



BISON
IST-2001-38923

*Biology-Inspired techniques for
Self Organization in dynamic Networks*

Abstract Models for Candidate CAS

Deliverable Number: D02
Delivery Date: June 2003
Classification: Public
Contact Authors: Geoffrey Canright, Andreas Deutsch,
Kent Engø-Monsen
Document Version: Final (January 30, 2004)

Contract Start Date: 1 January 2003
Duration: 36 months
Project Coordinator: Università di Bologna (Italy)
Partners: Telenor Communication AS (Norway),
Technische Universität Dresden (Germany),
IDSIA (Switzerland),
Santa Fe Institute (USA)

**Project funded by the
European Commission under the
Information Society Technologies
Programme of the 5th Framework
(1998-2002)**



Abstract

The BISON project seeks to develop self-managing systems which rely on complex adaptive systems (CAS) to do the management. The notion of CAS is not precisely defined. Here we review common definitions of the term, and then offer a BISON working definition of a CAS, in terms of many simple interacting agents. We then give a basic set of 'microscopic mechanisms' which we believe can be used to build generic CAS for BISON purposes. These microscopic mechanisms act at the individual-agent level, while network management functions (described in D01) emerge from these microscopic actions as collective phenomena. We give an analysis of five (idealized) biological systems in terms of our basic list of mechanisms; these hence represent five specific examples, motivated by biology, of how interesting collective behavior can arise from simple rules. Finally, we present some background material which compares the BISON CAS approach with related fields and concepts; and we discuss the notion of unpredictability for CAS, and its relevance for engineering with such systems.

Contents

1	Introduction	4
2	Working definition of a CAS	4
3	The emergence equation	9
4	Fundamental building blocks for agents	10
5	Microscopic mechanisms for biology-inspired CAS	12
5.1	Mechanisms	13
5.2	Analysis of five biological agents	14
5.2.1	Ants	14
5.2.2	Immune cells	14
5.2.3	Amoebae	15
5.2.4	Neurons	15
5.2.5	Viruses	15
5.2.6	Summary and table	16
6	Placing the BISON approach in a larger context	16
6.1	Mobile agents	17
6.2	Cellular automata	17
6.3	Artificial life	19
6.4	Process calculus	19
7	Engineering and emergence	20

1 Introduction

The purpose of this Deliverable is to establish the basic framework of the theoretical foundation for the BISON project. A great deal of the content here consists of definitions. Furthermore, much text is devoted to clarification of relationships between definitions adopted here, and those used in other fields. This kind of clarification is needed, because the definitions are of a nature that is still largely verbal and hence less than mathematically precise. The definitions themselves are useful and important: they represent a step towards higher precision.

Our most important definition here is that of a complex adaptive system or CAS (we take the plural also to be “CAS”). This is clearly a prerequisite to building any abstract *model* of a CAS. Section 2 is devoted to defining a CAS; the task is done slowly and carefully, as there is much variation and unclarity in the use of this term among researchers.

Section 3 then presents the key concept of *emergence*, in the form of a flow diagram (the “emergence equation”). The concept of emergence is embedded in our definition of a CAS (and virtually all others). Furthermore, emergence—ill-defined as it is—describes how BISON agents will perform useful functions on networks. Hence the concept requires attention.

Model-building begins in Section 4, with a set of abstract “fundamental building blocks”. These building blocks will be used to design agents, which in turn make up a CAS. They are represented in terms of state transitions, implying a view of an agent as a finite-state machine.

In Section 5 we elaborate on the microscopic (ie, agent-level) lifelike mechanisms which may be realized using the state transitions of Section 4. These microscopic mechanisms are then used as a way of analyzing the behavior of four distinct types of elementary biological “agents”: ants, immune cells, amoebae in a slime mold colony, and neurons. Hence Section 5 is the “fruit” of this deliverable, as represented in its title: Section 5 presents abstract models for real biological CAS; but the same building blocks may then be used, in novel combinations, to build new candidate CAS which may perform, collectively, useful functions.

Thus the foundation of the BISON approach is laid at the end of Section 5. The remaining sections then seek to clarify loose ends. In section 6 we identify explicitly “swarm intelligence” as the appropriate term for the BISON approach, and discuss a number of fields which are related to this approach. The aim here is to clarify the relationships—similarities and differences—between these approaches.

Finally, in Section 7 we revisit emergence, in order to offer a brief and general exposition of the tension between emergence and unpredictability on the one hand, and engineering and predictability on the other. BISON aims to engineer using emergence. Section 7 offers, in very general terms, an answer to the question: How is this possible?

2 Working definition of a CAS

The notion of “complex adaptive system” is widespread. However, there are few if any precise definitions of the concept; and furthermore, there is not a consistent (albeit imprecise) consensual definition of this phrase. That is, the words in this combination mean different things to different people. Here, for example, is a definition given by David Pines, based on a 1998

workshop [13] at the Santa Fe Institute:

complex adaptive systems can be regarded as a **collection of information-gathering entities (agents)** which

- **respond to the environment**
- **respond to one another**
- segregate information from random noise
- compress regularities into a model
- **modify their internal characteristics**—i.e., adapt to improve their performance of desired tasks

Typically, complex adaptive systems:

- possess intrinsic non-linearities which can lead to either negative or positive feedback
- **display emergent (self-organized) behavior**
- are unusually sensitive to initial conditions (“frozen accidents”)
- are rarely capable of finding an optimal state; instead, get “stuck” in local minima which are often “history” dependent.
- Finally, intervention in the affairs or behavior of a complex adaptive system, often gives rise to **unexpected consequences**.

Certain parts of the above description are placed in bold text here. These are the items that we find to be common to most or all definitions of a CAS. Those items that are both underlined and in bold text will be retained in our own definition. Because of the lack of consensus, we call the BISON definition of a CAS a “working definition”.

It is worthwhile discussing the highlighted terms above in some detail.

Collection of information-gathering entities (agents): An *essential* aspect of the CAS concept is that the global (collective) behavior is built upon the action of many small “parts”. “Agents” is a common name for these parts, and we will use that name here. The “information-gathering” aspect of these agents is essentially captured by the next two items.

Respond to the environment: An obvious ingredient, at least for BISON purposes. That is, one can imagine interesting types of CAS which are completely isolated from any environment; but this type is not useful for BISON.

Respond to one another: It is also essential that the many agents interact. In fact, here we will impose a further requirement on the interaction: it must be **nontrivial**. We illustrate the meaning of this term, in this context, by defining a trivial interaction. If the collective behavior of the N agents, acting in concert, is qualitatively (if not quantitatively) the same as that of one agent acting over a time span which is N times as long, then the interaction is “trivial”: it is

"brute parallelism". The notion of "nontrivial interaction" is thus closely tied to the notion of "emergence" (below); but it is more precisely defined.

Modify their internal characteristics: This may have several meanings. The words are directed towards concepts like learning, adaptation, and evolution. Here we seek some higher level of precision, as follows. We postulate that the agents have "states", in the standard sense of the term. The agents can then change their state, in response to the environment, to other agents, or to both. This ability to change state implies an ability to change behavior; and change of behavior represents a form of adaptation. A simple example may be taken from real ants, which may shift their state from one task (food gathering) to another (nest cleaning), based on inputs coming from the colony (environment + other ants).

Another form of adaptation is evolution. In biology, the time scale for adaptation by evolution is typically much longer than that for adaptation by change of behavior. We can think of the genetic material as embodying the set of "rules" for the organism, and the particular *activation pattern* of the genetic material as the state which yields the behavior. (Actually, other activation patterns are also highly relevant aspects of the state—for example, the activation pattern of a nervous system, when such is present. This discussion is thus highly oversimplified.) Mutation then involves a change in the genetic material ("rules"); and evolution acts, via selection, to induce changes in the population distribution of different rule sets. However, evolution as adaptation (via guided random search in the genome space) can occur on short time scales as well. An example [6] is seen in the immune system, when B-cells are stimulated (by a good match to an antigen) to explore very rapidly their genome space—seeking an even better match.

For the BISON working definition, the "internal characteristics" of an agent are its state structure: the states themselves, plus the transitions between them. This state structure represents (in the language of the previous paragraph) the "rules" of the CAS. And again two types of adaptive response are possible—state change (as noted above), which follows allowed transitions in the state structure, and allows behavior change; and change in the state *structure* itself. The latter is the analog of mutation for a BISON CAS; and, clearly, if mutation is allowed for a CAS (and some form of selection is also present), then adaptation by evolution is also possible for this CAS.

We take both kinds of adaptation to be possible, in principle, for a BISON agent. However we expect to rely most heavily on the former—that is, on adaptation via change of behavior, which arises from a fixed set of rules. The reason for this is that BISON agents are intended to act "on-line", in real time, in real network systems. Therefore there is a large, negative premium placed on maladaptive mutations. That is, adaptation by random mutation plus selection can be too dangerous for engineered systems. One has only to think of computer viruses to appreciate this point. Hence, if it is to be used at all, real-time evolution in a reliable, engineered system must be, at least, highly constrained, so as to hold the space of mutations away from known, or unknown, dangerous subspaces. And given the current state of the art, the space of state structures with unknown, hence possibly dangerous, behavior, is huge.

Having stated this caveat, we recall the example noted above of fast evolution in the immune system. This is an example of "safe, fast mutation" with obvious benefits for the organism. Possibly, CAS mutation mechanisms inspired by this example may be useful for supporting some of the functions to be addressed in the BISON project. For instance, searches in "information space" are in a formal sense similar to searches in "antigen space"—and so one can imagine

that fast searches in information space may be aided by fast mutation of the search pattern (the “antibody”, eg, keywords).

Display emergent (self-organized) behavior: The notion of *emergence* is a central element in the definition of a CAS. Hence our discussion of emergence will be thorough. We seek to (at least partially) define the term here—and will offer further discussion of emergence in Sections 3 and 7 below.

The many agents that together constitute the CAS have “microscopic” (internal) rules as discussed just above. However the behavior of *practical* interest is the *collective behavior* exhibited by the CAS as a whole. Specifically, for BISON purposes we want the collective behavior of the CAS to yield, in whole or in part, the functions to be implemented on the network. We can then implicitly define emergence by stating that the collective behavior *emerges* from the interaction of a large number of agents, each following its own microscopic rules.

This implicit definition has as yet little content. We can flesh it out a bit, by pointing out two (related) aspects of emergence, which are common to virtually every description of emergence. First, the interaction should be nontrivial, in the sense defined above. Otherwise there is no emergence—there is simply brute parallelism—and the system may be adaptive, but it is not complex. Secondly, emergence is almost invariably required to include an element of *unpredictability*. For instance, the last three items in the Pines definition above describe various types of unpredictability.

The first of these three, “sensitivity to initial conditions”, is an excellent, concise definition of chaos in dynamical systems. This type of unpredictability is understood and even quantified (via Lyapunov exponents, which can be both calculated and measured): in short, it is science. The second type mentioned by Pines (getting “stuck” in local minima) is very much like the first. In fact, a ball, moving over a crumpled surface under the influence of Newtonian mechanics, friction, and gravity, will get “stuck” in a local minimum; but *which* local minimum is eventually selected can be highly sensitive to initial conditions.

And yet there is still a third statement of unpredictability: that the “behavior of a complex adaptive system often gives rise to **unexpected consequences**”. This is a very broad statement of unpredictability, which may be of any type whatsoever. And yet this type of statement of unpredictability is ubiquitous in discussions of emergence. That is, emergent behavior is *defined* to be unpredictable—but the unpredictability itself is not understood or explained (as it is with classical dynamical chaos)—it is simply asserted to be present.

For example, here is (part of) a definition of a CAS, due to Kevin Dooley [3]:

A CAS behaves/evolves according to three key principles: order is **emergent** as opposed to predetermined, the system’s history is irreversible, and the system’s future is **often unpredictable**. The basic building blocks of the CAS are agents. Agents are semi-autonomous units that seek to maximize some measure of goodness, or fitness, by evolving over time.

Here we see the unpredictability in the definition (highlighted in bold). As another example, we quote from the home page of the journal *Emergence*:

... emergent patterns seem to be unpredictable and nondeducible from the components as well as irreducible to those components ...

This statement captures both the nontriviality concept stated above, and the “generic” unpredictability that is commonly found in discussions of emergence. It is important to note here that “generic” unpredictability—that is, unpredictability of global behavior, which is neither bounded, nor understood—is quite distinct from that unpredictability inherent in systems with stochastic behavior. Many useful and (globally) predictable systems are built upon stochastic behavior of components, say at the microscopic level. (See Section 7.)

We believe that “generic unpredictability” is not a useful part of any definition of how collective behavior emerges from microscopic rules. In the absence of any explanation (of the type given for classical chaos, or even of some weaker type), such statements are not helpful to understanding. Furthermore, given that BISON is interested in *engineering* with CAS, “generic” (not understood, and hence not bounded) unpredictability is highly undesirable. For these reasons, the BISON working definition of a CAS retains the word “emergence” to denote the qualitative distinctions between the collective behavior and the microscopic rules. That is, emergence in the BISON sense implies a nontrivial interaction among the agents. However we make no reference to unpredictability in our working definition.

It is clear that most interesting CAS (having a nontrivial interaction) will exhibit interesting types of unpredictability. We have only to look at one of the most simple types of CAS imaginable, namely, near-neighbor cellular automata on a line, with two states at each node [4]. Here one already sees not only behavior (class 3) that is strongly reminiscent of classical chaotic dynamics—and hence unpredictable in a similar sense. One also sees “class 4” cellular automata, whose behavior is so complex that it is believed to be Turing universal—which means, in this context, that there is no effective way to predict the behavior of such systems, other than by explicitly running them. This is an even stronger form of unpredictability; and such unpredictability may well arise in the CAS considered in the BISON project. However (again) we find no good reason to build this unpredictability into the working *definition* of a CAS.

Finally, we add an element to our working definition, which is not present in the above. That is, we require that our **agents are simple**. Our reasons for this requirement are twofold. One reason is practical: for some purposes (such as ad-hoc networks), limitations in battery life, processing power, storage, and/or bandwidth may make the use of highly sophisticated agents impractical or impossible. Our second reason derives from the spirit of “swarm intelligence”. That is, BISON will focus on systems for which the intelligence resides in the collective behavior, and *only* in the collective behavior. The agents themselves are to be unintelligent. Otherwise, the intelligence itself is not emergent, although the collective behavior is.

This requirement may be the only one distinguishing the BISON definition from that of the COIN system [14]. In the latter, each agent implements a reinforcement learning (RL) algorithm. In particular, each agent seeks to optimize, via RL, its own ‘local utility function’. The result of a ‘well-factored’ COIN system is that the action of many, concurrent, ‘selfish’ RL agents gives a collective behavior which optimizes a ‘world utility function’. One might then say that the agents themselves have the same sort of intelligence that would be expected of a global controller: they can estimate their environment, they can measure the payoffs of their behaviors, and they can modify their behaviors accordingly. Since the COIN definition does not specify the RL algorithm, it includes in principle very simple algorithms, such as searching in strategy space without estimating the environment. In this sense, the COIN definition overlaps our working definition of a CAS. However, sophisticated learning algorithms which place a signif-

icant intelligence in each agent are ruled out by our definition. Other than this one difference, the COIN prescription for a system of multiple interacting agents (see the list on page 6 of Ref. [14]), and the BISON definition of a CAS, are essentially the same.

It is worth noting that the COIN approach, like BISON, is aimed at inverting the emergence equation. In the COIN approach, the inversion is accomplished by ‘factoring’ the world utility function, ie, by finding a set of local utility functions that, when selfishly optimized by agents, yield good results for the world utility. In this regard, BISON is also distinguished from the COIN approach: BISON begins with an ‘agnostic’ viewpoint with regard to how the inversion is to be accomplished.

It remains to specify what is meant by the term “simple” here. Tentatively, we take our agents to be finite-state machines (supplemented by some computational capacity at the nodes). This already bounds their complexity: for a finite-state machine, the complexity may be quantified, simply by counting their states. This is only a tentative definition. Further precision in what is meant here by ‘simplicity’—and further clarification of the rationale for demanding ‘simple’ agents—may be expected to arise from experience with the other workpackages of BISON.

We summarize the above discussion with a concise statement of the BISON working definition of a CAS:

- **many agents,**
- **sensitive to their environment, and**
- **interacting via local rules;**
- **the agents are simple, and**
- **the interaction yields nontrivial collective behavior**

The new word “local” here simply reflects the reality that communication is never instantaneous—so that agents reach beyond nearest neighbors only by going through the latter. Thus, since each hop takes time, agents reach nearest neighbors first.

3 The emergence equation

The purpose of this section is to amplify and clarify the kind of dynamic structure that BISON seeks to understand and to exploit. In the previous section, emergence was used to name the process by which collective behavior arises from many interacting microscopic agents. Here we summarize that idea with a diagram:

$$\left(\begin{array}{c} \textit{set of} \\ \textit{microscopic} \\ \textit{mechanisms} \end{array} \right) \equiv \left(\begin{array}{c} \mathbf{abstract} \\ \mathbf{CAS} \end{array} \right) \xrightarrow{\textit{emergence}} \left(\begin{array}{c} \mathbf{collective} \\ \mathbf{behavior} \end{array} \right) = (\textit{function})$$

We focus on this diagram in order to make clear that there are two levels of behavior under study. At the microscopic level we have agents, with well-defined rules of behavior embodied in their state structure. The collection of such agents—which interact with one another,

and with their environment—constitutes a CAS. The “swarm” of agents then will exhibit various forms for collective behavior, which emerges from the interactions among the microscopic rules. The goal of BISON is to build CAS which give rise to *useful and predictable* collective behavior, which may be exploited to realize all or part of one or more desired network functions. These network functions are also a form of behavior; and it is an important goal of BISON (see D01) to define this behavior as precisely as possible.

Hence we seek precision at both ends of this “emergence equation”. The CAS is well defined at the microscopic level; and the desired functional behavior is also sharply defined. All the uncertainty then lies in the emergence arrow.

Thus the emergence equation offers a useful, albeit abstract, description of the BISON project, seen as a whole. The principal missing ingredient is of course the network on which the dynamic behavior is to take place. That is, the emergence equation focuses on behavior. The behavior of course takes place on a substrate, the network. We argue however in D01 that the network structure is not a passive boundary condition for the behavior. Instead, agents can actively seek to influence or even determine the network topology—which in turn affects the behavior of the agents. Thus, in this view, the network topology is implicitly included in the emergence equation, with its explicit focus on behavior.

One might say that the intent of BISON is to “run” the emergence equation, many times—sometimes, with the right-hand side specified, and the left-hand side to be determined. A further object of BISON is to *invert* the emergence arrow—at least, implicitly and heuristically—because it is the *right* side that is given, for technological problems. Hence BISON will seek, as it generates many different instances of emergent behavior, to extract some coarse forms for regularity in the relation between the left side and the right side. Any such regularity—again, almost certainly heuristic—will aid in the inversion of the arrow.

Note that BISON’s grounding in biology will be extremely useful in this regard. Biology represents very many “runnings” of the emergence equation. And, in many cases, we know much about both the left-hand side (the microscopic mechanisms) and the right-hand side (the global behavior). Hence an understanding of emergence, as displayed by (eg) the four biological systems selected for the BISON project, will provide insight towards the general goal of inverting emergence. We begin this task in the next Section.

4 Fundamental building blocks for agents

Recall from the working definition of the BISON CAS that an essential role is played by the “internal characteristics” of the agents. For BISON, these internal characteristics are the agents’ state structure—which consists in turn of the states of the agents, plus the transitions between the states. Since this structure determines the behavior of the agents, it can be viewed as providing the fundamental building blocks of the CAS.

Agents are assumed to be of different types (A, B, C, . . .). Different types are distinguished by different state structures (states, and transitions between states). We take each agent type to have a finite set of internal states (a, b, c, . . .). Our idea is that in a CAS collective behaviour emerges as a (macroscopic) result of microscopic interactions of simple individual agents. We believe that we can compose new types of CAS by choosing a subset out of a small set of el-

elementary agent building blocks. We are not restricted in our choice, since the selected subset need not correspond to any known biological system. Our goal then in this section is to present a short list of fundamental building blocks for agent behavior. The list of building blocks presented here is not intended to be a complete set; it is preliminary. It will be likely expanded in the course of further work. After introducing this set, we will (in the next Section) analyze selected biological systems, characterizing them in terms of this list of microscopic mechanisms.

In our attempt to construct a list of possible agent transitions, we took inspiration from “biological agents”, in particular biological cells or organisms. We have analyzed what kind of elementary (microscopic) “activities” biological entities as cells (organisms) can perform. We believe that a good starting list (possibly to be extended in the future) of microscopic agent mechanisms is the following:

- (1) $Aa \implies Aa$ (survival, positive selection, memory)
- (2) $Aa \implies 0$ (death, degradation, negative selection)
- (3) $A \implies AA$ (mitosis, proliferation)
- (4) $Aa \implies Ab$ (response, interaction, processing)
- (5) $A \implies AB$ (signalling, diffusive and nondiffusive)
- (6) $A \implies A$ (mobility)
- (7) $A \implies B$ (mutation)

In the next section we will give biological interpretations of these microscopic mechanisms. Here we offer a few, more general, comments on our choice of mechanisms, and on the language used to describe them.

- (1) Here the event is that ‘nothing happens’. This implies that the agent survives. We also say that this absence of change represents memory. The definition of memory can be subtle and nontrivial. However, we believe that this is the case for *collective* forms of memory—those arising from the interaction of many agents. At the microscopic level, in contrast, we interpret ‘memory’ to mean simply the retention of state. For instance, after an agent changes state, due to input from the environment, or from another agent—in fact, for any reason—the retention of that state is its memory. Hence, mechanism (1), in which ‘nothing happens’, actually represents survival of the agent type A , and at the same time its memory (retention) of state a .
- (2) We use a single symbol 0 to represent death of an agent. It may be thought of as a state with no outgoing transitions, and no interaction with the environment; or as disappearance of the agent.
- (3) Proliferation is a mechanism lying outside the realm of behavior of a finite-state machine or FSM. We take it for given that such can occur. The states of the two copies are not specified here; they need not be the same.

- (4) This is simply a state transition of agent type A. Clearly, a great deal of complexity may be built into the state structure of the agents. This complexity is somewhat hidden by the simple notation shown here. That is, an enormous variety of behavior, and range of complexity, is implicit in the simple fact that the agents' behavior is represented by their states and state transitions.
- (5) Here an agent emits a different type of agent (in contrast to proliferation). Again, the states of the two agents A and B are not specified here. The descriptive term 'signalling' is motivated by biological ideas, and may perhaps be replaced with a more general term.
- (6) Here the agent retains its type, but moves to a new location. This behavior, like proliferation and signalling, is outside of the repertoire of a standard FSM. (See Section 6.4 for a discussion of the process calculus—which also seeks to embellish isolated automata with features such as interaction and mobility).
- (7) In our language, mutation is change of state *structure*, rather than change of state. Thus, under mutation, it is the agent type which is changed, rather than its state. As discussed in Section 2 above, we are inclined to take a conservative (ie, risk-averse) approach to evolution; but it may play a role in some successful application of agent-based CAS. Hence we include the microscopic mechanism of mutation here.

Finally we comment on some slight formal redundancies in the set given here. In principle, change of state structure can be modelled by a state transition of a machine with a large state space. That is, the $A \implies B$ transition (7) (change of agent type) may also be viewed as a change of state [mechanism (4), $Ca \implies Cb$], by letting the agent type C have all the state structure of types A and B, plus allowed transitions between them. Let us call such a state structure—embodying, in some reasonable sense, the state structures of more than one type of agent—a 'compound' state structure. If we accept the existence of such structures, then it follows that, since the final states are left unspecified in our proliferation mechanism (3), we could in principle represent the signalling transition (5) as if it were simply proliferation of agents with compound state structure.

In short: if we allow compound state structures, then (4) can represent (7), and (3) can represent (5). Hence, even though our list is quite short, we see that not every mechanism on the list is *logically* necessary. We retain all seven mechanisms as distinct mechanisms, based on our belief that it is most convenient to do so.

5 Microscopic mechanisms for biology-inspired CAS

Now we start with the above list, and from it build up lifelike models of real biological systems. The aim is not to demonstrate completeness of our set, but rather to show how much one can obtain from it.

5.1 Mechanisms

We go again through the list in order; the aim here is to point out the existence and importance of these mechanisms in biology.

- (1) $Aa \implies Aa$ (survival, positive selection, memory). At the agent level this mechanism represents survival [as opposed to (2), death]. Survival is the persistence of organisms, and as such is an important biological mechanism—positive selection. A similar phenomenon occurs on a much shorter time scale for immune cells, which are subject to a fast form of mutation and selection. Survival also in this case implies positive selection. Finally, as argued above, persistence of state is a simple microscopic form for memory. This kind of memory may be thought of in biological terms as stability: life depends on the persistence of certain structures and states, as well as on change of state.
- (2) $A \implies 0$ (death, degradation, negative selection). Biological systems heavily exploit death, at many levels of scale: from cells through whole organisms, species, even ecosystems. Also, we will wish to remove agents when they are not needed. In short, whenever there is proliferation, there must also be death.
- (3) $A \implies AA$ (mitosis, proliferation). Agents can have “children”. Here the children are identical to the parents; and only one parent is needed. This simple form of proliferation is most like mitosis in biology. Proliferation with variation (not represented here) involves mutation, sex, or both. We have mutation on our list. Sexual reproduction is not on our list, as it requires (at least) two different agents in the initial condition. It may of course be added if we find it to be useful.
- (4) $Aa \implies Ab$ (response, interaction, processing). This simple mechanism is nothing more than a state transition. If it occurs as a result of an input from another agent, or from the environment, then this transition represents, respectively, interagent interaction, or sensitivity to the environment. It can also be part of the internal dynamics of the agent. In this latter form it is a type of information processing. Biological examples are obvious: a neuron changes state in response to input from another neuron; an amoeba changes its own level of cAMP output based on detection of cAMP in the environment; or, an amoeba changes its own level of cAMP output based on an internal sense of hunger.
- (5) $A \implies AB$ (signalling, diffusive and nondiffusive). Here a new agent is born, but of a type which differs from that of the parent. Our motivating example is signalling. An agent (ant, immune cell, amoeba) can emit a signal, which takes on a life of its own after emission. That is, the signal then follows its own rules. The signal is thus a different type of agent. If it is mobile then it represents a diffusive signal such as the cAMP emitted by slime mold cells. The ant pheromones are typically nondiffusive.
- (6) $A \implies A$ (mobility). This is the only item on our list which refers to a space in which the agents can move. In biology we think of the space as being the world, or a restricted (say, two-dimensional) version of the world. In some cases (such as ants following trails) the ‘world’ may even be simplified to the point of being a network of nodes and links. It is important however to keep in mind that the space studied in BISON may be quite unlike any space experienced in biology. For instance, overlay networks can form structures

which correspond to a four- or higher-dimensional space, or even structures which do not resemble a D -dimensional space for any finite D .

- (7) $A \implies B$ (mutation, evolution). Adaptation via evolution—a collective phenomenon, characterizing populations—depends on mutation, which is a microscopic mechanism taking place at the level of a single agent. The importance of mutation in biology is obvious. We repeat a point made earlier: for BISON purposes, we expect that ‘online evolution’ is a mechanism to be used with great care. One promising application is ‘fast, safe’ evolution, as seen for example in B-cells, proliferating and mutating rapidly in order to home in on an antigen. Here the mutation is safe as well as fast, since it is confined to a specialized region of the genome.

The above is a translation of the simple, abstract symbolic list of Section 4 into relevant biological behaviors. In the following, we give a breakdown of the behavior (at the individual-agent level) for five biological examples. The analysis is carried out with respect to our fundamental list.

5.2 Analysis of five biological agents

It is helpful to characterize biological systems by their elementary interactions. In the following we focus on five selected biological systems: food search by ants; immune response; the aggregation of slime molds; information processing in a neural network; and epidemic spreading via viruses. In each case we focus on individual “agents”: respectively, ants, immune cells, amoebae, neurons, and viruses. Obviously, we have idealized both the agents and their behavior.

5.2.1 Ants

Ants seek food, and mark trails to/from the food with pheromones. Ants searching for food produce a non-diffusible pheromone, possess an individual memory, and are able to migrate. Also, they are influenced by the pheromone traces left by other ants (4). Thus, this simplified, single-task view of ants’ behavior can be described by elementary activities (1), (4), (5), and (6).

5.2.2 Immune cells

In general, there is a cellular and humoral immune response [8]. Lymphocytes move around in the body (6) and can recognize (foreign) antigens that have similar receptor patterns; i.e. a sufficiently large affinity between antigen and the immune cell triggers (4) an immune response. Immune cells then proliferate (3) and release antibodies (5) that have exactly the receptor pattern that matches the antigen. In addition, there are devices (germinal centers) in which mutated immune cells are produced (7) that might have an even better affinity with the attacking antigen. These cells are subjected to a selection process, so that the less fit cells die (2). And, of course, survival and memory (1) play a crucial role: the fittest cells survive, and retain information about the antigen. Accordingly, the immune response can be characterized as embodying all of the elementary agent actions (1)–(7). It is interesting that later attacks by the same antigen induce a faster immune response. This shows that the immune system possesses a memory

that is probably not present at the individual immune cell level, but is likely instead a *collective* phenomenon arising from the specific coupling of the microscopic agent interactions.

5.2.3 Amoebae

In the slime mold (*Dictyostelium discoideum*) [2] the individual agents are amoebae. Upon starvation, populations of those amoebae switch from independent (autonomous) agent movement to collective behaviour (aggregation, fruiting body formation, sporulation). Amoebae move around (6) and feed. As long as they get enough food, they do not cooperate. Upon starvation (4), some amoebae produce a signal (cyclic AMP, or cAMP) (5). The signal can diffuse, and reaches other amoebae. There is a relay mechanism, i.e. if an amoeba receives a sufficiently strong signal, it will also begin to emit cAMP (4). Additionally, it will move up signal gradients (chemotaxis) (4),(6). Accordingly, the amoebae aggregate (a collective behavior). Aggregation centers serve as nucleation spots for fruiting body formation. The fruiting body can release a large number of spores, which drift over distances much larger than individual amoebae could crawl. When spores land, they can differentiate into amoebae which start the life cycle again. The net result is that the whole colony has dispersed into new regions with possibly better food supply. If we focus only on the aggregation behavior, we find that aggregation is the result of coupling the following elementary agent activities: environmental response (4), signalling (5), and migration (6). Finally, since the amoebae have state, we must give them the ability to retain their state (1). Hence we get the same list as that for food-gathering ants: (1), (4), (5), (6). The fact that the signal here is diffusive is however a significant difference.

5.2.4 Neurons

We have an abstract view of the neuronal system and neglect many details. The agents are the neurons themselves – immobile cells connected in a more or less static network. Neurons communicate via highly localized signals across membranes (synapses). Slightly simplifying, we term this signalling ‘non-diffusive’. There is memory at the microscopic (one-cell) level, in that the neurons have a state structure. (There is also memory in the neural network as a collective effect, which arises from the microscopic mechanisms.) The signalling affects the state of the receiving neuron, and so we have response.

5.2.5 Viruses

Finally we include a simplified version of a virus. The virus is either alive or dead, but otherwise has no state; we treat this rather degenerate form of state structure as stateless. Hence it has no memory. It is however mobile; it can proliferate; and it can die. With these simple mechanisms, a virus can infect a large system, sometimes very efficiently. The death of the virus may be a response to the environment; however, given the minimal state structure assigned to our idealized viruses, we see no reason to invoke mechanism (5) for these agents. Finally, a more powerful type of virus can mutate, and so resist negative selection; we leave that mechanism as an option.

5.2.6 Summary and table

We summarize the above discussion in Table 1, showing how to “build” a CAS that models a biological system from the basic building blocks. We split activity (5) (signaling) into two types—diffusive and nondiffusive—as there are good reasons to believe that this difference has significant effects on the collective behavior.

	Ants	Immune cells	Amoebae	Neurons	Viruses
Memory (1)	x	x	x	x	-
Selection (2)	-	x	-	-	x
Proliferation (3)	-	x	-	-	x
Response/interaction (4)	x	x	x	x	-
Signaling - diffusive (5a)	-	x	x	-	-
Signaling - nondiffusive (5b)	x	-	-	x	-
Mobility (6)	x	x	x	-	x
Mutation (7)	-	x	-	-	?

Table 1: Building a CAS from basic building blocks.

We emphasize that this table represents highly idealized models for real biological systems. However, it is not the goal of BISON to develop detailed, biologically faithful models. Instead we seek to model *behavior*. The claim of the above Table is then that the idealized microscopic behaviors listed in each column will give the respective idealized collective behaviors: finding trails to food (ants), recognizing and resisting antigens (immune cells), aggregating when needed (amoebae), and collective memory (neurons).

One can then hypothesize that the same idealized combinations may be used on networks to accomplish similar goals—for example, routing may be achieved using the ants’ combination. Furthermore, one has the freedom to test out other combinations than the five present here. Clearly, there is a large number of combinations! Hence the above Table does two things: it suggests ways of accomplishing specific functions using a CAS; and it suggests a large number of novel CAS, with as yet unknown collective functionality.

6 Placing the BISON approach in a larger context

For the purpose of comparison with related fields of work, the most appropriate descriptive phrase for the BISON approach is “swarm intelligence”. That is, in our language, swarm intelligence is the emergence of intelligent collective behavior from the activities of a swarm of simple, unintelligent, mobile agents. Sections 2 and 3 of this report describe this approach in some detail. Here we wish to briefly discuss some related approaches, in each case pointing out the similarities and differences with swarm intelligence. Our purpose is to extract some clarity from a rich and messy jungle of related ideas, approaches, and words.

6.1 Mobile agents

It seems intuitively reasonable to describe BISON agents as “mobile agents”. We have in fact explicitly endowed these agents with mobility. Here we test this intuition by comparing BISON agents with common definitions of mobile agents.

First we take the word “agent”. Clearly this word can have many meanings. As noted by Franklin and Graesser [5], the word “agent” can take one of two, rather basic, meanings. It can mean “one who acts”; or it can have the more explicit meaning of “one who acts on behalf of another”. Much work with software agents is directed towards the latter purpose: we humans (or other agents) send agents out into the network to accomplish tasks for us. BISON agents fit only the former, more general definition: they can act—they are “autonomous”. This word in turn lacks a precise definition; but it implies a degree of control beyond that associated with typical programs. An autonomous agent can choose what to do, when to do it, and (if mobile) where to go. Other agents may make “requests” of an autonomous agent, but the latter retains the freedom to execute the request, or not.

Jennings and Wooldridge [9] offer a definition of “intelligent agents” that is very close to the BISON agent definition of Section 2. Their intelligent agents are: responsive to the environment, proactive (roughly, autonomous), and “social” (responsive to other agents). Although, at the current state of common understanding, it is purely a question of semantics, we would drop the label “intelligent” from such agents. After all, ants, amoebae, and immune cells—and even very primitive, abstract (and arguably unintelligent) agents—have all of these properties. In our opinion, a preferred description of such agents would be, “agents which can, collectively, exhibit intelligence”. A better application of the phrase “intelligent agent” is then to those personal agents, acting in a secretary-like fashion on behalf of a human, which are endowed with considerable intelligence for such tasks as searching, evaluating, and negotiating.

Finally, we come to the term “mobile”. Again we find two common but distinct meanings for the phrase “mobile agent”. The minimal definition is that state is transported from node to node, along with protocols which allow the arriving state to initiate a thread of execution at the newly-arrived-at node. That is, the autonomy is mobile, as well as the state; but the executable code is not transported. The alternate definition is then that both code and state are transported. BISON agents will likely conform to the former, minimal definition of mobility. The reason for this is practical: we expect BISON systems to have few distinct types of agents, which are enabled on most or all of the nodes of the network. Hence the least-cost (in bandwidth) solution is likely to involve replicas of the agents’ code at each node.

6.2 Cellular automata

Recalling again the working definition of the BISON CAS, the essential role is played by the agents possessing “internal characteristics” (a state structure) which includes the states of the agents, plus the transitions between the states. Agents are assumed to live on a discrete space (the network). We also assume that they act in discrete time, and have a finite set of internal states (with types A, B, C,...). The state transitions may be viewed as being induced by the rules of a cellular automaton. So our working definition of a CAS matches the definition of a classical CA [4] nearly completely.

Let us review a standard definition. A CA is a 4-tuple $\{L, S, N, F\}$, where

- L is a regular lattice of cells
- S is a finite set of states. Each cell $i \in L$ is assigned a state $s \in S$.
- N is a finite set of neighbors, indicating the position of one cell relative to another cell on the lattice L (N can, and usually does include the reference cell itself).
- F is a function $F : S^{|N|} \rightarrow S, \{s_i\}_{i \in N} \rightarrow s$ which assigns a new state s to a cell depending on the state of all its neighbors, indicated by N .

The evolution of a CA is defined by applying the function F synchronously to all cells of the lattice L .

The key features of a cellular automaton are discreteness of time and space, the finite state set for each cell, and each cell being connected to a finite number of neighbors. Within these constraints, there are many possible variations of the standard CA definition. In probabilistic CA, F is not deterministic (and not a function in the mathematical sense), but probabilistic, i.e. $F : S^{|N|} \rightarrow S, \{s_i\}_{i \in N} \rightarrow s_j$ with probability p_j , where $p_j \geq 0$ and $\sum_j p_j = 1$.

Further variants have been considered in which the neighborhood is not uniform, but different for different cells. If the state space is not finite, one should not speak of CA, but of coupled map lattices [10]. In asynchronous CA, the updating is not performed synchronously but sequentially (e.g. in a fixed or random cell order). And, finally, a CA can be defined on any graph topology—not just on regular lattices.

Taking a CA to be defined in this broader sense (discrete state space, discrete space and time, local update rules) appears to give a definition which includes a CAS composed of BISON agents (FSM) acting on a network. In fact, the definitions of a BISON CAS and of a CA are close; but there are some important differences:

- a. Unlike the typical CA, we are not only interested in CAS operating on a regular lattice. Instead, other network topologies are of great interest (as described in detail in Deliverable D01).
- b. Since BISON agents are mobile, and can proliferate and die, the number of agents at any given node can vary over time. This means that the state space of each node is not constant over time. In fact, it is the *product space* of the state spaces of the agents present at the node at any time. It is presumably possible to regard this complication as a much bigger (but time-independent) state space for each node, which spans all expected combinations of numbers of different types of agents. However such a large state space is so unlike that typically studied in CA research that it represents a qualitative difference.
- c. The BISON “environment” for the agents (ie, the nodes on the network) is not as simple—nor as passive—as that for a CA. That is, we will retain the ability to store and process real numbers at the nodes. The agents in turn can react to this environment. Hence the combined system cannot be viewed as a CA. If we ignore complication b., then we can incorporate the behavioral capability of the nodes by considering the system as a hybrid combination of a CA and a coupled map lattice (on a graph which is not a lattice).

In summary: there are interesting parallels between discrete-state mobile agents on a network, and a CA on the same network. However, there are also nontrivial differences. It is possible that the similarities between the two will offer insights into the behavior of the former; but no literal carryover of results is likely to be possible.

6.3 Artificial life

Artificial life [11] is the study of manmade processes/structures with lifelike properties. Given this broad definition, the CAS studied by BISON are examples of artificial life. The main difference is one of motivation: the bulk of Alife research is aimed at understanding life. BISON, in contrast, is focused on applying abstractions of living systems. This difference is of course not so large, since application demands some understanding, and understanding naturally suggests application.

6.4 Process calculus

Automata theory [7] is the formal theory of the behavior of an isolated process—that is, an automaton, obeying explicit rules, and interacting only with a read/write tape. Process calculus [12] is a formal extension of this theory, to the case where multiple processes are active, and communicating with one another. That is, process calculus is the formal theory of interacting processes. Further extensions allow for the processes to be both interacting and mobile [12].

Thus process calculus promises to give the abstract foundation for a theory of mobile agents. And, just as automata theory offers a thorough and rigorous description of the behavior of isolated processes, process calculus seeks to describe the behavior of multiple, mobile, interacting processes. To quote Milner [12]:

A central question we shall try to answer is: when do two interactive systems have equivalent behavior, in the sense that we can unplug one and plug in the other—in any environment—and not tell the difference? This is a theoretical question, but vitally important in practice. Until we know what constitutes similarity or difference of behavior, we cannot claim to know what ‘behavior’ means—and if that is the case we have no precise way of explaining what our systems do!

Thus the goals of process calculus have much in common with those of BISON: understanding the behavior of abstract processes termed ‘mobile agents’, via precise description. In this context, it is interesting to compare the ‘primitives’ list for the ambient calculus (a form of process calculus, due to Cardelli and Gordon [1]) with the BISON primitives list of Section 4. The Cardelli/Gordon list includes:

- restriction
- inactivity
- composition

- replication
- ambient
- and three types of 'action'.

The first restricts the 'scope' of a name, and in so doing plays a role in describing usable links between processes. In other words, links defined by restriction are used to define in turn the network topology; the latter is not an independent entity in this theory. 'Inactivity' corresponds to our 'death'. Composition allows parallel running of processes, ie, multiple agents in one 'ambient' (site). This concurrency may be the result of mobility (arrival of a new process), or of replication. Replication is equivalent to our proliferation. 'Ambient' (here a noun) defines location. Finally, the three 'actions' support changes in ambient/process relationships, hence mobility.

Thus there are strong similarities between the ambient-calculus primitives list and the BISON primitives list. In spite of this fact, it is not clear that process calculus will be very useful for the BISON project's goals. Process calculus seeks a high level of precision and rigor in its description of interacting systems. The price for this is almost certainly an inability—at least, for the near future—to describe anything like emergence. In fact, it appears that the current level of achievement of process calculus is far from that needed to describe systems of the complexity that BISON will study. Thus, these two lines of investigation (BISON and process calculus), so similar in their starting point, will likely diverge strongly, due to the tension between the goal of *performance* in the BISON case, and the goal of *rigor* in the latter case. Nevertheless, as study of the BISON primitives and their implications continues, connections to the various forms of process calculus should also be studied.

7 Engineering and emergence

The purpose of this last section is to revisit the issue of predictability/unpredictability. This issue was discussed at some length in Section 2. The focus of the present discussion is on *engineering*, with its concomitant demands for predictability. The question is then how to reconcile the kinds of unpredictability associated with a CAS, with the need of engineering practice for reliability. In short: how does one engineer with emergence?


In a sense, the answer is the same answer that has held for as long as humans have been engineers. This answer is: the art of engineering is to *bound* the unpredictability.

This is in fact the whole answer. The remainder of this section is then included simply for clarifying these ideas. In so doing we will also place the BISON approach in a broad context, in which the main themes are predictability and its lack, in science and engineering.


Classical mechanics has been the science underlying the engineering of machines for centuries. In fact (as is often the case) the engineering preceded the explicit articulation of the science. But in any case, the two worked together in a satisfactory manner for centuries, completely unaware of what would happen in the 20th century—namely, that classical mechanics would be *proven* to have an intrinsic unpredictability in its workings.

Here we belabor the obvious, but the point should be clear. The unpredictability was present all along, and engineers were not fazed by it: they held it within acceptable bounds.


This reference to mechanical systems makes the main point; but it is far from a perfect analogy for the BISON approach. In particular, there is no emergence in such systems: their collective behavior is predictable, the behavior of each component making up the whole is also predictable, and finally, the manner in which the collective behavior derives from the behavior of the parts is known—even understood in great detail. In short, the machine is characterized by:



Consider instead a gas in a box. Here the collective behavior is also predictable, via the laws of thermodynamics. Unlike the machine, however, the microscopic behavior is unpredictable. That is, we have:



Here by 'brute parallelism' we mean something like the ideal gas law: that the pressure-volume product is proportional to the number of molecules in the gas. Hence the interaction is trivial (by our definition)—even though there is microscopic unpredictability—and there is no emergence. One can of course engineer with such systems also. But the gas, like the machine, fails to capture the notions of a CAS and emergence. These ideas involve:



Here the first line—microscopic predictability—does not preclude agents acting stochastically; the agents' behavior is determined entirely by design, and any stochastic components have bounded unpredictability. The important difference between this picture, and that for the machine, lies in the intermediate "via" arrow. For the machine, there is little or no uncertainty in that arrow: the behavior is, in principle, predictable (within bounds) at all levels of resolution, from the microscopic to the macroscopic. For the CAS, there is large uncertainty (at the current state of understanding) in the "via" arrow—it is still largely a "black box", which is typically realized in practice by computer simulations. The microscopic level can have known laws; and we *select* those CAS for which the macroscopic behavior also has bounded unpredictability. But

we lack the ability to predict, and sometimes even describe, the behavior at intermediate levels of resolution.

This last picture is the emergence equation, in slightly altered form. It is the aim of BISON to exploit this "equation" as a phenomenon, towards the goal of engineering functions on network systems. Biological systems (not mechanical devices, or thermodynamic systems) are the proof by existence that such is possible. The collective behavior of biological systems is "predictable enough"—to allow survival, reproduction, and evolution. It is then an *assumption* of the BISON project that one can engineer with emergence, such that the collective behavior is "predictable enough". This assumption is supported by the enormous existence proof of biology.

The problem with this approach (as noted above, and in Section 3) is that the laws of emergence are not known. Hence there is an important (and large) element of unpredictability in the emergence arrow itself—in contrast to the arrow for the mechanical system (classical mechanics) and for the gas (thermodynamics). It is *this* unpredictability which is presently far too large for many engineering purposes. Hence a principal goal for the BISON project—within the limited realm of implementing certain functions on networks—is to reduce the uncertainty in the emergence arrow itself. Finally—as has been true historically, with other technologies—the more modest goal of achieving "sufficiently bounded unpredictability", without knowing the laws giving rise to the macroscopic behavior, is also possible and desirable.

References

- [1] Luca Cardelli and Andrew D. Gordon. Mobile ambients. In *Foundations of Software Science and Computation Structures: First International Conference, FOSSACS '98*. Springer-Verlag, Berlin Germany, 1998.
- [2] P. N. Devreotes. *Dictyostelium discoideum*: A model system for cell-cell interactions in development. *Science*, 245:1054, 1989.
- [3] Kevin Dooley. A complex adaptive systems model of organization change. *Nonlinear Dynamics, Psychology, & Life Sciences*, March 1997.
- [4] D. Farmer, T. Toffoli, and S. Wolfram, editors. *Cellular automata: Proceedings of an interdisciplinary workshop*. North-Holland Physics Publ., Amsterdam, 1984.
- [5] S. Franklin and A. Graesser. Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents. In *Intelligent Agents III. Agent Theories, Architectures and Languages (ATAL'96)*, volume 1193, Berlin, Germany, 1996. Springer-Verlag.
- [6] Steven A. Hofmeyr. Introduction to the immune system. In *Design principles for the Immune System and Other Distributed Autonomous Systems*, pages 3–26, New York, July 1999.
- [7] J.E. Hopcroft and J.D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, Reading, Massachusetts, 1979.
- [8] C. Janeway, P. Travers, J. D. Capra, and M. J. Walport, editors. *Immunobiology: the immune system in health and disease*. Garland Pub., 1999.

- [9] Nicholas R. Jennings and Michael J. Wooldridge. Applications of intelligent agents. In Nicholas R. Jennings and Michael J. Wooldridge, editors, *Agent Technology: Foundations, Applications, and Markets*, pages 3–28. Springer-Verlag: Heidelberg, Germany, 1998.
- [10] R. Kapral. Chemical waves and coupled map lattices. In K. Kaneko, editor, *Theory and applications of coupled map lattices*, pages 135–168. Wiley, West Sussex, UK, 1993.
- [11] Christopher A. Langton, editor. *Artificial Life III*. Addison-Wesley, Reading, Massachusetts, 1994.
- [12] Robin Milner. *Communicating and Mobile Systems: the π -calculus*. Cambridge University Press, Cambridge, UK, 1999.
- [13] David Pines. Designing a university for the millenium: A Santa Fe Institute perspective. Technical Report 98-09-083, Santa Fe Institute, Santa Fe, New Mexico, USA, March 1998.
- [14] David H. Wolpert and Kagan Turner. An introduction to collective intelligence. Technical Report NASA-ARC-IC-99-63, NASA, Ames, Iowa, USA, March 2000.