Wiley.

Checkland, P. 1994. Systems theory and management thinking. *American Behavioural Scientist* 38(1): 75-91.

Checkland, P. 1992. Systems and scholarship. *Journal of the Operational Research Society* 43(11): 1023-1030.

Checkland, P., and J. Scholes. 1990. *Soft Systems Methodology in Action*. New York, NY: Wiley.

Checkland, P. 1981. *Systems thinking, systems practice*. New York, NY: Wiley.

Clements, F.E. 1916. *Plant Succession: An Analysis of the Development of Vegetation*. Pub. 242. Washington, DC: Carnegie Institute of Washington.

Conway, G.R. 1985. Agroecosystem analysis. *Agricultural Administration* 20:31-55.

Conway, G.R., and E.B. Barbier. 1990. *After the green revolution: sustainable agriculture for development*. Toronto, ON: Earthscan Canada.

Dubitzky, W., D. Bell and J. Hughes. 1996. Generic, object-oriented case-knowledge representation scheme, and its integration into a wider information management scenario. *Expert Systems* 13(3):219-233.

Gauthier, L., and R. Guay. 1998. Using object-oriented database management technology in agricultural decision support software. *Canadian Agricultural Engineering* 40(3):219-226.

Gauthier, L., and T. Néel. 1996. SAGE: An object-oriented framework for the construction of farm decision support systems. *Computers and Electronics in Agriculture* 16:1-20.

Holling, C. S. 1999. Introduction to the special feature: just complex enough for understanding; just simple enough for communication. *Conservation Ecology* 3(2):1. Available online at: <http://www.consecol.org/vol3/iss2/art1> [Accessed 20 January, 2000].

Holling, C.S. 1995 What Barriers? What Bridges? In *Barriers and Bridges to the Renewal of Ecosystems and Institutions*, ed. L.H. Gunderson, C.S. Holling and S.S. Light, pp. 3-34. New York, NY: Columbia University Press.

Holling, C.S. 1994. New Science and New Investments for a Sustainable Biosphere, in *Investing in Natural Capital: The Ecological Economic Approach to Sustainability*, ed. A.M. Jansson, C. Folke, R. Costanza, and M. Hammer, pp. 57-73. Washington, DC: Island Press.

Holling, C.S. 1993 Investing in Research for Sustainability. *Ecological Applications* 3(4):552-555.

Holling, C.S. 1987. Simplifying the complex: the paradigms of ecological function and structure. *European Journal of Operational Research* 30:139-146.

Holling, C.S. 1986. Resilience of ecosystems; local surprise and global change. In *Sustainable Development of the Biosphere*, ed. W.C. Clark and R.E. Munn, pp.292-317. Cambridge, MA: Cambridge University Press.

Kay, J.J. 1991. A non-equilibrium thermodynamic framework for discussing ecosystem integrity. *Environmental Management* 15(4):483-495

Kay, J.J. 1984. *Self-organization in living systems*. Unpublished Ph.D. thesis. Department of Systems Design Engineering, University of Waterloo, Waterloo, ON, 458 p. Chapter 2 of this work available online at: <http://www.fes.uwaterloo.ca/u/jjkay/pubs/thesis/toc.html> [Accessed 20 January, 2000].

Kay, J.J., and H. Regier. 2000. Uncertainty, complexity, and ecological integrity: insights from an ecosystem approach. In *Implementing Ecological Integrity: Restoring Regional and Global Environmental and Human Health*, ed. P. Crabbé, A. Holland, L. Ryszkowski and L. Westra. Dordrecht, the Netherlands: Kluwer, (in publication). Also available online at: <http://www.fes.uwaterloo.ca/u/jjkay/pubs/NATO/> [Accessed 20 January, 2000].

Kay, J.J., H. Regier, M. Boyle and G. Francis. 1999. An ecosystem approach for sustainability: addressing the challenge of complexity. *Futures* 31(7):721-742.

Kay, J.J., and E.D. Schneider. 1994. Embracing complexity, the challenge of the ecosystem approach. *Alternatives* 20(3):32-38.

Kay, J.J., and E.D. Schneider. 1992. Thermodynamics and measures of ecosystem integrity. In *Ecological Indicators, Proceedings of the International Symposium on Ecological Indicators*, vol. 1, ed. D.H. McKenzie, D.E. Hyatt, V.J. McDonald, Elsevier, pp.159-182. New York, NY: Elsevier.

Khachiyan, L. G., and M. J. Todd. 1993. On the complexity of approximating the maximal inscribed ellipsoid for a polytope. *Mathematical Programming* 61:137-59.

Kosko, B. 1993. *Fuzzy Thinking: the New Science of Fuzzy Logic.* New York, NY: Hyperion.

Lovelock, J.E. 1979. *Gaia: a new look at life on earth.* Oxford, UK: Oxford University Press.

Lovelock, J.E. 1972. Gaia as seen through the atmosphere. *Atmospheric Environment* 6:579.

Lovelock, J.E., and L. Margulis. 1973. Atmospheric homeostasis by and for the biosphere: the Gaia hypothesis. *Tellus* 26:2.

Margulis, L., and J.E. Lovelock. 1974. Biological modulation of the Earth's atmosphere. *Icarus* 21:471.

Maturana, H.R., and F.J. Varela. 1980. *Autopoiesis: the Organization of the Living.* Dordrecht, the Netherlands: Reidel.

Popper, K. 1961. *The Poverty of Historicism*, 2nd. ed. London, UK: Routledge.

Popper, K. 1959. *The Logic of Scientific Discovery*. London, UK: Hutchinson.

Prigogine, I. 1980. *From Being to Becoming*. San Francisco, CA: Freeman.

Schneider, E.D. and J.J. Kay. 1995. Order from disorder: the thermodynamics of complexity in biology. In *What is Life: The Next Fifty Years. Reflections on the Future of Biology*, ed. M.P. Murphy, L.A.J. O'Neill, pp. 161-172. Boston, MA: Cambridge University Press.

Schneider, E.D., and J.J. Kay. 1993. Exergy degradation, thermodynamics, and the development of ecosystems. In *Energy, Systems, and Ecology*, vol. 1, Proceedings of ENSEC' 93, Cracow, Poland, July 5-9, ed. G. Tsatsaronis, J. Szargut, Z. Kolenda, Z. Ziebik, pp. 33-42.

Schneider, E.D., and J.J. Kay. 1994a. Complexity and thermodynamics: towards a new ecology. *Futures* 24(6):626-647.

Schneider, E.D., and J.J. Kay. 1994b. Life as a manifestation of the Second Law of Thermodynamics. *Mathematical and Computer Modelling* 19(6-8):25-48.

Thornton, S. 1997. Karl Popper. In *The Stanford Encyclopedia of Philosophy*, ed. E.N. Zalta. Available online at: <http://plato.stanford.edu/entries/popper> [Accessed 21 January, 2000].

Todd, J. 1994. *From Eco-Cities to Living Machines™*. Berkeley, CA: North Atlantic Books.

Todd, J. 1980. *Tomorrow is Our Permanent Address*. Hudson, NY: Lindisfarne Press.

Wolfram, S. 1984. Universality and complexity in cellular automata. *Physica D* 10:1-35.

Zeigler, B.P. 1990. *Object-oriented simulation with hierarchical, modular models: intelligent agents and endomorphic systems*. San Diego, CA: Academic Press.