

Digital Business Ecosystem

Contract n° 507953

Workpackage 18

DBE UML Profile

Deliverable 18.1

Report on DBE-Specific Use Cases



Information Society
Technologies

Project funded by the European
Community under the "Information Society
Technology" Programme

Contract Number: 507953

Project Acronym: DBE

Title: Digital Business Ecosystem

Deliverable N°: 18.1

Due date: 31/10/2004

Delivery Date : 30/11/2004

Short Description: A wide ranging and interdisciplinary discussion of some of the theoretical problems underlying the DBE research, with practical implications on system architecture, use cases, and research methodology in order to incorporate evolution, intelligence, gene expression and network effects in the Digital Business Ecosystems.

Partners owning: LSE (Paolo Dini and Evangelia Berdou)

Partners contributed: ICL, Soluta.net, STU, Sun, TCD, UniS-Krause, UBham

Made available to: All project partners and the EC

Versioning

Version	Date	Author, Organisation
1	14 Nov 04	Paolo Dini, LSE
2	25 Nov 04	Paolo Dini and Evangelia Berdou, LSE
3	1 Dec 04	Paolo Dini and Evangelia Berdou, LSE
4	21 Dec 04	Integrated comments from 1 st reviewer P Dini and E Berdou, LSE

Quality check

1st Internal Reviewer: Jon Rowe (University of Birmingham)

2nd Internal Reviewer: Angelo Corallo (ISUFI, University of Lecce)

3rd Internal Reviewer: Petri Rasanen (TCH, Tampere)

PMEB Approval

Place:

Date:

Signature(s):



This work is licensed under the Creative Commons Attribution-ShareAlike License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/2.0/> or send a letter to Creative Commons, 559 Nathan Abbott Way, Stanford, California 94305, USA.

Table of Contents

EXECUTIVE SUMMARY	4
1. INTRODUCTION	5
Implications of the fundamental assumption of the DBE.....	5
The constructive cycles of the DBE.....	8
Modifications to workplan.....	8
Assessment of progress toward stated goals.....	9
Report road-map.....	9
2. STATISTICAL MECHANICS, DYNAMICAL SYSTEMS, AND SELF-ORGANISATION	10
Entropy.....	10
Free energy, phase transitions, memory, and emergence	13
Equilibrium, meta-stability, and the breath of life	17
Coupling and non-linear dynamical systems.....	21
Self-organisation in biological systems	21
<i>Spontaneous behaviour</i>	21
<i>Memory</i>	22
<i>Feedback and state transitions</i>	22
<i>Scale effects</i>	22
<i>An example: Phase transitions in an ant colony</i>	23
3. COMPLEXITY IN THE SOCIAL SCIENCES	25
Complexity theory and the social sciences.....	25
The complex and the social: interdisciplinarity, communication and commensurability	28
Social complexity in the social sciences: leveraging existing accounts and methods.....	31
Conclusions	33
4. INFORMATION AND NETWORKS: TOWARD AN ARCHITECTURE OF SELF-ORGANISATION	35
Shannon's information entropy and physical complexity (Briscoe, 2004)	36
"The network is the environment"	37
Cycles in fully connected and scale-free networks	38
The issue of time scales	42
Conclusions	42
5. SCIENTIFIC MODELS AND UML.....	44
Codifying self-organisation	44
On Engineering, Language, and DNA.....	46
Reconciling DNA with MDA™	51
Seeking the theoretical foundations of models.....	53
Some considerations on UML.....	56
6. DBE-SPECIFIC USE CASES	57
Basic elements.....	57
Use cases that rely on Science outputs	59
<i>Evolutionary computation</i>	59
<i>Distributed intelligence</i>	59
<i>Adaptive Object Model</i>	61
Global communication diagram of the DBE	62
PHYSICAL, BIOLOGICAL AND INFORMATION SCIENCES BIBLIOGRAPHY	64
SOCIAL SCIENCES BIBLIOGRAPHY	65

EXECUTIVE SUMMARY

This report examines several theoretical questions that relate to the formalisation of self-organising and evolutionary behaviour inspired by biology into the software infrastructure and services of the Digital Business Ecosystem. This software is meant to support and facilitate e-business practices between SMEs in such a way that optimum alliances and value nets will form among them, leading to regional economic growth in various parts of Europe. The report, therefore, addresses self-organisation simultaneously from the very different points of view of biological and socio-economic systems, with the objective of extracting universal models that can then apply also to software systems. We present a brief review of the new field of Complexity in the Social Sciences and of how other, more traditional analytical approaches in the social sciences have addressed the creation of order in socio-economic systems. We argue how the merging of evolutionary and information theory to define information relative to its environment can serve as the first step in understanding the dynamic aggregation of increasingly complex structures in these three types of systems. We conclude with a brief discussion of the context in which distributed intelligence and evolutionary computation will be used in the DBE, and with a snapshot of the current state of definition of the DBE system as represented through high-level UML diagrams and use cases.

1. INTRODUCTION

The vision of the DBE Project rests on the assumption that the self-organising and self-optimising behaviour of biological systems can be exhibited also by software systems. From a scientific point of view, the project's main objective is to understand how this could happen, and the various scientific tasks and workpackages are designed to address this central question in several different ways. Before we set out to solve this challenging problem, however, it is best to examine this assumption and its implications more closely. There are three main implications or expectations:

- 1- The fundamental physical laws upon which biological self-organisation depends have analogues in software systems
- 2- The self-organising behaviour of software, if attained, may replace top-down software engineering methods currently in use
- 3- The digital ecosystem dynamics and the business ecosystem dynamics will somehow reinforce each other

Implications of the fundamental assumption of the DBE

The first implication is addressed in tasks S22-“Scientific models and UML” and S23-“Evolutionary and self-organisation use cases”, for which LSE is responsible and that are discussed in this report. The second implication is not realistic, and may in fact carry the biological metaphor too far. It is more realistic to say that some progress has been made in the project toward the peaceful coexistence of evolutionary computation methods and model-driven software engineering, as shown in Fig 1.1. “U SME” means a software user SME, which could be either a consumer or a provider of services. “SW SME” means a software developer SME, i.e. a provider of DBE services. “Code segments” and “digital proteins” refer to the long-term view of the DBE where we may be able to assemble parts of individual services. This point highlights the pervasive presence in the project of two opposite philosophical positions on the construction of order: one is bottom-up and based on the emergence of symbols and structures from sub-symbolic signals and behaviour; the other is top-down design that starts with a language or a structured model and ends with behaviour. A practical reason to reconcile them is firmly anchored in a DBE design requirement or “boundary condition”: for all the biological behaviour that's expected of the DBE, human users representing SMEs must still be able to describe themselves and express their business and service offers and requests through a symbolic, graphical, or natural language interface. Additional motivation comes from the sensory and motor controls of the mind, which provide a very good example of how these two opposite processes can usefully be reconciled. It is not a coincidence, therefore, that the second implication is connected to the development of a UML profile that is compatible with evolutionary behaviour and artificial intelligence.

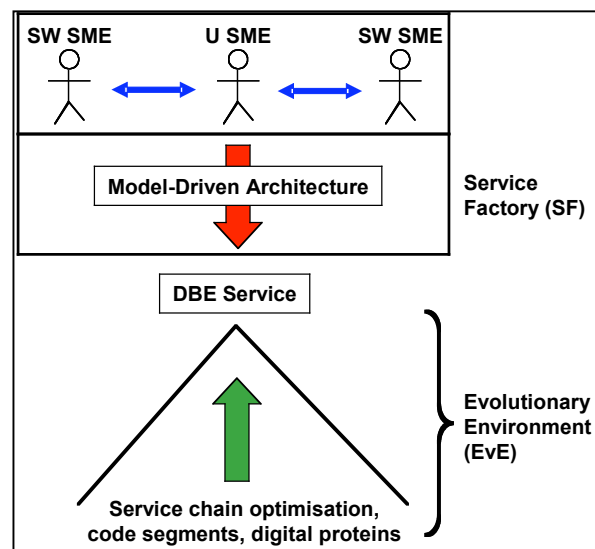


Fig. 1.1 Top-down and bottom-up construction of order in the DBE

The objective of the workpackage to which this report belongs is to develop a UML profile for the DBE. This requires capturing the desired dynamic and self-organising behaviour in the domain modelling constructs used to specify DBE infrastructural components and to some extent also the DBE software services. This objective is very ambitious partly because the problem to be solved is not very well defined. For instance, a theory of how biological systems achieve self-organising and self-optimising behaviour, i.e. how they manage to defeat entropy, is only now beginning to be developed. As a consequence, we are actually facing three problems: (1) understand the mathematical and physical basis of self-organisation; (2) apply these insights at a higher level of abstraction to postulate how software could achieve the same spontaneous order-creating behaviour; and (3) develop a language or semantic representation as a software engineering tool that incorporates the rules upon which such behaviour depends. The first two objectives are responsibility of LSE and are discussed here. The third objective is mainly covered by “C20-UML for DBE”, which is responsibility of Soluta.net and will be addressed in future deliverables.

The report also covers a review of the new field of Complexity in the Social Sciences, with a brief discussion of how other, more traditional approaches in the social sciences have addressed the creation of order in socio-economic systems. This workpackage aims to reach a quantitative understanding and subsequent formalisation of the dynamics of structure formation in physical and information systems, it does not pretend to develop a quantitative model of socio-economic growth from the complexity theory viewpoint. It is undeniable, however, that the formation of social institutions and economic systems can be rationalised as resulting, at least to some extent, from local interactions without conscious planning. Complexity theory tends to ascribe more autonomy to such “emergent” structures and systems than other currents of thought in the social sciences, so that, depending on where one draws the line, the process of structure formation in socio-economic and cultural systems appears more, or less, spontaneous. This line of thinking has been advanced by writers such as Diamond (1998) and De Landa (1997), and others discussed in Chapter 3, but a significant body of work has been done in the social sciences to study this phenomenon both before and after Complexity started being referred to as a discipline. From this point of view, over long time scales socio-economic systems have evolved spontaneously in different ways in different countries and parts of the world, building on pre-existing norms, over many thousands of years. Over short time scales, markets are a good example of self-optimising systems. For example, near-equilibrium at all scales, from local to global, between supply and demand could never be achieved through central planning (Kay, 2003).

Fig 1.2 attempts to capture the general idea as it may apply to the DBE. The figure shows how the DBE Services mediate business transactions and how these, in turn, contribute indirectly to the overall regional economy through a spontaneous process that is qualitatively similar to biological self-organisation in an ecosystem.




Fig. 1.2 Emergence of economic growth in the DBE

We therefore think that a review of how structure formation has been addressed in the social sciences literature will benefit us in three ways. First, socio-economic systems can serve as a source of empirical data for self-organising processes that complement the data from physical and biological systems. Second, compatibly with the interdisciplinary nature of the project, it is interesting and useful to compare from a philosophical and methodological point of view how different currents of thought in different disciplines have addressed the issues and the phenomena being grouped under the umbrella of Complexity Science. Third, this type of analysis will inform those tasks in the project that are concerned with the development of regional policies aiming to stimulate economic growth of the SME sector through the adoption of the DBE (B13-“Domain and Geographical Cluster Network Building”, B37-“Integration with Local Policies and Directives”, and B32-“The European Research Area and DBE Growth”).

Finally, getting to the third implication, it may at first sound suspiciously close to technological determinism. However, its bi-directionality combined with the above remarks opens other possibilities of interpretation. This implication or expectation is one of the main points made by the original discussion paper published by the ICT for Business Unit on the digital business ecosystem concept (Nachira, 2002), and strikes at the heart of the debates on the information society and the knowledge economy (Webster, 2002; Castells, 1996). Fig 1.3 proposes a possible integration of the main components of the DBE as three mutually reinforcing constructive cycles, each of which is driven by the business interactions between the SMEs. This integrated view is discussed in this report because the objectives of this WP require it to address most of the theoretical questions underlying the DBE and to find connections between them. Compatibly with the third implication, the discussion in Chapter 3 thus contributes to this unified view, working toward a unified theory of self-organisation and sustainability.

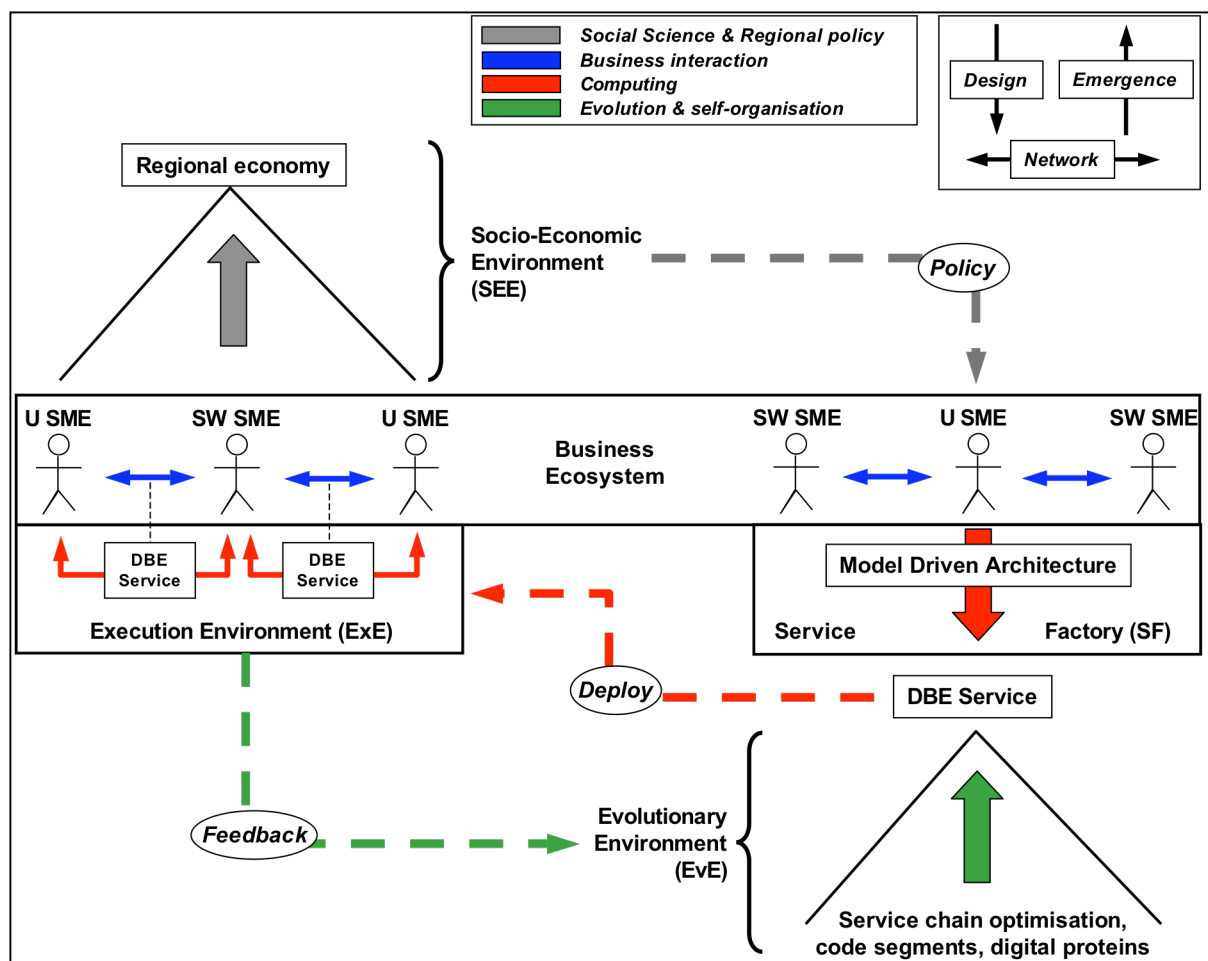


Fig. 1.3 The constructive cycles of the Digital Business Ecosystem

The constructive cycles of the DBE

The Business Ecosystem can be visualised as a two-dimensional plane (shown here in side-view) on which business networks of SMEs intersect each other in many different ways and at many different scales. Such networks could be correlated to geographical regions, business domains, or particular technological platforms. They may involve a small number or a large number of SMEs, they could involve large companies and multinationals, as well as business relationships with institutional bodies. They form the business ecosystem in the most general sense of the term. This is the engine that is driving all three cycles through “turbulent eddies” of business and market activity. “Emerging” from this activity is the Socio-Economic Environment (SEE) which supports the first cycle where empirical data from the local business ecosystems is formalised into regional economic policies to promote the growth of the SME sector and the adoption of the DBE (or ICTs in general). The second cycle springs from the same “turbulence engine” in the form of model-driven software engineering design and development, toward the description and implementation of software services in the Service Factory (SF). The SF is devoted to facilitating this task for SMEs, in order to raise the software development process from code implementation (resource-intensive, time consuming, expensive) to architectural modelling and specification (resource-light, fast deployment, lower-cost) (Dini et al., 2003). We will achieve this transition through business and service modelling and description languages (BML, SSL, SDL), recommendation systems, distributed intelligence and, to the maximum extent possible, automatic generation and testing of software from models. The deployment of the services in the Execution Environment (ExE) is effected by relying on the FADA¹ non-persistent storage layer, which is an http implementation for the Internet of Sun Microsystems’s proxy-based Jini technology. The third cycle also stems from the business interactions. User profiling provides passive and active feedback data that are correlated with BML and channeled to the Evolutionary Environment (EvE) to enrich the fitness data of services and models stored in the Knowledge Base. This stream of data and the explicit service requests by the SMEs represent the selection pressure used by the genetic algorithms to optimise chains of DBE Services, which are then tested and deployed to complete the cycle.

The pervasive presence of biological metaphors in all aspects of the project, combined with a clear emphasis on self-organisation and emergence as unifying concepts for the behaviour of the system, the infrastructure, and even the dynamics of interaction between SMEs, suggests that the DBE research is firmly aligned with the rubric of complexity science. This, however, is not exactly the case. The main reason is that while the interdisciplinary nature of Complexity is well suited to the aims of the project, it is too new and too heterogeneous to qualify as a “science”. Secondly, the very practical and applied objective to deliver a working software infrastructure before the end of the project dictates that we develop the system compatibly with existing technologies and design principles; we are deeply innovative but we cannot break with the past in a discontinuous way. Rather, we must find how to extrapolate from present software engineering practice toward an ambitious blending of ideas and design patterns from biology into software.

Modifications to workplan

In the original description of work for WP18 a social science component was not foreseen. The first author (Dini) changed institutions immediately before the proposal was submitted, when the text had already been written and the workpackages defined. It was therefore not possible to include in the original workplan a workpackage dedicated to the socio-economic aspects of the DBE research and housed at LSE. This research domain is therefore distributed in several workpackages, among which WP18. As stated above, because the tasks reported in this deliverable address some of the most fundamental and theoretical aspects of the DBE research, it was not difficult to redefine their scope in order to include a social science component with a theoretical focus. The initial results of this effort are discussed in Chapter 3 and will be developed further in future deliverables.

¹ Federated Advanced Directory Architecture, developed under FP5 project FETISH (<http://www.t-6.it/fetish/>)

Assessment of progress toward stated goals

We have made good progress toward the first goal, understanding self-organisation in Nature, at a conceptual level. We still have to develop a satisfactory mathematical theory or model for a self-organising dynamical system. The progress made in understanding self-organisation intuitively, however, is sufficient to give some indications for how this phenomenon may take place in software and information systems, our second goal. We believe that our suggestions and insights can aid and inform parallel research efforts of a more applied, engineering and empirical nature. In addition, they will benefit from a close scrutiny by and an interactive discussion with those partners in the consortium that are involved in similar theoretical efforts. We have also made some progress toward defining aspects of socio-economic systems that could benefit from a combined analysis from the mathematical and social science viewpoints, such as Social Network Analysis can provide. However, we were optimistic about how much progress we would be able to make toward our third goal. The connections between the theoretical basis of language and the theoretical basis of scientific models are sketchy and are discussed only briefly in Chapter 5. Consequently, how self-organising behaviour could be incorporated in the specification of a domain modelling language such as UML is not yet clear. We are estimating that reaching this third goal will require 6-12 months more than originally expected. Finally, the development of DBE-specific use cases is also behind. In Chapter 6 we go to great lengths to explain that “DBE-specific” is meant to emphasise the internal workings of the system and the interactions between its infrastructural components, which are characteristic of the DBE and not likely to be found in more traditional distributed software systems. It does not, therefore, refer to the more common meaning of “use cases” from the users point of view. Even with this in mind, the inner workings of the DBE infrastructure have taken longer than expected to be defined by the Computing team, and the contributions by the Science activities are likewise taking shape only now. So Chapter 6 does not attempt to capture more than a high-level view of the use cases specific to the DBE infrastructure.

Report road-map

The report is organised by focusing on the first goal in Chapter 2. The purpose of Chapter 3 is to open the discussion of self-organising systems as widely as possible, taking into consideration also socio-economic systems. While such a point of view is motivated by our desire to analyse and understand the dynamics of sustainable regional development, and the new field of complexity in the social sciences purports to achieve that, we are not oblivious to the huge theoretical and empirical challenges that it implies. Chapter 3 is therefore significantly more reflexive than the rest of the report, as we begin to construct the analytical foundations for how such an ambitiously inter-disciplinary research programme might be carried out. Chapter 4 then discusses mainly our second goal and provides some recommendations for the architecture of self-organising systems. Chapter 5 presents some initial ideas on the relationship between scientific models and domain models, as a first attempt to incorporate self-organising behaviour in the specification of the architecture and functionality of the software. The DNA forms the core of the concepts discussed in this chapter. Finally, Chapter 6 presents a high-level UML view of the DBE in its current state of development, with an emphasis on the use cases that make greater use of the output from the Science tasks.

Although this report pretends to address fundamental aspects of physical, biological, information, and socio-economic systems, it is written in an informal style and at a non-specialist level. This reflects the higher priority we have given in this workpackage, during the first year of the project, to building connections between very different disciplines over the specialised advancement of any one of them. We therefore hope that this report will contain understandable and useful information for all the stakeholders of the DBE project.

2. STATISTICAL MECHANICS, DYNAMICAL SYSTEMS, AND SELF-ORGANISATION

The purpose of this chapter is to touch on the main ideas and results from physics that underpin our current understanding of self-organisation in biology, which remains limited. A thorough and complete review is not attempted. Rather, we wish to identify potentially fruitful connections between physics and computer science that will warrant deeper study in forthcoming deliverables. The intended audience are other DBE partners who may need to understand the conceptual framework of self-organisation from a physics/biology point of view in order to follow and contribute more easily to the discussion of later chapters of how these ideas might apply to information systems.

We start with an elementary discussion of entropy, which we need to understand thoroughly if we are to have a chance at understanding the construction of order. A short description of phase transitions follows because of their potential relevance to emergence. Interactions, coupling, and non-linear coupling are discussed next. How these different classes of phenomena may conspire to give rise to self-organisation in biological systems is then postulated to conclude the chapter.

Entropy

Statistical mechanics studies systems of many particles that may or may not interact with one another and with other systems. This theory allows us to derive macroscopically observable and measurable equilibrium properties of these systems—for example, pressure, energy, heat capacity, etc—from a knowledge of the average (“statistical”) behaviour at the microscopic level.

The most important aspect of the dynamics of many-particle systems is that there is a definite time axis given by the increase in entropy. This is in contrast to classical mechanics of particles, which is symmetrical to inversion of the time variable in the governing differential equations. To explain the fundamental idea, take an insulated box with 10 identical free particles in the absence of gravity. Give the particles an initial random position and velocity such that the total kinetic energy of the system is thereafter constant. The initial positions of the 10 particles are all on the left-hand side of the box. Assume there is no interaction between the particles, thus the potential energy of the system is zero. The particles will bounce around and fairly soon will visit both sides of the box. To get a sense of how likely it is for the particles to find themselves again on the left-hand side at some future time, let's count the number of configurations for the three cases shown in Fig. 2.1. There is clearly only one way for the 10 particles to be on the left. Notice that this follows from the fact that we are not subdividing our position resolution grid beyond the 2 cells shown, left and right, thereby obtaining a so-called binary system. If we allow any 3 of the 10 particles to be on the right there will be 120 ways to achieve that:

$$\text{"10 choose 3"} = \binom{10}{3} = \frac{10!}{(10-3)!3!} = \frac{10 \cdot 9 \cdot 8 \cdot 7 \cdot 6 \cdot 5 \cdot 4 \cdot 3 \cdot 2 \cdot 1}{(7 \cdot 6 \cdot 5 \cdot 4 \cdot 3 \cdot 2 \cdot 1) \cdot (3 \cdot 2 \cdot 1)} = \frac{10 \cdot 9 \cdot 8}{3 \cdot 2 \cdot 1} = 120 \quad (2.1)$$

For 5 particles on each side there are 252 possible combinations. It should be pretty apparent that the case where the particles are distributed fairly evenly between the two halves will be observed most often while the case where there is only one or no particles on either side will be observed very rarely.

Entropy is introduced to provide a quantitative model of this behaviour. We treat the number of particles on the right of the box as a parameter that characterises a macroscopic state or constraint on the system. For each such macrostate we count how many configurations or microstates there are, thereby obtaining a function that we call the multiplicity. Boltzmann defined entropy as the logarithm of the number of microstates, so that equilibrium corresponds simply to the maximum of this function. This argument rests on the subtle point that while we have introduced a parameter to distinguish between different macrostates, each microstate, irrespective of the macrostate it belongs to, has equal probability of occurrence (“fundamental assumption of statistical mechanics”). This is because all the microstates we have discussed have equal energy and therefore in this isolated system they all qualify as accessible quantum mechanical states. So in a sense there is a hierarchy in the macroscopic

constraints we are discussing: the principal constraint that cannot be violated by our insulated box is its total energy, which is constant. The additional parameter we have introduced, the macrostate, is more a device for classifying the accessible states into different groups than an actual constraint. This classification, in turn, enables us to define the entropy function. The fundamental assumption then allows us to regard the largest such group as the most probable.




Fig 2.1 Entropy: 10 free particles in a box

The hierarchy mentioned reflects the relative importance of the first and second laws of thermodynamics. The first law is a statement of conservation of energy, while the second law is concerned with the direction of processes at a given energy. Our language, however, does not reflect this distinction as accurately as the mathematics. When we say “macrostate” we could be referring to either the amount of energy our system has or to the finer classification of its internal groupings for a given energy. We therefore need to pay close attention to the context. For isolated systems macrostates that correspond to the largest number of accessible states will also correspond to equilibrium and to the state of highest entropy. The Gaussian curve shown in Fig. 2.1 thus generalises to any isolated system. Addressing finally the time axis question, it arises from the simple probability argument that it is highly unlikely for a system that starts in a state of high entropy to evolve toward a state of low entropy. For physical systems composed of many particles, on the order of 10^{23} , “highly unlikely” assumes galactic proportions. This situation justifies the distinction between reversible and irreversible processes and the 2nd “law” of thermodynamics in general.

The condition of equilibrium identified as the state of highest entropy that applies to this example is valid for all systems that are insulated from their environment such that their energy is constant. The expression that quantifies the probability of an isolated system to find itself in one of its accessible states is referred to as the “microcanonical distribution.” Such systems are sometimes called microcanonical ensemble. Ensemble refers to the averaging effect provided by observing a single system for a very long time and that can be achieved equivalently by averaging across a large ensemble of identical systems. If we start with the fundamental assumption and denote the number of accessible states for a system of N particles $g(N)$, the probability that the system is found in any one of its accessible quantum states is given by the microcanonical distribution mentioned above:

$$P = \frac{1}{g(N)} \quad (2.2)$$

This distribution is perfectly uniform and in this context it is not very useful. It is better to speak of the probability of a given macrostate. To that end, we define a parameter similar to the number of particles in the right side of the box, that measures deviations from equilibrium, and that we call the “particle excess” (Kittel and Kroemer, 1980):

$$s = \frac{1}{2}(N_{left} - N_{right}) \quad (2.3)$$

(the factor of $\frac{1}{2}$ makes the algebra easier later) s varies from $-N/2$ to $N/2$. The number of ways in which the particles can be arranged for each value of s is

$$g(N, s) = \frac{N!}{N_{left}! N_{right}!} = \frac{N!}{(\frac{1}{2}N + s)! (\frac{1}{2}N - s)!} \quad (2.4)$$

while the total number of possible combinations for all values of s is

$$g(N) = \sum_s g(N, s) = 2^N \quad (2.5)$$

Either (2.4) or (2.5) is referred to as the multiplicity function, depending on what level of macrostate one means. Finally, the probability of the system being observed at a particular value of s is simply

$$P(s) = \frac{g(N, s)}{g(N)} \quad (2.6)$$

For instance, the probability that 3 particles (for which $s=2$) will be observed on the right of the box is

$$P(2) = \frac{120}{1024} \approx 0.117 = 11.7\% \quad (2.7)$$

While the probability that 5 particles will be observed on each side is

$$P(0) = \frac{252}{1024} \approx 25\% \quad (2.8)$$

When N is large, on the order of Avogadro's number or 10^{23} , Eq. (2.3) becomes unmanageable because the factorial function grows too rapidly. A good approximation has therefore been developed:

$$g(N, s) \approx g(N, 0) e^{-\frac{2s^2}{N}} \quad (2.9)$$

This is the equation for the Gaussian curve sketched in Fig. 2.1. For large N , this function has an extremely sharp maximum, which makes the equilibrium condition enormously more probable than any other. The probability that a particular macrostate will be observed is thus given by

$$P(s) = \frac{g(N, 0)}{g(N)} e^{-\frac{2s^2}{N}} \quad (2.10)$$

Free energy, phase transitions, memory, and emergence

Free energy, phase transitions and critical phenomena all rely on the existence of interaction forces between the particles of a statistical system. No such interaction forces exist between abstract digital entities. There are some aspects of these physical phenomena, however, that are relevant also for information systems, as we will discover in this chapter.

Free energy is a somewhat abstract concept that is difficult to absorb at first. It is analogous for systems of many particles to the mechanical energy of macroscopic mechanical systems (defined as kinetic + potential energy). In other words, it is most usefully thought of as the energy of a system that is “free” or “available” to do work. Unlike mechanical energy, however, it is generally not possible to obtain an absolute measure of the free energy of a system; rather, changes in free energy are usually calculated. As we might expect intuitively, if during a process the free energy *decreases* (the change is negative), then the process is *spontaneous*, whereas if the change in free energy is positive then the process is not spontaneous and work needs to be done on the system by an external agent in order to arrive at the desired end-point of the process. It is an interesting and important fact that a process that is spontaneous at a particular temperature may not be spontaneous at a different temperature.

In a system of many particles free energy is not made up of just kinetic and potential energy. It also measures changes in entropy. The Helmholtz free energy is defined as

$$F = U - TS \quad (2.11)$$

where U is the internal energy, T the temperature and S the entropy. This definition does not mean much in intuitive terms, but things improve if we think about the change in free energy during a process. The Helmholtz free energy applies to processes occurring at constant temperature and volume:

$$\Delta F = \Delta U - T\Delta S \quad (2.12)$$

U is the internal energy and contains both potential energy of interaction and microscopic thermal energy. Notice that specifying a constant-temperature process usually implies that the energy of the system is not constant, i.e. the system exchanges energy with a reservoir. If $\Delta U = 0$, then we recover the condition for equilibrium of the microcanonical ensemble, meaning that the minimum of free energy corresponds to the maximum of entropy.

There is another form of free energy, the Gibbs free energy, that applies to processes occurring at constant temperature and pressure:

$$\begin{aligned} \Delta G &= \Delta F + p\Delta V \\ &= \Delta U + p\Delta V - T\Delta S \\ &= \Delta H - T\Delta S \end{aligned} \quad (2.13)$$

where H is the enthalpy. Enthalpy plays the same role for systems at constant pressure that the energy U plays for systems at constant volume. The Gibbs free energy is generally more useful for describing chemical reactions since they usually occur at constant pressure.

Free energy is of fundamental importance for describing and understanding self-assembly. Protein crystallisation is a good example.² Take a super-saturated solution of protein molecules that deposit at the bottom and grow into a crystal. Why does this happen? One way to look at this is to say that the potential energy of interaction between molecules is higher when they are apart than when they “fall” together due to favourable electrostatic attraction. This would lead to a decrease in entropy due to

² I (Dini) am indebted for this example and many fruitful discussions to Dr Steve Durbin of Rightnow Technologies Inc.

their more ordered state, but the change in potential energy would be large enough to compensate for that and give a negative change in free energy. In fact, a completely different process appears to drive crystallisation. Each protein in solution is covered by a layer of water molecules that stick to it in whichever orientation happens to be more energetically favourable. This configuration has a relatively low entropy from the point of view of the water molecules. As the concentration of the protein molecules in solution increases, the fraction of water molecules that fall into this relatively ordered state increases. The super-saturated state, therefore, could be defined as the state where the increase in entropy of the water molecules due to their freeing themselves of the proteins is sufficiently large that it overcomes the energetically *unfavourable* binding of the proteins into a crystal. In other words, the free energy of the whole system decreases even though the crystallisation process by itself is *not* spontaneous in the absence of water.

Analysing this situation more carefully, we need to emphasise that the concept of free energy only makes sense when the particular system of interest is in thermal contact with a reservoir, such that the system energy is not constant. Otherwise we recover the microcanonical ensemble already discussed since the only thing changing is the entropy. In the case of the crystallisation, therefore, perhaps the correct way to think about it is that there are two coupled and interacting systems:

- The first system is comprised by the growing crystal plus a thin layer of proteins near the interface, with the water solution playing the part of the reservoir. In this case the potential energy of the crystal *increases* as it grows (since the proteins are actually *repelling* each other), absorbing energy from the water. Its entropy *decreases* as more proteins achieve a highly ordered state; thus we expect a sizable *positive* change in free energy—the process is not spontaneous by itself.
- The second system is comprised by the water, with the growing crystal playing the part of the reservoir. In this system a small part of the energy of the water is leaking out into the crystal but, more importantly, the entropy is increasing significantly due to the release of the water molecules into a state of high disorder. The free energy of this system, therefore, is definitely *decreasing*.

It would appear that the decrease in free energy in the second system is greater than the increase in the first system so that the crystallisation process can indeed occur spontaneously in a water solution.

In these examples of free energy some properties of the system are fixed (constraints, such as the number of particles) and some properties are allowed to vary (such as the energy of a sub-system). Minimisation of the free energy implies the spontaneous variation of one or more properties upon which the free energy depends until a minimum is reached. Free energy can be regarded as an example of a more general thermodynamic potential Φ that is minimised when certain parameters of the system such as the energy, the number of particles, the density, or the magnetisation are allowed to vary. Under certain conditions, the thermodynamic potential may exhibit more than one minimum when expressed as a function of a variable parameter to which we can give the generic name ϕ . This is the cause of phase transitions. Fig. 2.2 shows qualitatively the dependence of the thermodynamic potential on the parameter that is allowed to vary, at different temperatures (Bowley and Sanchez, 1999). In this case the order parameter could be the density of water that jumps to a lower value as the temperature crosses the boiling point. Figs. 2.3 and 2.4 show the same information in 3-D, with specific volume plotted instead of density (specific volume = 1/density).

If the phase transition takes place between the solid and liquid states then it acts as a memory mechanism. Familiar examples are the alignment of magnetic domains in permanent computer memory. The head adds just enough energy to the magnetic material to flip the domain in the desired direction, letting it “freeze” again to retain its state. A large number of memory phenomena in Nature are based on the same principle of freezing a desired pattern into a medium that is stable and will retain its state. Self-organising processes on long time scales, such as biological evolution, have perfected the ability to encode in the DNA not just static patterns but triggers for asynchronous instructions on how best to interact with the environment to extract energy to reproduce itself.




Fig. 2.2 Thermodynamic potential at different temperatures (Bowley and Sanchez, 1999)




Fig. 2.3 Visualisation of first-order phase transition (1)




Fig. 2.4 Visualisation of first-order phase transition (2)

The 3-D figures above show how the free energy changes as a function of the order parameter and temperature *for a constant pressure*. Thus, they correspond to the dotted line in the phase diagram of water in Fig. 2.5, below. If, instead, we follow the solid phase transition curve along the points 1, 2, 3 and 4, Fig. 2.6 shows how the two symmetrical minima disappear the closer we get to the critical point 4. Notice that at the critical point the free energy has a very flat minimum. If the surface is extended beyond Point 4 the potential will develop a quadratic shape with a single well-defined minimum.




Fig. 2.5 Phase diagram of water




Fig. 2.6 Visualisation of critical phase transition

The presence of a very flat minimum at the critical point means that the actual minimum is not well defined and large fluctuations in density are possible (and indeed observed). Another way to put this is that near the critical point the range of the correlated regions grows without bound. As Callen says,

This onset of long-range correlated behaviour is the key to the statistical mechanical (or “renormalization group”) solution to the problem. Because the large regions are so closely correlated, the details of the particular atomic structure of the specific material become of secondary importance! The atomic structure is so masked by the long-range correlation that large families of materials behave similarly—a phenomenon known as “universality”... (Callen, 1985)

The implications of universality are remarkable and in fact suggest that critical phenomena may transcend physical systems. In our discussion they are potentially relevant in two ways. Firstly, biological systems rely on the fact that the change in free energy between the end-points of a process is a function of temperature to turn on and off spontaneous chemical and self-assembly reactions inside the cell through small temperature fluctuations. Since the local fluctuations in temperature in the cytoplasm (the medium of cells) due to exogenous or endogenous reactions are extremely small, the system must be very carefully tuned around $\Delta G = 0$ for a significant transfer of control between different parts of the cell to be possible. Such transfer of control allows the events triggered by a particular gene becoming active to act, themselves, as triggers that will cause the next state change of the DNA. In this manner metabolic cycles are enacted entirely through spontaneous reactions. A corollary is that the DNA only holds part of the “instruction segment”. The other part is distributed in its environment. In fact, a change in the environment (a toxic substance) can change the metabolism in radical ways. Biological organisms, therefore, execute what we call a “distributed algorithm” in cooperation with their environment, as if organisms and their environment were part of the same “computer”. In the follow-up to this deliverable we will investigate the question of whether this whole dynamic requires the cell to be near a critical phase transition, as the $\Delta G = 0$ boundary would seem to suggest.

Secondly, at the critical point local interactions are amplified to propagate to the whole system, providing a possible amplification mechanism for communication of signals between widely different scales, which is one of the phenomena that can be associated with the concept of “emergence”. Inside the cell, such an effect would be desirable, for example, to communicate to all the ribosomes (millions per cell) that are working on the synthesis of a particular kind of protein to produce it together, before going on to the next task required by the DNA. The large numbers of proteins produced (either structural or functional, i.e. enzymes) give a high probability that a desired action coded in the DNA will be executed on a meso- or macro-scale. Thus, critical phenomena could be a mechanism to achieve parallel computation while simultaneously approximating deterministic coupling to macroscopic dynamics.

The reason these phenomena may be relevant to the DBE is that they appear to exploit statistical behaviour to achieve emergence. This over-used term can mean many things and also in this context could refer equally well to a significant reduction of degrees of freedom, to the synthesis of control signals from concentration distributions of chemical agents, or to the mediation of non-linear coupling between control centres of the cell. The important point is that the state of the DNA is affected by these chemical and statistical processes and, in turn, a change in state of the DNA causes a wave of signals to propagate back to influence and control the time evolution of the protein synthesis reactions. At this physical level, this scenario suggests that the important part is not the medium, it is the information. As a clear example of a similar dichotomy between the “channel” and the “message”, the analogue computers of the 1950s were able to perform various mathematical operations such as integration and differentiation of functions by relying solely on “analogous” electromagnetic interactions between inductors, capacitors, and resistors. The important part was not the manner in which the mathematical operations were performed, the important part was the mathematical operations themselves. Digital computers, in fact, were designed soon thereafter that could perform the same mathematical operations through elementary arithmetics. This equivalence mirrors the fundamental theorem of calculus, which states that a continuous integral is the limit of a discrete sum as the number of elements approaches infinity and their size approaches zero. Thus, we are motivated to analyse and brainstorm about how non-linear, cooperative control and parameter reduction is achieved in biological systems because we aim to extract the abstract model of how such dynamical systems are able to perform such a broad array of sophisticated functions whilst remaining stable. One of the aims of the tasks reported here (S-22 and S-23) is to design such an abstract model into the software of the DBE. But first we need to clarify what we mean by stability.

Equilibrium, meta-stability, and the breath of life

Linear systems in equilibrium are called stable if there is a (linear) return force that opposes their displacement from equilibrium. Any such system endowed with some form of inertia will have

acquired some velocity upon reaching equilibrium on its way back from its displaced position, and will therefore overshoot the equilibrium point. The return force will bring it back again, repeating the cycle and giving rise to an oscillation. All vibration and wave phenomena can be shown to arise from this simple mechanism (equilibrium + return force + inertia), which is invoked well beyond mechanical systems and, as metaphor, well beyond physical systems.

A simple harmonic oscillator by itself is exactly solvable. Two or more coupled oscillators (unless they are “weakly coupled”, in which case the solution is usually also possible) give rise to a wide range of behaviour, from orderly to random, that cannot in general be solved analytically. This is actually good. The most interesting behaviour of such (non-linear) systems can be observed when their parameters are tuned half-way between order and chaos. Quite rich and pleasing structures can be seen, and yet the constant trading of energy between the components causes the state of the system to meander quite unpredictably, generating ever new and surprising patterns. Although we talk about “trading energy”, it is a little recognised fact that physicists do not know what energy actually is. All we know is that it is conserved, that it is exchanged, and that it is transformed. And that when it is exchanged or transformed things happen, and we speak of the exchange or transformation of energy from one form to another as an energy “flow”, as when the chemical energy stored in natural gas flows through our houses in winter, heating them up, upon being transformed into thermal energy by the combustion process. Energy flow can also mediate an exchange or a flow of information, as mentioned above in connection with analogue computers. This observation forms an important link between natural systems and the DBE, in which the “breath of life” arises not from a flow of energy but from a flow of information. This concept will be developed partly in this chapter and partly in Chapter 4.

The application of the concepts of oscillations and stable equilibrium to biological systems is not so straightforward. We may be tempted to view muscles as springs, but in fact there is no “return force”. When we fall from 1 metre and bend our knees to absorb the shock we are not overshooting an equilibrium point, there is no tendency in our muscles to “spring back” to their original position. When they extend to slow our fall the muscles simply release the chemical potential energy of sugar, and when they contract to bring us back to a standing position they burn some more sugar. This chemical potential energy is never recovered as such, although it can be stored as gravitational potential energy of our body as, for instance, when we climb a mountain. All biological organisms ultimately rely on never-ending spontaneous chemical and physical processes to function, such as the release of energy from sugar. These processes are spontaneous precisely because they correspond to the approach to equilibrium. This happens both at the molecular scale in terms of the response of the atoms to electronic forces as they participate in chemical reactions, as well as at the macroscopic scale in terms of thermodynamic equilibrium as the most likely distribution of matter in a given volume (maximisation of entropy as discussed above). Although biological systems are always approaching or “falling” toward equilibrium, they never reach it. So the concept of equilibrium is still present when talking about biological systems, but in a more general way than how it applies to simple mechanical oscillators.

In this connection, it seems worthwhile to clarify that the not-so-well defined and very popular term used in complexity studies, “far from equilibrium”, actually means just what we said: it describes systems whose state is changing spontaneously because they are *approaching* equilibrium. Approaching equilibrium does not necessarily imply that they ever reach it. In biological systems, equilibrium is more closely associated with death than with life. Thus, living organisms manage always to be falling toward equilibrium without ever reaching it, and yet to retain some kind of stability.

We all know that as long as we don’t hit our head too hard, as long as we eat healthy food, and as long as we don’t catch a bad disease we will remain alive—until our DNA decides it’s time to turn off. Since the chemical potential energy released by food molecules is never recovered by the same molecules themselves but is traded for energy of new compounds, until sugar is finally converted into mechanical energy and heat, on the scale of the human body the flow of energy is uni-directional. We

can visualise this flow as energy that enters the organism carried by food molecules and trickles “down”, turning the wheels of the metabolism like sand falling out of a sand mill’s funnel turns the wheels of this popular beach toy for children.




Fig. 2.7 Sandmill as metaphor of the body’s metabolism driven by the energy of food, or of the biosphere driven by the energy from the Sun

On a smaller scale, however, as they fall toward equilibrium the many components of biological systems trade energy and information between them, in a direction that is “perpendicular” to the overall energy flow, affecting each other’s behaviour and causing the organism to meander through its own state space. This is as if the two wheels of the sandmill were connected and the motion of one affected the motion of the other, and their state together affected the state of the sandmill. In this respect and at this scale of description biological systems behave like dissipative non-linear oscillators of many degrees of freedom, where the trading of chemical potential energy between compounds is not unlike the trading of potential energy between springs. We could hardly call such systems “stable”. However, we do distinguish between “alive” and “dead”. This distinction means that biological systems operate within the analogue of a finite potential well of stability. Since they are never static and since they are never in equilibrium, a new term is called for. We are referring to this state as “meta-stability” or “generalised stability”.³

This section started with a discussion of stability but unwittingly uncovered another, more fundamental feature of complex self-organising systems. Insofar as self-organisation necessarily implies a change of state (or a succession of changes of state) there must be a cause for such state changes. Self-organisation refers to a process in which the end-state is more “organised” than the initial state, however we may wish to define what “organised” might mean,⁴ but says nothing at all about how such a change of state might come about. In other words, even if there is no engineer deciding how the pieces should be assembled, the pieces still need to move in order to assemble themselves. In physical systems the “cause” is provided by an energy flow and the fall toward equilibrium; the above discussion argues how the unidirectional energy flow from the Sun drives the biosphere into infinitely varied interactions that, through various memory mechanisms, give rise to ever-changing ecosystems. In the DBE we do not have such an energy flow (although something that has a similar effect is discussed in Chapter 4). The closest analogue to an energy flow in the DBE is an information flow since it too causes changes of state in the components of the system, and the analogue of the Sun as the source of this flow are the users. The analogue of the sandmill metaphor in the DBE, therefore, is represented by Fig. 2.8, which shows an early depiction of the DBE architecture that was developed during the proposal preparation as part of a discussion addressing the minimum requirements for a system to be self-organising.

³ The term “meta-stable” has already been defined in the dynamical systems literature in a manner that is related to the concepts discussed here. An alternative term found in the literature is “non-equilibrium steady state”, but it’s not clear if it refers exactly to the same phenomenon.

⁴ See Chapter 4.

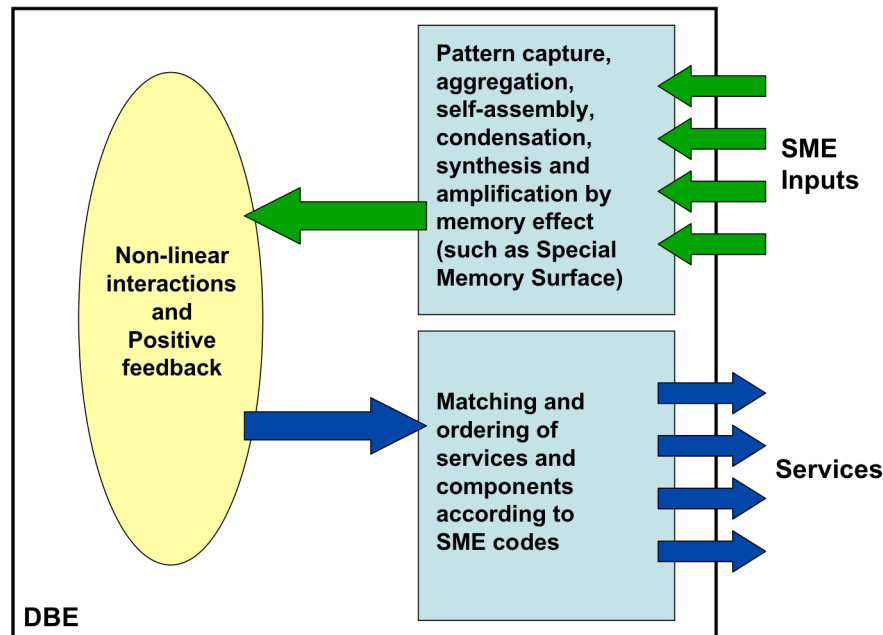


Fig. 2.8 Software structures arising from an information flow driven by the users

This discussion is partly inspired by De Landa (1997), who describes the formation of structures in the history of humanity as arising at the cross-roads of a flow of energy and a flow of information, and recognises a similar memory-based construction process when looking at human history through geological, genetic, and linguistic lenses. This however is not the right place for an in-depth discussion of this fascinating book. For our purposes it is sufficient to have understood that the main purpose of the DBE, to support interactions with and between its users, can in principle be harnessed as the source of both the “energy” that drives its internal state changes as well as the information that, through a passive memory mechanism, shapes these internal state changes into organised behaviour. What is still missing is an understanding of how meta-stability can be relied upon to “design” desired behaviour into open self-organising systems such as the DBE.

Engineering systems have traditionally been designed to stay close to equilibrium and certainly stable. Application of linear mathematical models to design have carried us a long way and complex machines like airplanes can now operate stably and predictably within a well-defined operational envelope. Business, economic, and social system, by contrast, defy exact quantification, or even satisfactory approximation through reductionist models that attempt to “capture” their essence. Yet, to varying degrees, they appear to function and we to function within them. This is more due to the fact that they tend to remain stable than to the fact that we have been, or will ever be, able to develop quantitative mathematical models that truly reflect how they work. Their stability is greatly dependent on history, i.e. on incremental additions to a previous system that was already working reasonably well. Thus the rich free-market economies of today can be traced to economic systems before the industrial revolution that already operated by similar principles of free trade, and to social and cultural institutions that co-evolved with these economic systems over a period of many centuries. Borrowing from John Kay’s recent book (Kay, 2003), economic development experiments such as India after WWII reinforce the importance of meta-stability. There, new technologies and economic policies from the West were brought into a system whose social and cultural institutions had developed along a completely different path to reach a completely different state of meta-stability. Injecting technologies artificially and planning the economy from the top did not work. After 50 years and after the co-evolution mechanisms of free markets have had a chance to work for the last 15, India is finally starting to show signs of the type of growth that had originally been envisioned.

The above example is only meant to highlight the “mechanical” and history-dependent aspects of complex systems. It is not a defence of unbridled free-market economics, whose shortcomings we are all familiar with. What matters here is that there are some basic principles of meta-stability and self-organisation that characterise both socio-economic systems and biological systems and that we will need to accept and learn to live with in a constructive way. Furthermore, if we truly want to develop evolving, adaptable, and self-optimising software, we need to understand how these ideas of meta-stability and self-organisation can be applied to the behaviour of software components and to the design of software architectures and algorithms.

Coupling and non-linear dynamical systems

At this point a discussion of low-dimensional non-linear dynamical systems, generally associated with “chaos theory”, would be logical and helpful. However, we hesitate to present it here because we have not been able to establish yet a direct and quantitative relevance to the DBE system. We will therefore only mention at a more qualitative level the aspects that seem most relevant. Non-linear systems of coupled degrees of freedom provide the clearest model for systems where no single actor can be assigned the role of “cause” or “effect”. Each variable or degree of freedom plays both roles. As we are beginning to understand, in low-dimensional smooth dynamical systems the transition between different states is modulated continuously by the exchange of energy between degrees of freedom. In an information systems context the same idea is expressed by saying that the transition between different states is effected discretely by the exchange of information between different components. In a dynamical system the physical constraints cause the energy to slosh back and forth between the degrees of freedom as would the water in a tub mounted on a seesaw. In a biological system such as the cell we have already hinted that the DNA and its environment are locked into a similar dance. In an information system we can abstract the same behaviour by programming each component to react to events raised by the others by changing its own state and sending back an event to the others. It is not difficult to imagine how such events could be interlocked in a sequence of mutually triggering steps. It is more difficult to imagine how such a real-time and context-dependent response could enact an engineering or functional specification or a business process, especially if the components in question were developed independently.

As a river follows the bottom of a valley, and is constrained by its geometry, so we need to learn how to constrain the desired behaviour of our software to follow the potential well of generalised stability that appears to guide non-linear processes with distributed control. The distinction between unstable and meta-stable states of non-linear systems has not to our knowledge been established in general. In any case, it would seem to require a solution, or some level of approximation, of the governing differential equations, which may or may not be solvable. In the forthcoming deliverables we will continue to investigate the symmetry analysis of non-linear systems of differential equations as a method that may be suited to detect ordered, meta-stable, and in some distributed sense programmable dynamical behaviour.

Self-organisation in biological systems

We have by now touched on a few mechanisms associated with self-organisation in different contexts: free energy minimisation for the self-assembly of crystals and molecules, the distributed algorithm of meta-stable non-linear systems for the autonomous running of the cell’s metabolism based on the interaction between the DNA and its environment, and biological evolution for storing successful structural and behavioural traits on the scale of many generations. These systems have several things in common, which we will now discuss briefly in order to derive a set of essential characteristics shared by all self-organising systems.

Spontaneous behaviour

Whether it is caused by electromagnetic forces (minimisation of energy), probability (maximisation of entropy), or a combination of the two (free energy), spontaneous physical behaviour at any scale can be considered an effect arising from a cause. This is a familiar and reassuring statement in its determinism. However, it is not as easy to accept that all the complex decisions made by a cat in hunting a mouse can be traced back to free energy minimisation in its brain as well as in its muscles.

We know it is true, but we suspect there is more to the story. In the same way, although we are aware that the most complicated software applications are reducible to bits, their structure and behaviour are defined and determined by higher-level constructs, composed of *groups* of bits. In the DBE we need to understand how such higher-level constructs could form spontaneously. In constructing similes with biological systems, we may recognise a similarity in the structural hierarchy. However, in information systems there is no internal source of energy; as we stated above the closest we can come to is a flow of information. Where could spontaneous behaviour in information systems originate from, therefore? This leads to a definition of self-organisation based on memory.

Memory

When the temperature of a multi-particle system decreases, the number of accessible states decreases, the kinetic energy of individual particles decreases, and groups of interacting particles settle closer to a low-energy, stable equilibrium configuration. Order wins over (Greek) chaos. As mentioned above, this is one of the fundamental mechanisms of memory. Self-organising systems use this mechanism as a ratchet, to add complexity incrementally with each interaction. In an inspiring 1969 book Edward De Bono (De Bono, 1969) presented a simple model of how thoughts are constructed by recursive application of the neurons' tiring effect and long-term decrease in activation thresholds, which conspire together to create a higher-level memory effect whose chemical and statistical mechanical basis we can afford to ignore. Looking "up" the scale hierarchy, the same principle applies again: there is a lot more to the story of how the mind works but all the functions of the mind can be reduced or traced back to the underlying simple mechanisms of the neuron. In this sense De Bono's book is a conceptual milestone, even if only the lowest rung in the theory of the mind. The book is important in this discussion for two more reasons. First, it provides a convincing scenario for how physical brain function merges continuously with the activity of the mind, without encountering a metaphysical boundary. This is relevant because in the DBE we too would like to understand how to extrapolate physical self-organisation to abstract information systems. Second, it introduces a definition of self-organisation that is quite general: the mind is a self-organising system because it is a passive memory system exposed to an environment that is the source of all patterns.

This definition provides an alternative to free energy minimisation as the cause of spontaneous behaviour. Like De Bono's model of the mind as a memory surface, information systems are also passive. The spontaneous behaviour originates *outside* of the system, in the environment, which is composed by their users. There is no energy exchange but the users provide information, patterns, and behaviour. So to obtain a self-organising information system we just need to give it a memory. Such a system is self-organising because the internal constructs form spontaneously and incrementally through repeated interaction with the environment. There is no plan and no outside control. Depending on the type of patterns being stored, we could be talking about an artificial neural network (sub-symbolic encoding of implicit behaviour in response to external inputs), or about an expert system (symbol copy and formal logic encoding of explicit behaviour). In the DBE we need to develop the ability to synthesise symbols and formal constructs from sub-symbolic learned behaviour.

Feedback and state transitions

The distributed algorithm is just one of countless examples of non-linear coupling and feedback in complex systems. The idea of the state of a part of a system affecting another part whose change of state, in turn, affects the first is easily generalisable to software systems. Whether or not the application of a non-linear model based on coupling and feedback to passive information systems with memory will lead to an AI system that can extract symbols from behaviour is still unclear, but this is one of the directions we are pursuing. Early examples are encouraging, as seen in methods such as Collaborative Reinforcement Learning.⁵

Scale effects

Whatever our model for self-organising systems will be able to do we know that it will have to be able to scale. This is both an unavoidable functional requirement as well as, potentially, an interesting non-

⁵ I (Dini) wish to thank Jim Dowling of TCD who brought this to my attention.

functional requirement. The argument hinges on the observation that as the number of interacting components or neurons increases, the feedback loops are not necessarily going to involve only pairs of nodes. Other configurations and interaction topologies may become possible or even preferred. In addition, self-organising systems in Nature tend to be organised in structural and functional hierarchies. Whether or not this is a consequence of scaling is not clear, but it would not be at all surprising if some global minimisation principle, such as minimisation of energy that determines the geometry of soap films, were at play. A possible effect of scaling on the architecture of self-organising systems will be discussed further in Chapter 4.

An example: Phase transitions in an ant colony

Having presented what we regard as the basic elements of a self-organising system (spontaneous behaviour, memory, feedback, and scaling effects), we are going to discuss a biological system that incorporates all of them in a manner that could be considered half-way between molecular self-assembly and our postulated vision of self-organisation in information systems: an ant colony. Many papers have been written on the topic, from the seminal essays of Dr Lewis Thomas (Thomas, 1974) (which formed one of the starting points of the popular recent paperback by Steven Johnson (Johnson, 2001)), to many research papers in the biology journals. We will discuss one in particular here because it provides a good example of the transposition of a model from physics to something that approximates an information system (Beekman et al., 2001).

Beekman et al. investigated experimentally the foraging behaviour of an ant colony as a function of its size. They found that below a threshold of about 600 ants the behaviour was disordered and inefficient, whereas above this threshold it was based on pheromone trails, well coordinated and efficient. This immediately makes us think of the question of critical mass in the DBE; namely, how many SMEs and how many software services are required to reach a self-sustaining digital business ecosystem? We don't know the answer to this question yet, but the ant colony provides some interesting clues.

The experiment investigated also how easy it was for the ants to find the source of food, a syrup feeder. This was controlled by varying the distance of the location of the syrup from the colony and was measured by counting the number of ants crossing a particular line 5cm from the syrup every 10min intervals. When the syrup is close to the colony, as the size of the colony increases the number of ants arriving at the syrup increases rapidly but continuously around the threshold size. This is equivalent to a critical or continuous phase transition. When the feeder is instead located far from the colony, the number of ants accessing the feeder jumps discontinuously at the threshold colony size, corresponding to when the ants begin to form pheromone trails. This is equivalent to a first-order or discontinuous phase transition. The experiment verified also the presence of hysteresis, meaning that, as the colony size was increasing, the colony remained disordered beyond the threshold, while, as it was decreasing, it retained ordered behaviour below the threshold before jumping down to the disordered state. The same type of hysteresis is observed at the phase transition between water and ice.

The reason these results are interesting and relevant to the DBE is that ants are not like molecules, they do not experience a feeling of attraction or repulsion from each other that is inversely proportional to the square of the distance between them. Thus their interactions do not involve minimisation of energy. Ants exchange information, and they do so indirectly through the pheromone trail. The stronger the smell the more likely an ant is to follow the trail, and the strength of the smell at a particular location is proportional to the number of ants that have walked on that same spot. Thus, a simple memory mechanism can have the same self-organising effect on a group of agents exchanging information as the minimisation of free energy that leads to the self-assembly of physical systems. Because the minimisation of free energy can lead to phase transitions, by analogy the same memory mechanism can serve to amplify local behaviour to the global scale, which is precisely what we observe in ant colonies. Paraphrasing Dr Lewis, an ant colony seen from 30 m in the air looks rather disturbingly like a large, black, *single* animal whose long tentacles slowly explore the ground around itself, that moves and migrates in response to floods and fire, that thinks!

This suggests that some of the adaptability of the DBE system and software on short time scales may be achieved through similar mechanisms, by endowing the software components (agents, services) with the ability to exchange information with each other and to recognise a “trail” of heavy usage through the SME business networks. This effect is not dissimilar to the spontaneous clustering of value chains and value nets into different business domains and industry sectors.

The definition of a space where trails can form and where software components can move, i.e. the ecosystem itself, is discussed in Chapter 4.

3. COMPLEXITY IN THE SOCIAL SCIENCES

This chapter reiterates and extends some of the points that have been made in other parts of this document regarding the links between biological, physical and social systems and contextualizes them within a discussion about the connections between the social sciences and the set of ideas that underlie complexity theory. More specifically, this chapter aims to open up a space for understanding the difficulties as well as possibilities of combining the two perspectives in an interdisciplinary framework that takes into account the conceptual depth and the methodological rigor of both these fields. For an integrated and intensively interdisciplinary project such as the DBE this is more than an epistemological issue. The success of the project depends significantly on the ability of its members to develop a common basis of understanding and a framework of action and intervention that cuts across science and social science and combines to a certain extent the strengths of the two.

This Chapter is structured as follows. First, we present some of the applications in the field of social sciences that clearly link themselves to complexity science. Central tenets of complexity theory such as the concepts of emergence and self-organisation have been appropriated during the last 20 years by researchers in the fields of economics, sociology, policy and regulation, historical studies, and organisation and management studies. This short review will serve as the basis for discussing in the second part of this chapter some of the problems that have emerged during the process of applying complexity theory concepts and methods in the context of the social sciences. There are two central, interrelated issues here. First, the extent to which complexity theory's concepts can be used metaphorically and secondly the ability of theories of social complexity to model social processes. These issues are connected with two more fundamental problems. The first one concerns the way that researchers from the two different communities communicate across disciplinary boundaries. The second one relates to the underlining commensurability of notions across the two fields and questions the suitability of complexity related ideas to model social processes.

The third part of the chapter refocuses the topic of the presentation on two lines of research that share many of the sensitivities of complexity theory with regard to evolutionary processes of change and emergence without explicitly taking on board its entire conceptual apparatus. The first line of research concerns the work of the economist Christopher Freeman whose work deals predominantly with the long run evolution of technologies and processes of incremental and radical change. The second line of research focuses on a perspective that has benefited enormously from advances in computing, statistics and mathematics, social network analysis. Although this perspective has been developing its own theoretical framework within the context of social and behavioural sciences for almost 50 years now, it has close methodological ties with complexity theory. The last part recapitulates and weaves together the arguments of the previous sections and highlights their relevance for the development of the DBE.

Complexity theory and the social sciences

To talk about complexity theory and the social sciences could mean that one assumes that the two fields are to a certain extent homogenous and coherent; that they are governed by certain underlining theoretical and methodological assumptions that give them an identifiable conceptual and epistemological basis. For different reasons, this cannot be further from the truth. Social sciences span a wide range of disciplines, from economics to sociology and psychology, each one with its own internal divisions and distinctive traditions. Complexity theory is an emergent field that is still searching for its centre of gravity. In their recent report on "Complexity and Behaviour in ICT Systems" Seth Bullock and Dave Cliff (2004) highlighted the little apparent consensus among the scientific community on definitions of complexity and emergence. This important issue is a considerable challenge DBE scientists have to grapple with.

In this section, however, we argue that the conceptual plurality that underlies the definition of complexity and emergence in the scientific community is mirrored and amplified in their appropriation within social sciences. Varying interpretations of complexity theory are being mobilized

here as part of distinctive programmes and agendas. Our aim is not to provide an exhaustive overview of the applications of complexity theory within the social sciences, but to illustrate some of the different ways that complexity related ideas have been mobilized in fields such as economics, sociology and organisation theory, highlighting along the way certain emerging patterns and trends.

One of the most notable applications of complexity ideas in economics has been made by Brian Arthur who has worked for a number of years in the Santa Fe Institute. Arthur has been predominantly preoccupied with the problem of increasing returns and network effects. Arthur argued that in contrast to conventional, equilibrium-oriented, economic approach, the more process oriented complexity approach allows one to study these phenomena as dynamic processes formed through random events and natural positive feedbacks. As he points out: “Complexity economics is not a temporary adjunct to static economic theory, but theory at a more general, out-of-equilibrium level. The approach is making itself felt in every area of economics: game theory⁶ [14], the theory of money and finance⁷ [15], learning in the economy⁸ [16], economic history⁹ [17], the evolution of trading networks¹⁰ [18], the stability of the economy¹¹ [19] and political economy¹² [20]. It is helping us understand phenomena such as market instability, the emergence of monopolies, and the persistence of poverty in ways that will help us deal with these. And it is bringing an awareness that policies succeed better by influencing the natural processes of formation of economic structures, than by forcing static outcomes”(Arthur, 1999:109).

In sociology the appropriation of complexity-related ideas has been frequently associated with attempts to revise and redefine the discipline’s scientific tasks and methods, a development that is often considered necessary in view of the effects and the processes of globalization. John Urry proposed that the modality of sociological concepts should be rethought along the lines of complexity and chaos theory (2000). Urry suggested that in order to survive sociology needs to become more expressive of movements and indeterminacies, two phenomena that are among the most important characteristics of complex systems. Urry’s proposition is indicative of a type of analysis pursued by a new generation of sociologists that lay emphasis on movements rather than on structures and positions and who aim to revise the concepts and the cognitive frames that are used to narrate individual experience and history (McLennan, 2001). Some examples that are akin to Urry’s work but do not include specific references to complexity theory include the work of Ulrick Beck and Zigmunt Bauman (2000;2000).

The search for underlying concepts on which new models of society can be built has also been part of a wider intellectual movement that has also appropriated in its agenda concepts of complexity theory, postmodernism. Among the intellectual characteristics of postmodernism are: a) the end of any belief in an overarching scientific rationality, b) the abandonment of empiricist theories of truth and c) an

⁶ 14. See K. Lindgren’s classic paper in *Artificial Life II*, C. G. Langton, C. Taylor, J. D. Farmer, S. Rasmussen, Eds. (Addison-Wesley, Reading, Mass., 1991). H. P. Young in *Econometrica* 61, 57, 1993; L. E. Blume, in *The Economy as an Evolving Complex System II*, 425, (Op. Cit.); B. A. Huberman, N. S. Glance, *Proc. Nat. Acad. Sci.* 90, 7716, (1993)

⁷ 15. R. Marimon, E. McGrattan, T. J. Sargent, *J. Econ. Dynamics and Control*, 14, 329, (1990); M. Shubik, in *The Economy as an Evolving Complex System II*, 263, (Op. Cit.); W. A. Brock, P. de Lima, in *Handbook of Statistics 12: Finance*, G. S. Maddala, H. Rao, H. Vinod, Eds. (North Holland, Amsterdam, 1995)

⁸ 16. T. J. Sargent, *Bounded Rationality in Macroeconomics*, (Clarendon Press, Oxford, 1993); D. A. Lane and R. Maxfield, in *The Economy as an Evolving Complex System II*, 169; W. B. Arthur, S. N. Durlauf, D. A. Lane, Eds. (Addison-Wesley, Reading, Mass., 1997); V. M. Darley and S. A. Kauffman, in *The Economy as an Evolving Complex System II*, 45, (Op. Cit.)

⁹ 17. D. C. North, in *The Economy as an Evolving Complex System II*, 223, (Op. Cit.)

¹⁰ 18. Y. M. Ioannides in *The Economy as an Evolving Complex System II*, 129, (Op. Cit.); A. P. Kirman in *The Economy as an Evolving Complex System II*, 491, (Op. Cit.); L. Tesfatsion in *The Economy as an Evolving Complex System II*, 533, (Op. Cit.)

¹¹ 19. P. Bak, K. Chen, J. Scheinkman, M. Woodford in *Ricerche Economiche* 47, 3, (1993); A. Leijonhufvud, in *The Economy as an Evolving Complex System II*, 321, (Op. Cit.)

¹² 20. R. Axelrod, *Am. Pol. Sci. Rev.* 80, 1095, (1986); K. Kollman, J. H. Miller, S. E. Page, in *The Economy as an Evolving Complex System II*, 461

emphasis on the fragmentation of experience and viewpoints. As Frank Webster points out: "Postmodernism is thoroughly opposed to each and every attempt to account for the world in these and similar ways, all of which seek to pin-point rationalities which govern change and behaviour. The presumption of Enlightenment thinkers that they may identify the underlying rationalities of action and change (which may well go unperceived by those living through such changes or acting in particular ways) is a focus of dissent for postmodernists." (2002:228). In addition to considering complexity theory itself as an indication of the decline of grand scientific narratives, postmodernists find in the processes underlying the emergence of complex phenomena such as the importance of local, initial conditions and bifurcation, a resonance with their ideas concerning the relativity, the locally specific character and the discontinuity of knowledge (Cilliers, 1998). As we are going to see in the next section, the often loose definition of scientific ideas and concepts that characterizes many postmodern attempts to apply aspects of the theory to the study of social processes has been often criticized by the scientific community.

An example of incorporation of complexity ideas in an agenda that lies at the other end of the modernist-postmodernist spectrum is provided by David Byrne's book "Complexity Theory and the Social Sciences" (1999). Byrne enrolls complexity-related ideas to the effort of revising the nature of the quantitative program of the social sciences. More specifically, Byrne attempts to set out the parameters of a "chaos/complexity"-founded programme in sociology and links this development to the collapse of linear and regression models in statistics and the necessity to rethink the design of traditional research tools, such as the social survey. Like Arthur, Byrne views complexity theory as part of a realist, positivist programme whose aim is not just to provide a better understanding of the social order, but also to assist in the development of policy interventions attuned to the nuances and interdependencies of complex social phenomena.

Complexity theory has been explicitly linked with policy making by several other researchers. Elliot and Kiel (1997) have explored the relevance of non-linear dynamics for social phenomena that are the objects of social policy. However, they also drew attention to the fact that headlong migration to the new tools offered by this approach can lead to the abandonment of the perfectly suitable, in certain cases, traditional statistical tools and techniques. Joseph Tainter has taken a more historically oriented approach and has examined the development of complexity in human societies with a view of improving the effectiveness of what he calls problem-solving institutions. Tainter considers complexity both as part of the problem and as part of the solution of our social and cultural evolution. He argues that: "Societies invest in complexity. In any system of problem solving, the initial strategies tend to be both effective (they work) and cost-effective (giving high returns per unit of investment). This is a normal economic process: simple, inexpensive solutions are adopted before more complex and expensive ones. So in the history of human efforts to feed ourselves, labor-sparing hunting and gathering gave way to more labor-intensive agriculture, which itself became more intensive as populations grew"(2000:6). Tainter's approach is indicative of a new type of historical analysis that aims to weave together accounts regarding the co-evolution of human and non-human actors and of human and non-human interaction events that cumulate through linear and non-linear, global and local processes. Some examples of this approach include De Landa's book *A Thousand Years of Nonlinear History* and Diamond's book *Guns, Germs and Steel*¹³. William McNeil sums up the agenda of this perspective by arguing that: "... it is time for historians to take note of the imperial role thus thrust upon their discipline by making a sustained effort to enlarge their views and explore the career of human kind as a whole, thus making human history an integral part of the scientific and historical worldview."(2001:1)

After economics, sociology, policy formulation, and historical studies, complexity notions and ideas have been also appropriated within the context of organisation and management studies. The view of

¹³ This kind of historical rationality is not entirely novel. French historian Fernand Braudel has been concerned with the importance of the material aspects of civilizations. Fernand Braudel, *Capitalism and Material Life:1400-1800* (London: Fontana/Collins, 1974), Fernand Braudel, *The Mediterranean and the Mediterranean World in the Age of Philip II* (London: Fontana, 1975).

organisations as complex adaptive systems has been widely adopted both by practitioners and researchers of the field, since it has been often mobilized within the context of many action research initiatives. The application of complexity theory in organisations is, therefore, predominantly guided by an agenda that aims to enable organisational innovation and change.

Donald McLean and Robert MacIntosh suggest that there are three major key insights from complexity theory that have informed their work on organisational transformation: “First is the notion that the structures, processes and procedures of an organisation can be thought of as being generated by a simple set of order generating rules...The second insight is that positive feedback applied to behaviour which is consistent with changes in these rules can drive an organisation from one state to another...The third insight is that in order for an organisation to become open enough to its environment to trigger a change in its order generating rules, it must experience far-from-equilibrium conditions.” (2003:151). Like a fractal pattern, the multiplicity of interpretations of complexity theory is also repeated within the context of organisational studies. Eve Mitleton-Kelly (2003) highlights ten generic principles of complex evolving systems that are applicable to social systems and organisations: self-organisation, emergence, connectivity, interdependence, feedback, far-from equilibrium, space of possibilities, co-evolution, historicity and time and path-dependence. Mitleton-Kelly provides several insights of how these concepts translate in terms of social processes and has in fact developed with her colleagues from the LSE Complexity Group an integrated methodology to study organisations as complex evolving systems (Mitleton-Kelly, 2003).

The appropriations of complexity theory in economics, policy, sociology, historical and organisational studies are often underlined by the desire to invent a new language for relating the macro to the micro, for talking about agency, structure and for the interdependence of such complex phenomena as globalization and localization. Within the context of the social sciences, complexity related ideas are often associated with endeavours to redefine modernity and to provide a more nuanced account of the evolution of societies and cultures that adopt a non-anthropocentric view of history. They aim to provide a new understanding of processes of historical and social change, by stressing the stochastic, non-linear character of events and interventions. This view implies that history and time are no longer seen as a separate, but as an integral dimension of events and processes. Complexity theory is often seen as the new science for the quest for patterns and discontinuities that aims to transcend the divisions and boundaries set by more traditional approaches.

The diversity of the conceptual landscape of complexity theory in the social sciences, that we attempted to paint with broad strokes in this section, creates the foundations for creative theoretical and empirical explorations across disciplinary boundaries. Both scientists and social scientists raise, however, significant objections about the theoretical and empirical value of many of these explorations. Part of the reason is that the mapping of concepts of complexity theory to social phenomena is far from straightforward. This problem is not only connected with the fundamental issue of the commensurability of natural and social sciences but also raises questions about the way that these two disciplines communicate.

The complex and the social: interdisciplinarity, communication and commensurability

In 1996 Alan Sokal, an American physicist, published an essay in the well-respected cultural studies journal *Social Text* titled “Transgressing the Boundaries: Towards a Transformative Hermeneutics of Quantum Gravity” (Sokal, 1996). The essay, which suggested that there are significant links between quantum mechanics and postmodernism, argued that the theory of quantum gravity has significant political implications and that physical reality is at bottom a social and linguist construct, an exaggeration of social constructivist ideas regarding the connections between the scientific, the social and the political. On the day of the essay’s publication Sokal announced in another journal that the article had in fact been a hoax and that he had employed scientific and mathematical concepts in ways that few scientists would take seriously. The debate that the hoax spawned, which became known as the “Sokal affair”, resulted in a series of discussions regarding relations between the natural and the social sciences and the standards of argumentation and scholarship in both disciplines.

Sokal's hoax illustrates in an extreme way some of the objections of the scientific community regarding the often loose appropriation of scientific notions and ideas within the social sciences and more specifically within cultural studies and postmodernism. Similar concerns have been expressed by scientists involved in complexity research (Bullock and Cliff, 2004), where the plurality and subjectivity of interpretations of self-organisation and emergence, that is frequently welcomed by many social scientists as a reflection of the inherent ambiguity and multidimensionality of these phenomena, is often considered as an issue that needs to be addressed by the elaboration of more objective and operational definitions. Debates within the scientific community regarding definitions of complexity are also echoed in discussions amongst social scientists regarding the way that these ideas are appropriated. Several researchers argue that these notions should not be used only metaphorically, in a heuristic, explorative way and that there should be an effort to operationalize them and apply them more systematically (Mitleton-Kelly, 1998). Although this approach appears to be more in synch with the natural sciences paradigm, the qualitative methodologies that are often employed within the context of social sciences are often considered inadequate for the exploration of mathematically demanding scientific ideas.

At the same time several objections regarding the appropriations of complexity theory are raised from the social sciences front. One of the main issues here is that theories of social complex systems are often poorly informed by existing sociological theories and ideas. Peter Stewart (2001) draws attention to the misuse of complexity theory within the field and argues that there are three major weakness of applications of complexity theory in the social sciences. First, complexity theory is often used as a meta-theory and is disconnected from relevant theories and debates in the social sciences fields (e.g. debates about power). This leads to overgeneralizations and oversimplifications. As he suggests: "My point is simply this: without recourse to the specificities of the relevant social fields, and without engagement in the full debates concerning the fields, a theorist is in a weak position to make major generalizations" (ibid: 332). Secondly, he argues that it might be misleading to consider that reality is entirely mathematical. He suggests that, if maintained rigidly in the field of social sciences the strong quantitative, mathematical dimension of many of the flavours of complexity theory can lead to positivism of numerical and spatial relationships that leaves out important dimensions of the social, such as the symbolic. His third objection is that theories of complex systems often regard society as a system, an idea which suggests that it is functional and coherent and lacks contradictions.

In his article 'Sociology's Complexity' Gregor McLennan (2001) connects the increasing popularity of complexity theory discourse with growing preferences of inter- and post-disciplinary over traditional sociological disciplinary approaches and argues that the complexity of sociological discourse deserves greater recognition. More specifically, although McLennan recognizes the potential relevance of complexity theory for sociological thinking, he seeks to confront suggestions that complexity theory reveals the traditional sociological mindset to be rigid and simplistic. Furthermore, McLennan points out that applications of complexity theory by sociologists like John Urry usually lack analytical precision and that concepts like emergence and self-organisation are "still in the mode of general ideas rather than modeled solutions" (ibid: 558).

Appropriations of complexity theory within the field of social sciences have been criticized then both by scientists and social scientists. However, not all explorations of society that mobilize complexity theories are conducted by social scientists. There is another side to the story and that concerns the modelling of social processes by scientists through the use of tools and methods closely associated with the development of complexity theory, such as computer simulations. Computer simulations have been widely employed as a means of exploring the properties of many kinds of complex systems. As Bullock and Cliff note: "For some complex systems researchers, simulations are a valuable source of additional empirical data on systems that are difficult or impossible to experiment on" (2004:25).

This is of particular relevance for the DBE, since this method is also employed for an examination of market efficiency and information (Maria Chli, ICL, S9-"Market Efficiency"). The application of

computer simulations in the social sciences has been regarded as a new way of doing science (Axelrod, 1997). However, it would be reasonable to assume that, as in the case of statistical modelling, the analytical and exploratory power of simulations would largely depend on the value of the theoretical models being employed. Simulations can shed light on current assumptions (Chli, *ibid*), but they can also perpetuate them or lead to misguided conclusions. Although the method has been employed with success to the modelling of cooperation and seems to offer a new way of exploring the relationship between structure and function (Lansing, 2002), it has yet to develop an effective and legitimate methodology, both from the mathematical and the social sciences standpoint.

A more important caveat regarding the more general effectiveness of non-linear computer simulations is given by Richardson who argues that: "Firstly, given the intractability of the emergent phenomena that occur within the simulation the analyst might not be able to provide any insights into the chain of events that led to a particular (modelled) system state. So, there is the possibility that the simulation itself might not offer any explanatory capability whatsoever even if the final state does indeed resemble 'real' systems' behavior(s). Potentially worse still is that, assuming that the Universe can best be considered a complex system, there are an enormous number of qualitatively different ways to model the same phenomena" (2002:3). Richardson attributes these difficulties to the phenomenon of equifinality, the idea that there are many different paths and trajectories that may lead to the same state.

The arguments presented so far are connected with two more fundamental issues. The first one concerns the way that researchers from the two different communities communicate across disciplinary boundaries. The second one relates to the underlining commensurability of notions across the two fields and questions the suitability of complexity related ideas to model social processes.

Bullock and Cliff have highlighted the institutional factors that make true interdisciplinary work, although popular in principle, hard to realize (2004). The structure of traditional knowledge institutions often results in the reification of disciplinary boundaries and, although promising in returns, the inherent risky nature of interdisciplinary research often translates in scarce funding opportunities. The fact that in spite of these difficulties complexity theory is becoming increasingly popular within social science means, however, that social scientists are finding other ways of coming in contact with complexity's basic notions and ideas.

Despite then barriers of communication set either through institutional factors or the frequently incompatible epistemologies of natural and social sciences, there appears to be some alternative channels through which complexity related ideas travel from one domain to the other. One of them is popular scientific accounts. Warren Smith and Matthew Higgins (2003) have examined their role in providing social scientists with a rich and promising source of materials and novel explanatory ideas. It appears, then, that in many cases the two fields communicate only indirectly, through a corpus of research and theory that is based on popular accounts regarding what complexity science is. Nigel Thrift (1999), on the other hand, suggests that the economy of concepts that have developed around complexity theory circulates the world through the three different but interrelated networks of science, business and New Age. Thrift highlights the iterative process of interpretation and dissemination by showing the different ways that each network has imported these concepts, processed them and re-exported them. In addition to its political dimensions, this observation draws attention to another important issue, which is that the relation between the natural and the social sciences is not necessarily one-way. By providing scientists with a rich source of data and a set of theories against which to test modelled social processes, social sciences can further their understanding of certain types of emerging phenomena and influence their definitions of complexity.

One of the key issues, however, is the extent to which there can be a constructive dialogue between researchers involved in complexity research from science and social science that respects the methodological rigour and theoretical depth of both these fields. In examining the frequently very different epistemologies, ontologies and methodologies that underlie these two different paradigms, we are bound to enter a discussion that has been going on, with different degrees of intensity, for

many, many years. In the philosophy of social sciences, theories, like complexity theory, that have included the social world in the natural order are usually termed 'naturalistic' (Hollis, 1994). This type of theory is frequently opposed to another group of social theories that are termed 'interpretative' or 'hermeneutic'. Such theories are underlined by the idea that the social world is inherently different from the natural order and must be understood from within rather than explained from without. Naturalistic theories assume that the same laws that govern physical reality will also underlie the social one, and that, although it may not be always straightforward, their respective mapping is therefore possible. Interpretative theories consider this impossible. What complicates things is that, as we have seen, appropriations of complexity theory within the social sciences do not fall neatly into one category or the other. Although most would be included by default in the 'naturalist' camp, some of them employ the ontologies and methodologies of the 'hermeneutic' paradigm. Such would be the case of John Urry's mobile sociology and of postmodernist interpretations. At the same time hermeneutic ideas regarding the constructed character of knowledge have found their way to the scientific realm by questioning assumptions regarding the picture of science as a realm independent of the enquiring mind.

This multiplicity of approaches, combined with the different agendas and amplified by the different standards of scholarship that are to be found in every field of study, create a series of barriers and links that connect and disconnect social sciences to and from complexity theory in ways that are often unproductive. In order to broaden and ensure a high level of dialogue between the two disciplines certain barriers need to be lifted and certain links need to be made explicit. This does not mean that the different communities of researchers need to settle epistemological and ontological questions that have occupied philosophers from the time of Enlightenment before they even start to consider cooperating. The development of compelling theories and models of social complexity that each community appears to require, however, will involve the willingness to engage with the notions of each field at a level that does not gloss over important conceptual differences and which involves a degree of understanding about the different research methodologies and traditions.

Social complexity in the social sciences: leveraging existing accounts and methods

In the previous section we argued that complexity researchers that wish to develop challenging theories of social complexity need to engage with social theory. As inherent properties of social processes, phenomena such as non-linearity, interdependence, and historicity have been highlighted and reflected upon by generations of social scientists. This section will suggest two certain specific lines of social science research that might add to the conceptual depth of the underlying complexity approach of the DBE project. These approaches share many of the sensitivities of complexity theory with regard to evolutionary processes of change and emergence without explicitly taking on board its entire conceptual apparatus. The first concerns the work of the economist Christopher Freeman, whose work deals predominantly with the long run evolution of technologies and processes of incremental and radical change. The second focuses on a perspective that has benefited enormously from advances in computing, statistics and mathematics, social network analysis. Both these perspectives are of particular relevance to the DBE, which not only aims to foster innovation and technical change and leverage the power of networks, but which is itself an object of an integrated Research and Development effort.

Freeman's work provides a good illustration of the vocabulary that social sciences have developed for discussing issues of technical and economic change, structure and agency. More specifically, Freeman has examined connections between technical change, institutional innovation and industrial growth primarily at a macroeconomic level. His work highlights the complex and untidy relationships that characterize these dynamics, and attempts to detect patterns and structures that underlie them. With his colleagues from the Science and Technology Policy Research Unit of Sussex university he has developed a taxonomy of technical change that distinguishes between the following five types of innovation: a) Incremental innovations, that occur more or less continuously and that are usually the outcome of inventions and improvements of those involved directly in the production process. Incremental innovations are crucial in the period that follows a radical innovation b) Radical innovations: these are discontinuous events that are distributed unevenly over sectors and in contrast

to incremental innovations they are usually the result of research and development activities c) new technological systems, the term refers to a cluster or constellations of innovations that are technically and economically interrelated d) changes of ‘techno-economic paradigm’: changes of this sort signify the most profound type of technological and economical transformation, since they alter the conditions of production, distribution and consumption for almost every branch of the economy (Freeman and Soete, 1987). Freeman believes that the transformations initiated by Information and Communication Technologies (ICTs) are of that scale and nature.

In examining relations between these types of change Freeman has argued that what matters more in terms of macro-economic effects is not the actual date of radical innovations, but the diffusion process that helps incorporate them to the realities of production, the everyday life of consumption and the actions that are being taken to ensure their match with an institutional framework that is usually resistant to change. As Freeman notes: “So far we have tried to insist that what matters from the standpoint of large-scale economic fluctuations is not so much the date of a particular basic innovation as a constellation of circumstances favourable to the exceptionally rapid growth of one or more new industries, each involving the combination of related inventions, innovation and economic and social changes. We would now insist furthermore on the vital importance of Shumpeter’s point about managerial and organisational innovations. These may often be as important as the technical changes for the growth of an industry or technology.”(Freeman, et al., 1982:70).

One of Freeman’s central propositions is that each paradigm change brings about considerable problems of mismatch and fragmentation with respect to skills, management styles and capital stock that give rise to problems of unemployment and adjustment. Until these issues are addressed, and the profound potential of the new technologies is realized through a series of process innovations and institutional changes, the full benefits of radical innovation and paradigmatic changes cannot take effect. In examining these structural crises Freeman makes it a point to note that he does not make any assumptions about their fixed periodicity or statistical regularity, but he stresses the fact that each paradigmatic shift is characterized by unique qualitative aspects that call for different investment strategies, government policies and institutional changes.

Freeman’s work is particularly relevant in considering DBE both as part of a wider complex of policies and interventions as well as an evolving technological framework and platform of innovation. Some of the insights that his work offers us are that the effects of changes that the DBE wishes to implement would depend on a number of associated institutional and regulatory policies. Their degree of fragmentation at a European, national, regional and even industrial level poses significant challenges to DBE’s deployment. At the same time, if we consider the DBE itself as a radical innovation, we should expect, according to Freeman, its success to depend on a cluster of incremental innovations. Although some of these are expected to take effect during the life-time of the project, through the take-up of the DBE platform by early SME adopters, it is more reasonable to expect that most of them would happen after the expiry of the project. This raises some interesting issues regarding the platform’s long-term sustainability. On the other hand, the DBE aims to work, at least at the early stages of its development, in concert with existing technologies and platforms, to function, in other words, to some extent as an incremental or process innovation. The balance between the incremental and the radical is yet another question that needs to be addressed within the context of the project.

Chapter 4 of this report examines the network architecture from an information theory perspective and makes the case that such a distributed system could also demonstrate self-organising properties. In this section we would like to argue that conceptualization of DBE as a networked information system involving different kinds of digital species interacting at different levels of abstraction could be productively complemented by a social networks perspective. It should be noted that issues of social networking are already being taken into account in several aspects of the project (e.g. through the selection of opportunity spaces and the elaboration of an SME recruiting strategy). These efforts could, however, benefit from a more systematic and theoretically informed approach that complements the information-theory and biologically-based aspects of the DBE as a network.

Social network analysis (SNA) is closely related to, but also quite distinct from the recent surge of interest in all kinds of networks (Barabási, 2002). More specifically, SNA shares many of the analytical tools of approaches that argue in favour of the pervasiveness of the network paradigm, but it has developed its own set of conceptual tools that apply to the study of social networks. What differentiates Social Network Analysis from other types of network analysis is, therefore, the special patterns and constraints that seem to underlie the development of social networks. In addition, the method, that has a long history in social and behavioural sciences, has emerged as an integrated scientific field concerned with the structural analysis of social interaction over the last sixty years. Although it has benefited enormously from advances in computer science, statistics and mathematics, the approach has its roots in anthropological studies that have first employed elementary network techniques in describing and studying relations of kinship. Gradually the perspective evolved and it currently incorporates theories, methods and applications. The perspective relies on three major mathematical foundations: graph theory, statistical and probability theory and algebraic models. As Stanley Wasserman and Katherine Faust point out: “The phrase “social network” refers to the set of actors and ties among them. The network analyst would seek to model these relationships to depict the structure of a group. One could then study the impact of this structure on the functioning of the group and/or the influence of this structure on individuals within the group” (1994:9).

Some studies that might be of particular relevance to the DBE concern dynamics of inter-organisational cooperation (Chung, 1996; Paulson, 1985), the flows and the bottlenecks of dissemination of information among social networks (Yamaguchi, 2002), and the structure of cooperation and trust (Buskens, 1993). As we are going to see in a following section certain Social Network theories, such as the “Small World” hypothesis (McCue, 2002; Watts, 1999), that suggests that one is connected to a new acquaintance by a surprisingly short chain, are already being taken into account. If this hypothesis holds equally for organisations, it could be one of the guiding factors of the evolution of the DBE business network and as such it should be taken into account in its regulatory and technical design.

If we weave together and follow through the implications of the previous arguments, we could conclude that SNA could be of benefit to the DBE in two ways. First, DBE aims to leverage existing networks of cooperation among users, software producers and regional catalysts. An understanding of the dynamics of these networks and of the structural constraints in scaling them could be of real value to the deployment and the sustainability of the initiative. However, as we know the DBE does not aim only to tap into existing industry and business networks; it also aims to extend and cultivate them by supporting a set of viable and productive relationships that will be of particular benefit to SMEs and regional economies. What is of particular interest here are the factors of social network dynamics that should be integrated in the regulatory and technical design of the DBE and the dynamics that will develop between these constraints and freedoms and the inherent dynamics of the networks of actors that will be brought to interact within its framework. SNA could provide a theoretical and methodological framework for mapping this complex interaction.

Whereas Freeman’s ideas present us with a body of theoretical and empirical work that can enrich our understanding of the co-evolution of technical, economic and institutional change and DBE’s position within these dynamics, SNA presents us with a set of concepts and techniques that can complement the informational and statistical aspects of the DBE topology. In the case of the first perspective the links are predominantly conceptual and can highlight the project’s multiple dependencies with institutional and policy aspects as well as draw attention to some of the dynamics of its own technical evolution as a platform for innovation. In the case of the second approach, the links to the DBE project are both theoretical and methodological.

Conclusions

This chapter aimed to provide the basis of understanding some of the difficulties of applying complexity theory in the social domain and to provide some suggestions as to the way the often conflicting paradigms of the complexity theory and social sciences could be productively combined within the context of the DBE.

In the first section we drew attention to the different applications of complexity theory within social sciences. Complexity related notions and ideas are being mobilized by different agendas and made part of widely disparate social science programmes that often stand in opposition with one another. All these different appropriations, however, seem to be characterized by a common desire to develop a new language and a new perspective for approaching old and new problems.

Scientists trained in natural sciences, mathematics and computer science are frequently apprehensive of the way scientific concepts are appropriated by researchers in the social sciences. One of their main concerns is that complexity theory is adopted primarily as a metaphor, a discursive strategy that fails to further our understanding of the phenomena under investigation and of complexity theory itself. Social scientists on the other hand are often concerned that the application of complexity ideas in social processes is frequently superficial and lacks the conceptual depth and richness that has been reached by corresponding ideas in social theory. The second section highlighted then the epistemological, conceptual, institutional and communicative barriers that appear to discourage interdisciplinary cooperation and at the same time pointed to some of the factors that contribute to the popularity and the wide diffusion of complexity-related ideas.

What are then the possibilities for deeply engaging interdisciplinary cooperation given:

- a) that it is almost impossible for social scientists to master the science behind complexity theory and of scientists to become familiar with the complexity of social theory?
- b) that it is unrealistic to expect that decades of epistemological debates could be settled in matter of a few years?

In this chapter we suggested that what seems to be important is the willingness to engage with the notions of each field at a level that does not gloss over important conceptual differences and which involves a degree of understanding about the different research methodologies and traditions. Intensive interdisciplinary projects like the DBE offer scientists from both fields a unique opportunity to cooperate on the basis of concrete and challenging problems that require the mobilization of the analytic power of both paradigms. The final section suggested two lines of social research that can add to the conceptual depth to the underlining complexity theory approach of the DBE project and that can help elaborate the formulation of actual solutions.

4. INFORMATION AND NETWORKS: TOWARD AN ARCHITECTURE OF SELF-ORGANISATION

Compared to physical systems, information systems are sparse, for two reasons. Firstly, the abstract nature of information systems makes the definition of a global metric space an unnatural imposition of structure that restricts the full range of possible relationships between digital entities. Depending on the level of abstraction involved, the proximity relationships between the same set of entities could be very different; for example, two nodes that are a few hops away at the physical layer may be many hyperlinks away at the application layer, and may be geographically located in different countries. The natural structure underpinning a discrete set of digital entities is a topological space, i.e. a network or graph. In such a more abstract space proximity depends on connectivity which, in turn, can depend on rules or criteria beyond spatial distance. Although the dynamics of the network topology is currently a topic of intense research in physics, computer science, and in complexity studies, in this report we are more interested in dynamical processes whose time scales are much faster than the topological evolution of the network.

Secondly, unlike physical systems, the interactions between digital components are not determined or influenced by energy exchange or minimisation. Although a small amount of energy is required to carry the bits, what changes the state of a digital system is the information carried by the interaction, not its energy. This was not case for the old analogue computers, but digital computer achieved an almost perfect decoupling between the information they hold and the underlying physical processes that enable them to hold it. Energy simply does not matter in the abstract space of information systems. To be fair, one could define a potential energy as a function of an arbitrary set of variables applicable to a particular domain, as is routinely done, in fact, in numerical simulations of dynamical systems. Such a definition, however, would not change the true nature of such an entity as information—albeit information designed to behave like energy within a formal system. Although in principle one could build a huge simulation of a whole ecosystem based on physical law, this would be hugely inefficient. Our principal purpose is to discover how information systems might exhibit analogous dynamical behaviour based solely on their own intrinsic properties.

Given the absence of a distance and of a potential energy of interaction, it becomes difficult to imagine how spontaneous processes based on the “fall” toward equilibrium, as discussed in Chapter 2, could take place in information systems. In other words, the spontaneous creation of order in information systems seems an even more formidable challenge than the spontaneous creation of order in physical systems. As in physics, we are confronted with entropy. As originally discussed by Shannon (1948) and Wiener (1948) in connection with communication networks, information is often lost, but it can never be gained. And in addition there are no attractive forces to trigger spontaneous aggregation.

And yet it does not take much insight to notice that socio-economic systems, which are made of heterogeneous “components” at many different levels of abstraction, have been assembling over long time scales in ways which are strongly evocative of spontaneous biological processes (Diamond, 1998). This suggests that self-organisation is not mediated solely by physical interactions embedded in metric spaces, it appears to transcend the physical. As discussed in Chapter 3, the most extreme philosophical position that evokes this point of view, although it predates complexity science or any attempts to apply it to the social sciences, is found in postmodernism. Roland Barthes, one of its best-known apologists, argues that objective reality does not exist, only the language that we use to interpret and describe our perception of reality exists (Webster, 2002). This view is clearly recursive, so that we soon realise that language does not interpret reality, but only other assemblages of signs and perceptions, which interpret another layer behind that, and so on ad infinitum. Even if we do not espouse the radical relativism of postmodernism, the plausible suggestion remains that our conception of reality is dependent on how we make sense of our environment. This provides a useful common ground between socio-economic systems, biological systems, and information systems.

We refer to the totality of symbols, signs and languages, visual, aural, olfactory or tactile, with which we codify our environment as “information”. Much work has been done since Shannon to define Information in a manner that could be used by communication and digital technologies while remaining compatible with the codified content (“signal”) they mediate for human consumption. We summarise the main concept here in order to introduce a discussion of network statistics that leads to a potential non-functional requirement for the DBE System.

Shannon’s information entropy and physical complexity (Briscoe, 2004)

As discussed in Chapter 2, Boltzmann’s entropy is a (logarithmic) measure of the number of states accessible (with equal probability) to a system:

$$S = k_B \log(g) \quad (4.1)$$

where k_B is Boltzmann’s constant, with units of energy/kelvin, and g is the multiplicity function. Shannon’s entropy is also a logarithmic measure. Given M distinct symbols and N symbols in total, if $N = M$ and each symbol appears only once, then the probability of each symbol is uniform and it is given by

$$p_i = \frac{1}{M} \quad (4.2)$$

In this case, for a binary alphabet, Shannon’s “information content” is given by the simple function (MacKay, 2003)

$$h = \log_2(M) = -\log_2 p_i \quad (4.3)$$

This is none other than the word size, in bits, needed to encode a set of M binary numbers (for $M = 512$, $h = 9$). For a source that outputs an infinite sequence of bits to communicate a finite set of symbols M , Shannon generalised the above function to express an average word length. The derivation is easier to see for a large but finite total number of symbols N ($N \gg M$). N_i is the number of occurrences of the symbol M_i , and therefore acts as a weight in the average:

$$H = \frac{\sum_{i=1}^M N_i [-\log_2(p_i)]}{\sum_{i=1}^M N_i} = -\sum_{i=1}^M \frac{N_i}{N} [\log_2(p_i)] = -\sum_{i=1}^M p_i \log_2(p_i) \quad (4.4)$$

What prompted Shannon to identify the average word length with the concept of entropy? Since physical entropy is a logarithmic measure of the number of accessible quantum states, Shannon made an analogy between a quantum state of a physical system and a symbol. Compatibly with the fundamental assumption of statistical mechanics that each quantum state (at the same energy) is equally probable, Shannon is implying that each symbol is generated with equal probability by the source. Although this is not strictly true (no reason to take the average if it were), it is true for the simpler situation leading to Eq (4.3), and the conceptual link is strong enough to justify keeping the name for the more general case.

As discussed in DBE Report D6.1 (ICL), this is explained by Adami et al (2000) by calling h and H *potential* information, which refers to the maximum number of symbols that can be coded by a word of length H (2^H , for a binary system). Adami then explains that, “...any arrangement of symbols [...] acquires the status of information only when its correspondence, or correlation, to other physical objects is revealed.” In other words, information must be grounded. We construct reality by selecting and retaining a subset of otherwise arbitrary symbols. This results in the opportunity to define information in relative terms, which is what biologists needed to do in order to provide a basis for

distinguishing between the DNA that corresponds to the organism (phenotype) and the redundant genes. Because those genes that represent a successful adaptation to an environment are selected and retained in the genome, we can discern the difference between information and noise by observing which gene patterns are replicated in a population, without any knowledge of the environment itself. The part of the sequence that codes for the environment is called “physical complexity”. Fig 4.1 summarises these concepts.

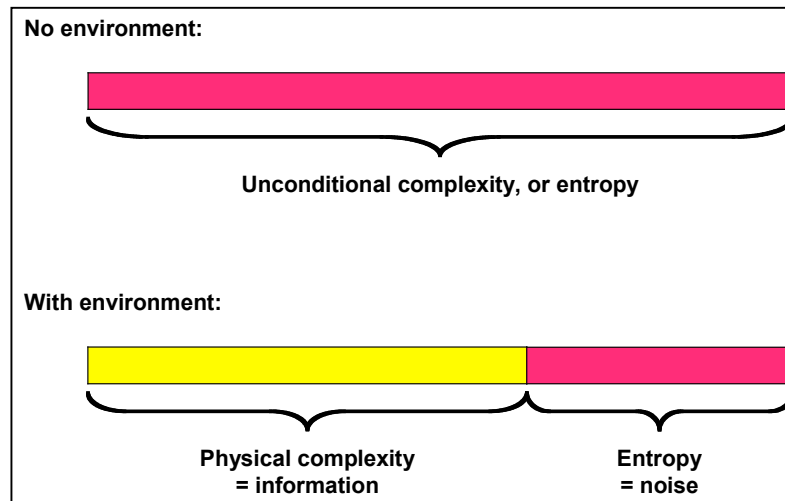


Fig. 4.1 Information defined relative to an environment (not to scale)

“The network is the environment”

In the Digital Business Ecosystem we would like the digital species to assemble and learn from their interactions with their environment, which is composed of other software components and of human users representing SMEs. We are therefore faced with the problem of defining the environment of each DBE “organism”. Such definition should ideally apply to different types of organisms living at different levels of abstraction, from the metering infrastructural services for accounting, to the DBE Business Services(Ferronato, 2004), reproduced in Fig 4.2 for convenience, to the SMEs themselves as holders of valuable business models. A graph or topological network approximates the idea of environment sufficiently flexibly to cover all these cases, regardless of the level of abstraction.

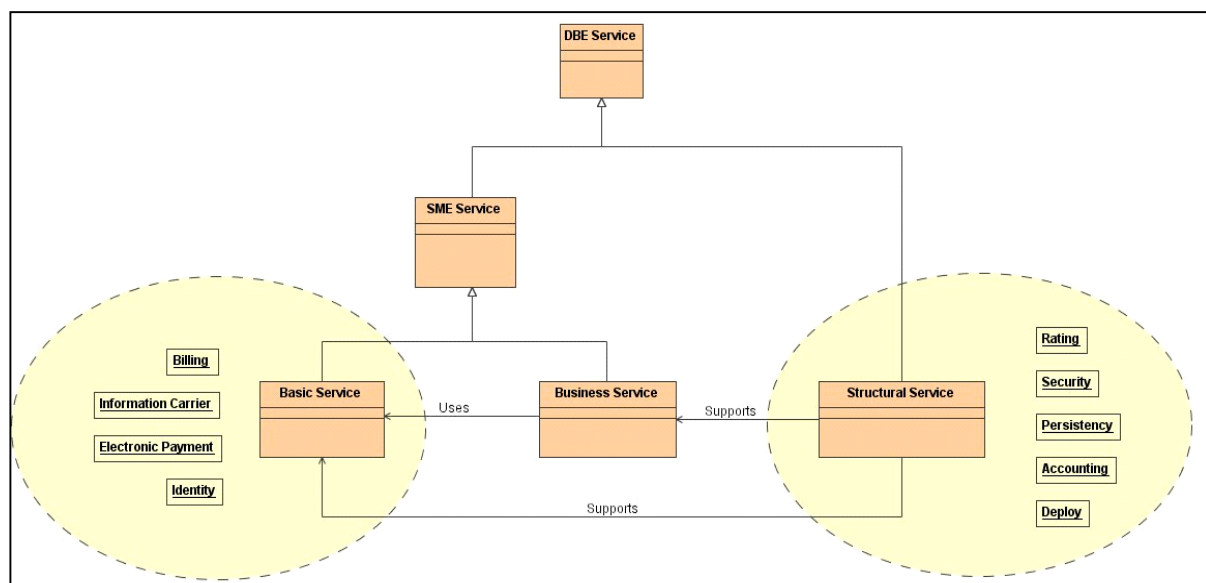


Fig. 4.2 Classification of DBE services (Ferronato, 2004)

Having defined information in relative terms, clearly the further construction and elaboration of the information into more complex patterns can only be meaningful if it is recognised as such by the environment, i.e. if it serves some function. In principle, complexity and function are concepts as arbitrary as information, but in our case we need to match specific requirements and usage scenarios arising from the SMEs in their various business domains. These provide the objectivity we need, which also serves as the selection pressure mediated by the network.

In Chapter 2 we built up the characteristics of self-organising systems from the physical point of view, generalising toward more abstract systems and ending up with the issue of scaling. In this chapter we have instead begun from a global abstract view of topological spaces and information theory, working toward the network as the unavoidable architecture of a distributed digital ecosystem. We know a few things that make or may contribute to making such a distributed system also self-organising. As discussed in Chapter 2,

The spontaneous behaviour of the system originates from its users. We can call it spontaneous as long as the users are using the system as a medium for pursuing their own ends, and centralised control of the system is not one of them. It is true that in an information system much of the activity is asynchronous, meaning delayed in time as the execution of a computer program can be. So even if the users do not explicitly set up automatic control measures (that would defeat the possibility of self-organisation), a certain level of random activity generated internally will originate from asynchronous commands. But over long time scales even programmed behaviour is reviewed and updated by human users based on changing human requirements.

Memory can be implemented in different ways, from a variable network topology, to the analogues of pheromones, to the weights of a neural network, all the way to human symbols. Something analogous to pheromones for migrating services¹⁴ is worth investigating further at the simulation level and as the actual DBE comes online. For example, Gerard Briscoe at ICL is investigating the migration history of services since it leads to links between habitats, the formation of new connections between habitats, clustering, and community formation in the habitats.

The non-linear coupling and feedback are much more difficult to quantify, model and predict, but given that as systems scale in size their statistical properties tend to become proportionally more important, we postulate that the statistics of cycles may provide the right intersection between non-linear effects and the inescapable law of large numbers.

In other words, the fourth aspect, scale effects, may be linked to the third, non-linear coupling and feedback. In this report we will briefly discuss cycle statistics and we will postpone the structural and functional hierarchies mentioned in Chapter 2 to the next deliverable.

Cycles in fully connected and scale-free networks

Several partners in the DBE are working on networks: UBham, ICL, and Sun, principally. The analysis reported here was inspired by interaction with all of these partners. We will look first at fully-connected networks since they provide an easier case with an exact solution for the cycle statistics. A network or graph is a set of nodes N some or all of which are connected to some or all of the other nodes. The degree of a node equals the number of links it has to other nodes. If all nodes are connected to all other nodes it is called a fully-connected network, whose nodes are all of degree $(N - 1)$. Random networks are constructed in a sequence of steps at each of which any two nodes are chosen randomly and joined, without any regard to existing links, until some criterion is met (such as average connectedness, or percolation across between the two farthest nodes). As shown in Fig. 4.3, the degree distribution of random networks has a peak somewhere between 0 and N ; the degree at which the peak forms is the “scale” of the network. Barabasi and co-workers (Albert and Barabasi, 2002) discovered that many networks (physical as well as abstract) have a scale-free degree distribution. These are networks where the formation of new links is influenced by the presence of previous links and is in fact more likely to occur on nodes that already have many links. The Internet is a famous example of a scale-free network. Finally, a cycle is a closed and non-intersecting (also called “self-avoiding”) loop of connections between nodes. The size of the cycle is simply the number

¹⁴ The *implementations* of the services do not actually need to move. What may migrate along the business networks are the Service Manifests (see Chapter 6), for which the statements above would still apply.

of links. Since a cycle must be a closed loop, the smallest cycle is made of two nodes and one link, the only case where a link is traversed twice. The largest cycle possible is of size N links.




Fig. 4.3 Degree distribution for different network topologies

The study of loops in graph theory has not received as much attention as other topological properties, probably because in communication networks loops are carefully avoided as they can give rise to undesired and unpredictable behaviour. This fact alone should be sufficient to tell us that there is something interesting about loops, since we know that in engineering traditionally we tend to avoid or suppress non-linear effects, while the dynamics of self-organising systems appears to be heavily dependent on non-linear behaviour and feedback. Our study of these effects is just at the beginning, but a property that seems worth investigating is the distribution of non-intersecting loops of size m in a fully-connected network of size N . The answer turns out to be remarkably simple:

$$M = \binom{N}{m} = \frac{N!}{(N-m)!m!}, \quad (4.5)$$

since choosing m independent links is equivalent to choosing m nodes, and choosing m nodes out of N possible nodes is indeed the function “ N choose m ”.¹⁵ In other words, the cycle distribution is given by the binomial distribution, just like the entropy of a binary system discussed in Chapter 2. As a consequence, the most probable loop size in a fully connected network of N nodes is simply

$$m^* = N/2, \quad (4.6)$$

as shown in Fig. 4.4 for $N = 10, 50$ and 100 . The case $N = 10$ is mathematically identical to the example of 10 particles in a box discussed in Chapter 2. This figure is included here to highlight how quickly the probability of loops of size m^* increases and how sharp this maximum becomes with increasing network size.

This measure for the cycle distribution does not distinguish between permutations of the set of nodes that define a cycle. All such permutations are considered equivalent to the same cycle. If we count also the permutations then the number of possible cycles corresponding to m^* is larger:

$$M_{\text{max with permutations}} = \frac{N!}{(N/2)!} \quad (4.7)$$

¹⁵ Thanks to Jon Rowe (UBham) for pointing this out.




Fig. 4.4 Probability distribution of different cycle sizes for three values of N

Let us assess the implications. Another way to describe this fact is by measuring the average return time (time = number of steps) of a packet to its starting node. If at every step it is able to choose randomly which node to jump to among the nodes it is connected to, it will take it N steps on average to return. This is easy to see if we recall that in a fully connected network each node is connected to every other node, so at every step the packet can choose from $N-1$ nodes, with equal probability. If, instead, it keeps track of which nodes it has visited and never jumps back onto one of them unless it is the starting node, thereby traversing a self-avoiding loop, then the above result says that it will take $N/2$ steps on average to return. The implication therefore can be expressed as a question:

For what reason might we want to constrain a packet to follow a random self-avoiding loop?

The appeal of this idea, when we first stumbled on it, was the binomial distribution and its obvious echo to the maximisation of entropy. However, maximum entropy by itself is not sufficient, which can be seen as follows. If we relax the constraint we are back at the case of N steps, on average, to complete a loop. The number of ways to choose N things from a set of N is $N!$, but if the packet can additionally visit each node as many times as it wishes then the number of possible loops (still of size N , on average) is even larger. In both cases these are numbers much larger than Eq. (4.5). Therefore, probability alone does not support the formation of self-avoiding loops.

As pointed out by Jon Rowe, just as we can constrain a packet to follow a self-avoiding loop, we could enforce different rules, leading to loops of different sizes. With each such rule we are adding some energy/work to the system, driving it toward more “unlikely” behaviour and away from a truly self-organising system. Looking at how evolution has addressed the problem of feedback in networks may shed some light on whether any such constraints are common, how they are enforced, and why they may be useful. At this point we are postulating the possible relevance of self-avoiding cycles solely based on the obvious importance of feedback in non-linear systems and on the strength of the binomial distribution. This point is clarified further by the following observations.

The Internet is not a fully-connected network, and neither are most other networks investigated by Barabasi. For a scale-free network a statistical study has been performed by Rozenheld et al (2004), who found that also in this case the distribution of cycles is nearly a Gaussian that tends to a delta function for large N . If the scale-free network is characterised by a degree distribution of this type,

$$P(k) \propto k^{-\lambda}, \quad (4.8)$$

where λ varies between 2 and 4, then for the deterministic scale-free network studied in that paper the most probable loop size is

$$m^* = N^{\frac{1}{\lambda-1}} \quad (4.9)$$

For $\lambda = 3$ this gives the very interesting result that the most probable loop size is $m^* = \sqrt{N}$. The reason this is interesting is because, whereas in the case of a fully-connected network of 1000 nodes, for instance, the most probable self-avoiding loop size is 500, for a scale-free network the size is on the order of 30. We may have a difficult time imagining a closed loop of information of 500 nodes travelling around in a network of 1000 nodes, no matter what the laws of probability may be telling us. But if the same Gaussian behaviour applies to a scale-free network and the most probable loop is of size 30, this suddenly starts to sound feasible.

Unfortunately we do not have at this time an estimate of the average return time for a random, unconstrained walk on a scale-free network of N nodes, but hopefully other partners (UBham, ICL, Sun, STU) will investigate this question in the near future. In the meantime we can postulate the following hypothesis. If we accept that closed loops mediate coupling between the inter-dependent nodes of a network, then in the scale-free case self-avoiding loops could be justified as supporting the right balance between maximum tightness of the coupling (small loop size, short return times) and maximum distribution of information (no repeated nodes). In other words, we postulate causal links between maximum entropy, optimum non-linear coupling (i.e. where “optimum” could mean “supporting meta-stability”), self-avoiding cycles, and scale-free topology. What remains to be seen is whether scale-free networks in natural systems and/or simulations support this claim.

We will also look at small world networks.¹⁶ These are more difficult to characterise because it only takes a few additional links to turn various kinds of networks into a small-world network. We will need to narrow down the characteristics of what we expect the DBE network topology to be, and we will also apply Social Network Analysis to strengthen our study of networks from the physical and mathematical point of view with the social science and anthropology viewpoints.

In conclusion, we suspect that the self-organisation of complex systems of many components or actors relies in part on the formation of closed information loops that provide both a mixing and a feedback function. The mixing acts to distribute information throughout the network, not unlike how vortices distribute momentum in a turbulent flow, whilst feedback provides coupling between the system’s actors. Such coupling may be important to support meta-stability, in the sense we defined, while preserving the maximum flexibility and dynamic response characteristic of non-linear systems.

The question of what kind of information, or data, would be flowing in such networks cannot be deferred further. A possible hypothesis is beginning to take shape and will be discussed in the next two chapters in the form of a distributed intelligence in the Evolutionary Environment (ICL). For the self-organisation of global behaviour arising from the individual interactions in such a topological space with memory we may need to postulate a function analogous to the sugar syrup for the ant colony. What is the incentive, the food, the energy that flows through the system and keeps it alive? A link to economics seems unavoidable. Is there a global cost function that we need to minimise? Or perhaps we need to maximise the individual utilities of the actors (DBE Services or SMEs)? These are opposite points of view on optimisation (investigated by Maria Petrou of UniS) that may turn out to be complementary both to each other as well as to Intelligence and Evolution.

¹⁶ This was suggested by Giulio Marcon, STU.

The issue of time scales

We have said that in information systems there is nothing equivalent to energy. This appears to be an accurate statement as regards interaction potential energy, although we saw that memory can serve as an analogous higher-level property that can lead to similar global behaviour. There is, however, another effect which has something in common with the effect of the mean energy or temperature of a physical system. As we saw in Chapter 2, as the temperature of a statistical system of interacting components (such as magnetic spin system) increases, the intensity of the random fluctuations increases also, leading to more spins flipping and a state of disorder. Conversely, as the temperature decreases the spins remain trapped in the configuration of minimum interaction potential energy. The high temperature limit is structureless and completely random behaviour, while the opposite limit is a frozen minimum-energy configuration of spins. In more complex systems, such as bacterial colonies, increasing the temperature accelerates their reproductive rate, which is why we refrigerate food. Because we can't define energy for an information system we can't define temperature either.¹⁷ However, there is an effect in information systems that is analogous to an increase in energy or temperature: an increase in CPU frequency or computing power.¹⁸ Regardless of what we want to base the network topology on (geography, business domains, etc) there will be areas of the network where the mean computing power is higher than in others. These will be able to process information more quickly. Likewise, and probably more importantly, areas where the coupling between the computers and the SME users is tighter will lead to a greater rate of e-business transactions. These areas are in some sense analogous to "hot spots" in the ecosystem. Since there is no interaction energy, such hot spots are not correlated with greater disorder, as would be the case in physical systems. Finally, from the point of view of network loops (either at the physical layer or at the application level), these spots are analogous to pumps that process and "energise" the information flow, pushing it along the loops as water in pipes.

From a global viewpoint the effect of a non-uniform distribution of CPU power across the DBE leads to different characteristic time scales for the dynamics. These may influence the formation of clusters and, if visualised on a geographical map, may provide useful information for the development of policies to address ICT adoption and the digital divide.

Conclusions

In current network technologies no time or effort is spent worrying about cycles. In fact, any time spent thinking about them aims to *remove and avoid* them. Thus cycles are not explicitly set up, and there currently isn't any underlying process by which this might be done. Our point of view is different. We are thinking about the spontaneous formation of structure as we observe in biological systems. For the DBE to be truly self-sustaining and self-optimising, we must understand how spontaneous, i.e. not directly planned, structure formation can take place in information systems.

We are in effect postulating that in large self-organising systems of interacting components we are likely to find, empirically, that cycles of information flow do in fact form, and that we have a way of predicting the predominant size of such cycles based on the cycle statistics of different network topologies. In this chapter and in Chapter 2 we have circumstantially supported this rather bold claim with a few apparently unrelated observations:

- 1- At the present time the only mechanism we *actually* understand that can achieve self-organisation is evolution and evolutionary algorithms. Evolution seems slow (in a relative sense) but, mainly, local; i.e. it is difficult to amplify the positive local effects of evolutionary algorithms to a complex system. We could do it through central control and explicit instructions, but the whole point of the DBE is to move away from central control. The intrinsic dynamics of the system are supposed to propagate the necessary information where it's needed, in the most efficient way possible and without direct control. The intrinsic dynamics are driven by the users interacting with and using the system.

¹⁷ Temperature is inversely proportional to the rate of increase of entropy with respect to unit energy increase. Thus, at low temperature the number of accessible quantum states increases very quickly with a unit addition of energy, whereas at high temperature the converse is true.

¹⁸ This idea was suggested by Gerard Briscoe, ICL.

- 2- Ant colonies provide a good starting point. They are closer to an information system than a cell is, and they self-organise through a memory mechanism. What we are suggesting is that, if we understand how ant colonies self-organise, we can try to relax the top-down control of software systems and, by applying appropriate rules and memory mechanisms, we can observe the system configuring itself in response to external stimuli to achieve a certain goal, such as foraging (achieving a certain goal could be modelled by the minimisation of a global cost function, a topic investigated in WP12 by Maria Petrou of UniS). Ant colonies are not however sufficient.
- 3- Report D6.1 clarifies the relative nature of information in ecosystems and the importance of the environment, which we identify with the network.
- 4- Self-organisation is a non-linear process, and non-linear systems are such mainly due to the coupling terms between the degrees of freedom that comprise them. This coupling manifests itself through feedback loops.

Based on these four points we still don't have enough information to decide whether cycles of a predominant size m^* will form purely as an effect without further consequences, or whether the formation of such cycles will then have a feedback effect on the self-organisation seen as relying on these four points, or even on the network topology. In other words, we do not know whether the characteristic cycle size is going to be a fifth factor needed to support self-organisation in distributed complex systems.

In the next chapter we begin to explore how the self-organising run-time behaviour of an information system might be incorporated into language, and a domain modelling language in particular.

5. SCIENTIFIC MODELS AND UML

In this chapter we discuss the role of DNA as a code that balances the concept of instructions (top-down) with self-organisation (bottom-up), since only in a specific environment does the code become active and acquire meaning. We take DNA as a metaphor for finding a similar balance between Model Driven Architecture™ (top-down) and the Adaptive Object Model and evolutionary algorithms (bottom-up) in the development and optimisation of software systems. Another way to say this is that whereas Chapters 2 and 4 are mainly concerned with understanding the self-organisation of run-time software systems, this chapter is mainly concerned with the program or model with which the run-time behaviour is specified. This leads to a discussion of the similarities and differences between scientific models and language. We close with some general considerations on the Unified Modelling Language (UML) developed by the Object Management Group (OMG).¹⁹

Codifying self-organisation

The motivation for pursuing Goal 2, understanding how self-organisation could take place in software, and therefore also Goal 3, codifying self-organisation, stems from the observation that, whereas evolution has taken billions of years to evolve complex organisms, such organisms live through self-organisation mechanisms *now*. Self-organisation is a complex mechanism that has required millions of years to develop, but that now works in real time (Fig 5.1). Thus, if we understand the theoretical and mathematical basis for this mechanism we may be able to build self-organising machines. If we can apply these principles to software, we may be able to build self-organising information systems. Chapters 2 and 4 suggest that some progress can be made in this direction. Self-organisation in biology, however, does not come as a “reusable module” that we might plug into different types of systems. It is difficult to separate the self-organisation of an adult organism from its inception and growth. Furthermore, it appears to implicate very large numbers of components organised hierarchically over a range of length scales of approximately 10 orders of magnitude. We are nowhere near being able to build physical machines of such complexity.




Fig. 5.1 Evolution and self-organisation

With software we have more freedom and fewer constraints. In principle we should be able to enforce both the developmental and the architectural requirements of self-organisation in the design of a software system. Chapters 2 and 4 focused more on the architectural aspects. An important observation we made was that in some cases the manner in which some species have harnessed the available molecular infrastructure and physical laws in order to further their successful adaptation to their environment has led them to *approximate* information systems at a macroscopic level, the ant colony being the prime example. In other words, whereas the individual ants may be driven by free energy minimisation, at the scale of the ant colony self-organisation takes place more through the

¹⁹ www.omg.org

exchange of information mediated by a memory mechanism than by interaction energy. This is very good news for us, it lends greater legitimacy to our second goal and gives us greater confidence that we may attain it by means of current technologies.

This chapter focuses more on the developmental aspects. DNA controls both the growth and development of an organism as well as its day-to-day metabolism. Although running the metabolism seems a little easier than understanding morphogenesis, both processes are controlled through gene expression, which we can then take as the interpretation of a code to perform or enact corresponding functions. As discussed in previous chapters this interpretation implicates the environment, leading to the term “distributed algorithm”. Distributed is meant both in space and in time, meaning that the “instructions” are to be found in different physical locations as well as that they become clear or active at different times. The instructions at the end of a metabolic cycle acquire meaning, and in fact become active, only if those at the beginning have been performed in the appropriate order and environment. At all points in the cycle the meaning is context-dependent, and the context changes as a function of time partly as a result of the instructions themselves (the essence of non-linear coupling).

This view of the DNA as an instruction register of the cell’s metabolism is taken from Kauffman (2000), who explains how different metabolic cycles can be considered equivalent to methods or functions. Each involves roughly 300 steps, each step consisting of a different state of the DNA and each state consisting of a different set of genes becoming active. The switching on and off of the genes sends signals both to the “protein factory” in the cytoplasm as well as to other genes that are either activated or de-activated. Kauffman talks about this in the context of a discussion on construction of order, marvelling at the fact that with 30,000 genes that can either be on or off the human DNA can be considered a binary instruction register with $2^{30,000}$ possible states. As shown in Fig. 5.2, a generous estimate of how many instructions may be necessary to run the human body (assuming that different metabolic cycles and cell types share no common instructions) gives a number of states that is microscopically smaller than the total maximum possible. This fact is in direct contradiction with the laws of probability, which is why we refer to the construction of order in biology as “defeating entropy”. As discussed in Chapter 2, the fact that even in the presence of mechanical and chemical perturbations the human body is able to remain within this extremely small set implies some form of stability, even if not in the elementary sense.

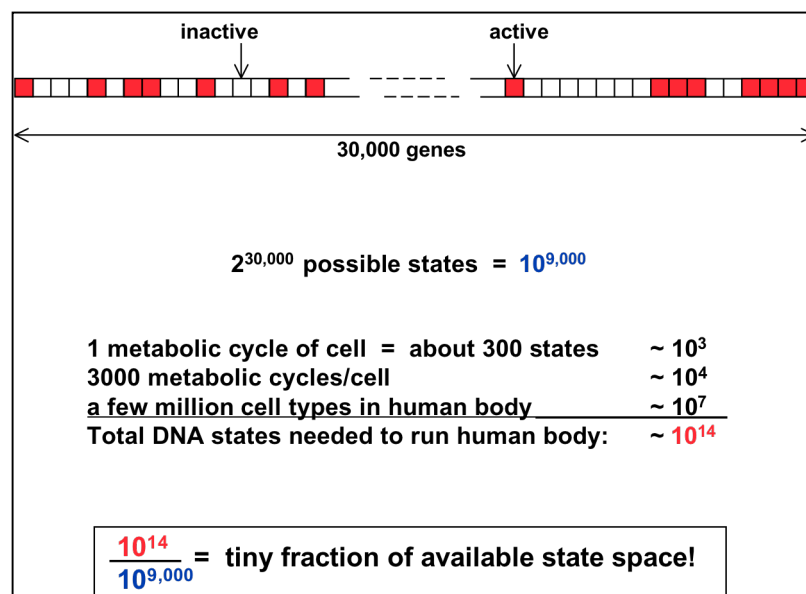


Fig. 5.2 DNA as instruction register (Kauffman, 2000)

On Engineering, Language, and DNA

When designing a bridge, the engine of a car, a wind-turbine blade, a manufacturing process, or a radio receiver, an engineer must be aware of the physical laws and constraints that determine the behaviour of the materials these various artefacts are made from, and of how the artefacts interact mechanically, chemically, electromagnetically, etc, with their environment. In other words, as shown in Fig. 5.3 in older kinds of engineering the relationship between structure and function is quite tight.

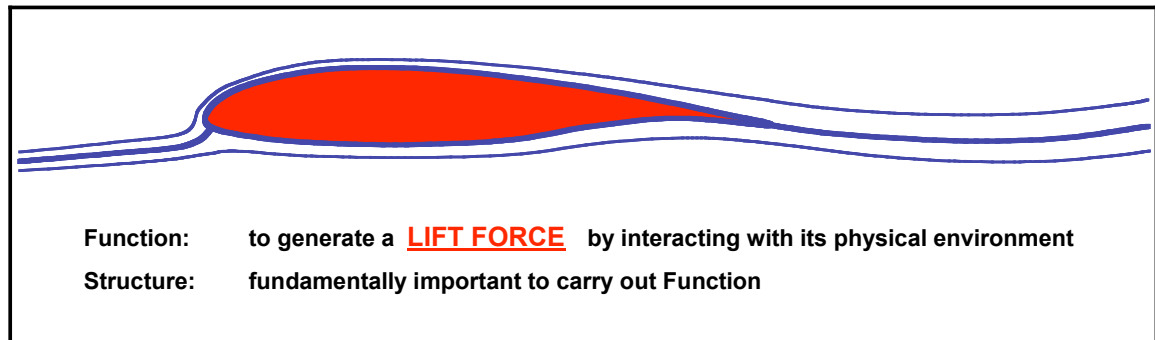


Fig. 5.3 Structure and Function in older kinds of engineering (wing section)

At the same time, the engineer must also be aware of the design requirements the artefact is supposed to satisfy and of the economic, social and environmental constraints. Costs, human resources, and timeline can be grouped under the general heading of Economic Constraints. As shown in Fig. 5.4, the result of the design process is a balance between the requirements on the one hand and these various constraints on the other. The arrows are meant to indicate a tension, as in a tug-of-war, that pulls the design process in different directions. Ideally, the engineering artefact should represent the point of equilibrium between these equal and opposing forces, although all-too-often social and environmental constraints are not enforced or are overlooked altogether.

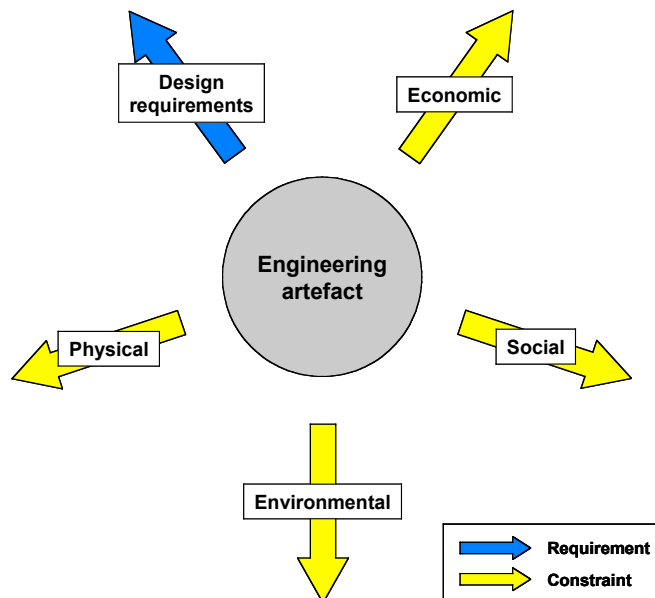


Fig. 5.4 Engineering design drivers as requirements and constraints

Unlike older kinds of engineering, in digital electronics something curious happens: whereas the user requirements and the economic, social, and environmental drivers remain as strong as for other kinds of engineering, there is significantly more freedom in deciding on the physical structure of the circuit. This is because the functional properties of the circuit are carried by its topological rather than its physical structure. As a consequence, the physical embodiment of the circuit can take many forms in order to fit in different enclosures, under different environmental conditions, etc, all achievable while

still fulfilling pretty much the same set of functions. Some physical constraints are still there: two high-frequency lines cannot be put too close to each other because they will interfere electromagnetically, too many high-power components cannot be crammed together because they will melt, etc. But these constraints are more at the component level and are less stringent at the circuit or board level. In short, electronic circuits are more abstract than other engineering artefacts and allow us to speak in terms of “levels”. Fig. 5.5 shows an electronic engineering artefact.

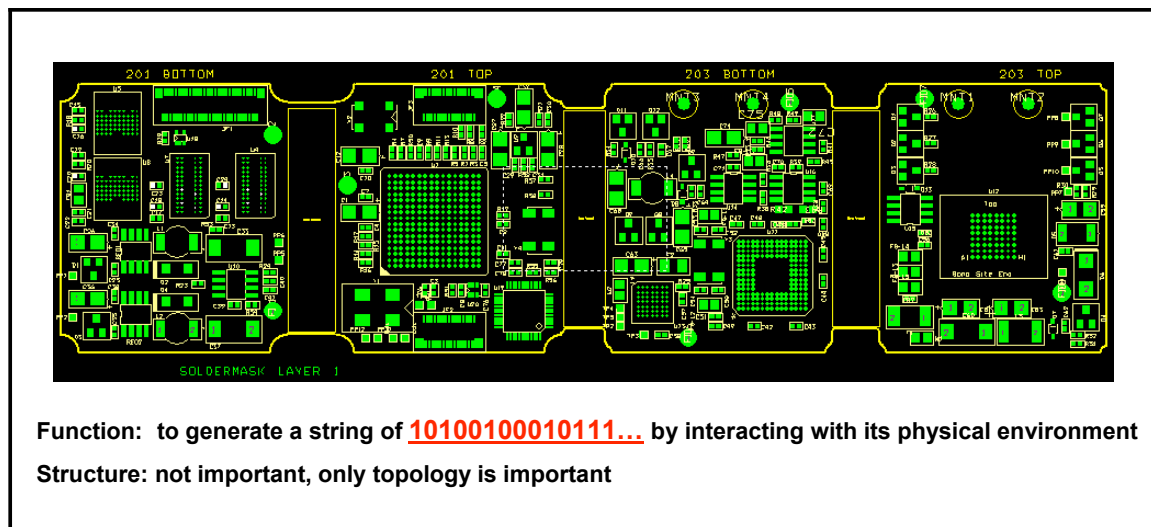


Fig. 5.5 Structure and Function in Electronic Engineering (wrist computer)

The abstraction of digital circuits is a direct consequence of their function. Rather than performing a physical function such as digging a hole or deflecting a flow of air, digital circuits perform an abstract function: they generate patterns of signals that on the physical plane are totally pointless, other than perhaps warming up the air around them. Their function lies in the meaning of the patterns of +3/0 Volts, which however are meaningful only to us. Fig. 5.6 shows the corresponding design drivers.

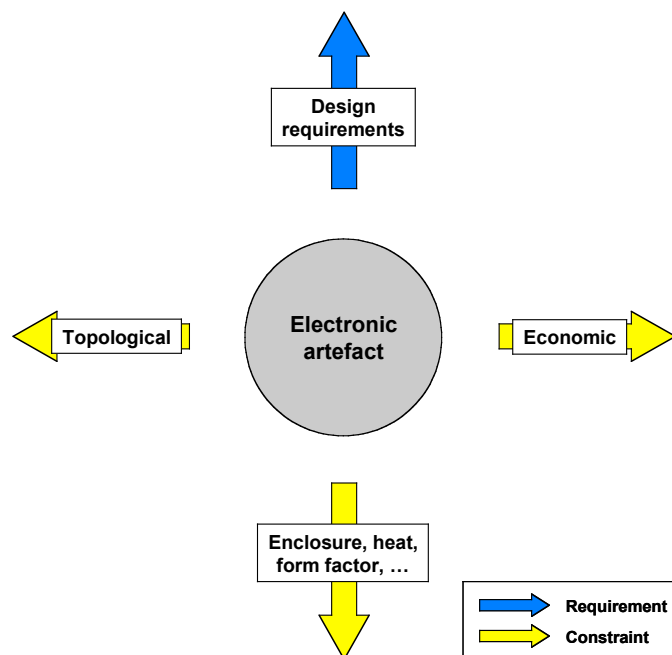


Fig. 5.6 Design drivers in Electronic Engineering

It is interesting and possibly important to point out that by operating on two levels, the physical and the symbolic, computers emulate coarsely the basic functional architecture of living organisms: short-

range interactions between components (molecules) at a small spatial scale that enable “emergent” processes at greater physical scales and higher levels of abstraction. The former tend to be internal to an organism and enable the latter, i.e. its interaction with the environment or with other organisms. For example, the break-up of sugar molecules (free energy minimisation) enables functions that are at a higher level of abstraction, such as the motion of a limb, the formation of a thought, or the utterance of a cry. Such higher-level functions only have meaning relative to the context with which they have co-evolved, and in this they are similar to human codes and languages.

All this allows us to claim that at higher levels of abstraction Function is dependent on meaning, which is relative, whereas at lower levels of abstraction Function is dependent on the physical context. In Biology, in fact, we know that Function is intimately dependent on Structure—and vice versa. For example, the metabolic function of a particular protein or enzyme is determined by its shape; change the shape and the protein becomes useless or “inactive”, as good as dead. Conversely, the structure and density of bones is greatly influenced by their use, as astronauts who spend long periods in orbit are well aware of. We wish to emulate certain dynamical aspects of biological systems in software; the dynamics of biological systems are based on interactions between components which, in turn, are based on a close interdependence between their structure and function. Thus, it seems plausible that the analogue to structure and function will play an important role in the dynamics of software. As odd as this notion may appear when we look at it from the point of view of Biology, the relationship between structure and function has in fact played a major role in the evolution of Computer Science since the beginning.

In biological systems the transition from physical context to meaning is quite smooth and gradual since it happens over many levels of the hierarchy. As research in the workings of the mind and cognitive processes of the last 40 years seems to indicate, there is no metaphysical boundary between the brain and the mind. Likewise, in software, many levels have evolved since its early days. Software originated around 1940 as sequences of instructions (“algorithms”²⁰) first for mechanical, then for electronic machines that performed the numerical calculations of complex problems in physics and engineering. Over the last 60 years algorithms have followed a progression of increasing abstraction: from binary machine code, to Assembler, to procedural languages, to structured programming, to objects, and now to Model-Driven Architecture and Executable UML. Each instruction at each level incorporates many instructions expressed in the language of the level below. This allows us to express desired behaviour of the software in terms of languages and concepts that are closer to human languages than to the binary number system or to sequences of voltage levels.

FORTTRAN²¹ was developed in the 50s to calculate the numerical solution of mathematical models of various types of problems in engineering and physics. When C++ started gaining acceptance many engineers were left in a state of confusion for several years. While Fortran had been developed by scientists and engineers in the 60s, the newer languages, especially object-oriented languages, were being developed by a new breed of scientists called “Computer Scientists” whose theoretical foundations lay not in physics but in the algebraic structure of language. Accordingly, the term “modelling” acquired a very different meaning to how it is used in science or in the social sciences, in spite of the fact that in all these communities it is associated with uncovering order and gaining understanding. In applied maths and physics order comes from the fundamental symmetries of the physical universe that show up as various conservation laws. In language a kind of order is found in the static relationships between concepts, such as inheritance and “property of”, that can be represented graphically as semantic networks and that tell us something about the structure of a particular conceptual domain. Semantics is the study of meaning in language. But, as we discovered in Chapter 4, meaning is not absolute, we derive the meaning of a concept based on its relationship to “nearby” concepts and its environment. Hence the strong dependence of meaning on conceptual structure or, in fact, “semantic structure”. The formalisation of semantic networks was given the

²⁰ From the name of an Arab mathematician who lived around 800 AD.

²¹ FORmula TRANslation

name of “computational ontologies” and the mapping of a domain through semantic networks and ontologies as “domain modelling”.

The evolution of Computer Science, Software Engineering, and computer languages in general has been determined first and foremost by engineering expediency. The increasing abstraction was motivated by the very practical goal to separate different functions into separate modules, thus facilitating software re-use and substitution of older modules with newer versions when they became available. Thus code maintenance, implementation efficiency, and robustness of the applications were the main motivators.

As we think of the abstraction process that started with digital electronics and has led us as far as executable UML, we get the impression that software is becoming increasingly pliable and that there is no limit to the level of abstraction algorithms and data structures can attain. In its broadest sense, we can see that Information Technology is indeed free from physical and material bounds, it can be used to represent anything, from business transactions, to movies, to chemical reactions, and it can do this in infinitely many different ways. And yet, we find that Computer Science and Software Engineering have evolved to *increase* the importance of the relationship between structure and function in software, creating constraints where none existed. Starting with a mapping of the user requirements through UML use cases, in fact, the underlying class structure of the application is gradually discovered (Fig. 5.7).

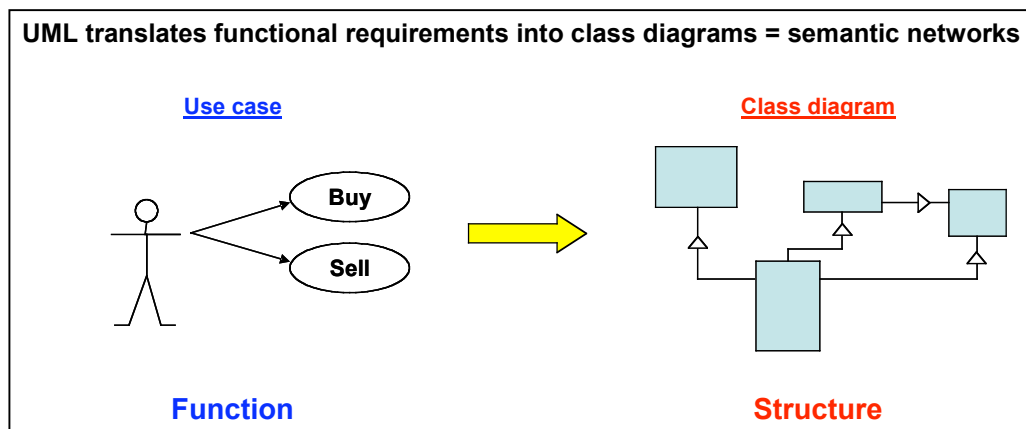


Fig. 5.7 Function and Structure in software engineering

Semantic structure is therefore intimately related to the functionality of the application, suggesting that we could draw an analogy with the relationship between structure and function in Biology. Rather than being constrained by physical laws and materials, software engineering is determined by user requirements, costs, and semantics (Fig. 5.8).²² The Semantic arrow in the figure is a requirement rather than a constraint because it represents the structural requirements needed to satisfy the functional requirements.

The standard software engineering approach works well if the requirements do not change over time, because once a domain has been mapped in terms of classes and semantic networks the relationships between them are assumed to remain fixed. The claim in Nachira’s paper (2002) is that in order to be truly useful to SMEs software services need to be able to adapt in real time to varying needs, without requiring a re-design with its very high associated costs. And that, in order for the software to be able to adapt to varying needs in real time, it needs to embody dynamical characteristics typical of self-organising and self-optimising biological systems. While it is certainly true that given the boundless flexibility of software such features could in principle be developed through arbitrary, abstract lines of thinking, taking inspiration from what we already know works so well (i.e. Biology) seems logical

²² The focus here is on software technology rather than on content, which is clearly influenced also by cultural, social and artistic drivers.

and sensible. Hence, in the development of the DBE our software engineering efforts will be subject to an additional constraint, the Scientific, as shown in Fig. 5.9. The user requirements defined by the SMEs are called “functional requirements”, while the real-time interaction and self-optimising characteristics of the software derived from biological systems are called “non-functional requirements”. Fig. 5.10 puts in perspective the different types of modelling and requirements in the DBE.

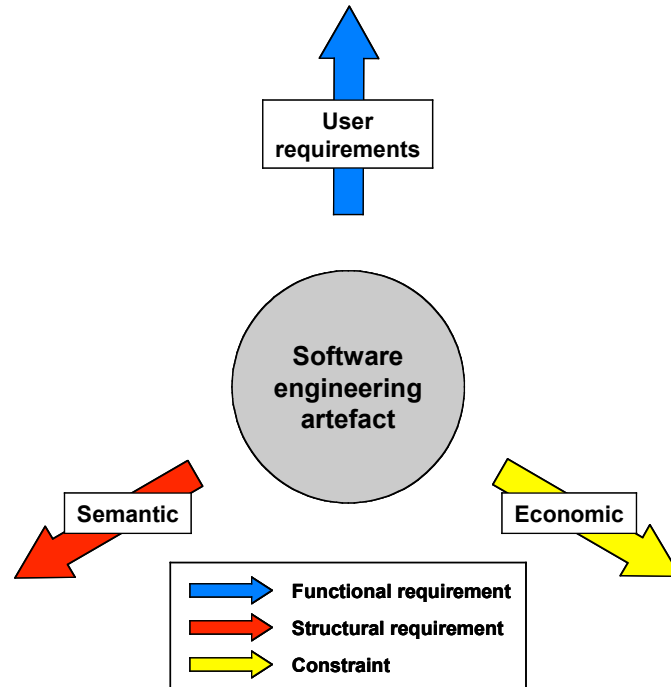


Fig. 5.8 Design drivers for software engineering artefacts

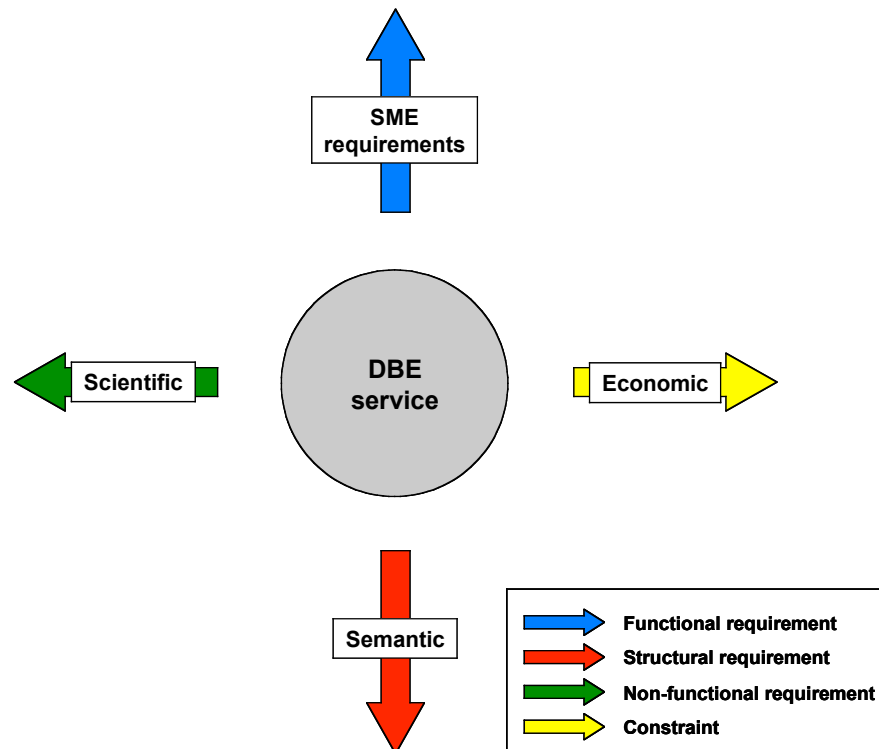


Fig. 5.9 Design drivers for DBE artefacts

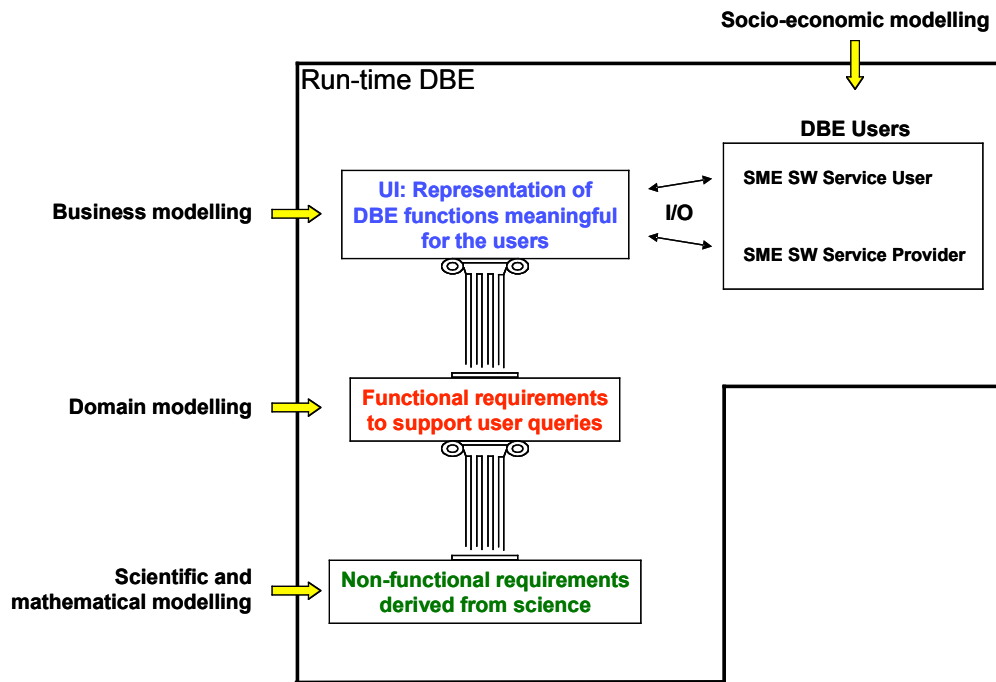


Fig. 5.10 Different kinds of modelling in the DBE

The discussion of the last five pages is finally summarised in Fig. 5.11, which allows us to begin to guess what the non-functional requirements might entail. The Dynamics column highlights the fact that whereas the functions of physical machines are carried out in real time, those of software can be executed at any time and at different speeds, depending on the processor. As applications have become increasingly sophisticated, however, event-driven and real-time applications have become more common. User-centred applications can already demonstrate that interesting behaviour happens at time scales comparable to the frequency of interaction with the users. The DBE needs to go further in this direction and mirror physical systems, where the dynamics is exclusively driven by interactions at corresponding time scales. We emphasise that this figure represents a long-term view of what we may be able to achieve, or at least what seems worthwhile exploring, specifically the creation of services from their DNA equivalent. For the present and for the foreseeable future MDA provides an appropriate framework for the creation of DBE services and infrastructural components.

The reason for coupling the dynamics of the DBE closely to the interactions with the users and between components will take some work to understand fully and to explain, but for now we can get a glimpse of what may lie ahead by looking at the central column, Representation of Behaviour. At the physical level, representation of behaviour is the behaviour itself. At higher levels of abstractions one can separate the representation of behaviour—the code—from the behaviour itself—the execution of the code. Both are in any case deterministic. The dynamic requirements of DBE software seem to point to the necessity to leave a significant degree of indeterminacy in the code, implying that the representation of behaviour embodied in the DNA is only partial. In addition to carrying genetic information between generations, DNA acts also as the code that “runs” the organism, in that it runs all the metabolic cycles of the cell. It can only do so, however, if it is in proximity not just of energy and water, but also of the correct molecules, proteins and enzymes that interact with it and with each other to carry out the various cell functions. The “algorithm”, therefore, is distributed among several structures. It could be useful to the DBE to design into the software components the ability to carry out different tasks in the presence of different stimuli, while retaining a long-term record of successful behaviour in the ecosystem.

Reconciling DNA with MDATM

The workings of the DNA are certainly impressive, but we need to question whether such a complex mechanism is relevant and desirable for the DBE. There are two reasons why the answer is likely to

be affirmative. First, it seems unlikely that we will achieve the self-organisation of run-time software components without having coded at least some aspects of such behaviour in their design, and DNA seems to have perfected how this codification is done. Second, the DNA's ability to respond to its environment in an algorithmic way points to the automation of at least some aspects of the software development cycle, which is one of the stated objectives of the Model Driven Architecture working group of the OMG (Frankel, 2003).

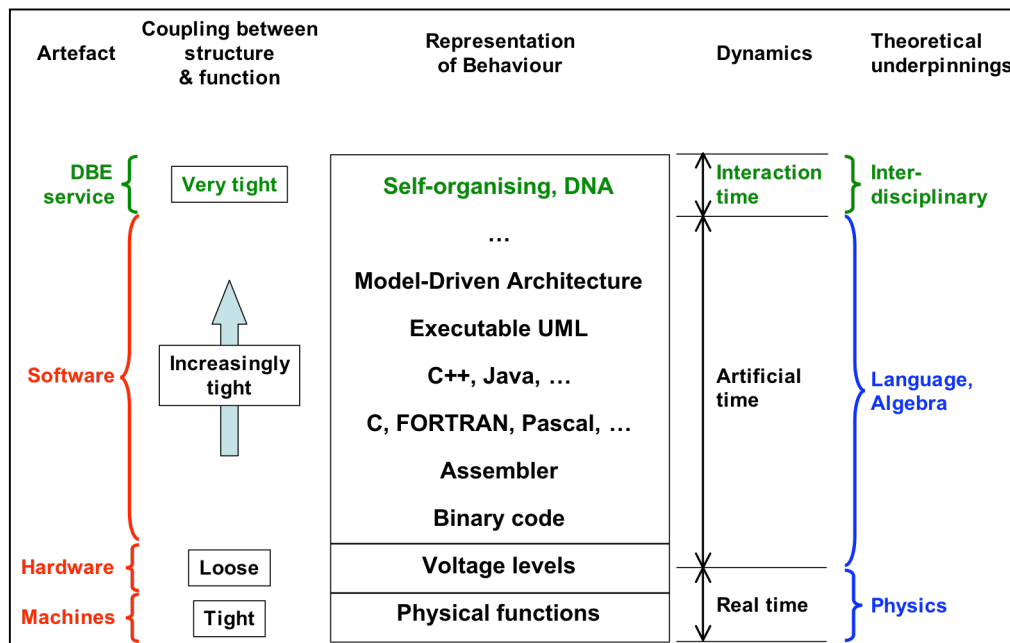


Fig. 5.11 Levels of abstraction for representing behaviour, long-term view

An overview of MDA is provided by Soluta in their reports (D21.1, D21.2), since they are in charge of defining the architecture of the whole DBE. Frankel (2003) provides a very good introduction. The main idea can be summarised by dividing software engineering data into four levels of abstraction: M0 is the data layer, M1 is the layer of models, M2 is the layer of languages (meta-models), and M3 is the layer of meta-languages (meta-meta-models). M3 is also referred to as the Meta-Object Facility (MOF). It is a meta-language that was defined a few years ago in order to define UML, which belongs at layer M2. MOF is self-describing, so that no additional abstraction levels are necessary. The DBE Consortium has agreed that all DBE artefacts must be MOF-compliant. The sub-divisions within each layer are somewhat subjective and not rigid, but an in-depth discussion belongs in the reports from the architecture and languages workpackages (WP21, WP15, WP14). What we can say here is that languages like BML (Business Modelling Language), SSL (Semantic Service Description), and SDL (Service Description Language) belong to layer M2 and are defined through an abstract syntax. Their implementations in UML or BSBR editors (Business Semantics of Business Rules, a new OMG working group) provide the concrete syntax that allows their utilisation to represent business models and service descriptions.

Even if we accept to identify the design and models of software components with their DNA, whose precise meaning will be defined in Chapter 6, there is a problematic implication. Whereas in Nature the “design” of different species is effected over long time scales by evolutionary processes that progressively specialise (in a relatively stable environment) the most successful genomes, there is no-one explicitly “designing” the DNA of a particular new species to be introduced in the ecosystem. This is shown schematically in Fig. 5.12. In the DBE, by contrast, we aim to balance the evolutionary optimisation of the software with their top-down model-driven design. The manner in which we have provisionally reconciled these two opposing requirements is to limit the optimisation to the

composition of service chains of *existing* DBE Services that are individually created and implemented by more traditional approaches.

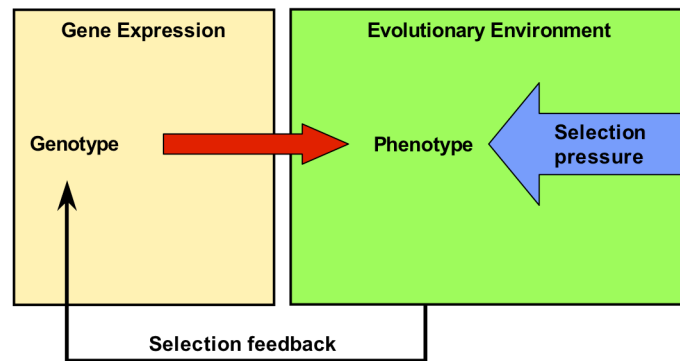


Fig. 5.12 Schematic of natural selection cycle

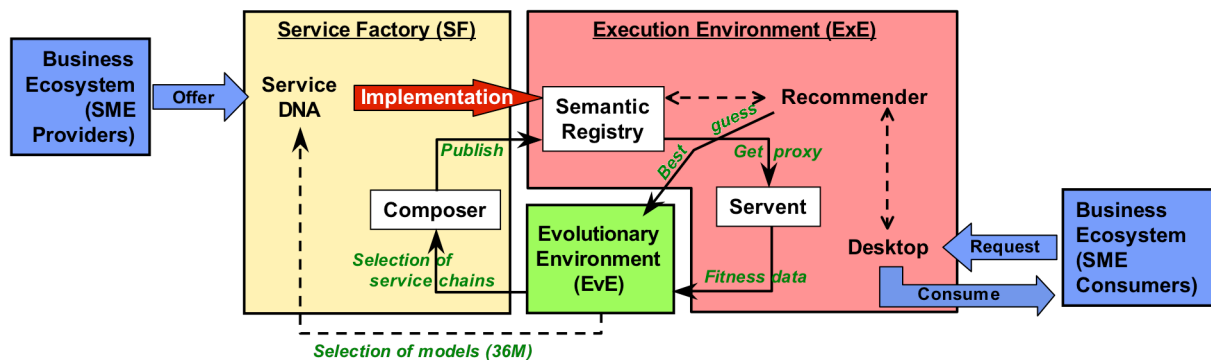


Fig. 5.13 Evolutionary cycle adapted to the DBE, 18-Month system

Fig. 5.13 shows how we have adapted the evolutionary cycle to the DBE. Four of the five environments of the DBE are shown along with some of the principal components. This figure captures the three main use cases of the DBE: creation and offer of a service, service request and consumption, and service chain optimisation. Each of these can be broken down in several others. The selection of models (Service DNA) based on real-time feedback from the performance of the phenotype will not be addressed until the later implementations of the DBE. For explanations of components please see D21.1, Architecture Requirements (Ferronato, 2004).

Seeking the theoretical foundations of models

The purpose of the above sections was to establish in what sense the high-level concept of DNA might apply in the DBE. We have not really begun to address gene expression, so clearly evolutionary algorithms are confined to the optimisation of chains of manually implemented software components. To achieve this in a manner that is useful to SMEs our current understanding of genetic algorithms is sufficient; the challenge is, rather, in how we perform a match between service interfaces, especially at the semantic level. In the long term, however, we want to keep our options open regarding code generation from models. This leads to the following questions:

- What kind of models are we talking about?
- How relevant are scientific models to self-organisation?
- How relevant are scientific models to the specification of self-organisation in software?
- How can we represent context-dependent causal relations between components in UML?
- Would we gain anything if we were able to do it?

Answering these questions will require a review of the main publications in the philosophy of language, from Russell's theory of types to formal semantics. We may be able to skip over human languages and pragmatics since we are interested in languages or codes that are somehow computable. At the same time, we do not want to be constrained to formal logics. We will also have to review the philosophical and theoretical foundations of scientific models. Whereas this will be addressed in the next deliverable, for now we can begin to say a few general things to help define the problem better.

We already hinted at the possibility of using state machines to represent mutually dependent components. State machines perform an action based on their current state, and this action can be programmed to affect another state machine. The state of the second machine will change as a result, and the new state can be programmed to trigger another action that affects the first machine. In this way the time evolution of two (or more) state machines can be coupled to each other. The n-tuple formed by grouping the states of the individual components can be regarded as the state of the overall system. The problem is that programming behaviour in this manner requires the knowledge beforehand of the range of situations or inputs that the components will experience, with the corresponding actions and responses. As long as the number of states and corresponding actions for each component is small the time evolution of such a system can be predicted and controlled deterministically, but as the number of interacting components, the number of states for each, and the number of possible actions increase the system's state quickly becomes unpredictable in practical terms.

The problem of combinatorial growth of the possible states of software systems as they scale in size is well-known and is, in fact, one of the main technical motivators for the DBE project itself. Looking at the universal success of domain modelling languages such as UML and at the growing popularity of Model Driven Architecture one might find it odd that such powerful software engineering tools are not able to ensure the stability of large software systems. The underlying reason is that domain modelling is most effective at modelling static relationships. Interestingly, the bulk of scientific research is bent on understanding and modelling time-dependent behaviour. Even in these two simple statements we can perceive the wide gulf that divides these two different ways of using the word "model". So what do we mean by model?

Domain models in informatics have strong connections to language because they are composed of symbols whose meaning to represent concepts is ultimately constructed through a social process. The model of a domain consists of a set of classes together with a set of relationships between the classes. This looks like an ontology, although a model is usually augmented with sequence and state charts to capture behavioural information.²³ There are three main characteristics of domain models that seem particularly relevant in this discussion:

- 1- Adaptive. Since they are built from requirements and concepts specific to the application domain, at the time of their definition they tend to take the "shape" required to perform the desired functions.
- 2- Invariant. The same class diagram captures static relationships that hold through the various functions of the application. Thus class diagrams extract the static structure from an application domain. Domain models are adaptive in finding this structure but, once they have found it, they are frozen.
- 3- Relative. Just like ontologies, the meaning of a concept in a domain model is strongly dependent on the concepts it is related to in the semantic network, where a semantic network can be considered the visualisation and instantiation of an ontology in a more specific domain. This reinforces our intuitive understanding of language and of information as discussed in Chapter 4.

The problem posed by the combinatorial growth of states in complex applications is analogous to the problem of the huge state space available to our DNA. Probabilistically, it is a constant marvel that we do not collapse into a wet mess of organic molecules on the floor, in the space of a few hours. Likewise, domain models can cope with a small number of possible behaviours, but become increasingly poorer at accounting for all possible behaviour in larger applications.

²³ This observation comes from a work-in-progress by Paul Krause (UniS).

Scientific models of static structures are similar to domain models. The information they contain, however, is not the geometry of the structures themselves (unless we are talking about architectural models, in which case it is). In science and engineering by “model” we usually mean the mathematical distillation of the *relationships* between the relevant concepts or entities of the domain in question. For a mathematical statement to qualify as a model it must apply to the whole domain equally well. For instance the mathematical model that describes the static relationship between the members of a steel bridge is

$$\sum F = 0, \quad (5.1)$$

meaning that the sum of the forces acting on every bolt, pin and rivet is zero. Obviously such a model is rather abstract and from it we cannot deduce the shape of the bridge without additional information. Eq. (5.1) is the condition of equilibrium. If we take into account the elasticity of steel the same equilibrium condition can be arrived at by searching for the minimum potential energy of the bridge structure. So this is the converse of the microcanonical ensemble of Chapter 2: there is no entropy to maximise and equilibrium corresponds to the minimum potential energy of the system.

There are many other kinds of scientific models, all of which involve quantitative relationships between the entities of the domain. If the entities change as a function of time then we may be able to say something about how their rates of change are related, which leads to differential equations. Whereas differential equations capture relationships between variables (and their derivatives) that change as a function of time, the relationships themselves hold for all time. So in this sense they still capture a static aspect of the systems they model. The time-dependent behaviour of the variables can be derived from these static structures through mathematical integration. Mathematical or scientific models, therefore, can refer either to the “governing” differential equations or to their solutions. More generally, the development of scientific models and theories progresses in steps, each of which is compatible with the available experimental data, and often to the *same* experimental data. In most cases models begin with the observation of some reproducible regularity in a physical system, which awakens our curiosity to find out the reasons behind it. The quest to find the cause of the observed effect fuels the development of models that as a rule become progressively more refined and more general.

One of the most well-known examples of this progression starts with the planetary astronomical data tabulated by Tycho Brahe in the second half of the Sixteenth Century. Half a century later Kepler recognised patterns among the data and derived two separate laws. These laws look more like correlations than actual explanations, from a modern viewpoint, but when they were proposed they did involve a significant elaboration of the raw data into concepts of significantly higher abstraction. An example is the rule that the radius joining the focus of the elliptical orbit (the Sun) to a planet sweeps a sector of constant area in equal time intervals, even though the planet travels at different speeds as it moves along its orbit. Not too long after that, Newton showed how Kepler’s 2nd Law can be derived from his own 2nd Law and by his Universal Law of Gravitation. Almost two hundred years later Hamilton showed how Newton’s 2nd Law is but a constraint that guarantees the minimisation of a variational integral called the Action, expressed as a time integral of the Lagrangian between fixed end-points. Although we can calculate the Lagrangian we do not truly understand what it is. However, we know that Hamilton’s principle is very general, in fact it can be considered the origin of Schrödinger’s equation of quantum mechanics and Feynman’s quantum electrodynamics, in addition to all of classical mechanics. The Universal Law of Gravitation was generalised in Einstein’s general relativity. At each step of this progression, if we look back the models seem simplistic and unimpressive, if we look forward we marvel at how such a generalisation might have been possible and we tend to ascribe to the more sophisticated model the status of TRUTH. In fact, the progression does not appear to end. In the case of Mechanics, as of today we can’t really explain Hamilton’s principle in terms of more fundamental concepts and are forced to take it as an underlying axiom of much of physics. String theory claims to be able to unify all of physics, but the practical impossibility to verify its predictions experimentally may well condemn physical theory never to find closure.

From this very quick synthesis of modelling in physics we can conclude that there are some fundamental regularities of physical systems that, even if we cannot explain completely, we can still take advantage of to predict quantitatively the behaviour of variables of interest. Our ability to predict the behaviour of physical systems, however, runs into a wall when we come to non-linearly coupled systems of 2 or more degrees of freedom. While we can usually find the analytical solution of linear systems, there is no general theory of integration for non-linear systems. The closest we come to a general method is the search for symmetries based on Lie groups, which will be discussed in the next deliverable. Furthermore, since not all (systems of) differential equations are integrable, we will need to address also their integrability. These are current areas of research in mathematics, and may for this reason seem a bit daunting and difficult to apply as part of the DBE project. However, consistently with every other area of DBE research, we combine rigorous theoretical investigation with engineering intuition, empirical data, and business acumen to invent new solutions and new methodologies of design and development. Out of this interaction arises advancement in theory as well as practical applications.

For the purpose of positioning this discussion relative to other disciplines and philosophical traditions, to a physicist the development of models that are based on ascribing explanatory power to statistical correlations, and in fact Positivism in general, appears not too dissimilar to Kepler's laws: interesting and valuable from a pragmatic viewpoint, but not an explanation that goes in sufficient depth. The minimisation of free energy and equilibrium statistical mechanics in general play a similar role in the explanation of protein synthesis relative to the actual non-linear processes that make protein synthesis possible but that we do not yet understand.

In summary, the most difficult aspect of self-organisation, such as takes place in the cell and in the brain, we do not yet understand. Models based on global minimisation principles apply and are accurate in describing quantitatively the time-local relationships between derivatives of functions, but are almost useless in finding the functions themselves. We need a new kind of model. In this connection, there are two possible reasons to be optimistic:

- 1- through emergence and phase transitions complex systems achieve parameter or dimensional reduction and in some cases behave in a manner similar to information systems (ant colony)
- 2- there exists a methodology to analyse non-linear continuous systems based on a systematic search for symmetries that we have only hinted at in this report but that may be applicable to information systems as well

The interplay between the continuous/discrete and the physical/abstract descriptions of systems that these two points imply offers some hope that over the next 12 months we may be able to construct simulations to demonstrate generalised stability and the distributed algorithm of complex systems, at least, and a theory of self-organisation, at best.

Some considerations on UML

Based on this discussion it seems worthwhile to consider developing a UML specification language that incorporates intrinsically a time-dependent element. The easiest way to do this, again taking inspiration from the DNA, is to partition a domain model into active and inactive parts, and to set up appropriate triggers and interdependencies between the parts or between the parts and the environment, in order to render the class diagram adaptive in real time. This implies that the class diagram will not be a static structure anymore, it will change shape and evolve through a finite set of states. It also implies that we need a "gene expression" mechanism in order to generate updates to the run-time code in response to fast-changing requirements and conditions. A first step in this direction is suggested by the Adaptive Object Model, briefly discussed in the next and final Chapter of this report.

6. DBE-SPECIFIC USE CASES

A good starting point for this chapter is to clarify what we mean by its title. The term “use case” refers to a series of actions performed by human or non-human actors to enact a required function. Use cases are generally utilised to capture user requirements at the beginning of a software project. They automatically also serve to communicate these requirements between different teams working on the same project. They are especially useful as a communication link and common ground between the product definition and marketing team and the design and technical implementation team. As a consequence, use cases tend to be biased toward the representation of the interaction between the human users and the software system being developed. In other words, they are generally biased toward the functional requirements. The task of capturing the DBE functional requirements from the SME point of view falls on FZI: B9-“Use cases definition” and on TCH and ITA: B29-“Functional analysis and specification”.

“Use case” can also mean a more generic usage scenario not necessarily enacted by human users. It can refer also to how the internal components of a software application interact and communicate. Such use cases are biased toward capturing the non-functional requirements of the system, i.e. those technical requirements that enable the system to fulfil the functional requirements for the users. This report and this chapter are more focused on this second type of use cases. “DBE-specific” refers to the fact that these internal system processes are characteristic of the DBE, i.e. they capture in particular the self-organisation and evolutionary behaviour of the software.

As mentioned in the Introduction and in the previous chapter, we have not reached this point yet, we are behind relative to the expected outcomes in the original workplan. It is still useful, however, to provide a high-level UML view of the system. This is not a complete and detailed view; rather, it is intended to give an idea of how the system as a whole functions. After defining the basic concepts we present a brief discussion of the main parts of the system where we expect the science tasks to make a contribution, along with appropriate activity diagrams to capture the main use cases. At the end of the chapter we collect these facts and behaviours in a global communication diagram of the DBE infrastructure (that we call the DBE System). We emphasise that our purpose is not to review all the science tasks of the project. Some science tasks interact more with the business partners and will therefore not be discussed here since they do not affect the software infrastructure directly.

Basic elements

The reader is referred to D21.1, Architecture Requirements (Ferronato, *ibid*) for a more complete definition of the various system components. Fig. 6.1 shows the atomic unit of the DBE, the Service Manifest, which can be viewed as equivalent to an agent.

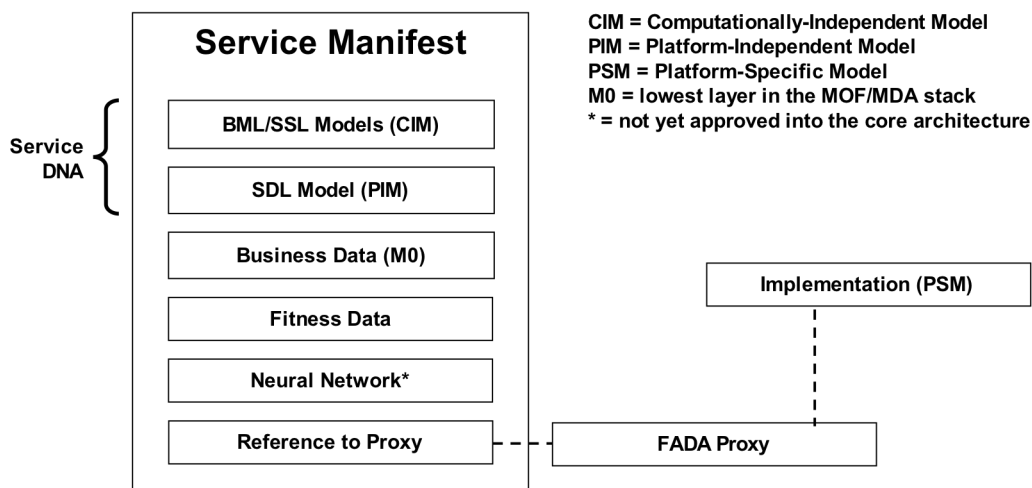


Fig. 6.1 The Service Manifest, the atomic unit of the DBE

As the Service Manifest is the focal point for the whole DBE, it is not surprising that a debate to define its nature and most convenient conceptualisation has been raging within the Consortium since the proposal writing stage. An alternative definition to Fig 6.3 is to say that the BML model and the SDL model represent a subset of the DNA that we can describe as the “species specification”. In other words, more than one SME and/or service may rely on the same models. According to this view, the SM as a whole can be considered equivalent to the DNA of an individual. This debate will be gradually resolved as our understanding of DNA improves and as we discover opportunities for mapping it to the structural and dynamic characteristics of DBE software.




Fig. 6.2 The three environments of the DBE System

The three environments of the DBE System are shown in Fig. 6.2. The arrows indicate dependence in terms of knowledge. Thus, the Evolutionary Environment (EvE) needs to know about the (state of) the Service Factory (SF) but not the other way around; the SF needs to know about the (state of) the Execution Environment (ExE) but not the other way around; and so on. A classification of the DBE actors is reproduced from D21.1 (Ferronato, *ibid*) in Fig. 6.3, with the addition for the DBE Business Service from Fig. 4.2.

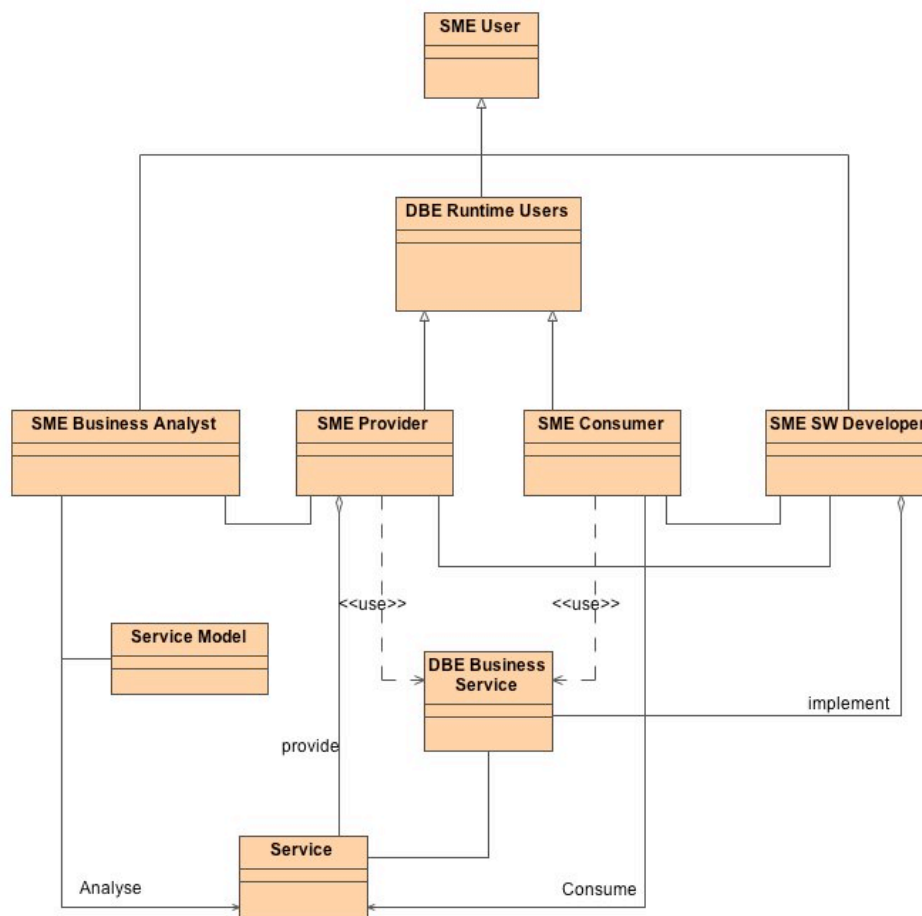


Fig. 6.3 A classification of DBE users (Ferronato, *ibid*)

Use cases that rely on Science outputs*Evolutionary computation*

As explained in Chapter 5, in the 18-Month implementation of the DBE we are applying genetic algorithms to chains of services. A group of partners has been working on the specification of the genetic algorithms (UBham), of the fitness function (STU), and of the architecture of the Evolutionary Environment (EvE) (ICL). Although this specification is far from complete we can show the first elements here. The preliminary EvE class diagram (Fig. 6.4) was originally created by Gerard Briscoe (ICL) and added to by Jon Rowe (UBham), Thomas Kurz (STU), and Jim Dowling (TCD), who introduced the Interceptor Framework. The interceptors are used also for the Accounting function (WIT) and for the User Profiling function (FZI). At the simplest level of explanation we can imagine one Habitat per SME, housing as many Population objects as there are requests. In reality, we will probably allow one habitat to span more than one SME, and each SME to hold multiple habitats. This is to accommodate the fact that the information and the models to optimise a service chain could be located at multiple, cooperating SMEs, or that one SME may do business in rather different domains. The former case provides a business context for developing distributed evolutionary algorithms that rely on distributed computing approaches such as GRID.

The Consume Service activity diagram is shown in Fig. 6.5. An important point that this diagram makes is that, following a service request, there are two strategies for fulfilling it. In the first use case the Recommender will use its own internal optimisation algorithms to find a best match among available single services or service chains (that it finds in the Semantic Repository). Simultaneously, the request and the best guess are passed on to the Evolutionary Environment that performs the optimisation of the chains. The component and integration testing (UniS-Krause) are closely connected to this process and affect the fitness of the services. Upon completion of the optimisation step the chain is passed to the Composer (TCD) that performs the binding and sends the completed chain as a new Service Manifest to the Semantic Registry in the Knowledge Base (TUC), as well as back to the Recommender that can now offer it to the SME user. Whereas the first use case is expected to happen in real time and is therefore more compatible with the “pull” point of view in service provision, the second use case does not have a time limit. When a solution is found by the EvE it is offered to the SME, enacting the “push” approach to service provision. The Recommender is used for both cases in order to present the same interface to the user.

Distributed intelligence

The Intelligence task in the DBE has had a longer gestation period. We understand evolution much better than we understand how intelligence works, so this is not surprising. ICL has proposed a daring vision of distributed intelligence that relies on neural networks either inside each service and/or outside, but within each Habitat, to accelerate what is described in D6.1 as the clustering of service chains. The corresponding component in the Habitat is therefore called the Clustering Catalyst. Two aspects are particularly interesting: 1) the services and the intelligence in the habitats will have the ability to learn over time, and 2) since the services move between different habitats this automatically leads to a distributed intelligence. In fact, the information and patterns the services will carry, even if at a sub-symbolic level, could represent precisely the information flow in networks that we discussed in Chapter 4. The dynamic nature of the networks will then open more possibilities for learning, clustering and adaptation at multiple scales. In other words, over time this approach has the potential to “bootstrap” more sophisticated forms of intelligence. It may also be able to address recognition of compatible interfaces between services and not just similarity between service chains, thereby helping the automatic composition process.

Interestingly, this idea can be presented from the opposite point of view. Namely, the same algorithms, data structures, and architecture can also be described as a method to augment the genetic algorithms, to speed up the selection and cross-over process in their convergence toward the desired goal. Seen from this point of view, as it is presented in D6.1, the evolutionary and the learning/intelligence processes become mutually reinforcing. We have not developed a use case or an activity diagram yet for these processes.

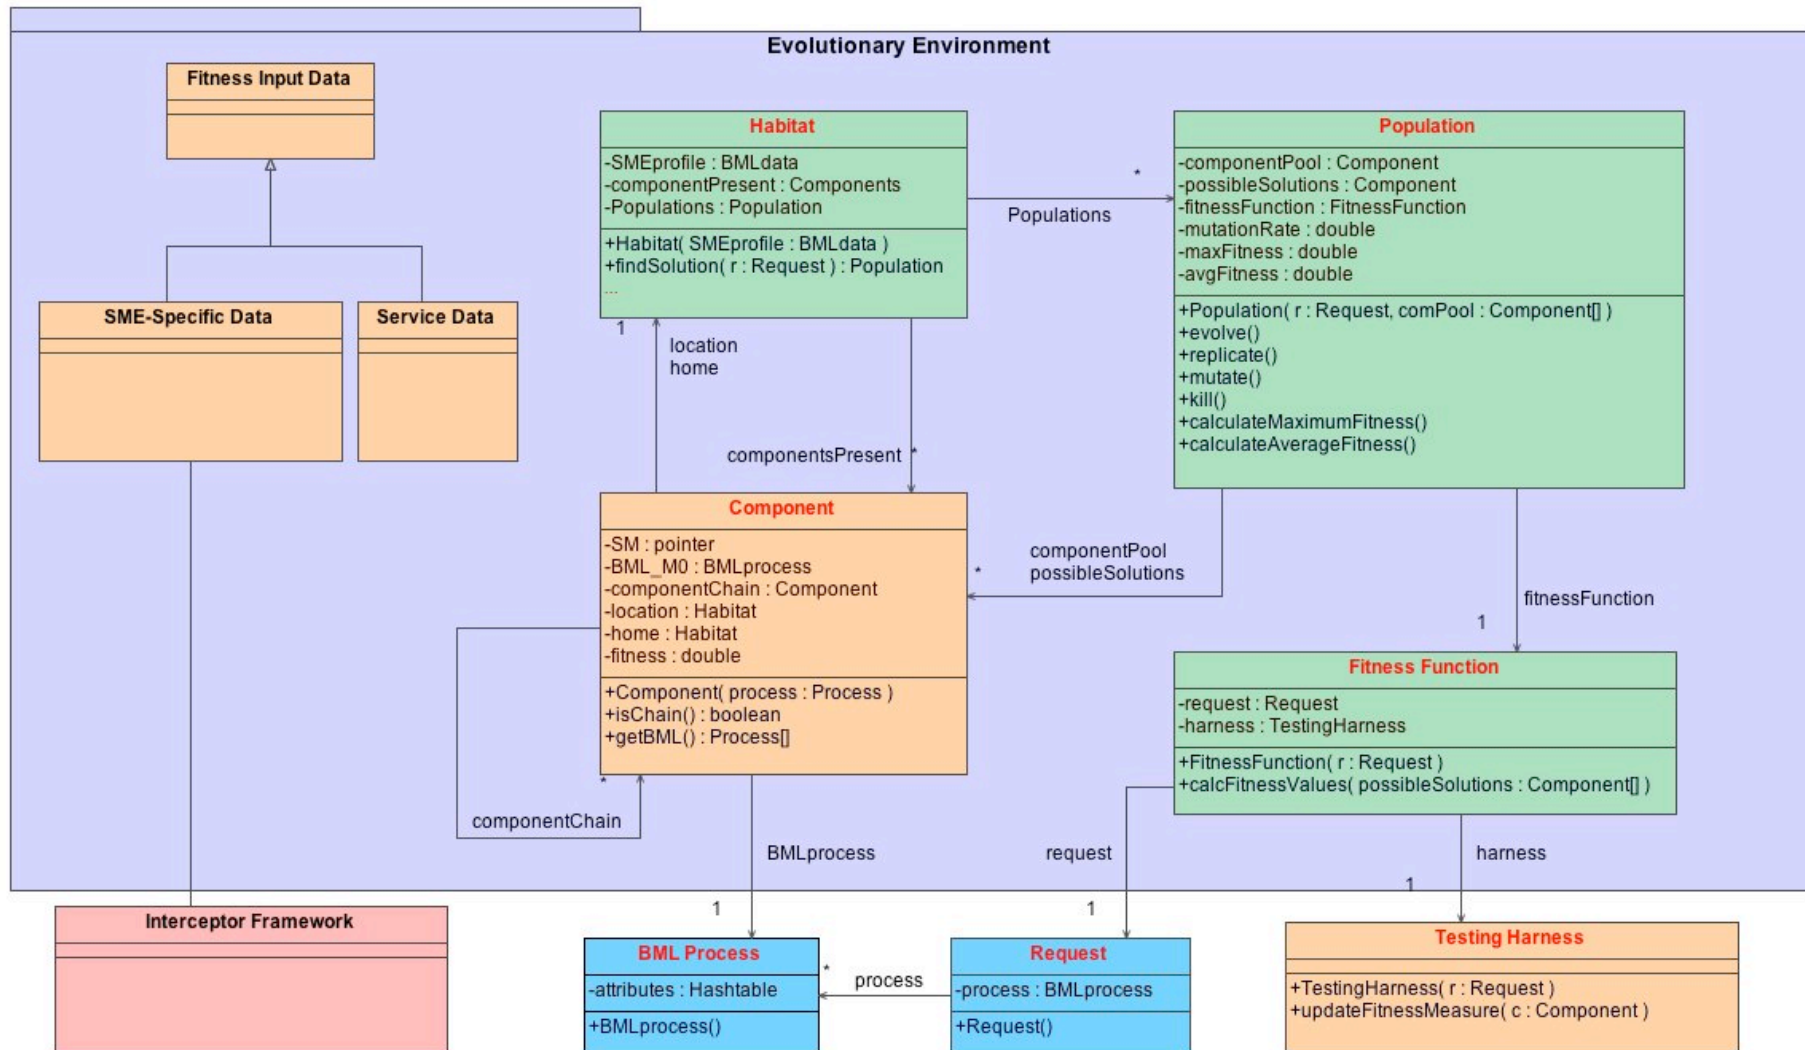


Fig. 6.4 Preliminary class diagram of the EvE (ICL, UBham, STU, TCD)

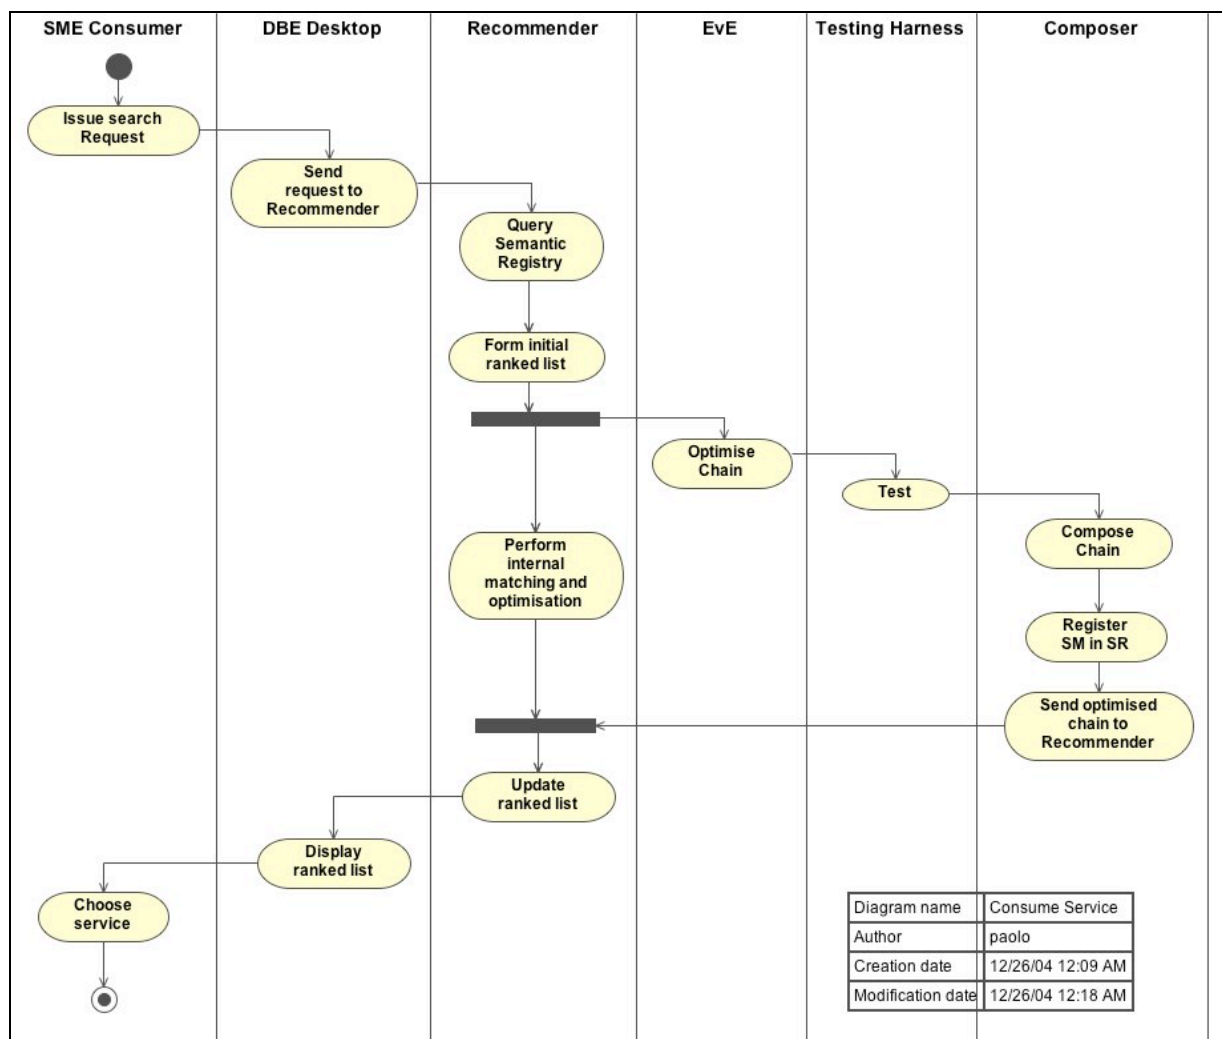


Fig. 6.5 The Consume Service high-level use case, including the evolutionary optimisation

Adaptive Object Model

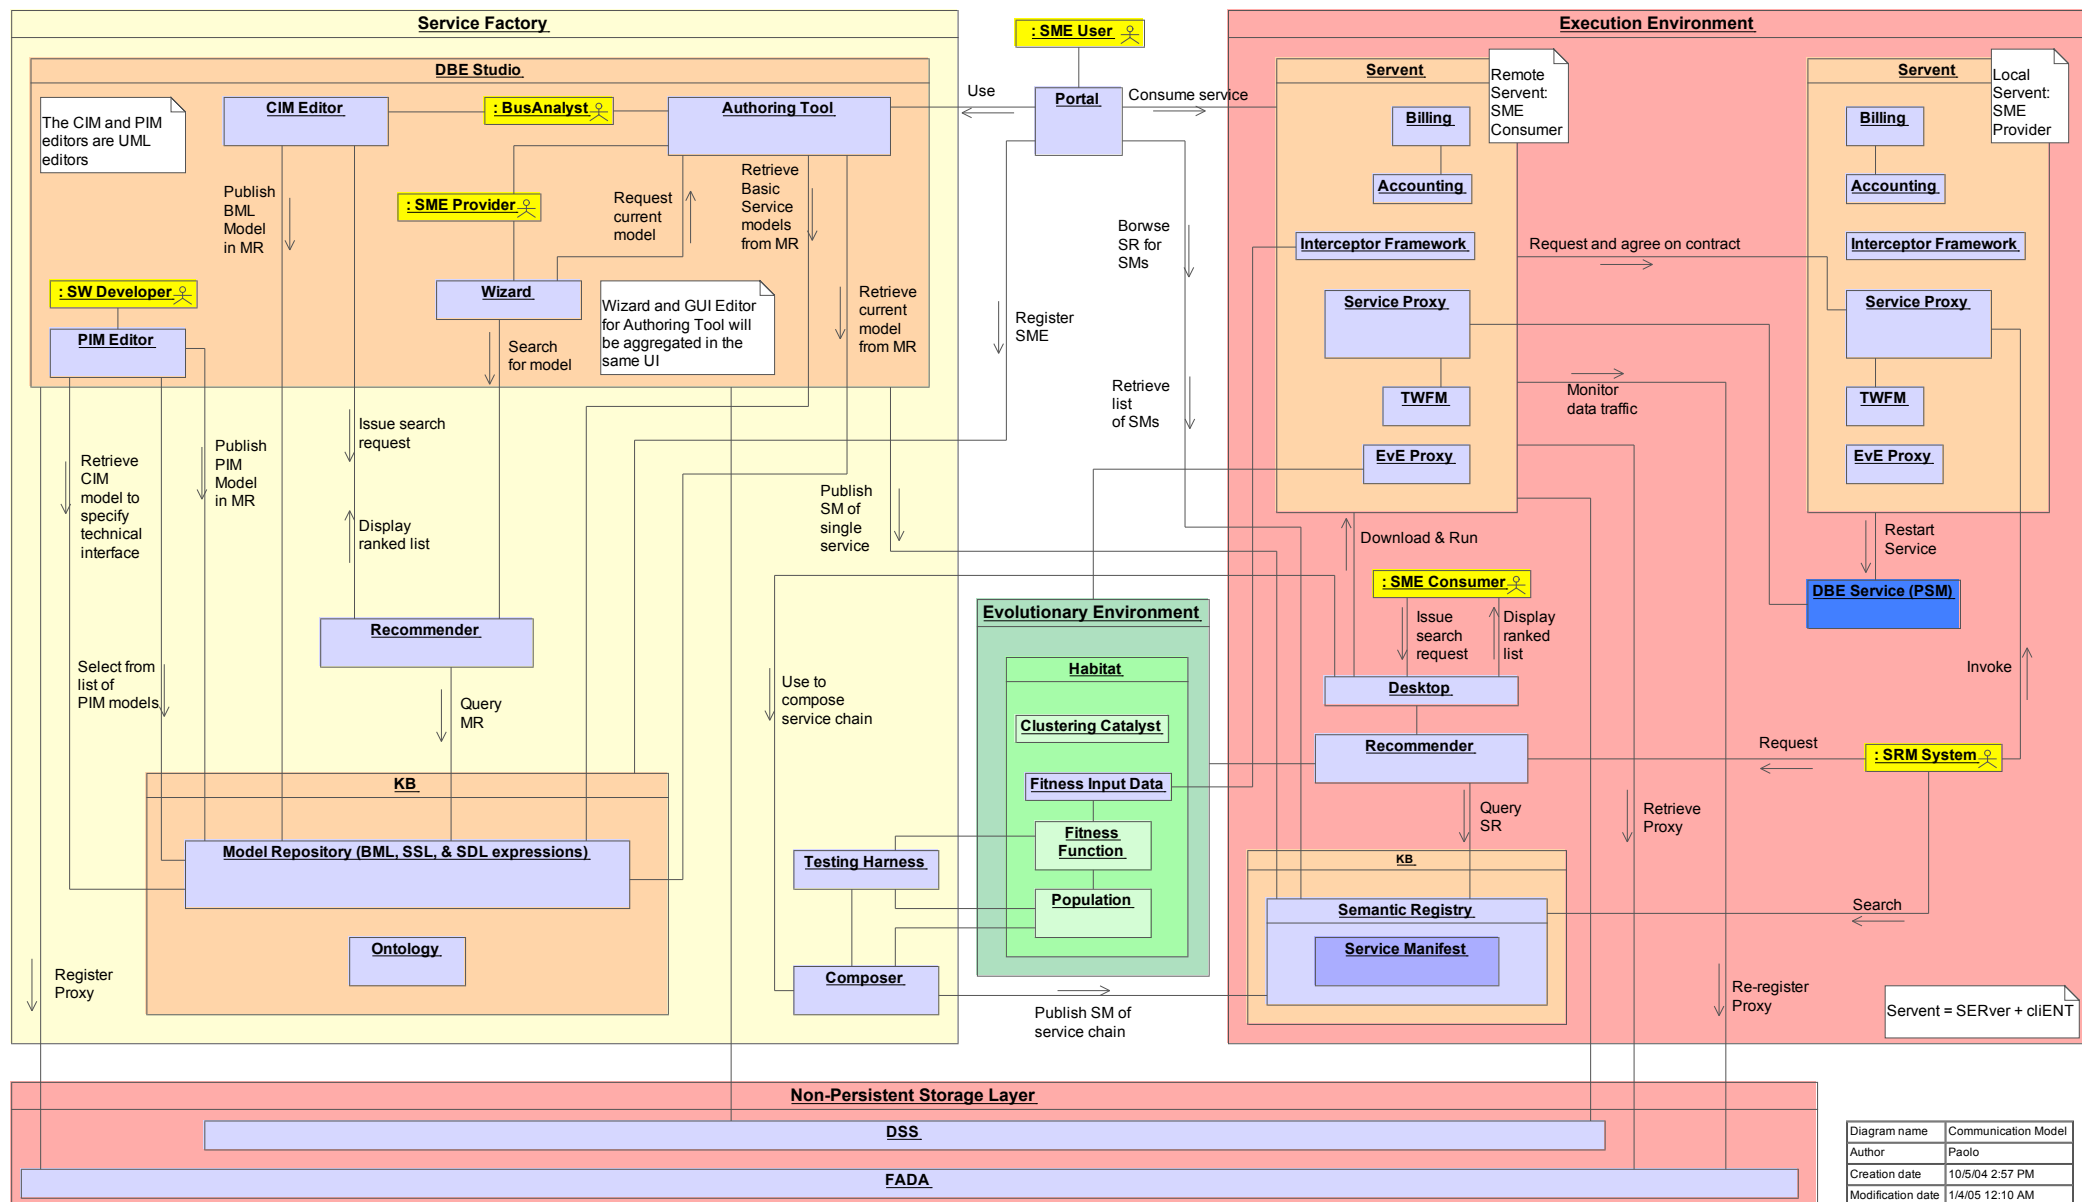
The main idea of the Adaptive Object Model²⁴ is to cast business rules as data rather than code. In this way they can be changed dynamically at run-time. Another way to describe this approach is to define abstract classes with the potential to acquire attributes and methods, but leaving them unspecified in the class implementation. When the class is instantiated, it can load the methods and attributes dynamically, subject to appropriate constraints. The result is greater adaptability to different environments. For example, an abstract and sparsely specified BML model could populate its classes with domain-specific methods by referencing a domain ontology at run-time. In the acquisition of information from its environment, AOM is conceptually similar to gene expression and the distributed algorithm concept. LSE is pursuing this line of research in collaboration with Claudius Masuch, STU.

In the DBE a mapping between AOM and gene expression is only just beginning. The goal is to achieve code generation similarly to xUML. However, our intent is to analyse gene expression carefully and to build into the method the most appropriate biological and self-organisation design patterns. The result will be a component that we are calling the Adaptive Service Generator (ASG), which will work in conjunction with the Intelligent System for the automatic creation of services from BML, SSL, and SDL models. A use case for this component is premature at this stage.

²⁴ www.adaptiveobjectmodel.com

Global communication diagram of the DBE

Finally, Fig. 6.6 is the general map of the system. This communication/collaboration diagram contains most of the links and a good number of the messages between the components. Most of the high-level use cases can therefore be found by following the links along different routes in this map. Although this figure can be a bit confusing since many use cases are collapsed into one diagram, we hope it will serve the purpose of helping the DBE stakeholders visualise the whole system at once, along with the principal components and interdependencies.



PHYSICAL, BIOLOGICAL AND INFORMATION SCIENCES BIBLIOGRAPHY

- Adami, C, Ofria, C, and Collier, T, (2000) “Evolution of Biological Complexity”, *Proceedings of the National Academy of Science*, USA, Vol 97, pp. 4463-4468.
- Albert, R, and Barabasi, A, (2002) “Statistical Mechanics of Complex Networks”, *Reviews of Modern Physics*, Vol 74, January.
- Badii, R, and Politi, A, (1997) *Complexity: Hierarchical structures and scaling in physics*, Cambridge.
- Barabási, Albert-László (2002) *Linked: The New Science of Networks*. Cambridge, Mass.: Perseus Publishing. Reprint.
- Beekman, M, Sumpter, D J T, and Ratnieks, F L W, (2001) “Phase transition between disordered and ordered foraging in Pharaoh’s ants”, *Proc Natl Acad Science*, USA, Vol 98 (17), August, 9703–9706. <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=55516>
- Berge, P, Pomeau, Y, and Vidal, C, *Order Within Chaos*, John Wiley, 1984
- Bowley R and Sanchez M, (1999) *Introductory Statistical Mechanics*, Oxford.
- Briscoe, G, (2004) “Entropy-Based Complexity Measure for the Evolution-Based Self-Organisation of Agent Populations”, DBE Report D6.1, October.
- Callen, H, (1985) *Thermodynamics and an Introduction to Thermostatistics*, 2nd Ed, John Wiley.
- De Bono, E, (1969) *The Mechanism of Mind*, Penguin (1990 Edition).
- Diamond, J, (1998) *Guns, Germs and Steel*, Vintage.
- Dini, P, Kuusisto, T, Corallo, A, Ferronato, P, and Rathbone, N, (2003) “Toward a Semantically Rich Business Modelling Language for the Automatic Composition of Web Services”, *Proceedings of the e-BRF 2003 Conference*, Tampere, Finland, 23-25 September. www.ebrc.info
- Ferronato, P, (2004) “DBE Architecture Requirements”, DBE Report D21.1, November.
- Frankel, D S, (2003) *Model Driven Architecture*, OMG Press.
- Hill, TL, (1986) *An Introduction to Statistical Thermodynamics*, Dover.
- Hoffman, A, (1998) *Paradigms of Artificial Intelligence: A Methodological and Computational Analysis*, Springer.
- Hook J R and Hall H E, (1991) *Solid State Physics*, John Wiley.
- Hydon, P, (2000) *Symmetry Methods for Differential Equations*, Cambridge.
- Johnson, S, (2001) *Emergence: The Connected Lives of Ants, Brains, Cities and Software*, Penguin.
- Kadanoff L P, (2000) *Statistical Physics: Statics, Dynamics, and Renormalization*, World Scientific.
- Kauffman, S, (2000) *Investigations*, Oxford.
- Kittel C. and Kroemer H, (1980) *Thermal Physics*, Freeman.
- Kohonen, T, (2001) *Self-Organizing Maps*, 3rd Edition, Springer.
- Kumar, S, and Bentley, P J, Eds, (2003) *On Growth, Form and Computers*, Elsevier.
- Ma S K, (1976) *Modern Theory of Critical Phenomena*, Perseus.
- MacKay, D J C, (2003) *Information Theory, Inference, and Learning Algorithms*, Cambridge.
- Mandl F, (1988) *Statistical Physics*, John Wiley.
- Pfeifer, R, and Scheier, C, (2000) *Understanding Intelligence*, MIT Press.
- Reif, F, (1965) *Statistical Physics*, (Berkeley Physics Course, Vol 5) McGraw-Hill.
- Reif, F, (1965) *Foundations of Statistical and Thermal Physics*, McGraw-Hill.
- Rozenfeld, H D, Kirk, J E, Boltt, E M, ben-Avraham, D, (2004) “Statistics of Cycles: How Loopy is Your Network?”, arXiv:cond-mat/0403536, Vol 1, 21 March.
- Shannon, C, (1948) “A mathematical theory of communication”, *Bell System Technical Journal*, July.
- Sornette, D, (2000) *Critical Phenomena in the Natural Sciences*, Springer.
- Thomas, L, (1974) *The Lives of the Cell*, Viking.
- Wiener, N, (1965) *Cybernetics: Or Control and Communication in the Animal and the Machine*, MIT Press (paperback edition of 1948 publication).

SOCIAL SCIENCES BIBLIOGRAPHY

- Arthur, Brian W. (1999) "Complexity and the Economy." *Science*. Reprinted in *The Complexity Vision and the Teaching of Economics*, D. Colander, ed., Edward Elgar Publishers, 2000 2 April, no. 284: 107-09.
- Axelrod, Robert. "Advancing the Art of Simulation in the Social Sciences." In *Simulating Social Phenomena*, edited by Rosario Conte, Rainer Hegselmann and Pietro Terna, 21-40. Berlin: Springer, 1997.
- Bauman, Zygmunt (2000) *Liquid Modernity*. Cambridge: Cambridge Polity. Reprint.
- Beck, Ulrich (2000) "The Cosmopolitan Perspective: Sociology of the Second Age of Modernity." *British Journal of Sociology* 51, no. 1: 79-106.
- Braudel, Fernand (1974) *Capitalism and Material Life: 1400-1800*. London: Fontana/Collins. Reprint.
- (1975) *The Mediterranean and the Mediterranean World in the Age of Philip II*. London: Fontana. Reprint.
- Bullock, Seth, and Dave Cliff. "Complexity and Emergent Behaviour in Ict Systems." edited by last accessed 10/11/04 available at: <http://comp.leeds.ac.uk/seth/DTi/report.pdf> Hewlett-Packard, 2004.
- Buskens, Vincent (1993) "The Social Structure of Trust." *Social Networks* 20265-89.
- Byrne, David (1999) *Complexity Theory and the Social Sciences*. London;New York: Routledge. Reprint.
- Castells, M (1996), *The Information Age*, Blackwell.
- Chung, Seungwha (Andy) (1996) "Performance Effects of Cooperative Strategies among Investment Banking Firms: A Loglinear Analysis of Organizational Exchange Network." *Social Networks* 18, no. 2: 121-48.
- Cilliers, Paul (1998) *Complexity and Postmodernism*. London and New York: Routledge. Reprint.
- De Landa, M, (1997) *A Thousand Years of Nonlinear History*, Zone Books.
- Elliot, Euel, and Douglas Kiel. "Non-Linear Dynamics, Complexity and Public Policy." In *Chaos, Complexity and Sociology*, edited by A. Eve Raymond, Sara Horsfall and Mary E. Lee, 64-78. Thousand Oaks, Calif.: Sage Publications, 1997.
- Freeman, Christopher, John Clark, and Luc Soete, eds (1982) *Unemployment and Technical Innovation*. Great Britain: Frances Pinter (Publishers). Reprint.
- Freeman, Christopher, and Luc Soete, eds (1987) *Technical Change and Full Employment*. Oxford; New York: Basil Blackwell Ltd. Reprint.
- Hollis, Martin (1994) *The Philosophy of the Social Science*. Cambridge; New York: Cambridge University Press. Reprint.
- Kay, J (2003), *The Truth about Markets: Why Some Nations are Rich and Others Remain Poor*, Penguin.
- Lansing, J.S. (2002) "Artificial Societies and the Social Sciences." *Santa Fe Institute Working Paper* 02-03-011.
- MacLean, Donald, and Robert MacIntosh. "Complex Adaptive Social Systems: Towards a Theory of Practice." In *Complex Systems and Evolutionary Perspectives on Organizations: The Application of Complexity Theory to Organizations*, edited by Eve Mitleton-Kelly, 149-65. Oxford: Pergamon, 2003.
- McCue, Brian (2002) "Another View of the "Small World"." *Social Networks* 24121-33.
- McLennan, Gregor (2001) "Sociology's Complexity." *Sociology* 37, no. 3: 547-64.
- McNeil, Arthur H. (2001) "Passing Strange: The Convergence of Evolutionary Science with Scientific History." *History and Theory* 4011-15.
- Mitleton-Kelly, Eve. "Complexity Research-Approaches and Methods: The Lse Complexity Group Integrated Methodology." In *Organizational Complexity*, edited by Kenskinen et al. Helsinki: TUTU Publications, 2003.
- . "Organizations as Complex Evolving Systems." Paper presented at the OACES Conferece, Warwick 1998.
- . "Ten Principles of Complexity & Enabling Infrastructure." In *Complex Systems and Evolutionary Perspectives of Organizations: The Application of Complexity Theory to Organizations*, edited by Eve Mitleton-Kelly. Oxford: Pergamon, 2003.

- Nachira, F, et al. (2002), "Toward a network of digital business ecosystems fostering the local development", Discussion paper issued by the e-Business Unit, European Commission, Sept. http://www.europa.eu.int/information_society/topics/ebusiness/godigital/sme_research/index_en.htm
- Paulson, K. Steven (1985) "A Paradigm for the Analysis of Interorganizational Networks." *Social Networks* 7:105-26.
- Richardson, Kurt A. (2002) "On the Limits of Bottom-up Computer Simulation: Towards a Nonlinear Modeling Culture." *Proceedings of the 36th Hawaii International Conference on System Sciences (HICSS'03)*.
- Smith, Warren, and Matthew Higgins (2003) "Postmodernism and Popularization: The Cultural Life of Chaos Theory." *Culture and Organization* 9, no. 2: 93-104.
- Sokal, D. Alan (1996) "Transgressing the Boundaries: Towards a Transformative Hermeneutics of Quantum Gravity." *Social Text* 46/47:217-52.
- Stewart, Peter (2001) "Complexity Theories, Social Theory, and the Question of Social Complexity." *Philosophy of the Social Sciences* 31, no. 3: 323-60.
- Tainter, Joseph (2000) "Problem Solving: Complexity, History, Sustainability." *Population and Environment: A Journal of Interdisciplinary Studies* 22, no. 1: 3-41.
- Thrift, Nigel (1999) "The Place of Complexity." *Theory, Culture & Society* 16, no. 3: 31-69.
- Urry, John (2000) "Mobile Sociology." *British Journal of Sociology* 51, no. 1: 185-203.
- Wasserman, Stanley, and Katherine Faust (1994) *Social Network Analysis: Methods and Applications*. Cambridge; New York: Cambridge University Press. Reprint.
- Watts, Duncan J. (1999) *Small Worlds : The Dynamics of Networks between Order and Randomness*. Princeton, N.J.: Princeton University Press. Reprint.
- Webster, Frank (2002) *Theories of the Information Society -2nd Edition*. London and New York: Routledge. Reprint.
- Yamaguchi, Kazuo (2002) "The Flow of Information through Social Networks: Diagonal-Free Measures of Inefficiency and the Structural Determinants of Inefficiency." *Social Networks* 16, no. 1: 57-86.