Project Number 001907

DELIS

Dynamically Evolving, Large-scale Information Systems

Integrated Project

Member of the FET Proactive Initiative **Complex Systems**

Deliverable D5.1.1

Desirable Lifelike Properties in Large-scale Information Systems

Start date of the project:	January 2004
Duration:	48 months
Project Coordinator:	Prof. Dr. math. Friedhelm Meyer auf der Heide Heinz Nixdorf Institute, University of Paderborn, Germany
Due date of deliverable:	December 2005
Actual submission date:	January 2006
Dissemination level:	PU – public
Work Package 5.1:	Novel Biological Metaphors for Information Systems
Participants:	Telenor R&D (Telenor), Norway Universita di Bologna (UniBO), Italy Universitat Pompeu Fabra (UPF), Spain

Authors of deliverable:	Geoffrey Canright (geoffrey.canright@telenor.com) David Hales (hales@cs.unibo.it) Ozalp Babaoglu (babaoglu@cs.unibo.it) Ricard V. Solé (ricard.sole@upf.edu) Sergi Valverde (svalverde@imim.es)

Abstract

This report comprises the complete D5.1.1 deliverable as specified for workpackage WP5.1 in Subproject SP5 of the DELIS (Dynamically Evolving Large-scale Information Systems) Integrated Project.

The essential goal of the DELIS project is to understand, predict, engineer and control large evolving information systems. It is becoming widely recognised that large-scale distributed information systems could benefit from certain desirable properties found within naturally occurring biological (living) systems. Examples of these properties are: self-organization, adaptivity, robustness, scalability, self-repair, etc. In this deliverable, we give an overview of living systems, in terms of both the *phenomena* of life and the *properties* exhibited by life. Both the phenomena and the properties, we find, are of interest to engineers wishing to build systems with lifelike properties[1].

[1]Most papers produced within DELIS are available from the DELIS website as DELIS Technical Reports. Where this is the case references are appended with the DELIS Tech Report number in square brackets. This indicates the paper was produced within the DELIS project, not some other project.

Contents

1 Introduction

It is becoming widely recognised that large-scale distributed information systems could benefit from certain desirable properties found within naturally occurring biological (living) systems. Examples of these properties are: self-organization, adaptivity, robustness, scalability, self-repair, etc. In this deliverable, we give an overview of living systems, in terms of both the *phenomena* of life and the *properties* exhibited by life.

We will look at the phenomena of life in a rather general way. The question, What is life? is of course not new; nor do we propose here to give a definitive answer. Instead our aim is more modest: we are seeking practical answers - and the demands of practicality (engineering) are somewhat different from those of pure science: rigor and finality are not absolutely required. Instead what is wanted is useful ideas for how to do things.

2 What is life?

2.1 Introduction

We work from the basic assumption that there is no fundamental difference between living and nonliving things. In other words, we reject any traces of vitalism [14] in our discussion. This position implies in turn that there is not a sharp boundary between life and nonlife: the definition of life cannot be as sharp as that for (say) a carbon atom. Instead, life will be viewed here as a collection of lifelike phenomena, which in turn exhibit lifelike properties [2]. In this section we present the phenomena.

2.2 Phenomena

Figure 1 presents a set of lifelike phenomena, presented in the form of a Venn diagram. As seen from the figure, we view these phenomena as having considerable overlap. Yet, at the same time, there is a degree of hierarchy in our picture.

We start with the two "largest" concepts. We say that life is characterized by two things: (1) persistence, and (2) evolution.

We comment on this choice. One might expect the fact that life is characterized by nonequilibrium phenomena to be a fundamental aspect. And in fact it is. However, equilibrium phenomena do not evolve. Hence we feel we have captured the aspect of being out of equilibrium in our term evolution. Furthermore, the two terms complement each other well. Very short-lived transient phenomena, which do not manage to persist over long time via self-maintenance and/or reproduction, will also not evolve. Hence we argue that persistence is needed for evolution. On the other hand, life on Earth has survived in the face of a dynamic and non-equilibrium environment. This means that, in general, the ability to evolve is also a precondition for persistence: crudely stated, the choice is change or die. Thus persistence and evolution are highly interdependent; and together they capture much of what we mean by life.

2.2.1 Persistence

We have placed three phenomena entirely within the realm of the phenomenon persistence. These three are: Homeostasis, regeneration, and development.

Homeostasis has of course to do with stability. Life is patterns that maintain themselves—often within narrow limits—in spite of environmental variation. Homeostasis is then this phenomenon of stable maintenance.

Within the concept of homeostasis we have placed several other phenomena. *Metabolism* is of course needed to maintain the form; here the idea is straightforward. Also, we have placed self/nonself

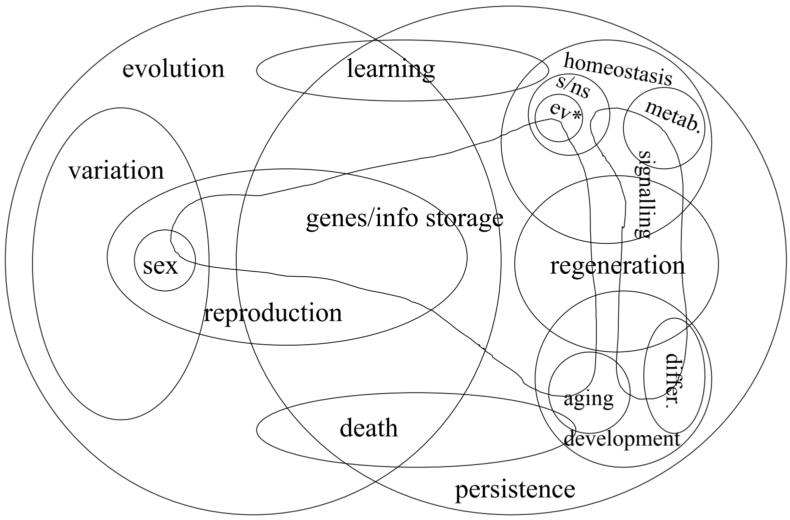

Figure 1: Lifelike phenomena, presented in the form of a Venn diagram. We view these phenomena as having considerable overlap. Yet, at the same time, there is a degree of hierarchy in our picture.

("*s/ns*" in figure 1) as a supportive phenomenon for homeostasis. This is of course a necessary function: the living organism must define and maintain a boundary between self and nonself—and the definition must be one that the organisms itself is capable of recognizing.

Inside the self/nonself function, we have placed the special form of evolution (*ev** in the figure) that is performed by immune systems [4]. That is, immune systems, when presented with a novel antigen, trigger an evolutionary process, in which the evolving populations are antibodies, and for which the evolutionary (mutation) rate is orders of magnitude higher than for other populations. This form of evolution is so different from "normal" evolution - both in purpose and in characteristics - that we have given it no overlap with the latter.

Regeneration has much in common with homeostasis. However the concept of regeneration presupposes that the organism or pattern has significantly deviated from its target state, and so must recreate or regenerate that state. Like homeostasis, this is a conservative mechanism—one that resists changes in form or pattern.

Development is a phenomenon of multicelled organisms. The cell is a primary modular unit of life (see *modularity* below); and one cell can become two by the process of mitosis. However, reproduction of multicelled organisms involves a process of development: the "copy" does not emerge full-blown from the "parent".

We have symbolized the strong similarities between development and regeneration by giving the two concepts an overlap in the figure.

Within the notion of development we have two further concepts. *Differentiation* is a crucial phenomenon for development: the multicelled, fully-developed organism is composed of many different cell and tissue types, which have emerged from the starting egg cell via differentiation. Also, we find that *aging* almost inevitably accompanies development: multicelled organisms age, and die.

Finally, we have placed the phenomenon of *signalling* in a rather unique position in the persistence bubble. Signalling (typically by use of chemicals, ie, molecules) is used in many ways in living organisms: to adjust metabolic rates; to mediate the immune response; and to guide the formation of patterns in the processes of differentiation, development, and regeneration.

2.2.2 Evolution

Here we use a very standard definition of evolution: change of form over a long time scale (many generations), mediated by genetic changes. Nevertheless we believe that this definition is not so narrow that it excludes application to engineered systems.

A principal mechanism for evolution is of course the phenomenon of *variation*. That is, without variation (specifically, of the genome), there is no evolution. The two obvious mechanisms for variation are then mutation (not shown explicitly) and sex. *Sex* is an efficient mechanism for mixing genetic material and so exploring the space of possible genomes. Sex is not however viewed here as having any necessary connection to persistence, since asexual reproduction is possible.

2.2.3 "Bridging" phenomena

We have shown four phenomena which bridge the two main concepts of persistence and evolution.

Reproduction is clearly a persistence mechanism of multicelled and single-celled organisms. However it also supports variation, since (i) natural selection works on reproducing organisms, and (ii) reproduction (via sex, and/or with mutation) gives genetic variation.

Genes are structures; the phenomenon or function which they represent is *information storage*. Genes support variation via mutation and sex; they are passed on during reproduction; they steer development, regeneration, and (likely) aging; their information handling properties are crucial for storing the self/nonself distinction; and genes are the medium in which fast evolution takes place during response to new antigens.

Death is the failure of persistence. However, death of the individual organism supports stability (persistence) of the population or species. In a similar way, death of cells supports persistence of the organism's form. Also, death plays a necessary role in evolution via selection.

Learning is included here to express the ability of some living organisms to find new responses to the environment without recourse to genetic variation. Learning with cultural transmission then becomes a form of evolution. Learning can also serve to support persistence of an organism or species. We will say much more about this aspect of life in Section 4 on human social systems.

3 Abstract aspects of life

Figure 1, and the accompanying discussion, constitute our answer to the question of what is life.

Now we recall a principal goal of SP5 of the DELIS project: we wish to study the problem of biology-inspired engineering. The motivation for doing so is clear: there are obvious advantages to be expected from being able to build (or grow) engineered systems with lifelike properties - for example, the ability to self-repair, to reproduce, or to learn [23]. Along with these potential advantages there are of course potential problems: lifelike systems may not behave in the way we expect, and they may not be easy to kill.

The discussion of the previous section in any event gives a nice menu of possible features for engineered systems. We believe that all of these phenomena can be engineered. Furthermore, engineering such phenomena with information systems is in most or all cases more easily - even much more easily - accomplished than with hardware. That is, we expect that implementing/engineering the phenomena and properties of life will be in general easier when only the informational aspects are included, while the more difficult material/physical aspects are not. (One need only consider for example the difference in challenge between implementing self-assembling overlay networks—eg, peer-to-peer systems—and designing and building self-assembling robots.)

In this section we present a short list of more general properties of life. We view these as being more abstract than the previous phenomena/ functions of Section 2—although we readily admit that the abstract/concrete distinction is not completely sharp. We group these abstract aspects into two categories: organizational principles, and "properties for success".

3.1 Organizational principles

We have selected three organizational principles which characterize living systems: modularity, hierarchy, and self-organization.

3.1.1 Modularity

Creative reuse of inventions is a generic aspect of living systems [20, 8]. It arises because living systems explore the space of possibilities by hopping from known solutions—and the hops (via mutation and/or sexual recombination) are typically small. The point is clear however: modules which prove to be useful are reused again and again. The outstanding example of course is the cell; but many others exist. In fact, reuse of known solutions is ubiquitous in life.

Furthermore, it is not necessary in this picture that the use to which the known "solution" is put is the same in the reuse as it was in the earlier application [8]. That is, modules can be reused, but with a differing function than before. This is true at all scales, from molecules through cells to whole organisms.

3.1.2 Hierarchy

It is a universal challenge to organize many subcomponents into a working whole. It seems clear that both engineered systems and living systems have made heavy use of hierarchical structures in order to handle this problem. If we quantify the size of multicelled organisms in terms of number of cells, it is clear that as the cell number grows, both specialization and hierarchy increase. Decentralized solutions simply fail when the cell number grows too large. The same principle may be observed in Internet routing systems [13], or in peer-to-peer file-sharing systems [18].

3.1.3 Self-organization

This property of living systems is self-evident: living systems achieve a high level of organization without external planning or control. However, this property is of great technological interest [16], since the vast majority of engineered systems have little or no degree of self-organization. Furthermore, the small minority of engineered systems that are self-organizing are proving the technological potential of the idea daily. Here we have in mind peer-to-peer systems. For example, the peer-to-peer telephone system Skype [19] builds a working Internet telephone system with only a low level of central control (mostly for authentication). In particular, the 'white pages' of the Skype system appear to be entirely self-organized.

In short: self-organization is (i) obvious for living systems, (ii) both possible, and highly interesting, for engineered systems.

3.2 Properties for success

Our final short list presents generic properties which are abstract, as in the previous section, but yet which have no reference to modes of organization. Instead, these three properties—adaptation, robustness, and scalability—describe aspects of living systems which enable them to 'succeed' in the game of life. Since there is a notion of success in these properties, they have direct relevance also for technological systems: they are desirable properties for engineered systems.

3.2.1 Adaptation

We have already implicitly described adaptation - at least partially - by presenting evolution and learning in Section 2. That is, systems which can evolve and/or learn can adapt. However, to fully define the concept of adaptation, one needs the further concept of success: the structure or behaviour which is learned or evolved must help increase the organism's (or species') likelihood of survival and reproduction; otherwise it is a change, but it is not adaptive.

Thus adaptation is change in the 'right direction' - towards increased survivability or functionality.

Technological systems with feedback are also adaptive. This kind of adaptation however involves neither evolution nor learning. Such low-level adaptation occurs of course also in living systems; it is in fact an aspect of homeostasis - which itself is a completely valid concept for engineered systems, for example those with feedback control loops and setpoints [10].

The technological frontier thus lies at the 'higher' level of adaptation: can a system which is designed for a given environment learn (or evolve) to be able to handle an environment outside the bounds of its original design? Such systems would of course be highly desirable.

We note in this context that both evolution and learning involve persistent change; and the change is stored as information in each case: in the genome in the case of evolution, and in nervous systems, books etc in the case of learning. The distinction between the two is primarily thus a difference in mechanism, and (thereby) in time scale. Coming to some hypothetical engineered system which adapts well to dynamic environments, we may have difficulty in assigning this adaptation to one or the other of these two concepts - evolution or learning. The adaptation will be clear from the behaviour of the system. But the mechanism, even if known in detail, may not readily fall unambiguously into one of these two categories. The point here is that both kinds of biological adaptation involve some form of information storage for the adaptive change. The two forms (genes and brains) are very different for biological systems; but the difference will not necessarily be clearcut for an adaptive engineered system.

3.2.2 Robustness

The Merriam-Webster Online Dictionary defines "robust" as "having or exhibiting strength or vigorous health". The key concept here is that a robust system can function well even in the face of some kind of *challenge* to its functioning: low temperature, a bacterial invasion, loss of a limb, etc. Clearly, a good homeostatic system - in particular the self/nonself system - and a good capacity for regeneration, will contribute to a robust system - biological or not. But, as with adaptation, robustness may be measured by functioning or behaviour, without necessarily any consideration of the mechanism or mechanisms responsible for the robust behaviour.

The notion of "challenge" is however somewhat problematic. If an organism finds a way to handle low temperature, for instance by modifying its behaviour, then it is adaptive by our previous definition; and it is robust also, if we include temperature variations in our definition of challenges.

The clearest difference between robustness and adaptation, at least as they are defined here, is simply that adaptation (by definition) involves a change in the adaptive system; while robustness simply requires that the system function well in spite of some environmental challenge. Typically there is no reference to change in the context of robustness; but on the other hand, it is not obvious that the definition of robustness is more useful if one excludes the possibility of change, focusing instead exclusively on the aspect of restoring of the pre-challenge form as well as function. If one does not do so, then there is clear overlap between the concepts of robustness and adaptation.

3.2.3 Scalability

The notion of scalability is popular in networked systems. Here the idea is clear: a network system has some number N of participating nodes; and a scalable system is one that functions well, even when the number of nodes N grows large. Typically (as noted above) such systems resort to the use of hierarchical structure, in order to cope with growing N.

Biological systems have also exhibited scalability. Living systems have found ways to function well, even as the number of cells in a multicelled organism grows very large [21] - and (with somewhat smaller large numbers) as the number of organisms in a cooperating society grows large [22].

Thus the concept is clear, as is the application to at least some kinds of engineered systems.

4 Social systems and life: differences and similarities

Much of the proceeding discussion identifies properties that hold on many scales, from single, simple biological organisms up to large societies of complex, learning biological systems. However, it is useful to ask the question as to in what way, if at all, the *human* social level of organisation differs from the non-human biological level. In this section we briefly consider this question with regard to desirable properties for information systems.

We will focus here on human social systems (HSS), although some animal systems have similar properties [2]. An important aspect of HSS is their ability (like biological systems) to both preserve structures - with organisations and institutions persisting over time - and adapt to changing environments and needs. What is of particular interest here is that fact that HSS can maintain their structures for thousands of years without any simple physical process of inheritance, yet also change very quickly. The evolution of HSS is not based on DNA, but rather on a complex interplay between behaviour, learning, and individual goals. Thus, in short, this section may be regarded as a considerable expansion of the simple "learning" bubble in Figure 1. We will present this expansion in the form of a list of phenomena characterising human social systems. Since we take an evolutionary view of social systems most of our points are a synthesis from key texts [3, 15].

4.1 Rapid Change

A feature of HSS is the speed at which re-organisations can occur. Revolutions in social organisation can take place within the life-time of a single individual. Hence although HSS often show stable patterns over long periods, rapid change is possible. The ability to respond rapidly would appear to be a desirable property in rapidly changing information system environments; however, for engineering purposes, one must ensure that such fast changes (unlike revolutions!) can be both predicted and controlled.

4.2 Non-Darwinian Evolution

HSS do not evolve in a Darwinian fashion. Cultural and social evolution is not mediated by random mutations and selection of some base encoder over vast time periods, but rather follows a kind of collective learning process - applied to shared knowledge. That is, the information storage media supporting the change by learning—and hence (as noted above), both the *mechanisms* for change and their *time scale*—are very different from those of Darwinian evolution. Individuals within HSS can learn both directly from their own parents (vertical transmission), from other members of the older generation (diagonal transmission), or from their peers (horizontal transmission). Hence, new cultural traits (behaviours, beliefs, skills) can be propagated quickly through a HSS. This can be contrasted with simple Darwinian transmission in which, typically, only vertical transmission of genetic information is possible (although, in the case of bacteria, horizontal transmission is also possible [17]). Although it is possible to characterise certain processes of cultural evolution based on the fitness of cultural replicators [3] (identifiable cultural traits - or memes [7]) it is important to realise such replicators are not physical - like DNA - but part of a socio-cognitive process - passing through human minds - and may follow many kinds of selective process [15].

4.3 Stable Under Internal Conflict

HSS exist because individuals need others to achieve their aims and goals. Production in all HSS is collective, involving some specialisation of roles. In large modern and post-modern HSS roles are highly specialised, requiring large and complex co-ordination and redistribution methods. However, although HSS may sometime appear well integrated, they also embody inherent conflicts and tensions between individual and group goals. What may be in the interests of one individual or group may be in direct opposition to another. Hence, HSS embody and mediate conflict on many levels.

This aspect of HSS is of course present in other animal social systems as well. However, the high complexity of HSS supports in turn a highly complex and dynamic pattern of internal stresses and conflicts.

4.4 Only Partial Views and Controversy

Although HSS are composed of goal directed intelligent agents, there is little evidence that individuals or groups within them have a full view or understanding of the HSS. Each individual tends to have a partial view (often the result of specialisation within, and complexity of, the system). Such partial views, often dictated by immediate material goals, may have a normative (how things "should" be) character rather than a more scientific descriptive one (how things "are"). Consequently, the ideas that circulate within HSS concerning the HSS itself tend to take on an "ideological" form. This ability to see what should be, rather than what is, is perhaps an ability in which humans excel over all other species—and furthermore, one which affects our ability to study ourselves. Given this, social theories are rarely as consensual as those from biological sciences. Thus, social theories include a high degree of controversy, and they lack the generally accepted foundational structure found in our understanding of biology. However, from an information systems perspective, such controversy is not problematic: we do not care if a given social theory is true for HSS or not; we only care if the ideas and mechanisms in the theory can be usefully applied in information systems. (This last point, of course, also holds for controversial theories of biological systems in general.)

4.5 Trust and Socially Beneficial Norms

In trying to understand the stability of socially functional behaviour, much work within the social sciences has focused on the formation and fixation of "norms" of behaviour. Many researchers working with Multi-Agent Systems (MAS) have attempted to create artificial versions of norms to regulate MAS behaviours - although much of these have not been based on theories from HSS (although see [6]). Certainly the establishment and stability of beneficial norms (such as not cheating one's neighbour) is a desirable property visible in all stable HSS [12]. This point (the existence and power of norms) is of course closely related to the previous point, which notes that norms can influence understanding and perception.

It is widely agreed that, in HSS, many observed behaviours do not follow the same pattern as would be expected from simple Darwinian evolution or individual "rational" behaviour - in the sense of maximising the chance of achieving individual goals. Behaviour is often more socially beneficial and co-operative or altruistic, generally directed toward the good of the group or organisation within which the individual is embedded. (We note the widespread appearance of altruistic behavior among many species of social mammals—such that, once again, we speak here of a difference in degree between HSS and other social animals.) Many theories and mechanisms have been proposed by social scientists for this kind of behaviour [1], with many of these formalised as computer algorithms; furthermore, several of these have already been translated for use in information systems [5, 11].

4.6 Generalised Exchange and Economics

Almost all HSS evidence some kind of generalised exchange mechanisms (GEM) - i.e. some kind of money. The emergence of GEM allows for co-ordination through trade and markets. That, is, collective co-ordination can occur where individual entities (individuals or firms) behave to achieve their own goals. It is an open (and perhaps overly simplified) question whether certain norms are required to support GEM or, rather, most norms are created via economic behaviour within GEM [9]. Certainly, the formation and maintenance of GEM would be an essential feature of any self-organised economic behaviour within information systems - currently many information systems work by assuming an existing GEM *a priori* - i.e. they are parasitic on HSS supplying the trust and norms required (such systems require trusted and centralised nodes before they can operate - they do not emerge such nodes in on-going interaction).

However, given that GEM exist, a huge amount of economic theory, including new evolutionary economics and game theory, can be applied to information systems. However, this specific focus is beyond the scope of this deliverable or DELIS SP5 generally. DELIS SP4 focuses on economics approaches.

5 Desirable lifelike properties in large-scale information systems

The above discussion presents a picture of life. Our overview spans from the most primitive, to the (arguably) most complex: human social systems. Given our rejection of vitalism, the picture we present here is, we claim, a picture of what kinds of phenomena and properties are possible in both living and nonliving systems. That is, we claim that all of the phenomena on our list may be implemented in technological systems, and in particular in large-scale information systems. In this final section, then, we ask, which of these properties are desirable?

We again break our discussion into two: first we address general living systems, and then specifically HSS.

5.1 Desirable properties from general living systems

When we consider our list of phenomena and abstract proeprties of life, and ask which of them are desirable for engineered information systems, two obvious short answers come to mind: (i) all of them; (ii) those properties called “properties for success”.

Each of these answers may in fact be defended. For example, in defense of the first answer, one may simply refer to Figure 1: all of these mechanisms, whose utility has been proven over millions of years of evolution, should be included in the engineer’s toolbox. This simple answer of course dodges the harder question, concerning the likelihood that a given mechanism or phenomenon can in fact be realized in a useful, practical, and (especially) *predictable* way in an engineered information system. However this question is in general quite hard, and is in any case beyond the scope of this deliverable. Hence we retain short answer (i): all of these mechanisms and properties are desirable.

Short answer (ii) has its own obvious logic. That is, the overriding important issue for engineered systems is that they should function well. And the property of functioning “well” is captured explicitly only in the three properties of Section 3.2. Therefore, one can argue that the other items on the list (mechanisms, and organizational aspects) are not desirable in themselves - only as a means to the end of good functioning.

This argument in fact then brings us back to the first one: if all these other items are in fact likely to be useful as a means to the end of good functioning, then they too are desirable - even if not as an end in themselves. Thus we conclude with our short answer (i): all of the phenomena/mechanisms in Section 2, and the organizational principles and properties for success in Section 3, are desirable objects for implementation in large-scale information systems.

In particular, everything on the 'persistence' side of figure 1 is not only of interest, but also relatively safe: building a persistent system, but one without the possibility of evolving, gives us a relatively predictable system. Such systems are safe - assuming that we can kill them, if such is deemed desirable. There is of course some risk that some such systems will be so successfully persistent, robust etc that we will not be able to kill them. This risk is real - just as it is with biological systems - and it must be addressed. We cite two obvious examples, which are already "living" on the Internet, and which many people wish to "kill": computer viruses, and peer-to-peer file-sharing systems.

On the other side of figure 1 we find the ability to take new forms of structure and/or behaviour in a response to a changing environment. This ability is also of interest for information systems; it is part of the vision of building self-managing, adaptive systems. However it is clear that the left side of the figure involves an even higher level of risk than the right side: evolving systems are less predictable; their behaviour is in fact guaranteed to be less predictable than the same system without the ability to evolve.

Furthermore, biological evolution has a more forgiving criterion for success than does technological evolution: in biology, persistence = success. Therefore, in order to seriously consider the use of evolving systems in technology, we must find reliable ways of implementing *technological* selection - ie, of killing effectively those trials which fail to meet our specifications. To date, this has been done, but only in offline systems. In fact, we regard the moving of evolving systems online as an outstanding technical challenge - one to be attempted slowly and with care.

We note in this regard that the one biological example of 'online evolution' which is closest to what we envision for distributed information systems is the one termed 'ev*' in Figure 1. Here [4], immune cells undergo rapid evolution, with selection pressure implemented in two ways: first, in the 'training' period of the immune system, cells which recognize self as nonself (giving thus false positives) are killed. Second, in the 'active' period of fighting an infection, cells whose antibodies show a good match to the antigen(s) invovled are rewarded by being triggered into a proliferative state—they have more 'children'. Thus, study of the immune system and its ev* mechanisms promises to give useful insight from biology towards the problem of designing safe, evolving, online digital organisms.

Our comments with regard to evolution are also applicable to systems that can learn. In each case there is an inevitable element of trial and error; and in each case, some feedback must be given, which enhances some choices and suppresses others. The feedback for evolving systems is selection, implemented via a kill mechanism. For learning systems, the feedback comes from some kind of reinforcement loop, with an evaluation of the behaviour built into the loop. In both cases - especially with large, distributed systems - one can imagine situations in which the feedback is not fast enough to prevent some unforeseen, "stupid", even catastrophic mutation or behaviour. Also, one can imagine situations in which the feedback loop is broken, because the kill mechanism or reinforcement mechanism somehow fails - while the ability to try new solutions remains. We imagine that the case of *cancer* is such a case: a destructive variation of a cell type, coupled with broken feedback—in particular, inadequate or lacking kill mechanism.

Finally we come to the more abstract properties of Section 3. Here we find less reason for concern. Specifically: we are confident that modularity and hierarchy will be useful (as they have already been in other contexts) in the context of large-scale information systems. Furthermore, the property of self-organization is clearly desirable: we have already argued for this point in the context of peer-to-peer systems such as Skype. Finally, our 'properties for success' - as noted by our short answer (ii) - are clearly desirable. In particular, we include scalability, which is (as noted above) particularly relevant for large-scale networked systems.

5.2 Desirable properties from human social systems

Here we discuss the properties of Section 4, but not in the order used in that Section.

We begin in fact with a property that appears to be *un*desirable. That is, we consider the notion of "partial views". It is of course a common aspect of self-organized systems that local 'agents' have only a partial view of the universe in which they live and act. This aspect of what we might call 'imperfect knowledge' is not particularly desirable or undesirable; rather, it is a totally acceptable byproduct of the phenomenon of self-organization. However, we have pointed out another type of imperfect knowledge, namely, the ability of humans to 'know' things which are not true. This is a byproduct of the (desirable) phenomenon of social norms—but we see no redeeming features in this byproduct. It is of course possible that a deeper understanding of human social phenomena will reveal, or at least suggest, benefits from such 'false knowledge'. Here however we deem this aspect of imperfect knowledge to be undesirable; while, again, we regard simply incomplete but correct knowledge as an acceptable byproduct.

Next we note that the phenomenon of social norms is in itself a desirable phenomenon. That is, we wish to be able to build information systems in which the agents may be constrained to act in ways that benefit the system as a whole. This will in general give rise—as noted earlier—to behavior which may be described as altruistic—another desirable property.

Connected with these two phenomena is the issue of trust. Trust is needed in those situations in which not all agents can be assumed to follow the same norms (ie, the same protocols). This uncertainty in fact is valid for virtually all situations in which anything of value is involved: cheating is a ubiquitous, norm-violating phenomenon. Thus, given the undeniable existence of cheating (ie, of failures of the basic norm mechanisms), one needs trust-building mechanisms in order to establish communities of agents with (more or less) consistent norms. Thus we add trust mechanisms to our list of desirable phenomena.

Social norms, altruism, and trust have all received considerable attention in the context of Internet information systems. Thus, their desirability has already been implicitly or explicitly recognized in the work of many researchers.

In the same spirit, 'digital money' has been recognized as a useful tool—both in research and in real-world implementation. We note that *all* of these related mechanisms—norms, altruism, trust, and money—are subject to cheating. However, implementatons of these mechanisms which are robust against cheating are viewed to be highly desirable. That is, the principal downside with these phenomena is that they may be broken or defeated by deliberate attack.

Now we come to phenomena which are more two-edged: namely, the non-Darwinian (ie, non-genetic) nature of cultural evolution, including more specifically its ability to use a variety of transmission mechanisms, as well as its ability to undergo very rapid change. We view each of these aspects of cultural evolution as being simultaneously desirable and dangerous, with regard to the good performance of engineered information systems.

The danger in each case comes from the problem of prediction. The ability to predict the engineered system lies at the core of every form of engineering practice; and we see a danger of *un*predictability connected with each of these two aspects of cultural evolution.

Consider first the possibility of multiple transmission mechanisms. We have already argued, in the previous subsection, that online evolution—even in the limited form of 'genetic' mutations—confers a degree of unpredictability that is problematic. Clearly, increasing the scope and rapidity of genetic/cultural change, beyond this limited form, will only increase the unpredictability of the resulting system. Both the example of bacteria with horizontal transfer (relatively slow, but unpredictable), and the example of human cultural evolution (fast, and unpredictable) support this statement. In other words: a society of agents with a considerable variation in protocols, who can then exchange, nondeterministically, pieces of their protocols with one another—either by 'direct' exchange, or by leaving written traces which play the role of 'books' in their culture—is likely to be as interesting and unpredictable as is a society of humans.

Now we consider the accompanying phenomenon of rapid change. As noted earlier, rapid change can be highly desirable, *if it is adaptive.* And so we come back to the same, fundamental problem:

how to build systems which (i) can change so as to adapt to new and even unforeseen environments, while at the same time (ii) can be *predicted* to change in a way that is viewed by the engineer as positively adaptive. We regard this problem as highly nontrivial.

Another view on this same problem comes from noting that many of the rapid changes that are observed in HSS may be best described as *instabilities*. That is: they occur quickly, they are difficult or impossible to predict or control, and they are in fact desired by few if any of the participants in the change. In this language, we would say that engineering practice involves steering systems such that they avoid those regions of their phase space in which instabilities arise. And the problem is then to build artificial information systems (perhaps including coupling to living humans) which have all the richness, adaptivity, and susceptibiltiy to rapid change that we see in real HSS—while at the same time keeping such systems away from the (inevitable) instabilities.

References

[1] Axelrod, R. (1984) *The evolution of cooperation.* N.Y.: Basic Books.

[2] Bonner, J.T. (2002) *The ideas of biology.* Minneola, N.Y.: Dover.

[3] Boyd, R. and Richerson, P. (1985). *Culture and the Evolutionary Process.* University of Chicago Press, Chicago.

[4] de Castro, L.N. and Timmis, J. (2002) *Artificial Immune Systems: A new computational intelligence approach.* Berlin: Springer Verlag.

[5] Cohen, B. (2003) Incentives Build Robustness in BitTorrent. Presented at the *1st Workshop on the Economics of Peer-2-Peer Systems*, June 5-6, 2003, Berkley, CA. Available at: http://www.sims.berkeley.edu/research/conferences/p2pecon/

[6] Conte, R. and Paolucci, M. (2002) *Reputation in Artificial Societies. Social Beliefs for Social Order*, Kluwer.

[7] Dawkins, R. (1976). *The Selfish Gene.* Oxford University Press, Oxford.

[8] Edelman, G. M. and Gally, J. A. (2001) Degeneracy and complexity in biological systems. *Proc. Nat. Ac. Sc. 98(24):13763-13768.*

[9] Edmonds, B and Hales, D. (2005) Computational Simulation as Theoretical Experiment. *Journal of Mathematical Sociology 29(3):209-232* [DELIS-TR-0122]

[10] El-Samad, H., Goff, J. P., and Khammash, M. (2002) Calcium homeostasis and parturient hypocalcemia: An integral control perspective. *J. theor. Biol. 214:17-29.*

[11] Hales, D. and Edmonds, B. (2005) Applying a socially-inspired technique (tags) to improve cooperation in P2P Networks. *IEEE Transactions in Systems, Man and Cybernetics - Part A: Systems and Humans, 35(3):385-395.* [DELIS-TR-0111]

[12] Hales, D. (2002) Group Reputation Supports Beneficent Norms. *The Journal of Artificial Societies and Social Simulation (JASSS) vol. 5, no. 4.* http://www.soc.surrey.ac.uk/JASSS/5/4/4.html

[13] Huitema, C. (1995) *Routing in the Internet.* Prentice-Hall, Inc.

[14] Graeme K. Hunter. (2000) *Vital Forces: the discovery of the molecular basis of life.* Academic Press.

[15] Lumsden, C. and Wilson, E. (1981). *Genes, Mind and Culture.* Harvard University Press, London.

[16] di Marzo Serugendo, G., Karageorgos, A., Rana, O. F. and Zambonelli, F. (eds). (2004)*Engineering Self-organising Systems: nature-inspired approaches to software engineering.* Berlin: Springer.

[17] Ochman, H., Lerat, E. and Daubin, V. (2005) Examining bacterial species under the specter of gene transfer and exchange. *Proc. Nat. Ac. Science 102:6595-6599.*

[18] Oram, A. (ed). (2001) *Peer-to-Peer: harnessing the benefits of a disruptive technology.* Sebastopol, CA: O'Reilly.

[19] Salman A. Baset and Henning Schulzrinne. (2004) An Analysis of the Skype Peer-to-Peer Internet Telephony Protocol. http://www.citebase.org/cgi-bin/citations?id=oai:arXiv.org:cs/0412017.

[20] Solé, R., Ferrer-Cancho, R., Montoya, J.M., and Valverde, S. (2003) Selection, Tinkering, and Emergence in Complex Networks. *Complexity, 8(1):20-33.*

[21] West, G. B. and Brown, J. H. (2005) The origin of allometric scaling laws in biology from genomes to ecosystems: towards a quantitative unifying theory of biological structure and organization. *J. Exper. Biology 208:1575-1592.*

[22] White, R. D., Kejzar, N., Tsallis, C. and Rozenblat, C. (2005) Structure and dynamics, submitted. Available at: http://www.eclectic.ss.uci.edu/~drwhite/pub/citysize/paper_31_5.pdf

[23] The BISON project has addressed the question of "nice properties" exhibited by living and nonliving systems. See for example the BISON Deliverable D04 at: http://www.cs.unibo.it/bison/deliverables/D04.pdf; and also see Ozalp Babaoglu, Geoffrey Canright, Andreas Deutsch, Gianni Di Caro, Frederick Ducatelle, Luca Gambardella, Niloy Ganguly, Mark Jelasity, Roberto Montemanni, Alberto Montresor. (2005) "Design Patterns from Biology for Distributed Computing". *Proceedings, European Conference on Complex Systems, 2005 (ECCS05).*