

**FET FP7 Workshop of 21-22 April 2004 on
New Directions for ICTs in FP7: Grand challenges for basic research**

**Thematic Session 3
Systems Research**

Chairman: *Gilles Kahn*, INRIA, FR

Rapporteur: *Bob Malcolm*, Ideo, UK

Participants:

Ozalp Babaoglou, University of Bologna, IT

Paul Bourgine, Ecole Polytechnique, FR

Mark Buchanan, UK

Luca Cardelli, Microsoft, UK

Kemal Delic, HP, FR

Yves Demazeau, INPG-IMAG, FR

Paolo Dini, London School of Economics, UK

Gus Z. Eiben, University of Amsterdam, NL

Manuel Hermenegildo, University Polytechnica Madrid, ES

Nicos Karcanias, City University of London, UK

Miroslav Karny, Czech Academy of Science, CZ

Vladimir Kucera, Czech Technical University, CZ

David Lane, University of Modena, IT

Patricia Longstaff, University of Syracuse, USA

Wendy Mackay, INRIA, FR

Michel Morvan, ENS Lyon, FR

Andrea Nicolai, T6, IT

Corrado Priami, University of Torino, IT

Jakka Sairamesh, IBM, USA

Mark Shackleton, University of Lancaster, UK

Joseph Sifakis, INPG-IMAG, FR

Moshe Sipper, Ben Gurion University, IL

Mikhail Smirnov, Fraunhofer Fokus, DE

Paul Spirakis, CTI, EL

Gabor Vattay, Eötvös Lorand University, HU

Richard Walker, University of Napoli, IT

Martin Wirsing, Technical University of Munich, DE

Participants from the European Commission:

Ralph Dum, Loretta Anania, Leonardo Flores Anover, Fabrizio Sestini, Franco Accordino,

Teresa De Martino, Paul Hearn, Antonella Karlson, Panagiotis Tsarchopoulos, Jose

Fernandez-Villacanas (DG INFSO)

Table of Contents

<i>Executive Summary</i>	2
<i>Introduction</i>	4
<i>Context</i>	4
<i>Vision</i>	5
<i>Research challenges</i>	5
1. <i>Inferring systems models from incomplete information</i>	5
1.1 <i>Challenge</i>	5
1.2 <i>Rationale</i>	6
1.3 <i>Research priorities</i>	6
2 <i>Human-Responsive Systems</i>	7
2.1 <i>Challenge</i>	7
2.2 <i>Rationale</i>	7
2.3 <i>Research priorities</i>	8
3 <i>Foundations for Software-Intensive System Development</i>	8
3.1 <i>Challenge</i>	8
3.2 <i>Rationale</i>	9
3.3 <i>Research priorities</i>	9
4 <i>Control and Integrated Design of Complex Systems</i>	10
4.1 <i>Challenge</i>	10
4.2 <i>Rationale</i>	10
4.3 <i>Research priorities</i>	11
5 <i>M3P – Many Level, Many Unit, Many Purpose Systems</i>	11
5.1 <i>Challenge</i>	11
5.2 <i>Rationale</i>	12
5.3 <i>Research priorities</i>	13
<i>Conclusion</i>	14

Executive Summary

Large scale systems like communication networks, large databases and software systems, the Internet, large distributed control systems, businesses and the global economy, are examples of huge, interdependent open information-processing systems with behaviour that is increasingly difficult to predict and control. Modelling, simulation, design and control of such large scale systems in technology, business and the sciences are major issues to address in the coming years and were the subject of discussions held in this thematic session.

Present and prospective developments were discussed as a basis for understanding how to address the challenges of building systems that are robust, resilient, dependable and secure, exhibit multi-purpose functionality, and guarantee operation in mission critical tasks. It was recognised at this session that interdisciplinary research drawing on results from complex systems research is now essential if we are to establish new paradigms within which to address the challenges ahead. Five research directions were identified.

The first research direction addresses the need to *infer system models even when only inconsistent and incomplete information is available* about their functioning and interactions. We need to develop techniques for inferring the dynamics of complex systems, the laws governing their interaction, and ways to describe their behaviour, in order to simulate many systems for which there is at present insufficient direct knowledge. Such systems occur especially in ecology, medicine, molecular biology, certain technological-information processes such as the internet, and systems in management, finance, and economics whose behaviour is very dependent on the human in the loop.

The second research direction addresses the design of *human responsive ICT systems* that integrate well with humans and adapt well to human needs. It was recognised that today we can no longer treat ICT systems as separate from their human users and others affected by them, and that an ICT system and the context in which it is deployed together form a system. Research is required to establish new design principles that accommodate the changing needs and desires of human participants in complex systems, rather than treating them as outside the system, and rather than presuming to know what they need. The research would need to acquire better understanding of human behaviour - especially regarding group man-machine interactions – from which generative theories should inform the organisation of architectures that support and sustain participation, including participative co-design of the systems themselves.

Three complementary research directions address the need to ensure that systems built have the properties we demand. All of them insist on the need to underpin the formal description of such systems.

The first of this set of research directions is to develop underpinning *'foundations' for software-intensive systems*. The aim is to make a fundamental leap in the scientific basis of software engineering technologies to capture evolution and dynamics, selfish interests of individual entities, various levels of bounded rationality, learning aspects and self-emerging behaviour, in a strict, yet tractable way. Advances required include new algorithmic techniques for distributed systems and property-aware compilation and implementation techniques for non-functional properties such as security, safety, scalability, resource optimisation, quality of service, and efficiency, in order to be able to guarantee these non-functional properties.

The other two research directions in this set, namely *"Control and Integrated Design of Complex Systems"* and *"Many level, Many Unit, Many Purpose Systems"* are closely related. It was recognised that for systems comprising many interacting elements we currently lack techniques for modelling that enable prediction of the types of structures, the dynamic behaviour, and the properties that are likely to emerge at higher, collective levels of activity. A key characteristic of many such systems is that they involve components with uncertain operation and interactions, which are at least partially unpredictable and often not yet formalised in algorithmic terms. A number of biological, social, management and economic systems exhibit these properties, yet they must continue to function adaptably, malleably and resiliently, in the face of such unpredictability. The ambition is on the one hand to establish engineering guidelines that draw inspiration from complex natural, social, technological and economic systems and on the other to establish an underpinning framework of formal or mathematical techniques. Together, they should enable us to find cost-effective solutions to problems that cannot be solved with current techniques.

Introduction

This report summarises the main grand challenges for basic research in Complex Systems, discussed during a workshop held in Brussels on 21-22 April 2004 on “*New Directions for ICTs in FP7: Grand challenges for basic research*” organised by IST-FET.

Based on position papers submitted in advance, a panel of European expert scientists and engineers from a broad range of complex systems-related fields proposed and debated different scientific points of view and anticipated the likely evolution of current research areas.

The experts explored the new research directions required to address system complexity, and the challenges of building systems that are robust, resilient, dependable and secure, exhibit multi-purpose functionality, and guarantee operation in mission critical tasks.

By way of introduction, each expert gave a very brief description of themselves and the main thrust of their position paper. After some discussion in plenary about the themes and structure for subsequent discussions, the session self-organised into sub-groups that established and characterised five research directions, described in the sequel.

Context

System complexity, deriving from the interactions of large numbers of highly interdependent system components, is becoming an obstacle to further progress in the design and deployment of large scale systems. Communication networks, large databases and software systems, the Internet, large distributed control systems, businesses and the global economy, are examples of huge, interdependent, open, information-processing systems with behaviour that is increasingly difficult to predict and control.

No less important, the huge amount of data to be handled today in many areas of science and the increase of information flow in novel communication and computation systems – like the next generation Internet, the Grid, and a myriad interconnected nomadic wireless devices – confront us with unprecedented challenges.

Modelling, simulation, design and control of such large-scale systems in technology, business and the sciences are major issues to be addressed.

The session had therefore two principle foci:

1. Design and modelling of complex information and communication systems, and
2. The modelling and simulation of complex systems in other domains, such as biological systems and social systems.

Today however, ICT systems can no longer be considered in isolation: they are embedded in larger, and typically complex, technical and social systems (that may themselves be considered information and communication systems). So the session took account of the need to consider also ‘conventional’ ICT systems in the context of the wider-scale systems in which they must operate.

The experts also considered the impact of awareness of complex systems behaviour, and emerging results from complexity science, on the fundamentals of computer science and software engineering.

Vision

Recent scientific studies on the principles of functioning of complex systems – like living organisms and ecosystems, as well as societies – suggest that lessons on efficiency and reliability can be learned from such systems. Ideally, new system theories that are informed by the scientific study of such systems in other domains will identify architectural, algorithmic and functional foundations of systems that incorporate vast numbers of adaptive and interacting elements.

In turn, such theories will inform and underpin:

- The design of scalable ICT systems with novel functionalities and substantially higher autonomy, performance, and reliability.
- The effective use of IT systems in modelling, simulating and ultimately managing highly complex processes (for instance in biological or social processes), thus helping to understand their complex behaviour and making their complexity more ‘tolerable’.

Prospective examples include *autonomic computing and communication* that embody self-configuring and repairing computation; *self- and context-aware communications*; and *self-modifying software*, with reduced cost of building and maintenance.

Scalable tools for analysing and understanding large amounts of data from living systems could provide the basis for a new era in the life sciences. And in social systems the capacity to handle large amounts of data will open a new era of social simulation and prediction.

Research challenges

From the wide range of research directions suggested in the position papers of the experts, the participants at the session identified five priority challenges for research.

The first two, on ‘*Inferring systems models from incomplete information*’ and on ‘*Human-Responsive Systems*’ were quite discrete. The remaining three all addressed the aim for complex man-made systems to have the properties that we desire. ‘*Foundations for Software-Intensive System Development*’ seeks a transformation of software engineering techniques so that the properties of such systems can be guaranteed. The other two – ‘*Control and Integrated Design of Complex Systems*’ and ‘*Many level, Many Unit, Many Purpose Systems*’ – both seek to learn from complex natural, social, technological and economic systems: the former to establish engineering guidelines that draw inspiration from such systems; the latter to establish an underpinning framework of formal techniques for analysis and reasoning about such systems.

Since these challenges were developed by sub-groups of the workshop, operating independently, there is some overlap in the associated research priorities that they identified.

1. Inferring systems models from incomplete information

1.1 Challenge

The ambition is to be able to model and simulate many systems for which there is at present insufficient direct knowledge. Such systems occur especially in ecology, molecular biology, certain technological-information processes (internet, etc), systems with the human in the loop (management, finance, economics) and in medicine. To achieve this will require techniques for inferring models – dynamics, laws governing their interaction and dynamics and describing their behaviour – that may be used as the basis of simulation, even in the absence of full information and in the presence of inconsistent and partially false information.

1.2 Rationale

An implicit assumption in modelling of systems, as they occur in nature, society or technology, has always been that we know, or could in principle know, the individual elements and their interactions in great detail, or we know the evolutionary process that underlies the changing mechanism. For many complex systems in the real world this is definitely not the case: we lack a full description of the relevant components of the system; there is uncertainty in the values of parameters of the relevant components of the system; and there is incomplete information about their interactions (interconnection topology) and about the influence of the system's environment.

The explosion of dynamical data is now, for the first time, providing a basis for reconstruction of the dynamics of highly complex systems. But the research community faces a fundamental challenge in dealing with the volume and form of this data, in managing its storage, and in using it to build real understanding of the system being modelled. Deriving useful information and then knowledge from vast sets of data is prerequisite to development methods and techniques for solving real life problems in many domains.

The construction of long-term databases on complex systems must begin now and must be supported by methodologies and tools for managing and extracting information from them. Such databases must include information on the dynamic evolution of systems, not only on their structure. This means that data will be programs inferred from incomplete data and queries will act upon the semantics of the programs. Visualisation tools must be equipped with (stochastic) run-time support showing the execution of programs according to the quantities driving the interactions.

The information revolution has vastly improved our ability to gather, store and process data. As a result, contemporary science in many areas is struggling to deal with an overwhelming explosion of raw data. This is true in biology (in the so-called 'post genome' era) and in modern communications networks, as modern technology gathers data automatically, offering a rich storehouse of information on the dynamics of these systems. Progress on this challenge will offer Europe an ability to harvest this information explosion to the benefit of basic science, as well as in the service of important technological applications.

1.3 Research priorities

Research is required in order to:

- Describe systems at different levels of 'granularity' (accuracy), e.g., from a qualitative to quantitative level. Methodologies for robust computations on models with uncertainty are essential for breakthroughs in this area.
- Enable the extraction of meaning from heterogeneous, incomplete, inconsistent, partially faulty data sets. This involves development of concepts and computational tools that allow the systematic transition from data to information and then knowledge.
- Infer models to describe the dynamics of complex systems from incomplete and inconsistent data sets, and repositories for these descriptions. This will require development of experimental as well mathematical methods for the description of systems. An example would be the classification of topologies (spatio-temporal structures) in such systems.
- Enable automated maintenance of system models, drawing on data-mining of public databases of relevant information, or even from the open literature. This requires elaboration

and application of formal compositional description techniques so that new knowledge can simply be composed with existing knowledge in order to refine a behavioural model.

The suggested research requires contributions from a rich mix of disciplines - computer science, mathematics, physics, statistics and design of experiments, control theory, information, systems identification and communications systems. And, in addition, contributions would be required from the targeted disciplines of biology, the social sciences, and so on.

A first step would be to make all the data from the present information explosion available to the entire scientific community in an appropriate form. For this we need to develop a platform for achieving and guaranteeing the free sharing of data. This is a social challenge as well as a technical challenge.

2 Human-Responsive Systems

2.1 Challenge

The participants envisaged a new era of human-responsive systems that actively consider the user's perspective, whether designing, implementing or executing complex systems. Such systems will significantly augment human capabilities, rather than replace them with technology.

We can no longer treat ICT systems as separate from their human users and others affected by them. An ICT system and the human context in which it is deployed together form a system. By shifting the emphasis to increased human productivity, people will no longer be viewed as simply input to a complex system, but rather as the object of system development.

2.2 Rationale

On-going advances in technology have produced a stream of innovations, resulting in technology that is cheaper, better and faster. However, instead of a corresponding increase in human productivity, users are often over-whelmed and less productive. By recognizing that complex systems incorporate both human and technological components, we aim to ensure that such systems are responsive to people and that they increase human productivity.

Today's systems have reached the limits of current tools and techniques for supporting human-computer interaction. As we begin to explore the new possibilities offered by ubiquitous computing and mixed reality, which merge physical and virtual information, we will need to move beyond the desktop metaphor and take advantage of the full range of human sensory and motor capabilities.

Progress toward this vision will enable us to take optimal advantage of both human and technological strengths, enhancing human skills, such as judgment, handling unpredictability, social interaction and decision-making, while profiting from technological strengths, such as calculation and information management. The participation of human users in decision-making will then increase, resulting in increased overall effectiveness and greater satisfaction.

This human-responsive approach reinforces Europe's existing leadership role in participatory design, by demanding greater participation of users in the design, implementation and operation of complex systems.

2.3 Research priorities

We need new design principles that accommodate the changing needs and desires of human participants in complex systems, rather than treating them as outside the system, and rather than presuming to know what they need.

New processes, measures and techniques are required to ensure that the human element of complex systems is considered as an important aspect of design. Specifically, research is required in order to enable development of:

- Multi-disciplinary design processes for system designers to enable them to respond to human needs and capabilities when designing and controlling complex systems.
- New design patterns that allow complex systems to be responsive to the human participants within them.
- Technological tools and interaction techniques that increase human capacity for handling complex data and making effective decisions.
- Metrics for measuring the effectiveness of complex systems in both human and technological terms.

We also need better understanding of human behaviour in order to develop generative theories that should inform the organisation of architectures to support and sustain participation, including participative co-design of the systems themselves.

This research requires multi-disciplinary design teams who are able to cross boundaries and address design questions from various perspectives. The disciplines to contribute are the social sciences, including psychology, sociology and anthropology; computer science; software engineering; hardware engineering; and various design disciplines such as industrial engineering, graphic design and architecture.

Significant progress should be achieved within 5 to 10 years. However, in order to effect a paradigm shift that values human as well as technological performance within a system, we will need to make changes in the design process, in the interaction paradigms, and in the metrics used to determine success. For example, we need a measure of the simplicity of a solution relative to the complexity of the problem and a measure of effectiveness and productivity of human decision-makers.

3 Foundations for Software-Intensive System Development

3.1 Challenge

A fundamental breakthrough is sought for the scientific basis of computer science and software engineering technologies so as to capture evolution and dynamics, selfish interests of individual entities, various levels of bounded rationality, learning aspects and self-emerging behaviour, in a strict, yet tractable way.

In particular, we need to understand how the attributes of systems that are often regarded as ‘non-functional’ and that have typically been treated as emergent – such as security, safety, scalability, resource optimisation, quality of service, maintainability, flexibility, and efficiency - might be designed-in, and guaranteed by construction.

This would drastically improve the ability to meet industrial and consumer expectations regarding quality and functionality for software-intensive systems, and to reduce the effort required for their development.

3.2 Rationale

Software technologies are of crucial importance for the development of all advanced products and services in modern economies. The relative weight of software in the value of these products and services is constantly expanding to the point where it is a strong factor for differentiation and competitiveness. More than 59% of software in Europe is developed internally by enterprises. Software development is thus at the heart of European productivity and competitiveness, the fundamental Lisbon objectives.

Competitiveness in software intensive systems is starting to depend heavily on the capacity to produce high quality, complex software at a reasonable cost. This objective contrasts with the relative low level of the techniques currently in use for software development, debugging, and composition. A major step forward is urgently needed to meet the requirements arising from the following technological trends:

- Sensors, embedded systems (avionics, automotive, space), wireless telecommunication (mobile phones and wireless services), and pervasive computing.
- Large-scale distributed systems (Internet, with emphasis on security and resource awareness).
- Large-scale software systems and applications (telecommunications, air traffic control, enterprise operations, banking and e-commerce with emphasis on security, interoperability, and compositionality).

Maintaining and extending Europe's lead in some sectors such as embedded systems and telecommunications and services, improving the productivity of European industry (which, as mentioned, relies highly on internal software development), and acquiring the lead in other software-related sectors requires significant advancements in software development technologies in order to enable the production of very high-quality software with the least effort possible. Europe is in a unique position to make a contribution in this area by building on its strong traditions for foundational work and applying them to systems problems.

3.3 Research priorities

Research is required for:

- Guaranteeing non functional system properties (i.e., security, safety, scalability, resource optimisation, quality of service, etc).
- Devising new, high-level paradigms and languages for programming that encompass distribution, mobility, dynamic evolution, and that take into account non-functional properties.
- Developing new algorithmic techniques for distributed systems, taking into account non-functional properties.
- Conceiving and implementing modelling theory, methods, and tools for complex software systems, encompassing heterogeneous levels of abstraction.
- Property-aware compilation and implementation techniques, especially for non-functional properties, with automatic verification.
- Establishing the foundations for component-based development, especially theory, methods, and tools for building of software intensive systems that are correct by construction (compositionality, component-based techniques, modularity, etc.).

4 Control and Integrated Design of Complex Systems

4.1 Challenge

The aim is to develop modelling and computational tools that will enable understanding of emergence of complex types of behaviour in systems, and then permit the formulation of integrated design strategies and techniques for the synthesis of complex processes, the structuring of their organization (centralized-decentralized-hierarchical etc), the design of measurement and control structures, and finally their control design.

Such complex systems come from many diverse domains and disciplines: they may be the result of a design process; they may have been formed at a given instance; or they may be the result of an evolutionary process. They include traditional complex engineering processes, energy systems, financial systems, new technological processes (micro-, nano-systems etc), biological and social sciences, and management systems, including systems for the intelligent management of communications and computation applications (the Internet, mobile networks, etc.) It is important to note that complexity in such systems is not due to our inability to act upon them, but the result of their properties and history.

4.2 Rationale

For systems comprising many interacting elements we currently lack a framework of formal or mathematical techniques for modelling with which to predict the types of structures, the dynamic behaviour and properties that are likely to emerge at higher, collective levels of activity, or to clarify the conditions that will support such emergence. In other words, we lack a basic theoretical understanding of how properties of elements and element interactions at the micro level produce system properties at the macro and aggregate level.

Contemporary science has crossed an important threshold. In many areas, science has progressed to the point where we know the individual elements that make up important systems, as well as the nature of the interactions between those elements. This is true in materials science, micro-nano systems, for example, and increasingly in biology. It is also true for large scale engineering systems (power, chemical processes etc), communications and computational systems, where we have designed the components and engineered their connections, as well as integrated management problems (from a firm to networks of suppliers etc). But we now face the task of understanding how complex functions emerge in such collections of interacting elements and how to design or redesign such systems to achieve desired behaviours. Complexity is the next great barrier to technological progress and research to improve understanding and management of complexity is timely and challenging.

At the same time we now have the tools that enable researchers to attempt to breach this barrier. In particular, we are now able to carry out powerful and realistic simulations of complex systems. While simulations themselves cannot yield formal insight directly, they do make it possible to explore the behaviour of complex systems in an experimental way, complementing and verifying theoretical studies and thereby making fundamental progress more likely. Hence the emerging field of integration between simulation techniques and verification techniques could enable new insights towards the solution of the challenge.

Furthermore, theory for modelling dynamic, distributed and mobile complex systems is now well developed and can be exported to other areas: the major field of application seems to be that of biological systems or artificial bio-inspired systems.

Fundamental insight into the functioning of complex systems and the relevant modelling, simulation and computational tools will offer many economic benefits, as many economically important technologies must overcome the barrier of complexity in moving toward the next

generation. Furthermore, the emerging conceptual framework and tools will allow better management of resources and will enable us to control technological and biological processes, which sometimes risk becoming autonomous under the pressure of progress.

4.3 Research priorities

To meet this challenge, research is required to:

- Develop modelling tools based on logical and mathematical insights that will enable researchers and designers to understand the process of emergence and the nature of related system properties. They will then be able to explore their potential for tackling problems of integrated design (synthesis, design of measurement-control structures, organizational issues etc) and the consequences for stability, controllability, observability (reconstructibility) system efficiency (performance) and resilience, etc.
- Establish new computational paradigms and primitives to model and control complex systems. These should hide the complexity and provide tools to assess qualitative as well as quantitative properties affecting the structure and the behaviour of these systems.
- Inform the development of a sophisticated, freely available and easy to use platform for the computational simulation of complex systems and tools for evaluating the emergent properties and carrying out design.

Achieving these breakthroughs will require expertise from mathematics, operational research, computer science, statistics, physics, control theory and information theory. Knowledge from particular application fields is also required. In particular, expertise in biology, economics, management and social science will be required so that the resulting insight might be applied to important applications.

To realise this vision, participants stressed that it is best to focus first on relatively ‘simple’ complex systems, (that is, systems for which we understand in detail the nature of the individual elements and the interactions between them) and then specify clusters of generic (discipline independent) problems such as problems emerging from technological processes, before moving to biological, man-machine, and social systems. The idea is to first tune the formal framework over known systems so as to reproduce known behaviour, and then use it in a predictive fashion to control and design much more complex systems.

The workshop thought that significant progress on this challenge could be achieved within 5 to 10 years. The nature of application areas influences the rate of progress, since many of the required tools will be domain-specific.

5 M3P – Many Level, Many Unit, Many Purpose Systems

5.1 Challenge

The participants proposed a new paradigm for computational systems that draws inspiration from complex natural, social, technological and economic systems, and on existing modelling work (swarm computing, evolutionary computing, cellular automata, formal modelling methods for hybrid and heterogeneous processes etc.). It does not however limit itself either to studying the factors that make these systems successful or to simply exploring what emerges from such systems (as in *ALife*, for instance). Rather it aims to define engineering guidelines that will enable us to find cost-effective solutions to problems that cannot be solved with current techniques.

Examples include:

- Enterprise systems, value-chain networks, telecommunications networks, energy distribution networks, traffic networks, air traffic control systems, and complex industrial plant that are resilient, adaptive systems in the face of unpredictable failures and attacks.
- Naturally resilient, self-repairing systems for peer-to-peer telecommunications, data sharing and knowledge management.
- Large scale simulation models for biological, social, economic and artificial systems and support systems for decision makers operating with or inside such systems.
- Technological support for effective decision-making by competing and collaborating agents in political systems (e-democracy), industrial companies (e-management and value-chain management) and other large organizations (e.g. the internet).
- Mission-critical emergency warning systems including systems for ecological, epidemiological, space, industrial and security applications.
- Knowledge capture (growable) and data-mining in very large information spaces where individual ‘agents’ obey a central purpose but have access to only part of the relevant data.
- Novel approaches to ‘hard’ computing problems (e.g., weather forecasting, molecular modelling, hard optimisation problems, and scheduling problems) based on very large numbers of interacting units (perhaps confined in a small physical space – ‘a computer in a glass’).
- Evolutionary computing systems that self-organise into subpopulations and sub-layers.

5.2 Rationale

A key characteristic of many systems is that they involve components with uncertain operation and interactions which are at least partially unpredictable or hard to formalize in algorithmic terms (e.g., complex interactions with human users, interactions with the physical world, interactions with human-designed complex systems). Systems should be adaptable, malleable and resilient in the face of such unpredictability.

A number of biological, social, management and economic systems exhibit these desired properties. These systems are made up of many units, have a given organisational form of management and information structures that can be seen as performing ‘computations’ with local - and possibly global - performance criteria. These systems solve a class of problems that resist attacks by traditional mathematical methods and computing techniques. At present we do not know how they do it.

Although many natural, social and economic systems have been modelled by researchers (e.g., in *ALife* modelling) and a large research community has dedicated itself to the construction of Multi-agent Systems, the majority of modelling approaches suffer from one or more of the following defects:

- The number of units or agents (usually hundreds or thousands) in the system is many orders of magnitude smaller than the number of units (millions, billions, trillions) in the natural, social or economic processes the system is attempting to emulate.
- Models and systems are often organised at a single level. Even where systems do consist of multiple levels, these are deliberately designed. Only very rarely do they emerge naturally

from lower-level interactions as occurs in natural, social and economic systems. Emergent behaviour is critically dependent on architecture in a way which is not considered at present.

- In many models, such as those used in multi-agent systems, the knowledge they contain is explicitly programmed. As a general rule, current models have no way of incorporating new knowledge during their evolution or development.
- Many models (such as those produced by *ALife* modelling) have no “purpose”: while this kind of model may well exhibit interesting behaviour this has no ‘use’; there is no way to program the model, even when we define programming in the broadest possible sense.
- System models and design in the past have been done without economic, business or social objectives, and this has resulted in inflexible systems that have to be constantly reprogrammed or rebuilt in order to satisfy newer objectives. The levels of programmability have always been an issue, and almost every ICT system implemented anywhere has to be reprogrammed to handle new kinds of applications, new entrants, and new economic demands.
- Most models in the past have focussed on homogenous systems, where agents or agent-like environments only function in homogenous environments in relative terms. In reality, systems are very heterogeneous, for example, ICT systems in large enterprises are multi-vendor, with very different behavioural characteristics internally, but externally have some level of programmability that could be utilized to create order.

However, there are already in circulation many concepts and techniques which could contribute to M3P – Multiple Agent Systems, complex system theory, analyses of specific complex systems, systems biology, agent-based economics, swarm computing, evolutionary computing, large scale P2P applications, distributed systems. It is time to attempt a synthesis.

The development of computing technology means it is now technologically feasible to produce very large numbers of the highly sophisticated components required for implementation of the paradigm and with the Internet it is possible to practically test concepts involving thousands, tens of thousands or even millions of computational units

Foreseeable parallel developments in nano- and bio-technology offer the prospect of novel possibilities for implementation.

Work in this area would provide European industry with a head-start in a new area of hardware and software that is likely to be of great importance. This kind of head-start, similar to the US head-start in personal computing or the European advantage in 2nd generation cellular telecommunications, can translate into long-term strategic advantage.

At the same time, the systems made possible by the M3P approach could facilitate decision-making at all levels of European society. This would be especially important if Europe took the lead in the adoption of such systems (as the US took the lead in take-up of previous generations of computing technology).

5.3 Research priorities

Several research priorities and breakthroughs are required to meet this challenge:

- A theory must be developed to guide the design of M3P systems by formulating a set of generic conditions under which a real many-units purposeful system has to operate (typically, limited computational communication and computational resources) and devising a methodology that allows optimisation of their elements under these constraints.

- Ways must be found to ‘steer’ systems comprising very large numbers (perhaps millions, billions or even trillions) of autonomous heterogeneous computational units that exhibit spontaneously emergent levels of organization so as to satisfy some objective (social, economic or business).
- Approaches must be found to design systems so that they are aware of the environment they are operating in: this awareness will provide better intelligence to react accordingly.
- Approaches must be found for optimization so as to enable such systems to perform under tight physical constraints (e.g. limitations on the complexity of individual computational units, limitations on communications among units etc.)
- A design paradigm for the design of evolutionary systems must be established. The purposes of M3P systems will not be completely expressed as an explicit program. Like natural systems, M3P systems will evolve and develop, and in the process they will automatically incorporate information from the environment. This ability will provide them with flexibility and malleability unheard of in current multiple-agent systems. In particular such systems might evolve their structure: at least partially they will be self-assembling, self-programming systems. Structural methodologies and methods for aggregating behaviour will also be required.
- The inspiration coming from biology and sociology must be formulated in rigorous terms and operationalised in methods and tools: metaphor is not enough.
- The large dimensionality of such systems and their heterogeneous nature requires new methodologies for reducing the designing burden to manageable levels. Design problem decomposition, sequencing of design, decentralised and hierarchical control will be challenging issues in this new paradigm.

This wide range of research will require contributions from mathematics and computer science, statistical physics, complexity science, information theory, control theory and optimisation, theoretical and experimental biology, economics, sociology, materials science, nano- and biotechnology. A multidisciplinary approach that goes beyond the traditional computer science is required where modelling, control theory, operational research and optimisation have to be deployed.

Evolving a new computational paradigm can take decades, as shown by the history of the current Von Neumann architecture. Nevertheless, participants at the workshop expect that results of practical significance can be obtained within 5-10 years.

Conclusion

The thematic session identified a clear need for radical thinking about how we should design complex systems, recognising that ‘complex’ means much more than merely ‘complicated’.

While the Holy Grail of complexity science is to understand how to engineer micro-properties so as to cause desired macro-behaviour, the first step – acknowledged in several of the independently devised ‘research challenges’ – is to harness the power of existing computational technology to simulate such complex systems, so as to gain insight into their operation, to inform and test man-made complex systems, and to explore the effects of man-made intervention in existing systems.

The ramifications are extensive. At one end of the scale, our whole approach to understanding our ‘requirements’ will require revision to move away from ‘specification’ of ‘the system to be

procured' (with, maybe, vague hopes for flexibility) to take proper, rigorous, account of the need for evolvability. At the other end of the scale, we may expect fundamental changes in the nature of computer science, as it moves from dealing with 'the bits' to dealing with the qualitative expectations for the effects of those bits.

The phrase 'paradigm change' is much over-used. But it is difficult to find an alternative characterisation of the outcome of this workshop. Europe must seize the opportunity – not shrink from the uncertainty that this new way of thinking will create.