



Complexity and Emergent Behaviour in ICT Systems

Seth Bullock¹, Dave Cliff
Digital Media Systems Laboratory
HP Laboratories Bristol
HPL-2004-187
October 27, 2004*

E-mail: firstname.lastname@hp.com

complex systems,
emergent
behaviour, ICT,
DTI, foresight

Information and Communication Technology (ICT) practitioners are now readily able to create systems of such interconnected complexity that predicting the effects that small changes (such as minor component failures) will have on overall system performance may become very difficult or perhaps impossible. The notion that system-level behaviour "emerges" from parallel nonlinear interaction of multiple components in ways that are difficult or impossible to predict is explored in this document with reference to the UK's ICT investments and assets. We conclude that while it is true that there are currently limits to our ability to understand the ICT systems that we are capable of creating, nevertheless there are ways forward, including new ways of structuring and approaching software engineering, and teaching IT. This 25,000-word report is a briefing document commissioned by the Foresight Programme within the Office of Science and Technology of the UK Government's Department of Trade and Industry. Its findings are independent of government and do not constitute UK Government policy.

* Internal Accession Date Only

¹School of Computing, University of Leeds, UK

Approved for External Publication

© Copyright Hewlett-Packard Company 2004

Complexity and Emergent Behaviour in ICT Systems*

Seth Bullock
School of Computing
University of Leeds, UK
seth@comp.leeds.ac.uk

Dave Cliff
Hewlett-Packard Laboratories
Bristol, UK
dave.cliff@hp.com

Executive Summary

- The management and design problems facing modern ICT practitioners are critically concerned with ensuring reliability, usability, robustness, efficiency, effectiveness, security, and evolvability in the interconnected ICT systems upon which societies and economies increasingly rely.
- As our world becomes an ever more interconnected place, so-called “systems” ideas and perspectives become increasingly important. A central issue is the emergent behaviour of complex systems.
- In complex systems, non-linear interactions between component parts give rise to high-level “emergent” organisation that is not straightforward to explain. Workable definitions and measures of complexity and emergence remain elusive, however.
- Issues that sometimes appear to place complex systems beyond science and engineering (subjectivity, unpredictability, inexplicability) are often misconceived. Complexity is more than mere complicatedness and is not equivalent to unpredictability. Emergence is not equivalent to inexplicability. Complex systems and emergent behaviour are valid topics of scientific enquiry.
- Understanding adaptation in complex systems is a key problem; one that motivates an attention to biological systems and processes, because many biological systems exhibit characteristics desirable in engineered systems.
- A wide range of diverse stakeholders are implicated in complexity research: academics with varied backgrounds, industrialists from across many sectors, and many different funding agencies.
- The profile of the complexity science community is largely unknown; however, it is clear that few “complexity researchers” have enjoyed explicit or extensive training in complex systems ideas and methods.
- Achieving and maintaining effective interdisciplinary collaboration, especially between stakeholder groups, is identified as key to successfully mobilising this diverse and fragmented community.
- Academic complexity research that is relevant to ICT is extremely diverse. The UK is a key player in certain aspects of this field.
- We review here relevant industrial activity, examining the research of large ICT companies, large non-ICT companies, and small & medium-sized enterprises.
- Open research agendas in complex systems are many and varied, and include: fundamental research into characterising complexity and emergence; practical work on building effective simulation methodologies; theoretical exploration of the effects of spatiality, coupling, development, adaptation, and evolvability in complex systems. More specifically, open-standards software methodologies, the new scientific study of networks, resource management for grid computing, evolutionary and adaptive computing, network security, usability and interoperability, peer-to-peer systems, and even nano-scale ICT engineering are *all* domains in which complex systems research is likely to have (or continue to have) a significant impact.
- Open challenges that must be faced by the complex systems community include overcoming institutional and cultural obstacles to interdisciplinarity and industrial involvement in complexity research. For the UK in particular, most computer science undergraduate degree programmes currently have a manifest lack of formal training in complexity ideas and techniques: especially simulation modelling methods, experiment design, and statistical data analysis. This is an omission that should be rectified as a matter of urgency. Computer Science graduates should actually be trained *as scientists*.
- The UK is well-placed to meet these challenges and potentially to benefit enormously from a combination of its high-quality complex systems research and its innovative ICT industry.

*This report was commissioned by the Foresight programme of the Office of Science and Technology. However, its findings are independent of government and do not constitute government policy. Copyright © 2004 Hewlett-Packard.

1. Introduction

1.1 Background

It is not controversial to observe that we are living in an increasingly joined-up world. Processes of globalisation continue to drive ever-increasing levels of interconnectedness, encouraging greater population mobility and resultant economic and cultural exchange. What may be less apparent is that for such flows of goods, services, and capital, at regional, national, and international levels (and even for international terrorism and consequent homeland security issues) the many faces of modern globalisation have been underwritten by innovation in information and communications technology (ICT).

Indeed, ICT crucially underpins these activities at all levels: from manufacturing, logistics, retail, finance, and entertainment to healthcare, education, government, transport, and the provision of utilities such as water, electricity, and gas. As these ICT systems have experienced a largely unplanned and unregulated increase in interconnectedness they have begun to suffer from attendant growth in problems of design, management, maintenance, and decommissioning. As the effects of introducing a new component, making a new connection, or (more drastically) attempting to improve a protocol or impose a control, continue to become increasingly difficult to predict or understand, how can we ensure reliability, usability, robustness, efficiency, effectiveness, security, and evolvability in the interconnected systems that we all increasingly rely upon?

For instance, institutional inability to swiftly and effectively deal with the catastrophic 2001 outbreak of foot and mouth disease cost the UK economy approximately £8bn (National Audit Office, 2002). A poor understanding of the causal interconnectedness of livestock transportation, disease behaviour, and vaccination, coupled with bureaucratic delays and poor co-ordination, created a rapacious epidemic. Hundreds of farmers lost their livelihoods, millions of animals were destroyed, rural and tourism industries were decimated. Last year's high-profile infrastructural power failures on both US seaboard, in Italy, and in south London are evidence of similar problems in the context of electricity transmission — a technology more than a hundred years old. With respect to ICT, as the scale and interconnectedness of required systems explodes, the ability to deliver working solutions tailored to particular requirements has stalled — software development stagnates, interoperability erodes, and reliability and robustness can no longer be guaranteed. Recent reports suggest that the UK government has “squandered” £1.5bn on failed IT projects since 1998 (Arnott, 2003).

Consider the implications for the widely-shared vision of a near-future Europe permeated by a new kind of com-

putational infrastructure: a pervasive, ad hoc, decentralised, ubiquitous, dynamic computing network subserving commerce, government, health care, transport, manufacturing, education, and so on. Effectively and efficiently managing this kind of architecture and the critical processes that are supported by it is the overriding challenge faced by today's ICT research communities. Guarding against systemic failures of the kind exemplified by foot and mouth, or sudden catastrophic blackouts caused by the effects of minor component failures rippling through an electricity network, will be crucial to Europe's success and quality of life.

How did we arrive at this situation, and how does this challenge compare to those that have been faced by the ICT industry in the past?

1.2 History

Over the past fifty years, the dominant mode of commercial provision of information technology (IT) and computing equipment has undergone a number of significant “step” changes, and history shows that those changes occur roughly once every ten years. The 1960s saw the sale of physically huge “mainframe” computers, housed in air-conditioned rooms, which created a large and profitable global business for a previously inconspicuous company called International Business Machines Corp. (IBM). In the 1970s, mainframes were increasingly supplanted by the more robust, compact, and affordable “mini” computers, where disruptive technology advances were exploited by newly-formed companies such as Digital Equipment Corp. (DEC), who developed their PDP and VAX range of minicomputers to the point that DEC came to threaten IBM's dominance of the industry; and long-established electronics companies such as Hewlett-Packard (HP) also commenced commercial production of their own minicomputers. Approximately a decade after the initial commercialisation of minicomputers, advances in semiconductor technology (i.e., the development of single-chip microprocessors) made the provision of even cheaper and smaller single-user “personal” computers (PCs) economically viable, and again new companies such as Apple, Compaq, and Dell came to exploit this technology advance and in doing so these new arrivals threatened incumbents such as IBM and DEC, the latter of which was eventually acquired by Compaq.

In the 1980's, advances in computer communications and local-area networks allowed PC-sized computers to be connected together for cheap and reliable data interchange; this not only allowed expensive resources such as laser-printers or modems to be shared by a number of users, but also created a hardware base on which new forms of operating system could be developed and deployed, such as the client-server model of distributed network computing. Again, new companies achieved prominence or dominance by exploiting these disruptive tech-

nology advances, most notably Sun Microsystems, whose marketing slogan “the network is the computer” neatly summarised the vision motivating that transition.

Then, again roughly ten years later, the most recent transition occurred. This one was facilitated at least in part by a global shift toward deregulation and adoption of digital technology in the telecommunications industry: the mid-1990s saw the widespread adoption of internet communications protocols and document mark-up standards, which fostered the creation and truly astonishing growth of the World Wide Web. Again new companies, such as Cisco Systems, profitably exploited the demand for the newly disruptive technology. The entwining of information technology with communications technology has become so deeply embedded and accepted that the acronym “ICT” is now a commonplace replacement for “IT”. The irrational exuberance in the equity markets which swelled and then burst the web-enabled dot-com bubble saw the elimination of many new ICT companies with more fanciful business plans, and the following economic downturn in the ICT industry also led to a period of consolidation marked by the merger of industry giants HP and Compaq.

Within the computer industry, there are now clear indications that the next transition is currently underway. Major global computer companies such as HP, IBM, and Sun have all announced current or imminent commercial offerings which represent step-changes in the provision of computing, of equal or greater magnitude to the previous step-changes outlined above. The underlying vision offered by these three vendors is strikingly similar, despite the differences in terminology used by each. Sun talk of their “N1 Strategy”; IBM talk of “on-demand” and “autonomic” computing; while HP refer to “utility computing” and “adaptive enterprise” systems. The common vision is to create centralised compute facilities housing many thousands of computer “nodes”, where each node is itself a very powerful PC connected to a high-speed high-bandwidth network within that facility. Sub-networks consisting of some number of nodes and the necessary network connections between them can be assembled and disassembled remotely, via software commands rather than requiring human operators to physically alter switches or wiring. Thus, users can access these vast parallel computers remotely, via high-speed high-bandwidth telecommunications links, uploading software applications and hardware requirements, and receiving the results back down the telecomms link.

For most users, the cost of owning and operating such a facility would be prohibitive, but by sharing the usage of the facility over a large number of users, and thereby running all the nodes at close to 100% capacity for close to 100% of their lifetimes, significant economies of scale can be exploited. Thus, the cost of constructing and

maintaining the facility is spread over all users over the lifetime of the facility, and by spreading the costs in this way the users gain access to huge increases in computer power per unit cost. This new mode of computing provision is increasingly referred to as “utility” computing, because the intention is that use of the centralised facilities is accessed via a transmission network and charged to the users on a metered per-second or per-hour basis, in much the same way as other utilities such as electricity, water, or gas are supplied. Any one provider of such utility computing facilities would probably build multiple centres around the world, and there is likely to be more than one such provider, so (providing interoperability standards and protocols can be agreed) the current view is that there will soon be a federated global network of utility computing fabric, providing a hardware base for new modes of commercial provisioning of computer resources. This is unlikely to *replace* desktop or laptop PCs, but the availability of low-cost access to easily reconfigurable super-computing systems may change many businesses, either by lowering existing barriers to entry, or by enabling the application of new approaches which are currently too compute-intensive to be practicable.

However, there is a growing suspicion that traditional engineering techniques may not be well suited to the construction of artefacts as complex as federated global networks of massive utility computing facilities containing many tens of thousands of computing nodes. In part this is because traditional decompositional “divide and conquer” engineering techniques result in hierarchically organised designs, with a single overall executive controller situated at the top of the control hierarchy. Apart from the obvious vulnerability that loss of this central node presents, having such a single point of overall control introduces some problems which rapidly grow in severity as the number of nodes reporting up the hierarchy increases: simply gathering and routing all the necessary data up through the hierarchy can swamp the available network bandwidth and can often introduce significant delays. These delays, compounded by any noise or uncertainty in the data, can lead to incorrect actions or commands being issued by the master control node. Recognition of these problems has led to research aimed at developing reliable and robust *decentralised* control mechanisms, which could be *distributed* across large computer facilities, thereby reducing the delays and eliminating the single point of vulnerability, but still giving coherent overall behaviour.

These problems are not limited to the construction of utility-scale computing facilities. As just about every industry in advanced economies rushes to take advantage of the commercial and economic possibilities offered by current ICT systems, so every such industry is increasingly having to face problems caused by the unforeseen

behaviour, or poor scaling, of networks of interconnected components.

1.3 Complexity and Emergence

It seems unarguable that the key challenge facing modern ICT is the management of a transition from systems comprising many relatively isolated, small-scale elements to large-scale, massively interconnected systems that are physically distributed yet must remain secure, robust, and efficient. We in the UK are already surrounded by systems that are attempting to achieve this transition: from e-government and the digital NHS, multinational infrastructures and virtual universities, to peer-to-peer communities, grid computing and e-science, mobile, amorphous and pervasive computing, ad-hoc networks, and mass-scale RFID-tagging.

As noted above, it is widely suspected that traditional decompositional (divide-and-conquer) engineering approaches will not scale to solve this kind of problem because they are geared to the production of modular, hierarchical, and ultimately centralised command-and-control regimes. Consequently, engineering large-scale, integrated ICT systems can be a (hap)azardous and wasteful enterprise. Is there any alternative? The behaviour of certain naturally-occurring systems suggests that there may be. Examples of such systems include animal brains, immune systems, colonies of ants and other social insects, and even economic markets, all of which comprise enormous numbers of simple elements that combine to achieve sophisticated, robust aggregate behaviour.

Indeed, the natural world is full of systems that, at one level of analysis, can be described as being composed of many components that are individually “simple” and that interact with each other in relatively “simple” ways, often only directly influencing neighbouring components, yet simultaneously, at another level of analysis, are able to exhibit some “complex” overall system-level behaviour. In broad terms, it is those systems that exhibit this “emergent” globally-complex-behaviour-from-simple-components that we refer to as “complex systems”; and the prospect that the system-as-a-whole needs to be considered as something more than simply the sum of its parts is perhaps the most basic articulation of the notion of “systems thinking”, which stands in contrast to the traditional reductive/decompositional “componential” concepts and viewpoints that have dominated engineering and science for most of the last two hundred years. These notions of emergence, complexity, and systems thinking are revisited and expanded upon in the rest of this document.

A large class of natural complex systems are of particular interest because they exhibit attractive aggregate (emergent) properties that allow them to adapt to changing circumstances in an efficient and effective manner,

despite lacking any central authority or control responsible for this ability — i.e., they adaptively self-organise. As a result, such systems can be both extremely robust to perturbation and also behaviourally agile: properties that ICT engineers would like to design into their technological systems. This class of systems is generally referred to as complex *adaptive* systems (CAS), and the natural world is packed with them. Here are four examples:

- An individual nerve cell (i.e., a neuron) is a relatively simple device. In brief, a neuron is a cell body with plant-like branches emanating from it. It integrates electrical impulses received on its “input” branches over some period of time, and if that integral reaches a sufficiently high value then it fires an impulse down its “output” branches, which typically connect to the input branches of many other neurons. Yet attach enough neurons together with the right connectivity, and expose them to the right environmental stimuli for long enough, and the result is a brain capable of coordinating the sensory inputs and motor outputs of an adult animal, sufficient for the animal to survive in an unpredictable and unforgiving environment until it can meet a mate and reproduce. For a social animal as complex as a human, successfully mating and then raising the resulting offspring can require the brain to generate extremely subtle and sophisticated outputs over many years. More’s the pity.
- Although the behavioural repertoire of an individual ant is relatively simple, and its behaviour may indeed appear to an untrained eye to be largely random, colonies of ants can nevertheless be observed to engage in collective behaviours that are globally coherent and highly effective at achieving some goal. A widely-quoted example is the path-laying and path-following behaviour exhibited by many species of ant while they forage for food. As ants walk over the ground, they can excrete trails of inert marker-chemicals called *pheromones*, which once deposited will decay over time as they evaporate and diffuse. Ants returning to the nest from a source of food will typically lay a pheromone trail that other ants may then follow on their outward journey from the nest to the food, and ants tend to prefer trails with stronger pheromone trails to weaker ones. Different ants may follow different paths back from the food, or may occasionally randomly deviate from an established pheromone trail, but the shortest paths will always tend to have the freshest pheromone deposits, and hence will attract the most ant traffic, which will help to reinforce these as the paths with the strongest concentration of pheromone deposit. Thus, without any central synchronisation or control, colonies of foraging ants will quickly and reliably find the short-

est navigable path from food-source to nest, and if that path ceases to be navigable, then the colony will dynamically adjust to find a new shortest navigable path, again without central synchronisation or control.

- In a population of animals we can characterise each individual animal as a simple device attempting to survive long enough in its ecological niche for it to reproduce and create viable offspring; but different niches place different demands on their inhabitants. Over a large number of generations, the compound effects of random genetic variation plus directed selection (i.e.: Darwinian natural selection’s “survival of the fittest” and/or sexual selection’s “survival of the prettiest”) can create populations or species of organisms that are exquisitely tailored to their ecological niche, despite the absence of any centralised controller or designer.
- In economic markets, it is often the case that the actions and interactions of individual traders can be characterised in simple terms, yet the overall market dynamics that arise from the trader interactions can be subtle and sophisticated, in ways that are hard to relate back to the underlying simple interactions of the traders. For example, in commodity markets, buyers typically want to purchase units at the lowest price they can get and sellers want to offload at the highest price the market can bear. Bringing such rapaciously motivated traders together in the appropriate market institution (that is, making them interact via the right type of auction mechanism) can yield highly desirable overall market dynamics; i.e., the markets can consistently show a rapid and stable convergence of transaction prices to the market’s underlying theoretical equilibrium price,¹ despite none of the traders having any prior knowledge of what the equilibrium price actually is, and despite their being no central auctioneer or coordinator. That is, the traders in the market collectively discover the best price for the transactions, without any central control, and despite the fact that they are all acting out of raw self-interest.

Almost all such naturally occurring systems are biological in origin, and increasingly researchers in advanced engineering are turning to biology for inspiration: over the last twenty years, a growing number of researchers from disparate academic disciplines have studied complex systems in general, and CAS in particular. Exam-

¹In a nutshell, the equilibrium price for a particular commodity in a market is the price that best matches the quantity demanded by the buyers with the quantity supplied by the sellers; if transactions consistently take place either above or below the equilibrium price, then respectively either the buyers or the sellers are consistently being ripped off.

ples of natural complex adaptive systems that have inspired modern ICT engineering techniques include: the human immune system, which is being used as inspiration for new anti-virus approaches to computer security (e.g., Sana Security [I-066]; de Castro & Timmis, 2002); animal nervous systems, which have inspired the development of artificial neural networks (parallel distributed computing techniques that can perform powerful statistical computations for data recognition and classification, Rumelhart & McClelland, 1986); insect colonies, which have inspired “swarm intelligence” algorithms for, e.g., data mining (Bonabeau, Dorigo, & Théraulaz, 1999); evolving populations, which have been the inspiration for new kinds of evolutionary search and design algorithms (Goldberg, 1989; Mitchell, 1996); and even markets or auctions, which have inspired new computer resource-allocation systems in which ideas from free-market economics are used to automate the dynamic matching of the supply of scarce computer system resources (such as CPU time, memory space, or network bandwidth) to the demand for those resources from some number of users or applications, while the quantities of resource supplied and demanded vary in real time — so-called “market-based control” (Huberman, 1988; Clearwater, 1996).²

In response to the growing need for scientists and engineers skilled in dealing with such systems, in recent years a number of top UK university computer science schools or informatics departments have started to offer successful postgraduate courses in natural computation [A-001], evolutionary and adaptive systems [A-002], bio-inspired computing [A-003] and in biosystems-based multidisciplinary informatics [A-004]; and other universities are known to have similar initiatives under development.

Thus, systems concepts such as complexity and emergent behaviour are increasingly being recognised as centrally important to meeting current urgent engineering goals. However, as yet there is no established design methodology capable of rigorously supporting their use in industrial-scale ICT engineering.

1.4 Overview

In this document we present a review of the state of research into these topics, informed by interviews with leading complexity science practitioners and other multidisciplinary scientists, core ICT researchers, ICT industrialists, representatives of relevant funding bodies, policy makers and educators. Nevertheless, the resulting overview will necessarily be selective due to limitations

²In June 2004, The UK’s Engineering and Physical Science Research Council (EPSRC: the primary government funding agency for computer science research) announced a £1.6m grant to a consortium of researchers working on the application of automated optimisation to exactly this problem. The consortium involves researchers from the universities of Southampton, Liverpool, and Birmingham; and from industrial research labs operated by HP, BT, BAE Systems, and IBM.

of space, the biases of us the authors, and the particular interests of those consulted.

We will first discuss issues arising from the diversity of terminology and conceptual frameworks generated by complexity research, before assaying the large number of interested parties contributing to the evolution of these ideas and the complicated relationships between them. Subsequently we will outline some of the relevant leading-edge complexity-related research and development activity currently being undertaken and the open research questions that remain. Finally, an assessment of the challenges to future progress in this area is presented, before the paper concludes.

Before proceeding, an extended metaphor employed by Lord Robert May, current president of The Royal Society and former chief scientific adviser to the British government, will serve as a *précis* for the document as a whole. In his opening talk at a recent meeting³ on complexity research, May suggested that, like most scientific activities, the state of complexity research could be assessed by comparison with a path stretching from Tycho Brahe's increasingly accurate *observations* of planetary motion, through Johannes Kepler's discovery of *patterns* in this data, to Isaac Newton's eventual derivation of general *laws* that could account for these patterns and more besides. By analogy with this sequence, May asserted, complexity science can be regarded as on the cusp of a transition from the Brahe stage to the Kepler stage. That is, complexity researchers are just beginning to discover and describe coherent patterns in the increasing volume of accurate data gathered from complex systems, but they are still mostly "stamp collecting" and have some way to go before the formulation of a general framework of "laws" with which to account for these patterns. This position resonates with the opinion of many observers, and one goal of this paper will be to identify ways in which the progress required can be encouraged.

2. Conceptual Landscape

In this section, we describe the conceptual diversity running through the complexity literature and explore some of the key issues for understanding how concepts of complexity and emergence can be applied in an ICT context.

It is widely acknowledged that the notions of complexity and emergence employed across a wide range of communities are poorly and multiply defined, if they are defined at all. Informally, the term *complex* can be a synonym for challenging, interesting, complicated, or just large, while *emergent* is often used to convey the surprising, inexplicable, or mysterious nature of a system's behaviour. One might hear that "as software developers attempt to deliver ever more complex systems, our diffi-

culty in guarding against unwanted emergent behaviour becomes an increasingly pressing concern". Where attempts have been made to formalise these kinds of notion, a plethora of definitions, interpretations and perspectives have been generated.

In a recent selective review article, Adami (2002) distinguishes between several *approaches* to defining and measuring complexity: computational vs. statistical, structural vs. functional, sequential, hierarchical, etc. He further describes a handful of existing *definitions*, from "algorithmic complexity", "Kolomogorov complexity", and "minimum description length", through measures of fractal dimension and entropy, to "effective measure complexity" and "effective complexity", before introducing his own candidate: "physical complexity". An exhaustive review would have had to consider upwards of three dozen distinct but inter-related attempts at defining complexity (Horgan, 1995) drawing upon information theory, computer science, statistical physics, evolutionary theory, and so on.

The motivation driving all this activity is to provide an account that successfully underwrites an intuitive notion of complexity, refined in the following manner. Grouped at one end of the complexity continuum are static, regular or random systems. These are "simple" organisations with straightforward aggregate behaviours (e.g., periodic motion in a swinging pendulum, or Brownian motion in a gas). At the opposite extreme are complex systems that are much more difficult to understand, predict, control or explain due to the "entwined" nature of their multiple components, which limits the success of a standard divided-and-conquer approach to explanation. Of course, much hinges on whether these notions of explanatory "straightforwardness" or "difficulty" can be formalised in some way. Until they are, to claim that one system is more complex than another is simply to claim that we currently find it harder to explain (Edmonds, 1999).

Formal measures of predictability or irregularity such as Kolomogorov complexity⁴ cannot help us here as they distinguish ordered systems from disordered systems. That is, they attribute low scores to homogeneous or static systems, higher scores to those exhibiting some regularity or periodicity, even higher scores to complex or chaotic systems, and highest scores to completely random or irregular systems. From such a perspective, it is the *intermediate* systems that are challenging or interesting — the behaviour of completely ordered systems can be explained through traditional methods, while that of completely disordered systems can be explained statistically through the law of large numbers.

In summary, while complex systems often exhibit system-level organisation that is interesting and some-

³ *Complexity Science & 21st Century Issues*, London School of Economics, March 25-26, 2004.

⁴The Kolomogorov complexity of a system is, crudely, a measure of the length of the shortest algorithm that captures the system's behaviour.

times attractive (e.g., the robust, efficient, adaptive behaviour of an ant colony), it is not clear how this high-level systematic behaviour is brought about by interactions between the system’s components (e.g., the behavioural tendencies of individual ants). Such system behaviour can reasonably be termed *emergent* insofar as it “emerges” from the system’s low-level atomic interactions in a non-trivial manner. But in discussions of the meaning and utility of the notion of emergence, much hinges on the nature of this “non-triviality”.

As Clark (2001) points out, when the collective behaviour of a system derives straightforwardly from the contributions of each individual (as when a collection of small weights tips a balance) then we gain little by tagging this behaviour emergent. However, reserving the term emergence to describe only systems that are currently unexplained or perhaps systems that are somehow inherently inexplicable “robs the notion of immediate scientific interest”. Clark considers four “prominent” attempts to pin down the notion of emergence: collective self-organisation; un-programmed functionality; interactive complexity; and incompressible unfolding.

It is beyond the scope of this paper to fully review all four here, but predictably each of them drive at different ways in which the opacity in the relationship between the low-level behaviour of a system’s components and the high-level aggregate behaviour of the same system might come about. Simplifying the picture somewhat, to the extent that a system’s low-level interactions are *non-linear*, a successful account of their impact on global system properties will be increasingly involved. This non-linearity comes in many flavours, tending to occur when a system’s interactions are multiple, ecologically embedded, non-additive, inseparable, heterogeneous, interactive, asynchronous, lagged, or delayed.

At one extreme, simple, homogeneous, linear interactions between identical particles give rise to aggregate properties that are relatively straightforward to characterise (e.g., temperature, pressure, weight, etc.), and perhaps do not even deserve to be classed as emergent at all. At the opposite extreme are systems for which there may never be an adequate theory accounting for the global effects of local interactions. The simplest description of such “uncompressible” systems is a full account (perhaps a simulation) of the interactions involved. Predicting a protein’s three-dimensional structure from its linear amino acid sequence was once suggested by Marr (1977) as one possible example of such a problem.

Again, for complexity scientists it is the intermediate systems that are most interesting: systems that exhibit regularities or systematicities in the relationships between different levels of description, but where these relationships are not straightforward. One might term the behaviour of these systems as *moderately* emergent (see Figure 1).

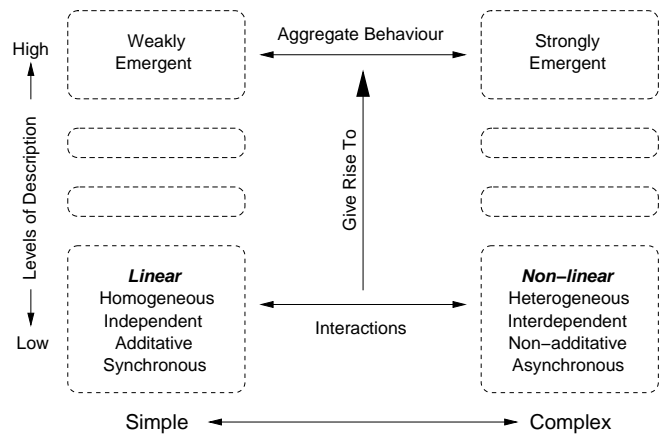


Figure 1: The aggregate, high-level behaviour of a system is likely to be strongly emergent when it arises from low-level interactions that are non-linear. By contrast, where these interactions are linear, aggregate behaviour may only be weakly emergent. In such cases, the term emergent may even appear superfluous. Systems may of course lie between these two extremes in which case their behaviour may be moderately emergent.

It is worth noting that many of these ideas are not new. On the contrary, they stretch back at least to the middle of the last century and have been repeatedly invented, challenged, discredited, and rediscovered in movements such as systems biology, cybernetics, systems science, catastrophe theory, artificial life, and complexity science. Given this history, can we expect significant progress any time soon? Before addressing this question, we will consider a number of problematic issues that can often arise in discussion of how ideas of complexity and emergence can apply in an ICT context. These are: the plurality and subjectivity of complex systems ideas and definitions; the unpredictability of complex systems; the distinction between complexity and complicatedness; and the phenomenon of adaptation in complex systems.

2.1 Plurality

The fact that there is little apparent consensus on definitions of complexity or emergence, and that different approaches are hotly debated and contested, is sometimes taken to reflect poorly on the field. However, we might expect diverse communities to arrive at multiple definitions of concepts as wide-reaching as these. The formulation of a single, tightly defined concept to replace the current plurality of ideas may simply not be possible or desirable. On the other hand, it might reasonably be expected that they share a “centre of gravity” that can be explicated (with some effort). Increasing interdisciplinarity should accelerate this explication, as it exposes researchers to multiple approaches both current and his-

torical, and discourages isolated activity. Indeed there is some evidence that recent treatments of complexity are more sophisticated in this respect, and more integrative as a result (Clark, 2001; Adami, 2002).

2.2 Subjectivity

If concepts of complexity and emergence are inherently subjective (“behaviour is emergent if it surprises *us*”; “a system is complex if *we* find it hard to understand”) then this limits their scientific utility. If researchers and practitioners cannot find more objective or operational definitions of these words, complexity science will be severely damaged, if not entirely undermined. The identification of non-linearity, a well-characterised formal notion, as the central root of complexity ideas has begun to account for the surprising, counter-intuitive and challenging nature of complex systems. The fact that, despite advances in understanding complex systems, they *remain* counter-intuitive does not imply that complexity and emergence are necessarily subjective or non-scientific concepts.

2.3 Predictability

Complex systems are often described as unpredictable. This is sometimes taken to imply that they are untrustworthy, or that it is inadvisable to employ a complex mechanism in an engineering context. For instance, a complex-systems approach to managing stock control might be vetoed on the grounds that stock control must be reliable and hence the system responsible for it must be predictable.

In what sense are complex systems unpredictable in a way that simple systems are not? It is true that complex systems are sometimes stochastic. A simulation of a foraging ant colony might model each ant as an essentially random process biased by local environmental cues. However, simple systems may also be stochastic. A model of an ideal gas may assume that molecules orient randomly after collision. In both cases, the low level entities that comprise the system are stochastic, and in both cases the statistical or aggregate behaviour of the system is predictable or regular. In the case of an ideal gas, we can predict properties like average pressure and average temperature with confidence despite being unable to predict the trajectories of individual molecules, while in the case of the ant colony, we can predict (*a posteriori*) the average rate at which different sources of food will be depleted despite being unable to predict the trajectories of individual ants.

In terms of stock control, it may be perfectly acceptable to make use of a stochastic process (a robot arm that picks items from a bin might take variable time; the exact retail outlet that a particular stock item will be sent to might be indeterminate) so long as the statistical properties of this process are reasonable and appropriate.

What distinguishes simple from complex systems is not that the aggregate behaviour of the former is regular and predictable, whereas that of the latter is irregular, unpredictable, random, or even chaotic. Rather, whereas explaining the statistical, aggregate behaviour of simple systems is relatively straightforward, doing the same for complex systems is much harder. We have a well-established canon of statistical techniques, modelling approaches, and well-understood simplifying assumptions that can be deployed in the former case, but not in the latter.

It is this difficulty in forming rigorous causal characterisations of the aggregate behaviour of a complex system (rather than the *absence* of regularity or predictability in this aggregate behaviour) that is the more legitimate barrier to adopting complex-systems approaches in an ICT engineering context. Without a grasp of how the system’s configuration gives rise to its aggregate behaviour, it is difficult to understand how a small change to that configuration will affect the system’s aggregate behaviour. This inability is compounded by the fact that while small changes to the configuration of simple systems tend to result in proportionately small changes in aggregate behaviour (as when more gas molecules are added to a container), similarly small changes to the configuration of a complex system (whether stochastic or deterministic) may have anything from a very minor to a very major impact on the character of the aggregate behaviour (e.g., removing a few ants from a colony may have little or no impact, or it might possibly destroy the colony’s ability to forage). The mismatch between these issues and the needs of traditional approaches to quality assurance in engineered systems is stark: in the absence of theoretical or analytic quality-assurance techniques for engineered complex systems, the obvious exhaustive empirical approach is to make all possible small changes to the system under test, noting which changes have no effect, which have negligible effects, and which (if any) have a major effect. The problem with such an approach is that the number of change-tests that need to be performed can grow hideously large very quickly.⁵

One way in which this sensitivity to small changes reveals itself in complex systems has been termed “path-dependence”. Particular local events in a system’s history may strongly influence the subsequent global state of the system, as when a rat population learns an aversion to a type of food, or a social custom sweeps through a teenage population, or positive feedback amplifies an initial random fluctuation in a stock price. These trends and tendencies are evidence of a kind of systemic “mem-

⁵The combinatorics are poor because of the need for exploring the possible interactions between different small changes (i.e., performing change-tests involving different combinations or sequences of the individual small changes), and the need on each change-test for exploring responses under different test-environment conditions (i.e., controlling for variations in the system’s free parameters and/or for different initial conditions).

ory” that is not present in simple organisations such as a gas in a container.

Complex systems researchers sometimes rely upon relatively detailed simulation in order to construct a post-hoc pseudo-empirical account of the system’s aggregate behaviour, as when an ant colony model is simulated over a range of foraging tasks in order to estimate its efficiency or sensitivity to parameters. Unfortunately, even if our simulations were sufficiently accurate, we would have to collect results from every conceivable scenario in order to grasp the full implications of system configuration for system behaviour. Even this “grasp” will remain weak in the absence of analytical accounts relating low-level interactions to high-level behaviour.

What are the prospects for the development of such analytical techniques for complex systems, the “Kepler” activity alluded to by Lord May, as described in this paper’s introduction? An analogy with Kepler-like activity in chemistry may clarify the situation. The properties of a bar of gold at room temperature (heavy, shiny, solid) are predictable, but difficult to derive from knowledge of the properties of a gold atom. Moreover, adding a single proton to each atom in the bar of gold has a series of “surprising” and “unpredictable” effects on its aggregate properties. The bar of gold changes from a solid metal to a liquid “transition metal” (an isotope of mercury). While the periodic table of the elements organises and labels these surprising transitions, it does not, on its own, account for them. Complexity scientists are just beginning to build their own “periodic table”. For example, a strand of research papers are beginning to reveal commonalities in the impact of moving from synchronous to asynchronous update across a range of complex systems (crudely, intricate patterns that may be exhibited under synchronous update are often lost when synchrony is relaxed (Nowak & May, 1992; Huberman & Glance, 1993; Harvey & Bossomaier, 1997)); and researchers analysing networks of interaction in a diverse range of natural and artificial systems are finding surprising commonalities in the underlying network topologies and dynamics (Barabási, 2002). It will take truly “Newtonian” insight to make sense of the patterns that such a “periodic table of complexity” might exhibit.

2.4 *Complicated vs. Complex Systems*

Complex systems are often distinguished from those that are merely complicated. The distinction is made on the following grounds: while a complicated system (e.g., a car) may be difficult to understand as a result of the interactions between its many components, unlike complex systems, complicated systems eventually succumb to a divide-and-conquer decompositional explanation because they are inherently modular.

For example, the turning circle of a car (an aggregate property) is largely determined by the properties of its

wheels in relation to one another (i.e., simple relationships between a few sub-components). The ability of a car to withstand collision is also largely a result of properties of specific parts of the car designed to enable it to do so: bumpers, crumple zones, side impact bars, etc.

By contrast, the emergent properties of a complex system cannot be explained in this way. Hofstadter (Hofstadter, 1979, p.308–309) gives an example in terms of a multi-user time-share CPU system that begins to *thrash*⁶ when dealing with more than 35 users. Hofstadter jokingly asks why the computer engineers don’t just find where the number “35” is stored in the computer’s operating system, and change it to “60”. But there is no system module or parameter (or handful of parameters) that corresponds to a “thrashing number”. The system’s thrashing behaviour arises from a complex interaction between many or all of its subcomponents; in Hofstadter’s words: “The point is, of course, that there is no such place. Where, then, does the critical number (35 users) come from? The answer is: *It is a visible consequence of the overall system organization*” (Hofstadter, 1979, p.308, original emphasis).

While reaching an understanding of very large and intricate complicated systems may be time-consuming, it will not require the development of new modes of thinking or analysis. Divide-and-conquer will remain the correct approach to such problems. As such, decreasing a car’s turning circle should perhaps be a far easier task than fixing the network thrashing problem. All it requires is a relocation of the wheels or an alteration to the angle through which individual wheels are able to turn. In a well-designed car, the implications of these alterations will be easy to calculate due to the limited number of well-defined interactions between tightly-specified component modules.

However, altering a complicated system is not always straightforward. Despite the car’s apparent modularity, alterations to components or their configuration may have multiple and compound influences on a range of systemic properties. For instance, a shorter wheel-base may decrease a car’s turning circle, but simultaneously alter its aerodynamics, stability, aesthetics, etc. Indeed, for many complicated systems, it is very hard to exhaustively specify the inter-relations between components. Newly designed cars are still occasionally recalled as a result of some unforeseen problem arising from unplanned interactions between components that compromise the car’s safety or efficiency. Even entirely modular complicated systems, such as large pieces of well-engineered software, may exhibit “unprogrammed functionality” (i.e., emergent behaviour) when coupled with

⁶Loosely speaking, “thrashing” refers to the situation where a multi-user computer system spends almost all of its time deciding which user to deal with next and dealing with the swap-over from one user to the next, rather than doing what it is actually intended to do – i.e., doing proper computing for each of the users.

a (perhaps untrained and evolving) user community.

Conversely, truly complex systems may come to exhibit the properties of complicated systems. Consider the human body with its musculo-skeletal system; its separate organs of respiration, digestion, and reproduction; and its systems of circulation, regulation, and communication. Or an ant colony with its caste system of workers, soldiers, queens and reproductive males; or the partially independent exchanges and marketplaces that emerge in major financial centres. In each of these cases stable modular structure has arisen (emerged) as a result of complex interactions. Understanding the manner in which this type of emergent modularity comes about is a key challenge in complex systems science (Dawkins, 1986; Maynard Smith & Szathmary, 1995).

As such, while the complex-complicated distinction is often assumed to be clear cut, it is not. In particular, while ICT engineering typically involves building complicated systems, these systems are unlikely to be free from complexity. Simultaneously, while complex natural systems have not been designed to be modular, they often appear to exhibit complicatedness in the form of compartmental or quasi-hierarchical structure or functionality. While it is clear that much is to be gained from initially focusing on the study of uncomplicated complex systems, if we are to derive useful design principles from studying natural complex systems, the relationship between complexity and complication needs to be explored seriously rather than taken to be a distinction separating complex systems research from traditional engineering.

2.5 Adaptive Behaviour and Adaptation

While the engineering design and management issues presented by increasingly complex systems tend to cast complexity as a pressing *problem* that must be coped with, there is a sense in which complexity research also holds the key to potential *solutions*. Some complex systems, principally evolved biological ones, exhibit kinds of emergent behaviour that would be extremely attractive features of engineered systems. Where large-scale integrated engineered systems can suffer from brittleness, stagnation, swamping, or inefficiency, organisms and biological organisations balance robustness (resilience, stability, fault-tolerance, graceful degradation, etc.) against flexibility (evolvability, creativity, adaptability, agility, etc.) through processes of adaptation (evolution, learning, homeostasis, and habituation) in order to achieve a range of properties (self-organisation, self-regulation, self-repair, self-calibration, self-decommissioning) which are now widely recognised (under the banner “self-^{*}”) as desirable aims in future ICT systems [I-047, I-048].

Examples of such naturally occurring complex adaptive systems include single cells, organs, nervous systems, immune systems, whole organisms, insect colonies, evolving populations, families, crowds, markets, lan-

guages, and cultures. What makes these systems notable and relevant is that they solve some of the problems associated with large-scale complex engineered systems without recourse to sophisticated centralised control mechanisms.

ICT engineers can learn lessons from biology in several distinct ways. First, particular low-level biological mechanisms can act as inspiration for the design of system components. The neuron is a key example in artificial intelligence: its study having led to the development of neural computing systems (Rumelhart & McClelland, 1986; Haykin, 1998; Stone, 2004). Second, the configuration of engineered components can take inspiration from particular kinds of biological organisation. Insect societies are a key example, having been used as models for optimisation algorithms (Schoonderwoerd, Holland, Bruten, & Rothkrantz, 1997; Bonabeau et al., 1999). Third, biological processes such as learning or evolution can serve as inspiration for engineered artefacts and processes (Holland, 1975; Goldberg, 1989; Mitchell, 1996; Sutton & Barto, 1998; Arkin, 1998; Brooks, 1999). In each of these cases biologists have a reasonable understanding of what kinds of properties are associated with each biological system. However, large questions remain unanswered. In particular: what aspects of complex adaptive systems underpin their attractive aggregate properties? Which are contingent or coincidental and which are fundamental?

Advocating the use of biologically-inspired approaches to ICT engineering is not to say that biological systems are perfect. To assume so is a form of the so-called “naturalistic fallacy”: the assumption that the way a natural system actually *is*, indicates the way that similar systems *ought* to be. The fact that nature appears to solve a particular problem in a particular way does not make that solution right or optimal. Natural systems have not arisen to solve our problems. However, despite this, we are quite happy to work with them where their natural abilities, propensities or capacities can be turned to our advantage (“good bacteria”, brewer’s yeast, food crops, cash crops, domestic farm animals, sheepdogs and sniffer-dogs, even our human co-workers).

Indeed, the distinction between organism and machine has been eroding for some time. Advances in genetics and biotechnology are colouring our perception of organisms, casting them as controllable and machine-like in several ways. Simultaneously, our exposure to new kinds of decentralised technologies is increasingly encouraging us to treat machines as “organic”. Our interactions with the internet, peer-to-peer file-sharing, mobile telephony systems, and even the operating system of a modern PC are now founded on expectations of patchy functionality and performance coupled with deep impenetrability. We do not “rely” on these machines in the way that we rely on a spoon or a pair of scissors or even a car (where we

demand a strong sense of control). We rely on them in the way that we rely on a colleague or a community. Despite these developments, the signs are that significant cultural change will be required of engineers and other industrialists (and complexity researchers) before working with *bio-inspired* systems, solutions, and services is entirely acceptable.

3. Stakeholder Landscape

The diversity of ideas and approaches represented by research into complexity and emergent behaviour is compounded by the variety of (for want of a better word) stakeholder groups implicated in this activity. Here we describe the makeup of this stakeholder landscape and identify some of the consequent issues for complex systems ICT research. We consider three sectors: academia, industry, and funding bodies; before addressing some issues arising from their interactions.

3.1 Academia

A large number of established domains have begun to develop specialised, poorly defined, esoteric (and sometimes short-lived) sub-fields interested in exploring and exploiting complexity ideas. Despite this, only a small, transient, and fragmented community would label themselves “complexity researchers”. Figure 2 represents an attempt to further characterise the diverse activity in complexity research. In the absence of systematic studies exploring the character of complexity research, we have no objective data on the exact distribution of this activity, who is undertaking it, how effective it is, or how it is being funded. However, here we present some subjective impressions of the current situation.

The pyramidal form of Figure 2 reflects the relative abundance of activity concerned with domain-specific questions addressing relatively concrete issues in particular disciplines: biologists characterising the complexity of bird-song repertoires, or the structure of evolved neural network circuits; engineers studying the effects of ICT network topology on performance; town-planners simulating congestion in a traffic system; or national security advisers evaluating resilience to attack in a telecommunications network. The domains involved here are many and varied, spanning the physical and life sciences, various kinds of engineering, and the humanities. In addition to the disciplines listed across the bottom of Figure 2, linguistics, social science, biochemistry, and even art could be included. By comparison, there is a relatively small amount of “complexity science” activity addressing domain-general, abstract, theoretical issues, such as, what is the nature of adaptation, or emergence, or homeostasis, etc. Note that the actual distribution of activity is unknown, and that it alters to reflect researchers’ changing interests and changes in funding ini-

tiatives, rather than being fixed in some way.

Where do complexity researchers come from? What are their backgrounds? What are they trained in? Figure 2 depicts the influx and outflux of complexity researchers by a series of peripheral double-headed arrows below and to the right of the pyramid. Two types of intake are identified. First, a diverse range of domain-specific researchers, each trained in a single discipline become involved in complexity-related research questions relevant to their own discipline. This type of intake is supplemented by an influx of researchers trained in more abstract theoretical tools from physics, philosophy, biology, mathematics and ICT. Researchers also leave complexity research by similar routes, returning to the study of more mainstream domain-specific questions or alternative theoretical problems. Our informal survey suggests that the turnover of complexity researchers is relatively high, with few academics establishing a career in “complexity” *per se*. However, we can only guess at the motivations driving individuals to join and, perhaps more importantly, to leave this activity. The general absence of complexity modules and courses from university curricula preclude significant direct intake of “complexity scientists” from degree programmes. This ensures that researchers are rarely trained in the abstract theoretical tools and ideas of complexity theory. Rather these ideas and tools are obtained at post-graduate level, and tend to either be self-taught or taught by a self-taught PhD supervisor.

Within the pyramidal structure of Figure 2, the work of a single researcher cannot usually be represented by a single point. Rather, the fact that individual researchers pursue a range of questions implicates a number of scattered points (the small crosses in Figure 2). Migration, collaboration and communication across the space of research activities (represented by arrows within the pyramid) is limited by methodological and terminological barriers separating theoretical from empirical work (horizontal dashed lines) and separating disciplines from one another (vertical dashed lines). These barriers are enshrined in the mono-disciplinary departmental organisation of universities, funding bodies, professional bodies, etc., ensuring that there are also social, spatial, and political barriers separating different communities so that they rarely come into meaningful contact with one another. As such, a researcher’s initial training and point of entry into complexity science is a significant determining factor for their subsequent research work. Note that, ideally, barriers between disciplines should become less significant as increasingly theoretical questions are addressed. The extent to which this is actually true in practice is unclear.

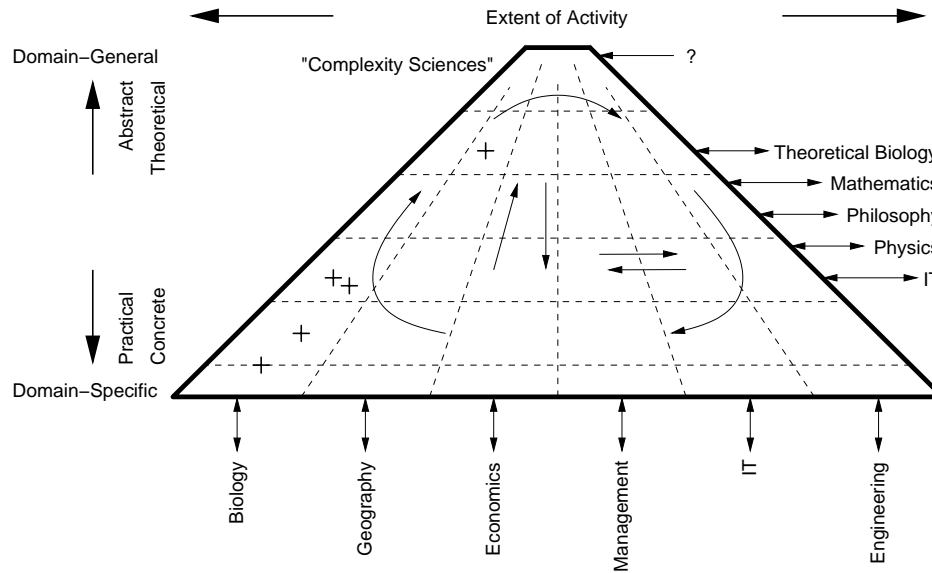


Figure 2: Schematic “pyramidal” representation of the diverse range of academic complexity activity. See text for explanation.

3.2 Industry

In addition to the complexity-related research taking place within industry (some of which is detailed below), industrial influence on academic complexity research comes in a variety of forms. From inputs to the strategic decision making that informs funding initiatives, to shaping the job market that strongly influences the makeup of undergraduate degree programs. Beyond core ICT companies, current interest in complexity-related research appears to span the transport, manufacturing, retail, finance, defence, aerospace, and utilities sectors. In particular, many research-led companies are actively involved in complex ICT research, including multinationals with their own significant in-house research and development activities. Interest from small to medium enterprises (SMEs) is less evident, although some early-adopter SMEs and start-ups in “bleeding-edge” areas (biotechnology, nanotechnology, software agents, etc.) have shown interest. Finally, governmental organisations are a significant driver in the form of initiatives within the NHS, MoD, etc. However, the manner in which these interests impact on academic research is unclear. While “coping with complexity” may be a widely-voiced industrial aim, there are certainly mixed opinions within industry concerning the applicability of academic complexity research to real world problems.

3.3 Funding Bodies

Across UK funding bodies, the range of activity supporting and promoting complexity research is extensive. In particular, recent calls from the EPSRC, and collaborative activities with the BBSRC, have targeted complex-

ity issues in the context of novel computation, e-Science, and the mathematics underpinning the life sciences. The European Commission, as part of their Information Society Technologies activity, has recently pursued various initiatives including extensive complex systems funding under the Future and Emerging Technologies program. UK charities such as the Leverhulme Foundation have also recently targeted complexity research, investing considerable funds in what are perceived to be projects with significant impact on quality of life in the UK

3.4 Interdisciplinarity Issues

It is difficult to satisfy or co-ordinate the varied and wide-ranging suite of interested parties outlined above. In particular, while there is wide agreement that genuine progress can best be achieved through effective interdisciplinary collaboration involving academics working closely with industrialists, it is unclear how best to foster this type of collaboration. A few of the issues relevant to the UK are briefly outlined below.

Communication: This is recognised as the key problem facing the complexity community, applying at all levels and particularly between stakeholder groups (e.g., between academia and industry).

Coherence: Currently there appears to be relatively little integration or co-ordination at the level of (i) academic research communities, (ii) national and international research funding initiatives, (iii) education and training opportunities, (iv) industrial representation.

Quality Control: Ensuring *rigour* at interdisciplinary boundaries is fraught with difficulty. Like an executive with two offices who can always claim to have been “in my other office” if someone comes calling, research ac-

tivity that attempts to connect two separate disciplines can end up satisfying neither. As a result, interdisciplinary research is often perceived as of dubious quality. This can lead to infighting as academics attempt to distance themselves from what they perceive as misguided or sloppy work. Even the long-established process of peer review is not an infallible safeguard of quality control: once a sub-field reaches a certain critical mass, where there are enough researchers (nationally or internationally) to form a circle of people all “patting each other on the back” (i.e., writing positive reports as grant-reviewers and paper-referees for each other), there is a very real danger that the community of practice thus created continues to perform research that is internally consistent relative to the norms of that sub-field’s community of practice, but which steadily diminishes in relevance to the reality that the sub-field was originally intended to address, or indeed to any reality at all.⁷

Transience: Without a natural centre of gravity, the turnover of ideas, personnel, and “brand” identities is rapid. For example, the first decade of research in Artificial Life (a multi-disciplinary effort to engage computing, biology and other disciplines in the simulation and synthesis of life-like systems) saw the inception of an explosion of slightly different “disciplines” aimed at applying simulation modelling to biological problems (“digital biology”, “computational neuroethology”, “synthetic ethology”, “synthetic behavioural ecology”, etc.). Very few of these labels are in common usage today, and many of the researchers who coined them have moved on to different questions or fields of enquiry. Relatedly, long-lasting complexity projects, research groups, institutes, or other structures are few and far between.

Resource Demands: Interdisciplinary research can be more resource-hungry than that adopting an established mono-disciplinary methodology. Time and effort are required to instigate and then maintain any effective collaboration, more so when collaboration spans disciplinary boundaries. Project researchers drawn from mono-disciplinary backgrounds must be trained and educated considerably before they can commence productive interdisciplinary work. Novel techniques, tools, and methodologies must be developed before useful results can be generated. With novelty comes risk, and the prospect of dead-ends, failed projects, and disappointing outcomes.

Career Pressures: The above-listed factors coupled with many institutional barriers to interdisciplinarity (see “Challenges to Progress” in Section 6 below) mitigate against a career at disciplinary boundaries. Currently, most of the risk of pursuing this type of research is borne disproportionately by young research students supervised by established, tenured academics. Given the

⁷No names no packdrill. Most academics will be able to name at least one field or sub-field (not their own, of course) in which this has clearly happened.

reality of pressures on modern academics, especially at the outset of their careers, it is perhaps remarkable that we see as much interdisciplinary activity as we currently do.

4. Current Activity

Continuing falls in the real costs of processor power, disk space and RAM, coupled with increasing telecommunications bandwidth and speed, have led to a situation where computational tools and techniques that were academic curiosities a decade ago are now readily deployable in a number of real-world problems of very high significance both to developers and to end-users of ICT systems. One such set of technologies involves complex adaptive systems (CAS) that are massively parallel, decentralised and distributed, and which are inspired to some extent by naturally-occurring adaptive systems such as immune systems, nervous systems, evolving populations, social structures, and economic markets. Examples include artificial neural networks and parallel distributed processing; evolutionary computation (e.g., genetic algorithms, genetic programming, and artificial immune systems); adaptive artificial autonomous agents for colony-based computation systems; and adaptive artificial autonomous agents as artificial traders for e-commerce, business modelling, and market-based control.

This section gives a selective and partial review first of some of the relevant major academic and non-profit CAS research institutes around the world, and then gives a similarly selective and partial review of the CAS research activities of commercial companies.⁸ Since the activity of academic groups is in general much more transparent and easier to find out about than industrial research and development, we have biased our coverage in favour of industrial activity, which we divide into three sections: large blue-chip companies in the ICT sector; large non-ICT companies; and small or medium-sized enterprises (SMEs) which range in size from around 300 employees down to start-ups and consultancies with employee numbers in single figures. For the large companies, we describe what research efforts are underway, and what products are offered. The SMEs are categorised by domain (consulting/software; hardware; business simulation; business operations systems; and entertainment/media), and details of each enterprise’s origin, customers and products are given.

Additional information on relevant activity within the UK can be found on the webpage of the Natural Computation Applications Forum (NCAF) an informal consortium of interested parties from British commerce, industry, and academia [I-044]. NCAF is currently chaired by

⁸As this report was commissioned by the UK Government’s Department of Trade and Industry, we have deliberately taken a UK-centric view here rather than attempt a comprehensive country-independent global survey.

Graham Hesketh, who works at Rolls-Royce’s Strategic Research Centre in Derby.

Recall also that this paper is addressing complexity and emergence with specific reference to ICT systems, and hence our review of academic activity concentrates on research in departments of computer science and electronic engineering; there are of course other such institutions housing talented and productive researchers that are omitted from our lists, and furthermore we know of related (but perhaps less immediately ICT-relevant) work in departments of mathematics, physics, biology, and social science, all of which could be included in a more comprehensive review, but which is beyond the scope of this document. For example, Chris Budd in the Mathematics Department at the University of Bath has recently been awarded a major grant to create the Bath Institute for Complex Systems: the funding (approx £1m) coming jointly from EPSRC and BBSRC sources.

As a final caveat, we note that although biologically-inspired or CAS-based approaches in computer science and electronic engineering can reasonably be described as novel, in the sense that they are non-traditional and relatively recent developments, they are increasingly being assimilated into the mainstream armoury of ICT engineers, in some cases largely replacing the previous mainstream methods. For this reason it is a somewhat arbitrary choice as to where to draw the line in deciding which companies should be included in the review of industrial activity presented here: there are probably many hundreds or thousands of companies of various sizes who now routinely use such techniques. Once again, the intention here is not to provide an exhaustive review but rather to provide an illustrative overview.

4.1 *Academic and Non-Profit Research*

Here we provide brief overviews of the major university and non-profit research centres where significant groupings of CAS researchers work. We have divided the review into two sections: research groups in the UK; and those in the rest of the world, which we deal with first. Both lists are presented in alphabetical order.

4.1.1 *The Rest of the World*

ATR Kyoto employed Tom Ray, developer of the award-winning *Tierra* and *NetTierra* artificial life systems [A-005].

Brussels (Université Libre de Bruxelles and Vrije Universiteit Brussel) has been a site of world-class research in artificial intelligence (AI), cybernetics, complex systems and systems biology for many years [A-006]. ULB is home to Marco Dorigo, who has pioneered research into new computer optimization techniques inspired by ant colonies. VUB is home to

Luc Steels, an early pioneer of biologically-inspired AI who has in recent years concentrated his research on the evolution of linguistic communication systems.

CalTech houses Chris Adami, responsible for the *Avida* artificial life system and also author of the first textbook on Artificial Life. Adami works within a nuclear physics group here [A-007].

Case Western Reserve University in Cleveland, Ohio, employs Randall Beer, who for many years has been applying CAS tools and techniques to better understand sensory-motor control and cognitive functioning in real animals and in synthetic systems.

École Polytechnique Fédérale de Lausanne

EPFL houses an Autonomous Systems Lab where lead researchers Dario Floreano and Aude Billard explore issues in biologically-inspired and CAS approaches to robotics, biology, and developmental psychology.

Los Alamos National Labs in New Mexico is nearby to the Santa Fe Institute (SFI; see below). Much of the research carried out at the Santa Fe Institute has been in collaboration with LANL, which has often employed SFI fellows [A-008].

MIT Artificial Intelligence Lab was recently reorganised under the title of “Combined Computer Science and Artificial Intelligence Laboratory”, the MIT AI Lab has since the mid-1980’s been responsible for major advances in biologically-inspired AI and robotics primarily under the guidance of its director Rod Brooks (see the discussion of *iRobot* in the review of industrial activity, below) [A-009]; with expertise in machine learning coming from Leslie Pack Kaelbling and from Paul Viola, and in evolutionary computation from Una-May O’Reilly. The MIT AI Lab also houses the *Amorphous Computing* group founded by Hal Abelson, Tom Knight, and Gerry Sussman, which draws inspiration from cellular and developmental biology in the design of new massively parallel computing systems, and which is also exploring using biological substrates (such as bacterial genetic regulatory networks) for new computing architectures [I-118].

MIT Media Lab researchers such as Pattie Maes, Mitch Resnick, Cynthia Breazel and Bruce Blumberg have generated considerable CAS-style research on autonomous software and hardware agents of various kinds [A-010].

New England Complex Systems Institute

(NECSI) hosts an annual international conference on Complex Systems, and also various workshops and tutorial courses. Coming to be seen by some as a *de facto* “Santa Fe Institute East” [A-011].

University of New Mexico employs Stephanie Forrest, a prominent researcher in evolutionary computation for many years, who has pioneered the development of computer security systems inspired by the actions of animal immune systems.

Santa Fe Institute (SFI) has been home over the years to many prominent complex adaptive systems researchers including Brian Arthur, Stuart Kauffman, and Chris Langton. For an excellent and informing account of the establishment and early activity of the SFI [A-012], see Waldrop's (1994) book, *Complexity: The Emerging Science at the Edge of Order and Chaos*.

University of Southern California employs faculty such as Michael Arbib, George Bekey, Maja Mataric, and Stefan Schaal, who pursue novel approaches to AI in the form of behaviour-based robotics and computational neuroscience [A-013].

University of Zurich has an Artificial Intelligence Lab run by Rolf Pfeiffer, which for many years been a centre for research into biologically-inspired artificial intelligence and CAS approaches.

4.1.2 UK Universities

University of Birmingham run a research group specialising in nature-inspired computation, led by Xin Yao [A-014]; commercial exploitation and dissemination of the expertise of this group is now handled by a special institute, the Centre of Excellence for Research in Computational Intelligence and Applications (CERCIA) created with the assistance of Birmingham's regional development agency Advantage West Midlands.

University of Edinburgh contributed strongly to the development of biologically-inspired AI during the 1980s and 90s through the work of roboticists such as Tim Smithers, Chris Malcolm and John Hallam; and computational neuroscientists such as David Willshaw. More recent work has included Barbara Webb's bio-mimetic robotics within the Institute for Adaptive and Neural Computation [A-015], and Simon Kirby's work on the evolution of human languages.

University of Essex recently recruited Owen Holland and Ulrich Nehmzow who are involved in a number of autonomous robotics research activities involving CAS concepts; Ricardo Polli is a noted evolutionary computation researcher; while Edward Tsang and Sheri Markose (among others) are involved in a new center for computational finance and economic agents [A-017].

University of Exeter has recently appointed a number of new faculty members who pursue CAS-related research: Martyn Amos is an expert in molecular and cellular computation; David Corne works in evolutionary computation; and Todd Kaplan is an experimental economist who wrote one of the most successful early trading-agent algorithms [A-016].

University of Leeds hosts Seth Bullock, Jason Noble and Netta Cohen who work on bio-inspired computing and biological modelling within the Biosystems group, and ran a recent EPSRC-funded 12-month research cluster on *Simple Models of Complex Networks* [A-018].

University of Liverpool Department of Computer Science has a BioComputing and Computational Biology group whose membership includes Ray Paton (specialist in cellular computing) and Peter McBurney (specialist in CAS approaches to multi-agent systems).

London School of Economics hosts the *Complexity Programme*, directed by Eve Mitleton-Kelly, which celebrated its 10-year anniversary this year [A-019].

University of Oxford features work on complexity in biological, cultural and socio-economic systems through the work of researchers such as Janet Efstathiou, Neil Johnson, Felix Reed-Tsochas and Gesine Reinert, who ran a recent EPSRC-funded 6-month research cluster on *Complex Agent-based Dynamic Networks* [A-020].

University of Southampton School of Electronics and Computer Science (ECS) has a recently-formed "Bio@ECS" research group which gathers together the 26 faculty members who have active interests in biologically-inspired complex adaptive systems and related issues. Notable researchers in this initiative are John Shawe-Taylor and Nigel Shadbolt.

University of Sussex has been a centre for complex adaptive systems research for roughly 15 years. Their *Centre for Computational Neuroscience and Robotics* is run by Phil Husbands and Mick O'Shea, and involves, amongst others, Ezequiel Di Paolo, Inman Harvey, Emmet Spier, and Adrian Thompson, all of whom work on adaptive biological and computational systems [A-021]. Harvey also co-directs Sussex's *Centre for the Study of Evolution* which explores natural and artificial evolutionary processes from a CAS perspective.

University of the West of England has long-standing research activity in autonomous robotics (Chris Melhuish), non-linear computational media (Andrew Adamatsky), and adaptive computation (Larry Bull) [A-022].

University of York has CAS-related research underway both in its Department of Computer Science, where Jim Austin leads a large group of neural network researchers and Susan Stepney works on artificial life and evolutionary computation; and also in its Department of Electronic Engineering where Andy Tyrrell, Julian Miller, and others work on biologically-inspired approaches to electronic hardware design and implementation, drawing inspiration from embryology and from evolution.

4.2 Industry: Global ICT Companies

Blue-chip ICT companies are generally big and consequently have diverse business interests, and the deployment of CAS techniques are not a dominant or central theme for any of them. However, some do have CAS-related products, and some ICT corporate research labs conduct work in this area. The following sections describe the research and products of (in alphabetic order): BT, HP, IBM, Intel, Microsoft, Mitsubishi, NEC, Oracle, PARC, Sony, and Sun.

Of these, Microsoft and IBM have the most clear CAS research agenda and have a number of CAS-related products; while BT, NEC, and PARC are more research-focused; Oracle is more product-focused; and Intel and Sun have relatively little investments in the area as far as we can determine. Sony and HP sit in the middle of this range, with some relevant research and some relevant products.

4.2.1 British Telecom

BT Labs has a long history of working in CAS-related areas, primarily from a perspective of applications in telecommunications, but also with significant involvement in funding more blue-sky research in several UK universities. British Telecom's Labs currently operate using the name BTEExact [I-022]. Core CAS researchers and team-leaders at BTEExact include Mark Shackleton, Fabrice Saffre, Paul Marrow, Sverrir Olafsson, and Robert Ghanea-Hercock. A recent issue of the BT Technology Journal [I-023] provides an overview of work at BT on CAS related areas. From the BTEExact website, it is apparent that this work is divided into a number of areas:

Complexity research [I-024]. This group seeks to develop new approaches to the analysis and the modelling of large and complex network systems using methods drawn from statistical mechanics, chaos theory and non-linear dynamical systems. They study the dynamic properties of various network systems with particular focus on aspects of complexity as it emerges in data networks, distributed filing systems and access systems. Particular emphasis is on the

modelling of rapid performance deterioration and its predictability within probabilistic frameworks.

Emerging technologies [I-025]. This group is interested in ad-hoc working, parasitic networks, futurespace (flexible workplace environments) and smart interfaces.

Evolutionary algorithms [I-026]. EOS is a software platform developed by BT's Future Technologies Group that is similar to Swarm [I-027], an agent-based modelling framework developed at the Santa Fe Institute. EOS supports research and rapid implementation of evolutionary algorithms, ecosystem simulations and hybrid models. It also supports fast prototyping of industrial applications that use these technologies.

Future technologies [I-028]: The Future Technologies group is interested in nature-inspired computing, new approaches to computation, mobile software agents (both malicious and benign) and information ecosystems,

Intelligent agents [I-029]: BT's Agent Research group is concerned with the development and analysis of sophisticated artificial intelligence problem-solving and control architectures for both single-agent and multiple-agent systems. Current research themes include multi-agent coordination and negotiation protocols, cooperative and non-cooperative multi-agent systems, organisational self-design, multi-agent adaptation of coordination strategies, computational economics as well as multi-agent building platform. This group has developed the *Zeus* agent-building toolkit

Intelligent business systems [I-030]. This group has done work on scheduling for repair personnel, exploiting information assets and competing in electronic markets. They have particular expertise in agent based systems, agent platform development, ontology management and self-organising adaptive systems. Currently the repair personnel scheduling system is the best example of BT leveraging this research activity in their business.

4.2.2 Hewlett-Packard

HP Labs have headquarters in Palo Alto, California, and a European base in Bristol, UK. The Bristol site is home to HP's Complex Adaptive Systems research group [I-036], founded by Dave Cliff [I-037]. Cliff's work for HP has included research on adaptive trader-agents and automated market-mechanism design [I-038], and although the results of this research are primarily intended for market-based control of complex ICT systems, a number of major institutions in the global financial markets

have taken a keen interest in this strand of research. Other work in the HP CAS group, by Matt Williamson [I-039], has explored CAS approaches to computer security [I-040; I-041] resulting in the development of “virus throttling”, a method that limits the spread of “malware” (mobile malicious code such as computer viruses and worms) [I-042]. Additional projects recently conducted by the HP Labs CAS group include the development of document-classification mechanisms inspired by the human immune system, evolutionary optimisation of print-layout problems, and automated design and optimisation of storage area networks (SANs) using methods inspired by ant foraging and by Darwinian evolution. Members of the HP CAS group work with a number of external companies on research collaborations (currently focusing on partnerships developing applications in financial markets and investment banking, and in the energy industry). HP’s Palo Alto Labs are home to the Information Dynamics Laboratory, led by Bernardo Huberman, where a number of CAS-relevant projects are also underway [I-043].

Other CAS-type research conducted at HP Labs includes work by Barry Shackelford on using evolved configurations of field programmable gate arrays in order to predict protein conformation; Ira Cohen’s work on Bayesian Networks [I-044]; Evan Kirshenbaum’s work on genetic programming [I-045]. As part of its *Open-View* product line, HP has since 1999 been offering “self-healing servers” [I-050], and more recently has been promoting a suite of *Adaptive Enterprise* technologies; both of which are broadly similar to IBM’s *Autonomic Computing* initiative, discussed further below.

4.2.3 IBM

IBM is involved in a number of CAS-related research activities: The Information Economies Group [I-009] worked for several years on trading agents, exploring techniques developed at IBM and also at HP Labs. Several of IBM’s researchers involved in this project, Jim Hanson, Raja Das and Jeff Kephart have all previously worked at the Santa Fe Institute [I-010]. The work of the anti-virus group [I-011] is also relevant, although that research appears to have terminated some time ago. IBM are also interested in novel ways of processing text and information [I-012]: for example the *Clever* project [I-013] explored collaborative filtering and hypertext classification for applications such as bookmark clustering. Other research areas include using optimisation to solve constraint based factory scheduling problems and the *Quest* data-mining group [I-014].

IBM are now deploying CAS techniques in several products. Their premier database product *DB2* now incorporates plug-ins that can perform image query-by-example [I-015] and audio query-by-example [I-016]. They also sell one of the premier data mining tools, *Intelligent*

Miner [I-017] which can be applied to data or text [I-018] for identifying and extracting business intelligence from data assets. *Lotus Discovery Server* [I-019] is a web-based knowledge management tool that provides text sharing, searching, classification and expertise location. Their *eLiza* self-healing servers [I-020] are also relevant.

In 2001, IBM announced their major *Autonomic Computing* initiative that encompasses both research and product development, and which [I-021] proposes to take biological inspiration in order to provide server systems that are self-organising, self-managing/self-regulating, and self-healing. IBM’s stated reason for using the term “autonomic” is because “[a computer system] must act like our autonomic nervous systems. It must provide an unprecedented level of self-regulation while hiding complexity from the user. And it will be a radical shift in the way we conceive and develop computing systems today. This will call for more than retooling old systems – autonomic computing calls for a whole new area of study.” [I-046].

Since launching their autonomic computing initiative, IBM have actively sought to foster global collaboration and cooperation on relevant work in both the academic and the industrial research communities. IBM have staged a number of informal meetings to encourage such interactions,⁹ and IBM researchers have also led the organization of the first international conference on autonomic computing (ICAC-04) [I-047]. Some quarters of the academic community have shown a resistance to describing their work as “autonomic computing”, in the belief that it is somehow too closely associated with the one big computer company; the phrase “self-star” computing is commonly used as a vendor-neutral term when organising conferences and workshops (e.g., [I-048].)

4.2.4 Intel

Intel, as far as we can see, has neither a large investment in CAS research and nor does it offer any CAS-based products. They have done some work on technologies for text processing, information extraction, retrieval and classification. They have aimed this work specifically at the Pocket PC platform. An Evaluator Toolkit [I-054] is available that uses a vector-classifier based on the vector space model for performing categorisation tasks [I-045].

4.2.5 Microsoft

Microsoft Research has been investigating a variety of CAS techniques. They have organised this work into three areas:

The Adaptive Systems & Interaction (ASI) group [I-001] is working on automated reasoning, adaptation and human-computer interaction. ASI also does research on

⁹For example, IBM hosted a one-day meeting on Autonomics at the IEE’s offices in London, in November 2003.

information retrieval and management, including work in automated text classification and clustering.

The Data Management, Exploration and Mining group [I-002] works on exploiting data mining techniques, i.e., applying statistical and machine learning techniques to detect patterns in databases.

The Machine Learning and Applied Statistics [I-003] group is focused on learning from data and data mining. By building software that automatically learns from data, they enable applications that perform intelligent tasks such as handwriting recognition and natural language processing, and help human data analysts explore and better understand their data more easily.

In addition Microsoft are also conducting research on collaborative filtering [I-004], natural language processing [I-005] and face detection and recognition [I-006].

The work on adaptive systems and interaction has led to the Microsoft Agent toolkit [I-007]. Work from the data management, exploration and mining work was recently incorporated into Microsoft SQL. Microsoft uses Bayesian networks (a machine learning technique) for several applications including the printer trouble-shooter in Windows, Office Assistant, and the Microsoft customer support line. They are also investigating other machine learning techniques such as Support Vector Machines [I-085] for use within their information management tool *Sharepoint* [I-008].

4.2.6 Mitsubishi

Mitsubishi Electric Research Laboratories (MERL) [I-059] are the North American arm of the central R&D organisation of the Mitsubishi Electric Company. MERL's artificial intelligence research has pursued very few projects that can reasonably be described as having CAS as a central theme. In collaboration with the Brandeis University DEMO Lab [I-060] one MERL researcher worked on evolutionary optimisation of 3D machines and mechanisms through simulation studies [I-061], but this project concluded some years ago.

4.2.7 NEC

NEC Corp. [I-062] operates research laboratories in Japan, USA, Europe, and China, the best known of which is their US operation: NEC Labs America Inc. which was created in November 2002 by the merger of the Princeton (New Jersey) based NEC Research Institute and the Cupertino (California) based NEC Computer and Communications Research Lab. Major areas of research interest with relevance to CAS include: Machine Learning (specifically, advancing the state of the art in support vector machines); Bioinformatics (experimental and computational approaches to discover new protein folds or to synthesise proteins with new folds); Robust and Secure Computer Systems (discovering new

'cognitive' learning techniques that are unsupervised, distributed, deal with dynamic environments, and learn on-line so that computer systems can adapt and evolve to better levels of robustness or security through autonomous learning and reasoning); and Broadband and Mobile Networks (self-organising *ad hoc* communication networks). NEC Labs' website states that "The primary focus is on technology research and early market validation in support of NEC's core businesses.", but little information is available regarding what inventions from NEC Labs have fed into NEC's products or businesses. The NEC Labs website includes a web-page entitled "Licensable Technologies" which, at the time of writing this document, contained the single sentence "Information to come". In academic ICT circles, probably the most famous output from NEC Labs is *CiteSeer*, a set of scientific-literature digital library tools that indexes research articles on the web, and which most notably performs autonomous citation indexing so that researchers can track how many times a particular paper is cited, and by whom it is cited [I-063].

4.2.8 Oracle

Oracle appears to be doing little in the way of CAS related research. However as it is the leading provider of database software, and because data mining is very strongly tied to databases, it is natural that they would have some interest in this area. Therefore Oracle do sell a number of products that use CAS-related techniques for business intelligence [I-034] both for performing on-line analytical processing, i.e., identifying credit card fraud in real time or offline data mining [I-035]. Their data-mining product, DARWIN, was originally created by Thinking Machines, the company responsible for the *Connection Machine* massively parallel computer that was a brief commercial success in the late 1980's.

4.2.9 PARC

PARC [I-031] (formerly Xerox PARC) is one of the most famous computer science corporate research labs in the world. A recent review [I-032] by Jacob Nielsen noted that PARC is the only corporate research lab to feature in the top three computer science research labs in each of the past three decades. However despite being a noted success from a research perspective, Xerox was not always successful in commercialising the research undertaken at PARC, and PARC was recently spun-off as an independent research facility. It is currently undertaking projects in three areas: *smart matter*, which encompasses micro scale devices and integrated systems; *knowledge ecologies*, which encompasses analysing document content; and *sensemaking and community networks & documents*, which encompasses image processing and mobile or wireless computing. Previous CAS work at

PARC included seminal work on market based control [I-033] undertaken by Bernardo Huberman and Tad Hogg who are both now research staff at HP Labs in Palo Alto.

4.2.10 Sony

Sony is primarily a manufacturer of consumer electronics with a product range including PCs, handheld computers, and their hugely successful PlayStation series of games consoles, yet they also have a well-publicised interest in consumer robotics. Early CAS related research at Sony was concentrated on AIBO [I-055], a robotic dog, which was initially demonstrated as a prototype in 1995. Recently Sony developed an impressive humanoid robot prototype [I-056, I-057]. Equipped with video cameras and seven microphones, it recognises faces and gestures, and has a 60,000 word vocabulary. For many years Sony has operated a dedicated research facility in Japan but more recently it also set up a European centre: Sony Computer Science Lab Paris [I-058]. The Paris Lab is investigating computational approaches to language and evolutionary linguistics that has led to a research project on robot language-formation called *Talking Heads*. The other main area of study in the lab is music, specifically the creation of adaptive listening environments. Luc Steels, an academic with a long history of world-class research in biologically-inspired robotics and artificial intelligence, is a prominent employee of CSL.

The AIBO dog robot has been made commercially available. However the commercial versions are slightly simpler than the research versions, and are comparatively expensive consumer electronics items.

4.2.11 Sun Microsystems

Sun, despite their well-publicised interest in technical and grid computing, have not been involved in CAS related research as far as we can determine.

4.3 Industry: Large non-ICT Companies

4.3.1 BAE Systems

BAE Systems [I-086] is a major European defence contractor, based in the UK. BAE Systems has its primary roots in the old British Aerospace company, but it now owns companies such as Vickers (submarines) and Royal Ordnance (tanks/artillery). They have a considerable history of research involvement in biologically-inspired CAS approaches to a variety of problems, from evolutionary design of components, through automated assembly in manufacturing via autonomous robots, to distributed and decentralised systems for minefield clearance and for battlefield asset management.

One notable non-ICT example of biological inspiration in BAE research is their exploration of microvor-

text aerodynamic permeable skins for variable-dynamics fixed-geometry airfoils. This work is inspired by the skin of sharks, which can alter its local hydrodynamic smoothness/roughness via muscular control of small scaly patches on the skin. BAE is experimenting with aerodynamic surfaces (aircraft wings) that incorporate a high-density mesh of very small vents that each exit on the outer surface of the wing. Compressed air can be pumped through the vents, creating small vortices that disrupt laminar airflow over the outer surface. Varying the distribution and magnitude of the compressed air flows to the vents alters the aerodynamics of the wing in desirable ways, without the need for any moving metal surfaces.

BAE also have an interest in CAS approaches to battlefield asset management. This is a response to a perceived desire in BAE System's customers for alternatives to AWACS-style¹⁰ centralised command-and-control systems. The desire is for alternatives that are distributed/decentralised, and potentially that also have a higher degree of autonomy, requiring less second-by-second control from human operatives. One exemplar application in particular that has received significant research attention (in academic labs, at least) is automated/autonomous clearance of landmines by swarms of cheap disposable (i.e., sacrificial) robots, where there is a danger that humans may be injured or killed.

Like HP, BT, and IBM, BAE Systems has an interest in self-healing, self-regulated distributed dynamic control; they are a partner in the recently-announced major EPSRC-funded project mentioned previously in this report (see Footnote 2). BAE Systems have also explored the use of teams of autonomous robots for assembly and repair of complex products (e.g. combat aircraft). This is necessary because computer aided engineering techniques can design very efficient airplanes but such designs often have the problem that the airplanes are very hard to assemble and/or to repair. One way to avoid this problem is to use autonomous robots in place of tool-wielding humans to perform the assembly and repair tasks. BAE Systems have also funded more abstract theoretical studies in areas such as ecology and game theory. Researchers at BAE Systems Advanced Technology Center with active interest in CAS approaches include Hector Figueiredo and Andy Wright.

4.3.2 NCR

The National Cash Register [I-087] company's primary business is supporting other companies when they interact with customers across the counter, by telephone, at a kiosk, at an ATM machine or over the Internet. As they have a clear business interest in consumer behaviour they are interested in areas such as data warehousing and

¹⁰AWACS: Airborne Warning and Control System.

data mining [I-088]. They have also done some work on agent based modelling of retail environments. In the mid 1990s, NCR set up a specialist research lab in London called KnowledgeLab [I-089]. KnowledgeLab carried out research on intelligent appliances such as bank-access via games consoles, and also (somewhat bizarrely) internet-enabled microwave ovens [I-090]. NCR also has a machine learning research group that has strong links with researchers from the Neural Computing research group at Aston University, UK.

4.3.3 Unilever

Unilever is a large multinational company that sells food, home-care and personal-care products with approximately 270,000 employees worldwide. They have several research facilities but the Adaptive Computation Group at Unilever Research Port Sunlight (UK) in particular has a history both in artificial intelligence research and in CAS-related research. This group has applied CAS techniques to several different areas including:

Product Design: There was a three year project at Unilever to develop software based on Bayesian Neural Networks and Genetic Algorithms to evolve product formulations for domains such as washing powder. These techniques were also used in the formulation of margarine, because the cost of its ingredients can vary drastically on a monthly basis requiring constant product reformulation in order to achieve cost efficiency.

Consumer Understanding: Unilever has a considerable investment in “Consumer Science” (studies of human-computer interaction and cognitive ergonomics). Unilever’s Adaptive Computation group has applied a number of CAS-related techniques to mining data obtained from consumers including graphical models and neurofuzzy modelling.

Natural Language Understanding: As Unilever is a multi-national company comprising of many localised operating units, there is considerable interest in automatic machine translation. There is also interest in using classification and searching techniques such as Latent Semantic Analysis in order to analyse textual information obtained from consumers discussing products.

Supply Chain Optimisation: Unilever has complex, highly diversified multi-national supply chains. This diversity can be advantageous, but increasingly major customers such as WalMart want very fast response to product stocking requests. This requires a high level of supply chain optimisation so Unilever has been investigating using agent based modelling in order to optimise supply chains.

Unfortunately there is little published material about this work outside Unilever, but one interesting paper discusses how evolutionary algorithms can be used to generate patents for bactericidal peptides (Patel, Scott, Bhakoo, & Elliott, 1998).

4.4 Industry: SMEs

There are a host of smaller companies with CAS-style product or service offerings. For the purposes of this report, the companies can be collected into five rough groupings by topic area. These are Hardware, for companies with an emphasis on physical (electronic and electromechanical) products; Business Simulation, where the focus is simulating the ramifications of strategic or tactical decisions in order to aid planning; Business Operations, which include companies offering on-line data processing systems for businesses; Entertainment, which is primarily animation and computer games; and Consulting/Software services. One again, the intention here is not to produce an exhaustive directory of relevant SMEs, but rather to list the activity of some notable exemplar companies.

4.4.1 Agorics

Company Type: Consulting and software development. **Location:** Los Altos, California. **Formation:** Founded in 1994 by Ann C. Hardy and Mark S. Miller, et al. **Target Industries:** Secure e-Business Solutions. **Customers:** Sun Microsystems. **Technologies:** Market-based resource allocation; secure e-mail systems.

At the time of their founding, Agorics [I-104] were centred on developing Market-Based Control (MBC) systems, working on a major project for MBC of network quality-of-service with Sun Microsystems. One high-profile Agorics co-founder was Mark Miller, who co-authored (with K. Eric Drexler) three papers on MBC (which Miller and Drexler referred to as Agoric Computation) in Bernardo Huberman’s influential 1988 edited collection *The Ecology of Computation*. On the current Agorics website, these three papers (and others on auctions and markets) are listed in their “Tech Library”, but the only current product offered is a secure corporate email management system. Thus it would appear that Agorics have either not been able to turn their MBC expertise into profitable products, or that they are unable to publicise any such success.

4.4.2 Apama

Company Type: Software development and consultancy company. **Location:** Cambridge, UK. **Formation:** Spin-off from research at Cambridge University. **Target Industries:** Financial, Telecoms, Mobile location based services, Customer Relationship Management, Supply

Chain Management. Customers: several investment banks, plus alliances with Oracle and Sun. Technologies: Distributed computing.

Apama [I-101], like Searchspace, provide on-line processing solutions that continually monitor streams of data for complex patterns of events, and provide real-time alerts when a match is found. Their software platform, the Apama Engine, indexes ‘monitors’ used to specify patterns of events to be watched for, and checks them against incoming data streams. The Apama Engine is a general-purpose technology; it can be applied to a range of applications. It is particularly well suited for applications in which data is constantly changing, where there is significant value in being able to react quickly to those changes, and where systems must operate on a large scale.

4.4.3 *Autonomy*

Company Type: Software development company. Location: Cambridge UK. Formation: Set up by Mike Lynch from Cambridge University. Target Industries: Telecoms, Energy, Public Sector, Technology, Life Sciences, New Media, Professional Services. Customers: Ericsson, Astra Zeneca, Unilever, ZkB, Pfizer, Novartis, Sonera, BAE, McGraw Hill. Technologies: Bayesian probability and Shannon’s information theory.

Autonomy [I-103] is a software development company that sells various products that perform text classification and search. Autonomy products can perform automatic categorisation, hyper-linking, retrieval and profiling of unstructured information, thereby enabling the automatic delivery of large volumes of personalised information. This can be used for enterprise portals, e-commerce, business intelligence, consumers accessing mobile or digital TV portals, knowledge management or customer and relationship management.

4.4.4 *BiosGroup*

Company Type: Consulting and software development company. Location: Sante Fe, New Mexico, US. European subsidiary, EuroBios, based in London and Paris. Formation: Joint venture between the Centre for Business Innovation of Ernst & Young (now Cap Gemini Ernst & Young) and Stuart Kauffman [I-105]. Target Industries: Food industry, airlines, automotive industry, energy, entertainment, financial services, army, manufacturing industry, telecoms. Customers: BT, Unilever, Procter and Gamble, Air Liquide, Cap Gemini, Ford Motor Company, Honda, Ivensys, SAP, Southwest Airlines, NASDAQ Stock Market. Technologies: Agent based modelling and multi-objective optimisation.

BiosGroup [I-106] claim to have pioneered the use of complexity science to solve complex business problems

and to now be the world leader in applying the techniques of this emerging science to large commercial applications. The Company’s European subsidiary, EuroBios [I-107] has locations in Paris and London, and was originally founded by Eric Bonabeau, who had previously published research on new optimisation techniques inspired by collective behaviour in ants. BiosGroup specialise in the techniques of agent-based modelling and multi-objective optimisation. They apply these techniques to supply networks and e-procurement, market analysis and planning, adaptive scheduling and routing, decision support, risk analysis and strategic and tactical simulation. In the past year, the US-based BiosGroup was acquired by NuTech Solutions Inc. [I-064], leaving EuroBios as an independent entity. Bonabeau remains as EuroBios’s Principal Scientific Advisor, but is now also Chairman and Chief Scientific Officer of a US company called Icosystem [I-065] which offers broadly similar solutions and services as EuroBios. Currently the President and CEO of EuroBios UK is Vince Darley, while the President and CEO of EuroBios Paris is Hervé Zwim.

4.4.5 *Cambridge Neurodynamics*

Company Type: Software development and consultancy company. Location: Cambridge, UK. Formation: Set up by Mike Lynch from Cambridge University. Target Industries: Security, Surveillance, Biometrics, Intelligence Gathering. Customers: Yorkshire Police Force. Technologies: Bayesian probability, neural networks, Shannon’s information theory and wavelet coding.

Cambridge Neurodynamics [I-093] is software development and consulting company founded by Mike Lynch (who also founded Autonomy) that specialises in the development of recognition systems. The company is organised into three groups: advanced systems, biometrics and witness. The Advanced Systems Group investigates the application of technologies such as Bayesian probability, neural networks, Shannon’s information theory and wavelet coding. They also work on problems involving pattern recognition, image and audio processing, texture analysis and synthesis and three dimensional processing and imaging. The Biometrics Group works on applying the techniques developed by the ASG to problems such as fingerprint recognition or face recognition for security systems. The Witness Group is concerned with security and surveillance. It produces solutions for car number plate recognition or very large scale (250+ cameras) surveillance systems.

4.4.6 *Forio*

Company Type: Software consultancy. Location: San Francisco, USA. Formation: Set up by two MIT graduates. Technologies: Business Flight Simulators.

Forio [I-095] is a company based in San Francisco that produces business simulations that allow managers to “learn through experience without the cost of experience”. Firstly they have created a simulation that investigates the competition in pricing between HP/Compaq and Dell [I-096]. They have also created a simulation called PDASim [I-097] that allows managers to manage a portfolio of PDA products across multiple product life-cycles. This simulation allows managers to learn how to use financial data to make pricing and product line decisions and experience how decisions can have consequences many years into the future. Forio maintain an interesting set of resources on simulation [I-098]. Their main business is writing simulations but they also host simulations, offer training via simulations, and partner with training or consulting organisations to deliver services to clients.

4.4.7 Gameware Development

Company Type: Entertainment software. Location: Cambridge, UK. Technologies: Game software architecture using neural networks, genetic algorithms and biochemical control.

Gameware Development is the company formed after a previous incarnation, CreatureLabs [I-114], went into liquidation. Gameware create games using a development environment conceived by Steve Grand (who now operates as a sole-trader inventor, under the name Cyberlife Research) that is a bottom-up, agent-oriented, software architecture incorporating powerful biochemical, neural network and genetics modelling systems, capable of creating artificial intelligent systems and believable life-forms. Creature Labs can produce life-like virtual organisms in a wide range of simulated environments, producing persistent environments that can be indefinitely upgraded and expanded by exchanging, creating and deleting the agents that make up the simulation. Gameware’s highest-profile project to date was providing the artificial life and 3D creature-modelling for a groundbreaking children’s TV format called *Bamzooki*, first broadcast by the BBC in early 2004 (with a new series currently in production).

4.4.8 GMAP

Company Type: Software development and consultancy. Location: Leeds, UK. Formation: Spin-off from University of Leeds. Target Industries: Retailing, Forecourt, Financial and Automotive. Customers: Ikea, Dixons, Asda-WalMart, Exon-Mobile, BP, Abbey National, Alliance and Leicester, Ford, Jaguar and Mazda. Technologies: Geo-mathematical and statistical models to represent given markets using geodemographics, consumer demand, travel patterns and competition data obtained from real-life sources.

GMAP [I-094] build simulators that help companies answer questions such as where to build their next superstore. They specialise in providing market intelligence and decision support solutions to global retail organisations that are looking to improve the efficiency of their network, their sales territories and their channels to market. These programs use geo-mathematical and statistical models in order to represent given markets using geodemographics, consumer demand, travel patterns and competition data obtained from real-life sources. This allows companies to determine “what is actually happening” within the market, “what might happen” given certain scenarios and “what should happen” within an optimal situation, enabling organisations to develop highly efficient networks and channels to market for each of their products.

4.4.9 Icosystem

Company Type: Software development and consultancy. Location: Boston, Massachusetts; Paris. Formation: founded in 2000 by Eric Bonabeau, noted swarm-intelligence and complex adaptive systems researcher. Target Industries: pharmaceuticals, energy, consumer packaged goods and software. Technologies: Agent-based business modelling for strategic scenario evaluation and optimisation.

Icosystem’s [I-099] business model is to act as a novel source of strategic analysis, exploration, evaluation, and exploitation for client companies that are typically market or sector leaders, and fairly large (i.e. Fortune 500 or similar). The Icosystem website stresses that Icosystem is not a consultancy company in the traditional sense, and that they “. . . avoid the temptation of trying to replicate anything that works with any and all companies in a sector. On the contrary, Icosystem does not normally work on similar projects for competing organizations”. Icosystem’s main mode of operation is to create a bespoke agent-based software simulation model of the relevant aspects of a client’s business, which replicates the behaviours of the key business participants and captures the interactions between them; then to use that simulation model for exploring “what if” questions concerning possible alternative strategies; followed by evaluation and exploitation of the most promising strategic alternatives that were identified in the exploration phase.

4.4.10 iRobot

Company Type: Robot development company. Location: Somerville, Massachusetts. Formation: Co-founded by Prof. Rodney Brooks, Director of the MIT Artificial Intelligence Lab. Target Industries: Defence, civilian disaster search/recovery, space exploration, industrial maintenance, domestic assistance, and

entertainment. Customers: DARPA, Hasbro, Universities. Technologies: Biologically-inspired autonomous robotics.

iRobot was co-founded by Rodney Brooks, Director of the MIT Computer Science and Artificial Intelligence Laboratory. From the mid 1980's to the mid 1990's, Brooks provoked a radical re-appraisal of established artificial intelligence (AI) practices in the design of control systems for autonomous mobile robots. In a series of papers (often describing working prototype robots) he argued, by reference to the evolved "designs" of the "controllers" (i.e., nervous systems) of simple creatures such as insects, that it should be possible to avoid the then-traditional complex centralised and hierarchical AI robot-controller architectures; replacing them with simple decentralised and heterarchical layered control architectures instead. In doing so, Brooks and his students made major contributions toward establishing the approach now known as biologically-inspired or behaviour-based AI, which today is widely (if still somewhat begrudgingly, in certain quarters) accepted as a viable alternative to the longer-established traditional approaches, and indeed is frequently seen as providing tools and techniques that are significantly better than those provided by the traditional methods.

The iRobot [I-092] company was set up by Brooks in order to commercially exploit his research on biologically inspired robotics. Their robots are designed for a variety of applications such as defence, civilian disaster, space exploration, industrial maintenance, entertainment, and domestic assistance. iRobot produce a standard platform for developers to create value-added robotic applications. They produce products such as *Coworker*, a wireless, mobile, remote telepresence platform that provides control of video, audio and movement through any internet browser without additional hardware or software, *Microrig*, a conveyance device for carrying sensor payloads and tools down into oil-well bores and *My Real Baby*, a toy aimed at young girls. Two of their more successful products are *Packbot* and *Roomba*. *Packbot* is a portable semi-autonomous tracked ground-vehicle for surveillance and access to hostile environments, used by recovery teams at the site of the two collapsed New York World Trade Center towers, and by US armed forces in the subsequent military action in Afghanistan. *Roomba* is a small autonomous vacuum-cleaner robot suitable for domestic homes, which sold in huge numbers in the run-up to Christmas 2003, no doubt aided by its receipt of a "seal of approval" from *Good Housekeeping* magazine.

As well as producing standard robots, iRobot is also engaged in a number of research projects such as DAMP, Distributed acoustic mobile positioning; Swarm, distributed programming of robots; Gecko, a wall climbing robot; and MUMS, the micro unattended mobility system. More details of these and other projects are

available on the iRobot website [I-092].

4.4.11 *Lostwax*

Company Type: Software development and consulting. Location: Surrey. Formation: Co-founded by Nick Jennings, a leading researcher in field of agent technology. Target Industries: E-Commerce. Customers: General Motors, British Airways, Egg Financial Services, City Index, MoneyXtra, Sony, Psion, Orange, Prudential. Technologies: Software agents.

Lostwax [I-111] is a software development and consulting company. Agent Technology is the latest approach to analysing, designing and building software. Agents create smart adaptable solutions that are particularly suited to complex; changing environments. Lostwax build e-commerce and portal systems on J2EE and .NET frameworks using agent architectures. They are members of two associations for agent technology: FIPA and AgentLink.

4.4.12 *Natural Selection*

Company Type: Software development and consultancy. Location: La Jolla, CA. Formation: L.J. Fogel, founder, invented Evolutionary Programming. Target Industries: Defense, Medical (diagnostic, and drug design), Finance. Customers: None named. Technologies: Evolutionary Algorithms, Neural Networks, Fuzzy Systems.

NSI [I-113] concentrates on the use of biologically-inspired algorithms for optimisation, signal processing, and control applications industry, medicine, and Defense. Their website is non-specific about client companies, stating that "Natural Selection, Inc. assists corporations and government entities in the areas of production and transportation scheduling, signal detection and pattern recognition, financial forecasting, mission and path planning, inventory control, bioinformatics, and computational drug design, as well as other applications that demand the rapid solution of difficult combinatorial and temporal optimisation problems" but without naming names. They were founded in 1993.

4.4.13 *Maxis*

Company Type: Entertainment software. Technologies: Business Simulators.

Electronic Arts [I-115] produce entertainment software, specifically the "Sim" family of simulations: The Sims, Sim City, Sim Coaster and Sim Golf. Although they no longer work on business simulations, they launched a spin-off enterprise called Thinking Tools Incorporated [I-116] that produced a simulator called Telesim [I-117] in 1994 that generated a lot of industry interest. Although TTI no longer exist, companies such

as BiosGroup, Forio, or IcoSystem are currently active and offering similar services.

4.4.14 *Norkom*

Company Type: Software development and consultancy company. Location: Dublin, Ireland. Target Industries: Telecoms, Financial Services and Insurance. Customers: ING Direct, HSBC Bank, Vhi Healthcare, British Airways, Esat Digifone, Actel Direct. Technologies: Predictive modelling.

Norkom [I-102], like Searchspace and Apama, provides software and services for on-line processing. However their software has a much stronger bias towards pattern recognition via predictive modelling, unlike Apama and Searchspace where the focus is very fast online processing. Norkom targets problems such as customer relationship management by creating predictive models that can be used to help client retention, cross-sell, up-sell, acquisition, churn and channel migration. They sell a software platform called *Norkom Alchemist* that allows integration with third party tools and models (e.g., data mining tools) as well as independent data (e.g., demographic data and other internal information systems etc).

4.4.15 *Sana Security*

Company Type: Software development and consultancy. Location: San Mateo, California. Formation: Start-up from research at University of New Mexico. Target Industry: computer security. Customers: IT service providers, global investment firms, nationwide multi-channel retailers, major government agencies. Technologies: Biologically-inspired computer security systems.

Sana Security [I-066] develops and markets “host-based intrusion prevention software” (HIPS) that protects computer systems from known and unknown attacks, with low and predictable operating costs. Sana was established to commercialize “Adaptive Profiling Technology”, developed by Sana’s founder Steven Hofmeyr who was a PhD student of Stephanie Forrest at the University of New Mexico, and Forrest is a member of Sana’s Advisory Board. Sana have also recently recruited Matt Williamson, the inventor of virus throttling, from HP Labs.

Sana claim that their first product, Primary Response, protects the broadest range of platforms and applications, and requires fewer resources to manage, deploy and scale. It does this by eliminating the need for constant updating and management by security experts. Sana state that Primary Response is the world’s only security software approach based upon the principles of the human immune system, and that the solution provides customers with out-of-the-box code-injection suppression, proactive vulnerability exploitation detection, and real-time prevention from attacks by worms and

hackers. Code injections are the largest class of vulnerability attacks on computer systems and applications. Deployments of Primary Response in organizations reduce or eliminate the need for frequent “patching” of fixes to flaws in security systems.

4.4.16 *Searchspace*

Company Type: Software development and consultancy. Location: London, UK. Formation: Spin-off from University College London Intelligent Systems Group. Target Industries: Banking, Security Broking, Telecoms, Securities Exchanges. Customers: Bank of New York, London Stock Exchange, Royal Bank of Scotland, Archipelago. Technologies: Machine Learning.

Searchspace [I-100] produce online processing AI-enabled infrastructure that already handles over \$400 trillion worth of business by automating fraud control, compliance reporting, risk management and customer service. They use a collection of “agents” known as *sentinels* that are capable of automating key business tasks – ranging from anti-money laundering detection, cheque fraud detection, real-time audit, debit card fraud, operational risk, individual customer profitability, price management through to contact management. Searchspace apply this technology within banking, security broking, telecommunications and exchanges.

For example within banking they currently have sentinels aimed at detection of: money laundering, cheque fraud, debit card fraud, credit card fraud, and payment fraud; and calculating operational risk and customer profitability. Searchspace argue that their sentinel approach delivers the best performance of any product on the market – and that this translates into bigger savings for fraud detection and comprehensive risk management for compliance. Searchspace sponsor research at University College London, Sussex University and the New Jersey Institute of Technology.

4.4.17 *Xilinx*

Company Type: Semiconductor company. Location: San Jose. Formation: Xilinx invented the world’s first field programmable gate array (FPGA) in 1984, one of the fastest growing segments of the semiconductor industry. Target Industries: Computer and electronic hardware. Technologies: Applying evolutionary algorithms to adaptable FPGA hardware.

FPGAs are adaptable integrated circuits, and are often used as basis of evolvable hardware experiments, where evolutionary computation techniques are used to design electronic circuits (Thompson, 1998). Xilinx have done some CAS-related work on using evolutionary computation techniques for fault tolerance and also for built-in self-test, a process that can be very difficult on complex chips.

5. Open Research Questions

For a field as diverse and dispersed as complexity research, there are too many live avenues of enquiry to explore them all in detail in this document. Rather than attempt to cover a representative sample, we identify some important themes and detail a small number of specific but idiosyncratic examples of upcoming complexity research relevant to ICT.

As outlined in the “Conceptual Landscape” section, one fundamental challenge confronting the complex systems community is the generation of consensually agreed-upon unifying theories of complexity and emergence, i.e., genuine “Newtonian” activity. As yet it is unclear what shape such theories will take, or how close to formulating them we may be. However, we can stipulate what a successful theory would achieve.

First, it would account for (rather than merely describe) what is common to complex systems across many domains, and as a result provide criteria for formally characterising (rather than merely identifying) complexity and emergence. Perhaps more importantly such a theory could underwrite a fundamentally stronger grasp of wider related processes such as adaptation, self-organisation, etc., thereby locating or reconciling complex systems with respect to simple linear systems, bringing their study within the remit of Kuhn’s “normal science” (Kuhn, 1962). It is perhaps only a minor exaggeration to compare the impact that such an account would have with that achieved by Darwin’s theory of evolution by natural selection, which similarly managed to organise a previously mysterious and a-scientific natural world under a single scientific account.

A more practical and perhaps more readily achievable goal centres on the formulation of a widely-adopted, robust, effective and legitimate simulation modelling methodology. Currently, simulation approaches often appear to be the only practicable option open to researchers exploring many kinds of complex system. While our ability to construct tractable mathematical models through simplifying assumptions is improving, simulation will remain an important tool for the foreseeable future. Unfortunately, as yet, despite the prevalence of simulation modelling across many branches of science and engineering, its role is remarkably poorly understood.

For some complex systems researchers, simulations are a valuable source of additional empirical data on systems that are difficult or impossible to experiment on (e.g., Ray, 1994). While one cannot literally re-wind and re-run the earth’s evolutionary history, or re-wire the neural systems of free-living mammals, or manipulate real large-scale social organisations such as stock exchanges, the internet, languages, religions, etc., one can perhaps simulate these scenarios (Bedau, 1998). However, while these simulations are eminently manipulable, that they

can act as sources of empirical data is extremely debatable.

How could one settle an empirical question by appeal to a simulation model of a system that we currently do not understand? Consider the question “*Is there life on Mars?*”. A simulation of Martian evolution would surely not convince one way or the other. We have insufficient data to adequately constrain any such simulation, and no established grounds upon which to base such a simulation. Less fancifully, consider *What impact would introducing traffic congestion charges to Leeds have?* Could we rely on a detailed simulation of Leeds in order to answer this question? To what extent would we have to accurately capture the decision-making processes of individual drivers? How could we be sure that our model accurately captured the way in which any initial impact of congestion charges on traffic behaviour would go on to influence the *subsequent* decision making of drivers, or retailers, or employers? More mundanely, what if TCP/IP were altered to cope differently with lost packets, what impact would there be on internet congestion? Have we the wherewithal to adequately simulate even this relatively limited scenario? Surely the same issues that complicated the automotive traffic model are equally relevant here?

By contrast, if I want to know whether my new desk can be manoeuvred up the staircase leading to my office, a simulation of the situation would seem to be a good way of settling the matter before purchasing the furniture. The difference between this example and those in the previous paragraph is that we have extremely high confidence in our ability to model the behaviour of medium-scale solid objects as they are rotated and translated in 3d space, but much less confidence in our ability to model veridically and in detail the complex systems mentioned above. Until we are able to increase this confidence, it is unclear to what extent simulation models can be relied upon to produce detailed predictions of real-world complex system behaviour (Silverman and Bullock, in press).

Less controversially, simulations are good sources of insights into our *ideas* about complex systems. Even a simulation of Martian ecology and evolution might be a useful way to shed light on our current assumptions concerning biological adaptation. Such simulations can serve as intuition pumps, proofs of concept, hypothesis generators, or “computational thought experiments” (Bedau, 1999; Di Paolo, Noble, & Bullock, 2000). How best to mobilise simulation techniques in order to achieve these ends is currently an under-explored question.

Relatedly, there is a growing need for an inferential statistics tailored to the analysis of complex systems. While descriptive statistics drawn from graph theory are being used to characterise complex networks, there is currently no equivalent to the enormously useful body of

inferential statistics that has become associated with the normal curve. We require an analogous suite of statistical methods to organise around systems that are better described by power laws than Gaussian distributions. In particular, measures of significance and statistical power are currently largely absent from accounts of the structure or dynamics of complex systems. Such tools would enable us to assess the extent to which a complex system is modular or hierarchical or decomposable in its behaviour or structure, to further identify these sub-components and characterise the coupling between them, and to compare these answers to what would be expected as a base-line for systems of the type we are considering (e.g., Tononi, Sporns, & Edelman, 1994; Seth & Edelman, 2004).

Aside from methodological questions, there are many general properties of complex systems that deserve more extensive consideration. For example, the role that *spatiality* plays in underpinning complex adaptive behaviour is poorly understood. Many complex systems exhibit spatial organisation that appears to bear crucially on their ability to remain robust, or to achieve sophisticated behaviours. For example, neural networks and cities both rely upon or exploit a number of spatio-temporally constrained processes (neurotransmission and neuromodulation; traffic and pedestrian flows), and employ spatial plasticity in order to organise adaptively (Changeux, 1993; Holscher, 1997; Fujita, Krugman, & Venables, 1999). More generally, we have little understanding of how to model *coupling* between complex systems, and between them and their environments (and those environments may of course themselves be composed of multiple coupled complex systems). Examples of such couplings include those between the internet and its user community; the health service and the pharmaceutical industry, or between ad-hoc wireless and fixed wired networks (Ottino, 2004). Moreover, the ways in which the structure and functionality of complex adaptive systems are constrained by processes of growth and *development* are also under-explored. Increased understanding of the way in which biological development is influenced by genes and environment could enable us to *grow* engineered systems rather than specify and fabricate them in explicit detail. Finally, the *evolvability* of complex systems (their ability to profit from small changes) is only recently beginning to be explored. Complex adaptive systems often appear to be able to avoid the lock-in or fragility exhibited by large engineered systems in favour of a balance between robustness and flexibility (Wagner & Altenberg, 1996; Kirschner & Gerhart, 1998). It is clear that there is considerable scope to exploit our understanding of all four of these issues in the design and management of engineered complex systems.

Turning to examples of more specific questions, complexity and emergent behaviour issues are relevant across

a wide range of ICT systems and at every level. In terms of strategic ICT management, large-scale development projects are running into complexity related problems that require circumventing. For example, how can an effective digital NHS be deployed and run? Here, issues range from the technological (effectively maintaining and integrating massive data sets) through socio-technological (data protection issues and security and trust) to social (overcoming cultural inertia in an entrenched user population). A key issue is discovering “open-standards” software development methodologies that can cope with the scale of the systems envisaged. How can massive pieces of software be engineered to deadlines in such a way that they remain under budget, on schedule, fit-for-purpose, usable, maintainable, extensible, robust and resilient, and so on? Here there may be scope for developmental approaches inspired by natural systems that exploit self-organisation, etc., in order to circumvent the same problems.

At an architectural level, opportunities are presented by the recently-documented apparent ubiquity of certain kinds of network structure (small worlds, or scale-free networks, for instance) in natural and technological contexts (Albert & Barabási, 2002; Newman, 2003). Researchers are already exploring whether these naturally existing structures can be exploited in ways that do work for us. For instance, BT are currently exploring whether the small-world structure of social networks can be usefully exploited by communications technologies.

Moreover, resource allocation issues are likely to become critical as infrastructural ICT networks grow and interconnect. How can we manage grid computing architectures such that they maintain quality of service without relying on centralised controls that will not scale to cope with large, dynamic, heterogeneous ICT systems. One approach suggested by complex systems ideas is to design a topology that supports efficient co-allocation of resources via the action of one or more decentralised artificial market-economies. Such a mechanism, it is argued, could manage resources effectively without the requirement of a central executive (Cliff & Bruten, 1998; Bye, Salle, & Bartolini, 2003; Davy, Djemame, & Noble, 2003).

Considering this kind of complex systems design problem more generally, how can we use automatic processes such as those employed in evolutionary and adaptive computing to configure or optimise the organisation of complex systems? Here we might exploit the problem-solving abilities of adaptive populations to discover solutions that are not only high quality, but robust to perturbations and easily tuneable. For example, designing attack-resistant topologies for computational infrastructures might be recast as a problem that does not require a single topology to be discovered (itself a difficult task), but rather a space of topologies that can be rapidly and

easily re-worked to resist new kinds of directed or undirected attack (Thompson, 1998; Layzell, 2001). This emphasis begins to shift the burden of design from a “right-first-time” stance to one of “continuous-redesign” (i.e., adaptation) in the face of changing circumstances. This relaxes traditional requirements of provably correct and robust designs that are required to solve fixed problems for all time, in favour of methodologies that are maximally responsive to failures that are considered an unavoidable fact of life (an approach sometimes termed “satisficing” rather than “optimising”, Simon, 1981). Similar approaches could perhaps be developed to cope with spam, worms, viruses, and system incompatibilities or conflicts, etc. Here the self-healing or self-repair abilities of natural complex adaptive systems are a direct inspiration.

Can ICT systems remain usable as their complexity increases? Workable interfaces, and interoperability between multiple complex ICT systems remains an outstanding challenge, especially given the rates of change in the functionality and make-up of those ICT systems and the variability and growing antipathy of user populations towards reading or consulting manuals for the ICT products that they use. Here the prospect of systems that are able to adaptively reconfigure themselves to cope with their changing environment (i.e., their partner systems and users) is attractive, but as yet untested. New forms of responsive, dynamic ICT support might exploit users’ growing reliance on organic on-line communities rather than fixed manuals or help pages.

Relatedly, how can we establish ICT software and middleware that efficiently support and encourage light-weight peer-to-peer communities, and in what contexts can they provide a workable substrate for commercial or business activity? Currently peer-to-peer activity is most closely associated with potentially illegal file sharing that swallows significant bandwidth. However, there may be scope for similar legitimate systems if problems of co-ordination, licensing and infrastructural load can be resolved (Saffre, 2004).

Finally, how can we effectively exploit continuing progress in the production of simple light-weight, low-power computational devices? Highly respected computer science researchers of long standing have been (for some years now) seriously considering the computational problems and opportunities arising from likely future developments in the cheap mass-fabrication of micro-electro-mechanical systems and nano-scale computer systems: extremely small, simple, slow, and unreliable computer processor units, each with its own very-low-range wireless communication transmitter/receiver capabilities, could be produced at extremely low cost per unit (e.g., literally ten-a-penny). Hundreds of thousands of such units could be randomly scattered over some space, and each unit could then be required to es-

tablish communication contact with nearby units and for the whole connected network of units to then cooperate on some globally coherent computation. All the current indications are that such *amorphous computing* networks are likely to require biologically-inspired self-organisation processes for development and ongoing regulation, and indeed the leading amorphous computing research group has vigorously pursued such biological inspiration from the outset [I-118].

6. Challenges to Progress

In this final section, we review some of the main obstacles that currently confront attempts to advance complex systems research in the UK. We review factors that prevent or limit interdisciplinarity, cultural barriers that curtail industrial uptake of results and ideas from complex systems research, and the limitations of current training for ICT complexity researchers.

6.1 Obstacles to Interdisciplinarity

Perhaps the key threat to the long-term prospects of complex systems research is the instability of interdisciplinary research. At many levels, there is an increasing recognition of the value of interdisciplinarity in collaborative research. In particular there is an increased awareness of particular disciplinary interfaces (bioinformatics, geographic information systems, etc.) which has led to the employment and development of relevant faculty, the establishment of research groups, and the targeting of specific funding initiatives. However, the oft-stated more general assertion that “interdisciplinarity is important” has yet to translate into many institutional structures that support and promote work that genuinely spans disciplinary boundaries.¹¹

First and foremost, it is widely perceived that the research assessment exercise (RAE), designed to rate the quality of UK university research activity and thereby encourage excellence across the board, has inhibited and restricted interdisciplinarity. As yet, the extent to which the next round of assessment will continue this trend or remedy it is unclear. What is apparent is that an approach that divides individual researchers into a limited number of units of assessment organised along disciplinary lines will tend to let genuinely interdisciplinary research fall between the cracks. This tendency could be redressed in future assessment exercises by instituting explicit mechanisms for assessing research regarded

¹¹Indeed, we note with sadness the recent demise of the School of Cognitive & Computing Sciences (Cogs) at the University of Sussex. Cogs was for many years the UK’s only *truly* interdisciplinary centre for research in artificial intelligence and the cognitive sciences, and from the late 1980’s onwards it became the nucleus of the UK academic community for biologically-inspired complex adaptive systems research. In general it takes many years to create an interdisciplinary institution as intellectually rich and fertile as Cogs was, but only a small amount of time to break one up.

to be “non-standard” by the criteria of a particular unit of assessment, and, crucially, by convincing the relevant academic institutions that these mechanisms will actually reward interdisciplinarity. Without *both* measures in place, university appointment panels will probably continue to regard the employment of interdisciplinary researchers as a risk.

Along similar lines, interdisciplinary grant proposals are perceived by academics to suffer from a higher than average risk of rejection as a direct result of current peer-review mechanisms. While it may be true that research funding bodies are improving their ability to judge the quality of interdisciplinary proposals, a widespread perception within complex systems research is that achieving funding for this type of work remains difficult. Even funding initiatives explicitly designed to target complex systems research are prone to misfire if panels are comprised of individuals who cannot work together to identify high-quality proposals that may appear to be very different in content from the panellists’ own research. Such panels can tend to be multi-disciplinary rather than interdisciplinary, in that they are comprised of domain specialists from several fields, rather than researchers who work at disciplinary boundaries and interfaces. As a result, consensus building can be particularly difficult. The same problem is presented by peer-review for mainstream journals and conferences, and departmental job panels.

The impact of review mechanisms on interdisciplinarity might be acceptable if it served to simply raise the bar, demanding higher quality of interdisciplinary activity than might be required of more standard job applicants, grant proposals, or research papers. However, there is no evidence that it is only high-quality interdisciplinary activity that survives peer-review. In fact, activity at interdisciplinary boundaries is perceived to be of very patchy quality, with sloppy, misguided or trivial work proposed, funded and published at higher-than-average rates compared to mono-disciplinary standards. In particular, the reinvention of established ideas and techniques, and the glib, content-free, exploitation of hype and fashion are frequently complained of both within the complexity research community and outside it. Indeed, it is these perceptions that fuel the tendency for panels to subject interdisciplinary work to increased scrutiny.

Opinions on how to address these problems differ, with some researchers believing that there are inherent difficulties in assessing routine interdisciplinary research (let alone ground-breaking work) that cannot be avoided, that such work will always be risky, slow, of patchy quality and that a degree of disenfranchisement is a fact of life for interdisciplinary researchers. While there is some disagreement over the accuracy of this picture, few dispute that the ability of institutional structures to cope

with and promote interdisciplinarity could be improved significantly.

6.2 *Obstacles to Industrial Involvement*

While industrial interest in exploiting complex systems ideas and methods may be growing, there remain significant cultural differences between industrial ICT engineering and the study of complex systems within academia. In particular, the real-world problems of specifying, designing, manufacturing and deploying systems that solve real problems and thereby save or make money for businesses are often perceived to not be the problems that academics are interested in solving. Complexity research can appear to focus on what seem to be arbitrary natural systems (e.g., avalanches on sandpiles) or mathematically abstruse problems (e.g., random networks, or Conway’s “game of life” (Gardner, 1970; Poundstone, 1985)) and to pursue questions with little relevance to the pressing problems facing industry; seeking a theory of everything, rather than theories of particular somethings. Conversely, industrialists can appear wedded to a traditional design and engineering paradigm that has little or no room for non-standard solutions, especially where those solutions rely on emergent (i.e., mysterious) properties of complex (i.e., impenetrable) systems. Several common properties of complex adaptive systems tend to raise concerns: a frequent strong reliance on the aggregate behaviour of essentially *random* elements; the lack of a single, central authority or executive; an inability to provably demonstrate robustness; and worst-case scenarios that can range from poor to catastrophic. These communication problems are surmountable, but they require considerable effort on the part of both communities.

6.3 *Training*

While there has been recent growth in complexity research, and also in the recognition of the relevance of complexity research to ICT, there has not so far been a proportional growth in related training and education opportunities to create the required human capital (i.e., appropriately-skilled researchers). Discussion of Figure 2 earlier in this paper highlighted the diversity of backgrounds experienced by complexity researchers, but also drew attention to the relatively unprepared state in which most enter the field. As far as we are aware, neither the notions of complexity and emergent behaviour, nor the tools and techniques with which complex systems can be studied, are included in standard computer science curricula at any level in UK education. While there are a handful of advanced undergraduate modules and a small number of masters degree programs that target complexity among a number of issues, these are in the minority. Where such ideas are actually encoun-

tered in ICT degree programs, their treatment is often brief and as a result necessarily superficial. There are few, if any, ICT-relevant complex adaptive systems textbooks or educational primers alongside the profusion of popular science treatments.

If the complex systems community are correct in arguing that issues of complexity and emergent behaviour will become increasingly significant factors for ICT in the near future, they should perhaps be increasingly concerned with both conveying these ideas to ICT practitioners in general, and explicitly supporting a new generation of scientists and engineers equipped to develop and extend those ideas. At minimum, the absence of complexity from core-curriculum ICT modules denies the average student both an awareness of the complexity-related problems that beset modern computational infrastructures, and also an ability to make informed choices with respect to the techniques and ideas being developed by complexity researchers. To the extent that these problems and ideas may rapidly become central to modern ICT practice, they should certainly not be taught as peripheral, optional, specialist, and esoteric.

As noted previously in this document, many significant factors mitigate against opting for an academic research career in interdisciplinary complex systems science. Trusting that significantly many bright researchers will continue to take this gamble is perhaps unwise. Moreover, even if sufficient numbers choose to pursue questions of complexity and ICT, the educational structures that are currently in place may fail them.

In particular, current educational trends are likely to increase the rate at which poorly-understood, complicated, detailed, simulation models are built and relied upon as “realistic” in some way. Three factors are involved. First, the above-noted lack of significant exposure to ideas of complexity and emergent behaviour at primary, secondary and tertiary levels results in a naivety with respect to the problems of modelling complex systems. Second, the increasing attention paid to ICT and programming skills across many diverse degree programs equips students with the relatively limited computing skills required to build a simulation, but not the other skills that are needed; this is a situation in which a little knowledge really can be a dangerous thing. Third, and most importantly, the continuing fall in numbers of students taking mathematics qualifications prior to their university degree, and the consequent erosion of both mathematics undergraduate programs, and mathematics content in ICT degrees is producing a math-poor student cohort unable to underwrite their simulation building activity with a relevant mathematical foundation, or to design appropriate experiments and controls, or to rigorously statistically analyse the results that their simulations generate.

Simulation modelling is an increasingly important

skill, and one that should be being taught carefully on many undergraduate courses. If this is to happen, it must not occur at the expense of teaching the mathematical skills that necessarily underpin their proper use. It has always been difficult to keep ICT education abreast of what is a very fast changing area. However, there is a real risk that failing to incorporate simulation modelling skills rapidly into computing curricula will lead to significant problems. The public’s perception of statistics as an illegitimate activity that is used by politicians, scientists, management, and other lowlife, to artificially support dubious lines of argument could soon be augmented by a similar attitude to simulation. Computational models are increasingly being used to predict and control real-world scenarios such as intervention in traffic systems, the ecological impact of genetically modified crops, the ability of our armed forces to achieve military targets, and similar examples in many other fields. Without the establishment of good practice in complex agent-based simulation modelling, “simulation” could rapidly join or even replace “statistics” alongside “lies” and “damned lies”.

7. Conclusion

The UK is very well-positioned to benefit enormously from a combination of its high-quality complex systems research and its innovative ICT industry. Innovation at the intersection between these two activities is likely to be crucial to the future prosperity of the world’s computing industry and the myriad activities that ICT underpins. In particular, the problems posed by increasingly networked systems demand complex systems approaches that must be developed quickly through collaboration between multiple disciplines and industrial sectors. Effectively facilitating this kind of collaboration is a major challenge facing institutions at all levels: government bodies and funding agencies; professional organisations; industrial research and development labs; and universities, both as centres of research and as providers of education and training.

Acknowledgements

This report was commissioned by the Foresight Directorate, part of the Office of Science and Technology, in the UK Government’s Department of Trade and Industry. The commission involved a consultancy contract that fixed the number of working days available for researching and writing this report. The present document is the result of the 20 days work that were commissioned; given more time, we could of course have written a much longer and/or more comprehensive report.

We would like to thank the following people for their suggestions and discussion of the ideas presented here: Prof Maggie Boden (School of Cognitive & Computing Sciences, University of Sussex), Dr Paolo Dini (Media & Communica-

tions, London School of Economics), Dr Ralph Dum (Future & Emerging Technologies, European Commission), Dr Mark Hylton (EPSRC), Dr Bob Malcolm (Managing Director, ideo), Dr Robin Mansell (Media & Communication, London School of Economics), Prof Eve Mitleton-Kelly (Director, Complexity Research Programme, London School of Economics), Dr José-Luis Fernández-Villacañas Martín (Software Technologies, European Commission), Dr Felix Reed-Tsochas (Saïd Business School, Oxford University), Dr Fabrice Saffre (BTExact Labs), Prof Luc Steels (VUB and Sony Europe), Dr Mike Wheeler (Philosophy, University of Dundee), Prof Sir Alan Wilson (Director General for Higher Education), and the University of Leeds Biosystems group.

Sections 4.2 to 4.4 are revised from a review co-authored by members of the HP Labs Complex Adaptive Systems group; Dave Cliff thanks his HP colleagues Dr Mark Butler, Dr Andrew Bye, Dr Steve Cayzer, Dr Paul Layzell, and Dr Matt Williamson for their help in writing the initial review and for their permission to re-use the material here. Cliff also thanks Martin Merry and Robin Gallimore at HP Labs for allocating scarce resources to the process of writing this report.

Finally we thank Dr Alex King and Dr Andrew Jackson of the DTI Foresight Directorate for their enthusiasm and encouragement, and for their comments on earlier drafts of this report.

References

- Adami, C. (2002). What is complexity? *BioEssays*, 24, 1085–1094.
- Albert, R., & Barabási, A.-L. (2002). Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74, 47–97.
- Arkin, R. C. (1998). *Behaviour-based Robotics*. MIT Press, Cambridge, MA.
- Arnott, S. (2003). Government IT projects squander £1.5Bn. *Computing*, <http://www.computing.co.uk/news/1139438>.
- Barabási, A.-L. (2002). *Linked: The New Science of Networks*. Perseus, Cambridge, MA.
- Bedau, M. A. (1998). Philosophical content and method of artificial life. In Bynum, T. W., & Moor, J. H. (Eds.), *The Digital Phoenix: How Computers are Changing Philosophy*, pp. 135–152. Blackwell, Oxford.
- Bedau, M. A. (1999). Can unrealistic computer models illuminate theoretical biology? In Wu, A. S. (Ed.), *Proceedings of the 1999 Genetic and Evolutionary Computation Conference Workshop Program*, pp. 20–23. Morgan Kaufmann, San Francisco, CA.
- Bonabeau, E., Dorigo, M., & Théraulaz, G. (1999). *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, Oxford.
- Brooks, R. A. (1999). *Cambrian Intelligence: The Early History of new AI*. MIT Press, Cambridge, MA.
- Byde, A., Salle, M., & Bartolini, C. (2003). Market-based resource allocation for utility data centers. <http://www.hpl.hp.com/techreports/2003/HPL-2003-188.html>.
- Changeux, J. (1993). Chemical signaling in the brain. *Sci. Am.*, 269, 58–62.
- Clark, A. (Ed.). (2001). *Mindware*. Oxford University Press.
- Clearwater, S. (Ed.). (1996). *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, Hong Kong.
- Cliff, D., & Bruten, J. (1998). Simple bargaining agents for decentralized market-based control. <http://www.hpl.hp.com/techreports/98/HPL-98-17.html>.
- Davy, S., Djemame, K., & Noble, J. (2003). The application of bioinspired engineering principles to grid resource allocation. In *UK Performance Engineering Workshop*.
- Dawkins, R. (1986). *The Blind Watchmaker*. Longman, Harlow.
- de Castro, L. N., & Timmis, J. (2002). *Artificial Immune Systems: A New Computational Intelligence Approach*. Springer-Verlag, London, UK.
- Di Paolo, E. A., Noble, J., & Bullock, S. (2000). Simulation models as opaque thought experiments. In Bedau, M. A., McCaskill, J. S., Packard, N. H., & Rasmussen, S. (Eds.), *Artificial Life VII*, pp. 497–506. MIT Press, Cambridge, MA.
- Edmonds, B. (1999). What is complexity? The philosophy of complexity *per se* with application to some examples in evolution. In Heylighen, F., & Aerts, D. (Eds.), *The Evolution of Complexity*. Kluwer, Dordrecht.
- Fujita, M., Krugman, P., & Venables, A. J. (1999). *The Spatial Economy: Cities, Regions and International Trade*. MIT Press.
- Gardner, M. (1970). Mathematical games: The fantastic combinations of John Conway's new solitaire "life". *Scientific American*, 223, 120–123.
- Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley, Richmond, MA.
- Harvey, I., & Bossomaier, T. (1997). Time out of joint: Attractors in asynchronous random boolean networks. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 67–75. MIT Press.
- Haykin, S. (1998). *Neural Networks: A Comprehensive Foundation* (2nd edition). Prentice-Hall, Upper Saddle River, NJ.
- Hofstadter, D. R. (1979). *Gödel, Escher, Bach: An Eternal Golden Braid*. Basic Books, New York.
- Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor. Reprinted by MIT Press, 1992.
- Holscher, C. (1997). Nitric oxide, the enigmatic neuronal messenger: its role in synaptic plasticity. *Trends in Neurosciences*, 20, 298–303.
- Horgan, J. (1995). From complexity to perplexity. *Scientific American*, 272, 74–79.

- Huberman, B. (1988). *The Ecology of Computation*. Elsevier, New York.
- Huberman, B. A., & Glance, N. S. (1993). Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences, USA*, 90(16), 7715–7718.
- Kirschner, M., & Gerhart, J. (1998). Evolvability. *Proceedings of the National Academy of Sciences (USA)*, 95, 8420–8427.
- Kuhn, T. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.
- Layzell, P. (2001). Hardware evolution: On the nature of artificially evolved electronic circuits. D.Phil thesis, School of Cognitive and Computing Sciences, University of Sussex.
- Marr, D. (1977). Artificial intelligence — a personal view. In Boden, M. A. (Ed.), *The Philosophy of Artificial Intelligence*, pp. 133–147. Oxford University Press, Oxford. Collection published in 1990.
- Maynard Smith, J., & Szathmáry, E. (1995). *The Major Transitions in Evolution*. Freeman, Oxford.
- Mitchell, M. (1996). *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA.
- National Audit Office (2002). The 2001 outbreak of foot and mouth. Report by the Comptroller and Auditor General, HC 939 Session 2001-2002: 21 June 2002.
- Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45, 167–256.
- Nowak, M. A., & May, R. M. (1992). Evolutionary games and spatial chaos. *Nature*, 359, 826–829.
- Ottino, J. M. (2004). Engineering complex systems. *Nature*, 427, 399.
- Patel, S., Scott, I. P., Bhakoo, M., & Elliott, P. (1998). Patenting computer-designed peptides. *Journal of Computer-Aided Molecular Design*, 12(6), 543–556.
- Poundstone, W. (1985). *The Recursive Universe: Cosmic Complexity and the Limits of Scientific Knowledge*. William Morrow, New York.
- Ray, T. S. (1994). An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artificial Life*, 1, 179–209.
- Rumelhart, D. E., & McClelland, J. L. (Eds.). (1986). *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 1. MIT Press, Cambridge, MA.
- Saffre, F. (2004). Complexity science and the future of middleware in a P2P world. Software and Complex Systems Workshop: Ninth IEEE International Conference on Engineering Complex Computer Systems.
- Schoonderwoerd, R., Holland, O., Bruten, J., & Rothkrantz, L. (1997). Ant-based load balancing in a telecommunications network. *Adaptive Behavior*, 5(2), 169–207.
- Seth, A. K., & Edelman, G. M. (2004). Environment and behavior influence the complexity of evolved networks. *Adaptive Behavior*, 12, 5–20.
- Silverman, E., & Bullock, S. (2004). Empiricism in artificial life. In Pollack, J., Bedau, M. A., Husbands, P., Ikegami, T., & Watson, R. A. (Eds.), *Proceedings of the Ninth International Conference on Artificial Life*, pp. 534–539. MIT Press, Cambridge, MA.
- Simon, H. A. (1981). *The Sciences of the Artificial* (Second edition). MIT Press, Cambridge, MA.
- Stone, J. V. (2004). *Independent Component Analysis: A Tutorial Introduction*. MIT Press, Cambridge, MA.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.
- Thompson, A. (1998). *Hardware Evolution: Automatic Design of Electronic Circuits in Reconfigurable Hardware by Artificial Evolution*. Distinguished dissertation series. Springer-Verlag.
- Tononi, G., Sporns, O., & Edelman, G. (1994). A measure for brain complexity: Relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Science (USA)*, 91, 5033–5037.
- Wagner, G. P., & Altenberg, L. (1996). Complex adaptations and the evolution of evolvability. *Evolution*, 50(3), 967–976.

Background Bibliography

- Axtell, R. and Epstein (1996) *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press.
- Bak, P. (1996) *How Nature Works: The Science of Self-Organised Criticality*. Copernicus Press.
- Barabasi, A.-L. (2003) *Linked: How Everything Is Connected to Everything Else and What It Means for Business, Science, and Everyday Life*. Plume Publishers.
- Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*. MIT Press.
- Brooks, R. (2002). *Robot: The Future of Flesh & Machines*, Allen Lane.
- Dreyfus, H. (2001). *On the Internet*. Routledge.
- Goodwin, B. (1994). *How the Leopard Changed Its Spots: The Evolution of Complexity*. Princeton University Press.
- Huberman, B. (2003). *The Laws of the Web: Patterns in the Ecology of Information*. MIT Press.
- Hofstadter, D. R. (1979). *Godel, Escher, Bach: An Eternal Golden Braid*. Basic Books.
- Holland, J. (1997). *Emergence*. Addison-Wesley.
- Kauffman, S. (1995). *At Home in the Universe: The Search for Laws of Self-Organization and Complexity*. Oxford University Press.
- Kauffman, S. (2000). *Investigations*. Oxford University Press.
- Krugman, P. R. (1996). *The Self-Organising Economy*. Blackwell.

- Ormerod, P. (1999). *Butterfly Economics: A New General Theory of Economic and Social Behaviour*. Faber & Faber.
- Resnick, M. (1997). *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds*. MIT Press.
- Sigmund, K. (1995). *Games of Life: Explorations in Ecology, Evolution, and Behavior*. Penguin.
- Waldrop, M. (1994). *Complexity: The Emerging Science at the Edge of Order and Chaos*. Penguin.
- Watts, D. (1999). *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton University Press.
- Winograd, T. & Flores, F. (1987). *Understanding Computers and Cognition: A New Foundation for Design*. Addison-Wesley.

Web Resources

- [A-001] University of Birmingham's *Natural Computation MSc*,
<http://www.cs.bham.ac.uk/study/postgraduate-taught/msc-nc/>
- [A-002] University of Sussex's *Evolutionary and Adaptive Systems MSc*, <http://www.informatics.susx.ac.uk/easy/MSc/>
- [A-003] University of Southampton's upcoming *Bio-inspired Computing MSc*, <http://www.ecs.soton.ac.uk/>
- [A-004] University of Leeds' *Multidisciplinary Informatics MSc*,
<http://www.informatics.leeds.ac.uk/msc-multidisciplinary-informatics/>
- [A-005] ATR Kyoto, <http://www.atr.co.jp>
- [A-006] Free University of Brussels, <http://pespmc1.vub.ac.be/VUBULB.html>
- [A-007] Caltech, <http://www.caltech.edu>
- [A-008] Los Alamos National Labs, <http://www.lanl.gov>
- [A-009] MIT AI Lab, <http://www.ai.mit.edu>
- [A-010] MIT Media Lab, <http://www.media.mit.edu>
- [A-011] New England Complex Systems Institute (NECSI), <http://necsi.org>
- [A-012] Santa Fe Institute (SFI), <http://www.santafe.edu>
- [A-013] University of Southern California (USC), <http://robotics.usc.edu/interaction>
- [A-014] Birmingham University, Natural Computation, <http://www.cs.bham.ac.uk/research/NC>
- [A-015] Edinburgh University, Institute for Adaptive and Neural Computing <http://www.inf.ed.ac.uk/research/ianc>
- [A-016] University of Exeter, Complexity Network, <http://www.dcs.ex.ac.uk/research/complexity>
- [A-017] University of Essex, <http://cswww.essex.ac.uk/Research/>
- [A-018] Leeds University, <http://www.comp.leeds.ac.uk/seth>, <http://www.comp.leeds.ac.uk/jasonn>,
<http://www.comp.leeds.ac.uk/netta>
- [A-019] London School of Economics (LSE), Complexity Research Programme <http://www.psych.lse.ac.uk/complexity/>
- [A-020] Oxford University, Complex Agent-based Dynamic Networks (CABDYN) research cluster
<http://sbs-xnet.sbs.ox.ac.uk/complexity/>
- [A-021] Sussex University, Centre for Computational Neuroscience and Robotics (CCNR)
<http://www.cogs.susx.ac.uk/ccnr/>
- [A-022] University of the West of England, Intelligent Autonomous Agents Laboratory (UWE), <http://www.ias.uwe.ac.uk>,
Larry Bull <http://www.cems.uwe.ac.uk/lbull/>
- [I-001] Microsoft Research, Adaptive Systems and Interaction, <http://research.microsoft.com/adapt/>
- [I-002] Microsoft Research, Data management, exploration and mining, <http://research.microsoft.com/dmx/>
- [I-003] Microsoft Research, Machine learning and applied statistics, <http://research.microsoft.com/mlas/>
- [I-004] Microsoft Research, Collaborative Filtering, <http://research.microsoft.com/users/breese/cfalgs.html>
- [I-005] Microsoft Research, Natural Language Processing, <http://research.microsoft.com/nlp/>
- [I-006] Microsoft Research, Face Detection and Recognition,
<http://research.microsoft.com/szli/FaceGroup/default.asp>
- [I-007] Microsoft Agent Toolkit, http://msdn.microsoft.com/library/default.asp?url=/library/en-us/msagent/agentstartpage_7gdh.asp
- [I-008] Support Vector Machines used in Microsoft Sharepoint, <http://www.informationweek.com/828/microsoft.htm>
- [I-009] IBM Research, Information Economies, <http://www.research.ibm.com/infoecon/>
- [I-010] Sante Fe Institute <http://www.santafe.edu/>
- [I-011] IBM Anti-virus group <http://www.research.ibm.com/antivirus/>
- [I-012] IBM Research Text and Information, <http://www.almaden.ibm.com/cs/k53/ir.html>
- [I-013] IBM Research, Clever Search, <http://www.almaden.ibm.com/cs/k53/clever.html>
- [I-014] IBM Research, Quest Data Mining Group, <http://www.almaden.ibm.com/cs/quest/>
- [I-015] IBM, DB2 Image Extender, <http://www-3.ibm.com/software/data/db2/extenders/image.htm>
- [I-016] IBM, DB2 Audio Extender, <http://www-3.ibm.com/software/data/db2/extenders/audio.htm>
- [I-017] IBM, Intelligent Miner, <http://www-3.ibm.com/software/data/iminer>
- [I-018] IBM, Intelligent Miner for Text, <http://www-3.ibm.com/software/data/iminer/fortext/>
- [I-019] IBM, Lotus Discovery Server, <http://www.lotus.com/products/discserver.nsf/>
- [I-020] IBM, eLiza Project <http://www-1.ibm.com/servers/eserver/introducing/eliza/>
- [I-021] IBM, Autonomic Computing, <http://www.research.ibm.com/autonomic/>
- [I-022] BTExact, <http://www.btexact.com>
- [I-023] BT Technology Journal Issue on Biologically Inspired Work, <http://www.bt.com/bttj/vol18no4>
- [I-024] BT Complexity Research, <http://www.btexact.com/projects/complexity.htm>
- [I-025] BT Emerging Technologies, <http://www.btexact.com/projects/lab7.htm>
- [I-026] BT Eos, <http://www.btexact.com/projects/eos.htm>
- [I-027] Swarm, <http://www.swarm.org>
- [I-028] BT Future Technologies Group, <http://www.btexact.com/projects/ftg/index.htm>
- [I-029] BT Intelligent Agents, <http://www.btexact.com/projects/agents.htm>

[I-030] BT Intelligent Business Systems, <http://www.btexact.com/projects/ibsr.htm>

[I-031] PARC, <http://www.parc.xerox.com/>

[I-032] Top Research Laboratories in Human Computer Interaction, <http://www.useit.com/alertbox/20020331.html>

[I-033] Market Based Control, Edited By Scott Clearwater, World Scientific 1996
<http://www.wspc.com/books/chaos/2741.html>

[I-034] Oracle, Business Intelligence Tools, http://www.oracle.com/ip/analyze/warehouse/bus_intell/

[I-035] Oracle, Data Mining, <http://www.oracle.com/ip/analyze/warehouse/datamining/>

[I-036] HP Labs, BICAS, <http://ginger.HP1.HP.com/research/bicas/>

[I-037] HP Labs, Dave Cliff, http://www.HP1.HP.co.uk/people/dave_cliff/

[I-038] HP Labs, ZIP traders http://www-uk.HP1.HP.com/people/dave_cliff/traders.htm

[I-039] HP Labs, Matthew Williamson <http://www.HP1.HP.co.uk/people/matwil/>

[I-040] Biological Approaches to Computer Security, <http://lib.HP1.HP.com/techHPubs/2002/HPL-2002-131.html>

[I-041] Virus throttling (HP Labs Technical Report) <http://www.hpl.hp.com/techreports/2003/HPL-2003-69.html>

[I-042] HP, Article on Barry Shackelford, <http://web.HP1.HP.com:8080/news/shackelford.html>

[I-043] HP, Information Dynamics Lab, <http://www.HP1.HP.com/research/idl>

[I-044] NCAF, Natural Computation Applications Forum, <http://www.ncaf.co.uk>

[I-045] Vector Space Model tutorial, <http://isp.imm.dtu.dk/thor/projects/multimedia/textmining/node5.html>

[I-046] IBM, Autonomic Computing, <http://www.research.ibm.com/autonomic/research>

[I-047] International Conference on Autonomic Computing (ICAC-04), <http://www.caip.rutgers.edu/~parashar/ac2004/>

[I-048] Self-Star: International Workshop on Self-* Properties in Complex Information Systems,
<http://www.cs.unibo.it/self-star>

[I-049] Modelling Disk Arrays Using Genetic Programming <http://lib.HP1.HP.com/techHPubs/2002/HPL-2002-20.html>

[I-050] HP Self-healing servers, <http://www.hp.com/hpinfo/newsroom/press/2003/030506c.html>

[I-051] Pharmacopeia, <http://www.pharmacopeia.com/>

[I-052] Compaq Research Publications, <http://src-www.pa.dec.com/CRL/publications/crl-papers.html>

[I-053] Compaq Software Download Site, <http://research.compaq.com/downloads.html>

[I-054] Intel DeveloperWorks, <http://www.intel.com/pca/developernetwork/>

[I-055] Sony, AIBO, <http://www.aibo.com/>

[I-056] Sony, EE Times article on Human AIBO, <http://www.eetimes.com/story/OEG20010111S0019>

[I-057] Sony, Pictures of Human AIBO prototype, <http://asia.cnn.com/2000/TECH/computing/11/22/sdr3.idg/>

[I-058] Sony, CSL Paris, <http://www.csl.sony.fr/>

[I-059] Mitsubishi Electric Research Laboratory, <http://www.merl.com/>

[I-060] Brandeis DEMO Lab, <http://demo.cs.brandeis.edu/index.html>

[I-061] Mitsubishi Electric Research Laboratory, Evolutionary Optimisation <http://www.merl.com/projects/evOpt/>

[I-062] NEC Labs, <http://www.nec-labs.com>

[I-063] CiteSeer, <http://citeseer.ist.psu.edu/cs>

[I-064] NuTech Solutions Inc., <http://www.nutechsolutions.com>

[I-065] Icosystem, <http://www.icosystem.com>

[I-066] Sana Security, <http://www.sanasecurity.com>

[I-085] Support Vector Machines, <http://www.support-vector.net>

[I-086] BAE Systems, <http://www.baesystems.com/overview/ourevolution.htm>

[I-087] NCR, <http://www.ncr.com>

[I-088] NCR Data Mining Services, http://www.ncr.com/services/dw_des_minser.htm

[I-089] NCR KnowledgeLab, <http://ncr.knowledgelab.com/>

[I-090] MIT Technology Review article on KnowledgeLab, <http://www.techreview.com/articles/carter0999.asp>

[I-091] Xilinx, <http://www.xilinx.com/>

[I-092] IRobot, <http://www.irobot.com/home/default.asp>

[I-093] Cambridge Neurodynamics, <http://www.neurodynamics.com/>

[I-094] GMAP, <http://www.gmap.co.uk/>

[I-095] Forio, <http://www.forio.com/index.htm>

[I-096] HP Compaq vs. Dell pricing simulation, <http://www.forio.com/pricing20010912.htm>

[I-097] PDASim, <http://www.forio.com/pdasim/index.jsp>

[I-098] Forio Resources on Simulation, <http://www.forio.com/resources.htm>

[I-099] Icosystem, <http://www.icosystem.com/>

[I-100] Searchspace, <http://www.searchspace.com/>

[I-101] Apama, <http://www.apama.com/>

[I-102] Norkom, <http://www.norkom.com/>

[I-103] Autonomy, http://www.autonomy.com/autonomy_v3/

[I-104] Agorics Inc, <http://www.agorics.com>

[I-105] Stuart Kauffman, <http://www.santafe.edu/sfi/People/kauffman/>

- [I-106] Bios Group, <http://www.biosgroup.com/>
- [I-107] EuroBios, <http://www.eurobios.com/flash/gb/eurobios.html>
- [I-111] LostWax, <http://www.lostwax.com/index.htm>
- [I-112] MASA, <http://www.animaths.com/>
- [I-113] Natural Selection, Inc. <http://www.natural-selection.com/>
- [I-114] CreatureLabs, <http://www.creaturelabs.com/>
- [I-115] Electronic Arts, <http://thesims.ea.com/>
- [I-116] Article on Thinking Tools Incorporated, <http://www.prweb.com/releases/1999/3/prweb7478.pHP>
- [I-117] SFI article on Telesim, <http://www.santafe.edu/sfi/publications/Bulletins/bulletin-spr95/4beyond.html>
- [I-118] MIT Amorphous Computing, www.swiss.ai.mit.edu/projects/amorphous