Computing and Informatics, Vol. 28, 2009, 369–391

DISCIPLINED EXPLORATION OF EMERGENCE USING MULTI-AGENT SIMULATION FRAMEWORK

Ognen Paunovski

South East European Research Centre (SEERC)
17, Mitropoleos str.
54624 Thessaloniki, Greece
e-mail: ogpaunovski@seerc.org

George Eleftherakis

City College
13, Tsimiski str.
54624 Thessaloniki, Greece
e-mail: eleftherakis@city.academic.gr

Anthony J. Cowling

University of Sheffield
Regent Court, 211 Portobello str.
Sheffield, S1 4DP, UK
e-mail: a.cowling@dcs.shef.ac.uk

Manuscript received 10 July 2008; revised 10 November 2008
Communicated by Jozef Kelemen

Abstract. In recent years the concept of emergence has gained much attention as ICT systems have started exhibiting properties usually associated with complex systems. Although emergence creates many problems for engineering complex ICT systems by introducing undesired behaviour, it also offers many possibilities for advance in the area of adaptive self-organizing systems. However, at the moment the inability to predict and control emergent phenomena prevents us from exploring its full potential or avoiding problems in existing complex systems. Towards this end, this paper proposes a framework for empirical study of complex systems exhibiting emergence. The framework relies on agent-oriented modelling and simulation as a tool for examination of specific manifestations of emergence. The main idea is to use an iterative simulation process in order to build a coarse taxonomy of causal relationships between the micro- and macro layers. In addition to the detailed description of the framework, the paper also discusses the corresponding verification and validation processes as important factor for the success of such a study.

Keywords: Complex systems, emergence, multi-agent modelling, simulation.

Mathematics Subject Classification 2000: 68U20, 68M15, 68U01

1 INTRODUCTION

There is a variety of systems which are perceived as complex. In the natural world cells, immune systems, nervous systems, ant colonies and many others can be viewed as complex systems. Similarly in the human world a wide range of cultural and social systems like families, political parties, companies, scientific communities, economical markets and many others are also complex systems. However, in recent years the study of complex systems has gained particular interest in the field of computer science. The main reason for this trend is the increase in the complexity of Information and Communication Technology (ICT). ICT systems are becoming more complicated, open and distributed [6]. In fact it is not difficult to imagine a future where automatic composition of billions of elementary web services will form more complex services. However, this will lead to enormous amounts of unplanned and unregulated interactions. These interactions will undoubtedly cause appearance of unexpected emergent behaviours which are usually associated with complex systems.

The importance of emergent behaviours in ICT can be viewed from two aspects. From the engineering perspective mastering the control of emergent phenomena can be very useful. Emergence is responsible for self-organization, self-optimization, adaptation and other beneficial properties encountered in complex systems. The utilization of these emergent behaviours in an information system can benefit the development and performance of the system making it highly available, scalable and robust. On the other hand, however, perhaps a more important concern for computer scientists and software engineers is the appearance of undesired emergent behaviour or so called "misbehaviour". Emergent misbehaviour can be viewed as unexpected behaviour with undesired effect which infringes the system's operation by diminishing its functional performance or by introducing behaviour which is different from the intended one. An example presented in [29] shows that complete degradation of the service can occur in fairly simple multi-tiered distributed application due to a small increase database server latency. This and similar examples show that it is vital to prevent appearance of emergent misbehaviour in ICT systems. However, since the phenomenon of emergence is inevitably linked with the complexity of the information systems, it cannot be simply avoided. Consequently there is a need to devise means for the development of correct systems which will guarantee (to some extent) that there will be no undesired emergents at runtime.

Nevertheless, given the stochastic nature of emergence it is practically infeasible to formally verify the appearance of emergent behaviour [35]. Therefore in recent years agent-oriented modelling and simulation have been suggested as a tool which can shed light on the problem [12, 13, 14, 23]. The idea is to model the components of a complex system as agents and use them in a simulation study. Nevertheless, so far there is no study which deals with the practicalities of constructing a disciplined way of analysing emergent formations. In this paper we address this issue by proposing a structured two-phase framework for empirical exploration of emergent behaviour through multi-agent modelling and simulation. The initial phase of the methodology addresses the verification and validation of the multi-agent model, while the second phase is an experimental process aimed at determining the causal relations between the micro level interaction and the visible effects of emergence at the macro level. The end goal of this process is to address the problem of analysing emergent behaviour in complex systems through a structured set of well defined activities and practices.

The rest of the paper is structured as follows. Section 2 offers an introductory discussion on the phenomenon and types of emergence, followed by a discussion on the different ways in which emergence is conceptualized in Section 3. The application of the multi-agent paradigm for modelling of complex systems is elaborated in Section 4. Section 5 provides a detailed overview of the proposed methodology followed by a discussion. Finally conclusions and future work are summarized in Section 6.

2 EMERGENCE AND TYPES OF EMERGENCE

The term "emergence" (from Latin "emergere") means "to become apparent", "to turn up", "present itself", "to appear" (Oxford dictionary). The whole idea behind emergence was popularized by Anderson in [1], where he argued that simple component interactions can give rise to complex phenomena which is more than a sum of the properties exhibited by the elementary components. A simple example of this is the liquidity in water. There is nothing to suggest, in a single H_2O molecule, that millions of molecules at room temperature have the properties of a fluid. Another example is the interaction between local weather patterns, which influences emergent formations like hurricanes, tornadoes, temperature inversions and other weather phenomena. Similarly stock market crashes can be viewed as emergent phenomena based on the interaction between traders on the stock market. The phenomenon of consciousness in the human brain follows the same basic principle of emergence. While one neuron is a relatively simple entity, the collective interplay of millions of neurons in the human brain can result in something much more than a simple sum of the neuron's properties and behaviour.

If we examine these examples in more detail, we can come to the conclusion that emergence describes a system where a global phenomenon arises from the local interactions between the individual (micro level) components of the system. However, due to the diversity and complexity of emergent phenomena, in natural as well as social systems, different sciences have focused on different aspects in the investigation of emergence. Consequently there is a variety of definitions [3, 4, 7, 19, 20, 23, 25, 34] used to describe emergence, but none is generally accepted. Nevertheless, in order to continue the discussion on emergence, there is a need for a working definition within the field of computer science. The authors of this paper adopted the definition proposed by Wolf and Holvoet in [12], where they view emergence as part of the system:

> "...when there are coherent emergents at the macro-level that dynamically arise from the interactions between the parts at the micro-level. Such emergents are novel w.r.t. the individual parts of the system."

In this context "coherent emergents" denotes orderly (logically or aesthetically) consistent effects (properties, behaviours, structures, patterns) which are product of the process of emergence at the macro (system) level, caused by interactions at the micro (individual, elementary) level. While the definition describes the basic principle behind emergence, it gives almost no insight in the particular manifestations of emergence. Therefore in order to understand the forms of emergence there is a need to differentiate and classify different types of emergent phenomena.

Very often the concepts of weak and strong emergence are used in order to differentiate between emergent phenomena [4, 8]. However, this classification, although relevant for philosophical discussions, is too general to be useful in the field of computer science. There is a need for more clearly defined classification structure like the one proposed by Fromm in [18], where he builds upon the classification for cellular automata proposed by Wolfram [37]. He distinguishes four primary classes (types I-IV) based on the causal relations of the phenomena. Furthermore, the classification follows a gradation in complexity. Class I contains the simplest emergent phenomena with a single feed-forward relation which can be found in engineered systems (e.g. the intentional design of a machine like a clock, computer program and so on) and systems exhibiting aggregated emergent phenomena (e.g. wave fronts in water, avalanches, cascades). Classes III and IV have the highest level of complexity. Class III phenomena have multiple feedbacks, both positive and negative. This type is common in open systems with high complexity and it is usually associated with activator-inhibitor systems (e.g. patterns in biological entities, stock market rush, prisoners dilemma) as well as evolutionary and adaptive systems (e.g. evolution of ecosystems, sudden scientific or mental revolutions and so on). Class IV, on the other hand, contains the emergence of completely new complex systems (e.g. culture, life). From an engineering perspective a particularly interesting case is type II emergence which encompasses systems exhibiting self-organization and other useful properties (type IIA), as well as emergent phenomena which are based on imitation and self-amplification (type IIB). The latter sub-type is responsible for the so called negative emergents like crashes and bubbles in the stock market, explosions of social unrest, buzz in the news and so on. In principle these are the kinds of phenomena which are perceived as misbehaviour in ICT systems. Consequently the framework proposed in this paper is primarily concentrated on emergent phenomena of this type (II).

3 DIFFERENT PERSPECTIVES ON EMERGENCE

Although emergence is one of the key issues in complexity research, the phenomenon is also a very active research topic in various scientific disciplines as well as a theme of philosophical discussions. In addition to exploring different aspects of emergence, different studies have adopted different views on the nature of the phenomenon, its relationship to the system and the observer, the structure of causal relations which influence it, the hierarchy of systems exhibiting it and so on. In this section we make an attempt to capture several perspectives on the nature and core ideas behind the concept of emergence as presented in different publications addressing these issues from computer science perspective. We start with subjective approaches and move towards more objective ones.

One very intriguing view on emergence is put forward by Bonabeau and Dessalles in [5]. In their view the key in understanding emergence is the observer rather than the system itself. Their argument is based on the notion that an emergent phenomenon can only be defined while the system is being observed. Consequently, their primary focus is on the relation between the observer (as a detector) and the system (being observed). They argue that the emergent patterns appear only when the observation takes place at the right level in the system's hierarchy. For example observing a city by walking through its streets cannot give an indication about its fractal structure which is visible in a picture taken from a satellite. This idea is closely linked with the so called "emergence of higher structures" theories, where the phenomenon of emergence is viewed as a property (pattern or behaviour) taking place at level L_h which is a result of processes taking place at a lower level L_1 and when the properties exhibited at L_h are impossible (or "difficult") to explain given the behaviour at L_1. Another aspect of the system-observer relation, addressed in [5], is the nature (capabilities) of the observer and the observation tools used. In this context, the authors do not consider emergence as an absolute property of the system, but as relative to the observer. This idea is supported by situations where a phenomenon that cannot be described or understood using a certain set of observers (and observation tools) can be detected and understood with the introduction of additional observers and/or observation tools. This implies that by introducing additional observers the description of the system can be simplified. Consequently, there is a perceived complexity shift in the observed system, which leads to the conclusion that "emergence is associated with a decrease in the relative complexity" [5].

The view put forward by Bonabeau and Dessalles is actually a "highly subjective" approach to emergence, where the main focus is on the observer and observation apparatus as main factors in the detection and definition of the emergent phenomena. Taken to the extreme this perspective on emergence correlates with Epstein's view expressed in [16] where he says:

> "To call something emergent is therefore not to say anything about the property at all, but merely to make a confession of scientific and mathematical incompetence."

In order to fully appreciate the depth of Epstein's statement, one may consider it through a presence of an ultimate observer which embodies all of the available scientific knowledge. In relation to this (hypothetical) all-encompassing observer, Epstein labels the concept of emergence as a "scientific mystery" (or a scientific unknown) which exists purely on the basis of the observer's limitations. In other words, according to Epstein, emergence exists due to the (current) unavailability of scientific mechanisms to detect, predict and control causal relations in certain complicated systems. Thus an emergent phenomenon as such is temporarily limited with the advance of the scientific mechanisms which will be able to resolve it. Although Bonabeau and Dessalles do not tackle this issue explicitly (since they do not assume an all-encompassing observer), the same reasoning holds because, like Epstein, they too perceive the observer (and/or the observing mechanism) as a dynamic entity which is subject to change.

However, in the case where a particular emergent manifestation is examined within a well defined observation setting, which is usually the case for practical studies of complex systems in the field of computer science, the discussed implications of the dynamic observer lose their relevance. Thus, in such an environment the emphasis is shifted from the observer to the behaviour perceived in the model. The basic idea behind this so called "emergence relative to model" is expressed by Rosen in [32]:

> "One way to define emergence is to call a behaviour of a system emergent when it can no longer be described by the model that described the system until then."

Building on this idea an approach proposed by Ronald et al. in [31] aims at devising a practical test for emergence based on three main elements: design, observation and surprise. In the proposed setting, emergence appears as a surprise from observer's point of view when the local behaviour of the system described in language L_1 during the systems's design is different from the language L_2 in which the global behaviour of the system is described at runtime and there is an non-obvious causal relation between the two languages. In this sense the element of surprise is basically a discrepancy between the intended and the actual (observed) behaviour of the system. In other words there is a "cognitive dissonance" between the observer's mental image of the system at design time and the observed behaviour of the system at runtime. The rationale for the appearance of the surprise, in Havel's view [22], can be explained as a gap in the observer's ability to interpret relations in the particular domain due to lack of understanding of the causal laws in that domain. Consequently, one might argue that an increase in the observer's understanding of a particular problem domain might lead to shift in the observer's domain horizon which will overcome the element of surprise in the particular context. This is in fact supported in [31] through an analysis of the flocking behaviour. Authors consider the flocking behaviour in birds to be an emergent behaviour when it first appeared in Reynolds's flocking simulations (see [30]) done in 1987. However, they argue that since then the flocking behaviour lost the element of surprise because the causal relations in the model have become well known. Thus, similarly to the view adopted by Bonabeau and Dessalles, the work of Ronald et al. [31] follows a "subjective" approach to emergence by signifying the role of the observer. Nevertheless, the view proposed by Ronald et al. incorporates pre-conceptions about the system's design, thus avoiding to rely only on observation as in the case of Bonabeau and Dessalles. This represents a significant shift in conceptualization of emergence, but at the same time limits the applicability of the approach to artificial systems.

Although the test for emergence proposed in [31] gives a vague description (from formal perspective), the work done by Freund et al. [17] shows that the basic concept guiding the emergence test can be practically applied in grammar systems. The authors illustrate this in [17] with an example where the simple sum of two finite grammars produces a finite language when applied in isolation, but at the same time generates an infinite language when the rewriting rules of the two grammars are applied in combination. In this case the surprise is due to the gap between the finite and the infinite behaviour exhibited by a system composed out of the same basic components. In this direction the authors formalize a theoretical framework for emergence in systems which can be described using string languages. In addition they also discuss how the defined formalism of emergence applies to various types of grammar systems.

Another attempt to define emergence in relation to grammar systems has been made by Kubik in [27]. Nevertheless, he completely refuses to acknowledge the observer's surprise as relevant in defining emergence:

> "...judging the behavior of complex systems on the basis of our subjective feeling of surprise is misleading and obscures better explanations."

Instead he combines Rosen's idea of emergence relative to model, with Bedau's notions of micro-macro relations [4] as primary postulates on which emergence should be formally defined. Furthermore, he supports the multi-agent approach to emergence advocated by Holland [23], where the behavioural rules of the agents and their mutual interaction are viewed as a way of describing the relations in a complex system exhibiting emergence. In this sense Kubik's view in conceptualizing emergence is the same as the view advocated by supporters of the multi-agent modelling and simulation approach. Nevertheless, his main goal is to move towards a more formal theory with well defined meaning. He proposes array grammar for this task, based on the notion that different grammars can be used to formally specify a multi-agent system [10]. Within this framework, Kubik defines emergence as an array obtained when isometric production rewriting rules are applied on the entire language. This is different from the sum obtained after the rewriting rules have been applied separately on all language parts, thus formally describing systems behaviour which is more than the sum of its parts. The major criticism on this approach is that without a restriction in the rewriting rules in some cases it might be impossible to determine whether a particular array is emergent. Nevertheless, this approach manages to avoid dealing with the subjectivity of observation and moves much closer to a concept where emergence is viewed as a property of the system. This is in fact the core idea behind modelling and simulation approaches to emergence, which is elaborated in more detail in the next section.

4 MULTI-AGENT MODELLING OF COMPLEX SYSTEMS EXHIBITING EMERGENCE

Many researchers [3, 9, 13, 14, 19, 23] dealing with the problems posed by emergence agree that an initial approach in understanding emergence should be done through modelling and simulation. One of the main reasons supporting this is based on the notion that emergence is "more than the sum of the parts". The main implication of this idea is that emergence can not be captured with a model of the system. According to this point of view an emergent phenomenon is only visible at runtime operation or through an animation of the system. This is due to the novelty introduced at the system level which cannot be deduced from the properties of the individual components. Consequently, the developer of the model cannot simply design a synthesis rule, but only use simulation as means of achieving the micro-to-macro aggregation. From a computational point of view, as defined by Darley [11], this means that a system of size n is emergent if the condition $u(n) \geq s(n)$ holds, where $s(n)$ is the optimal amount of computation for predicting the system's behaviour through a simulation and $u(n)$ is the amount of computation required in order to resolve the problem using some form of "creative analysis." In the manner described, the modelling and simulation approach adopts an objective approach in dealing with emergence.

In this context, the Multi-Agent Systems (MAS) paradigm is considered (at the moment) to be the most suitable way to model and simulate complex systems exhibiting emergence [14, 23]. There are several reasons behind such a claim. First of all there is a natural correspondence between the structure of complex systems and MAS. They both rely on many individual components (agents) in order to achieve their goals. Each agent in MAS is autonomous and able to interact in a stochastic manner with other agents. Moreover, there is no limitation on the interaction scenarios, which means that an agent is able to communicate indirectly on multiple levels by modifying the local environment, which arguably is the most common approach used for communication in natural complex systems. For example, ants communicate indirectly to each other by dropping pheromone which modifies the environment and serves as a guide towards food. Another correspondence between MAS and complex systems is the level of complexity. A MAS can achieve almost any level of complexity and thus a multi-agent system of specific complexity is essentially a complex system.

Nevertheless, this does not mean that any MAS is complex or possesses emergent behaviour by default. There are several properties common in MAS exhibiting emergence [24]:

Agent mobility or visible states for fixed systems – e.g. spatial repositioning for mobile agents.

Ability to influence the environment – e.g. chemotaxis, self-replication or other approaches for modifying the environment.

Ability to distinguish between groups and individuals – e.g. flocking of birds as a model composed of individual agents and groups (flocks).

In addition, it can be argued that the basic multi-agent modelling concepts – abstraction, decomposition and organization [24] correspond to the modelling of complex systems. The first notion is the ability to *abstract*. Basically, it denotes the ability of the agent models to simplify the representation of the system by hiding unnecessary complexity. Some properties of the model are emphasized while others are suppressed. This is very important when dealing with a system whose entities may be complex systems themselves. Otherwise it would be very difficult (if not impossible) to develop the complete model, at the same time it may require considerably more time and effort.

The second concept is *decomposition*. The idea behind it is to divide a complex problem into several smaller, more manageable components. Thus, each component could be examined and analysed in relative isolation. Nevertheless, the application of decomposition in systems exhibiting emergence is a very delicate issue. Decomposition of a system might diminish the emergent phenomena. For example, dividing a living entity into parts could result in a bunch of dead pieces. This is because "life", as an emergent phenomenon, relies on interaction between different components in the system and might not be a property of the component when examined in isolation. Nevertheless, this does not mean that decomposition is useless in the case of emergence, but rather that it is important to devise an appropriate decomposition strategy which will not influence the phenomenon.

The third concept discussed by Jennings and Woodridge in [24] is *organization*. Although the development of an individual agent is a relatively straightforward process, the intentional design of organizational structures formed by agents may be extremely difficult. This is primarily due to the dynamicity and unpredictability of interaction patterns within the system. This unpredictability of interaction creates problems from an engineering perspective since it diminishes the predictability of the system's behaviour. However, in the case where the MAS model is used in a simulation study, the uncertainty of the runtime dynamics is neither a disadvantage nor a problem for that matter. In fact, it offers the unique possibility to the investigator to gain insight of the possible behaviours that might occur in the modelled system.

Although the discussed principles and natural correspondence to complex systems make the MAS paradigm an intuitive approach for modelling complex systems exhibiting emergence, developing a MAS model of the system is not sufficient to understand the causality of a particular emergent manifestation. For this purpose there is a need for adopting a structured approach which will guide the animation of the model through an iterative experimental process.

5 FRAMEWORK FOR EXPLORING EMERGENCE

The following sections present a framework for a disciplined exploration of emergent phenomena using multi-agent simulation. The fundamental idea is based on an incremental increase in the understanding of the causal relations in the model under study. Section 5.1 presents a general overview of the two phases in the proposed framework, while a more detailed description of the activities and artefacts in each phase is elaborated in Section 5.2. The model verification and validation process is discussed in Section 5.3 followed by a discussion on the experimental phase in Section 5.4.

5.1 General Overview

The main object of study in the proposed framework is a multi-agent model of the system exhibiting emergent phenomena. The animation (simulation) of this model represents the main instrument to achieve the synthesis of the elementary behaviours and interactions to a macro level emergent behaviour. In addition, the simulation process is also used in order to generate the data needed for the analysis and delineation of the causes for the observable emergent effects. The overall process can be divided into two phases: a development and verification phase and an experimental phase.

- The **development and verification phase** is an initial phase which encompasses the development and verification of the model in respect to the expected behaviour. The role of the investigator, in this phase, is to make iterative refinements to the model in order to bring it closer to the desired state, thus continuously increasing the confidence in the behaviour of the model throughout the refinement. The refinement and evaluation of the model are an important part in the development because they provide an alternative to formal verification and validation of the model, which cannot be achieved for a dynamic complex systems with non-deterministic interaction [13, 35]. Once the model is validated the process can move to the second phase.
- The second stage of the investigation is an **experimental phase**. It is essentially an analytical process aimed at detection of invariants, interaction patterns,

local properties and other elements that can influence the macro effects of emergence. The goal is to define the possible causal relations which impact the observable emergents. The process relies on the investigator to form a testable hypothesis which will be evaluated in the next iteration. The observation of the model execution and the analysis of the gathered data should prove or disprove the hypothesis, thus increasing the understanding of the specific emergent phenomenon.

A phase, as defined in the framework, can contain several iterative cycles. Each cycle is composed out of predefined arrangement of activities (steps) and transitions. An activity defines the tasks (operations) which need be performed at a particular point, while the transitions define the output from one activity and the input to the next one.

5.2 Activities and Transitions

Figure 1 presents an overview of the activities and transitions envisioned in a cycle. The first two activities are performed only in the initial cycle of the verification and validation phase, since they deal with the initial system description and model specification. The activities 3-8 represent the elements of a single iteration. The description of the activities, transitions and the corresponding artefacts that follows is discussed through a herd dynamics case study example based on [21].

Activity 1: The initial transition is a jump from the real system to a corresponding theoretical description. Depending on the case study the most appropriate person to perform this activity is an expert in the field of study. For example when dealing with a natural system, like the herd dynamics case, the theoretical description should be performed by a biologist.

Activity 2: The second activity relies on the analysis of the system description which is provided in the initial step. The main goal in this activity is to derive a specification of the model properties. In this context the overall specification of the model should be captured in the following documents:

- The **agent type specification document** containing the description for each type (if more than one) of agent in the model. In the case of herd formation the document should provide the specification of the properties of an individual animal.
- The **environment and communication document** should describe the properties of the environment and the forms of interaction and communication in the model. This document needs to provide details of all explicit communication protocols (if any) and at the same time it should describe any indirect (or implicit) way of communication between the agents as well as the interaction between an agent and the environment. In the herd scenario example the communication is achieved through modification of the environ-

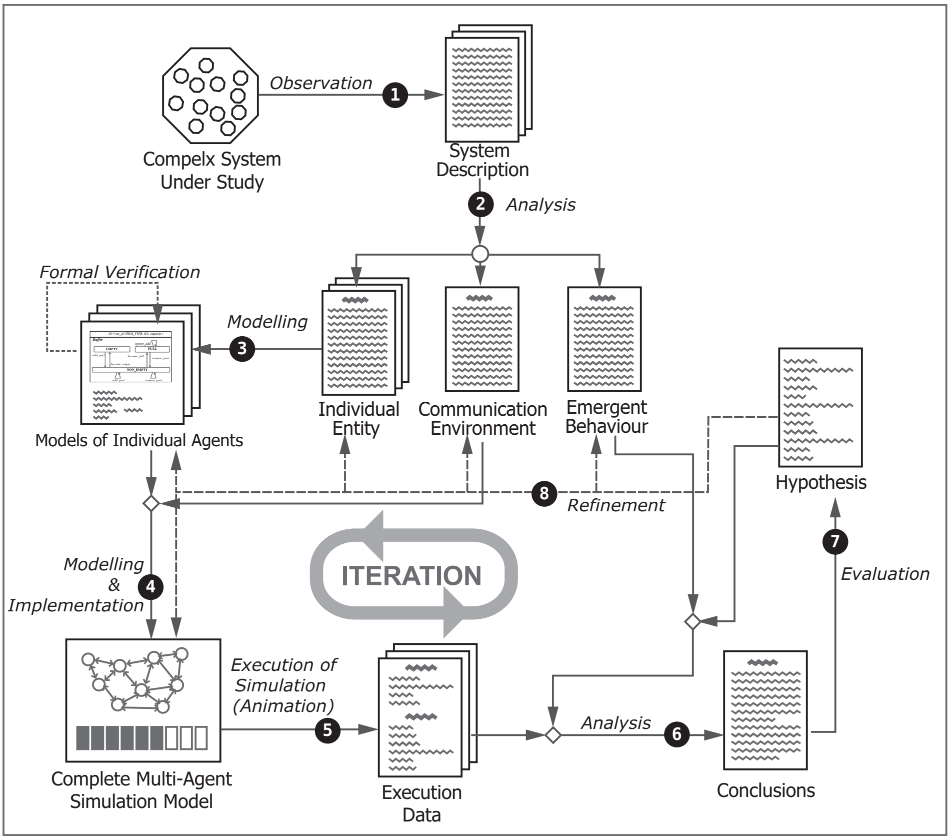

Fig. 1. Artifacts and transitions in the proposed methodology

ment by repositioning of the animals. The animal perceives the position of its neighbours and decides where to move next.

- The **emergence specification document** should focus on a clear (quantifiable, if possible) definition of the visible emergents which are expected to appear in the system. The characteristics of the emergent phenomenon could be defined through macroscopic variables as proposed in [12] or using other indicators. In the herd dynamics example (where the herd is the actual emergent phenomenon) such an indicator could be the level of herd cohesion in different groups.

Activity 3: Based on the specification created during activity 2, an appropriate model for each type of agent needs to be developed. A good practice is to use formal modelling language, in order to be able to verify the properties of the model and determine that there are no undesired discrepancies and errors in the models.

Activity 4: The model(s) of the individual agents should be combined with the appropriate representations (models) of the environment and communication in order to form the complete multi-agent model of the system. An important issue in this activity is the representation of the relation between the individual agent and the environment, which is supposed to enable non-deterministic multilevel interaction.

Activity 5: At this point the complete model should be transformed (implemented) in an environment which should allow animation of the model. The selection of a simulation environment may vary depending on the system under study. A visual animation of the model may prove very useful in cases like the herd dynamics scenario, where visual animation of the model can give insight into the herd formation process. In addition all of the data required for the analysis needs to be generated during the animation process. Therefore it is essential that parameters which define the model's behaviour, both at micro and macro levels, are recorded in sufficient details for later analysis. The data should present two views on the executed simulation.

- First containing quantitative measurements on the state of the micro-level at a particular time frame, offering traceability in terms of continuous model execution.
- The second type of data should provide insight into the properties of the macroscopic level. The most suitable way is to use global variables (preferably gradual rather than binary) which will describe aspect of the emergent phenomenon at a specific point in the execution of the model.

Activity 6: This step involves analysis of the gathered simulation data. The goal of the data analysis process is different in the two phases.

- In the initial phase (development and verification) the analysis of the data should be done in order to evaluate the current state of the model in respect to the desired model. Invariants in the data can be used in order to trace possible errors in the model.
- In the experimental phase, the analysis of the gathered data should be compared with the expected outcomes in order to test the hypothesis. The analysis should focus on evaluation of the stability of emergents in the model and detection of patterns by means of statistical and correlation analysis as well as detection of invariants. Based on the findings the investigator should derive conclusions about the behaviour of the model in the particular simulation run.

Activity 7: This activity involves evaluation of the findings in the previous step.

- In the development and verification phase, the goal of this activity is to identify elements of the model which cause discrepancies in the model behaviour.

Based on this the investigator can propose changes to the model in order to bring it closer to the expected behaviour.

- The main focus in the experimental phase is on the formulation of a hypothesis which will be examined in the next iteration. The hypothesis should be empirically testable in respect to the model. In addition, the investigator should define criteria according to which the hypothesis will be confirmed or refuted.

Activity 8: The final activity of the cycle can be viewed as an initial stage of the next iteration. It involves modification (refinement) of the model in order to refine the behaviour of the model or test a hypothesis. As can be seen in Figure 1, the modification can be done on all aspects of the model including the individual agent, the environment and communication. Additionally changes in the data reporting routines of the simulation environment may be needed in order to gather additional data.

The discussed activities represent a single cycle in an iterative process. The number of iterations is not fixed, but it should be sufficient to confirm or refute the hypothesis being examined. In this manner the proposed framework adopts an iterative experimental approach to exploring emergence in existing systems.

However, since the object of the study is a model of the real system rather than the system itself, there is a major concern that the developed model might not possess sufficient details, i.e. it might omit important factors or make wrong assumptions about the system. In this case the model would be potentially useless in respect to the goals of the study. This is the main reason why the verification and validation of the model is a very important part for the success of the study.

5.3 Verification and Validation of the Model

The verification and validation of the model is incorporated in the initial phase of the proposed framework. The main goal in this phase is to minimize the potential discrepancies between the real system and the model. In respect to this, several issues concerning validation and verification need to be addressed.

Figure 2, based on Sargent's work in [33], presents an overview of the modelling process and the associated verification and validation steps. The problem entity denotes the "real" system which is the object of the study, while the conceptual model is an abstract representation of the system which is developed during the modelling process. The validation of the conceptual model should determine whether the model corresponds to the real system for the intended purpose. Although the practical details of various validation techniques are beyond the scope of this discussion (for more information see [33]), there are several issues that need to be considered. In this context perhaps the most important issue is the selection of a representation technique. There are several aspects that need to be taken into account when selecting a representation technique [2]:

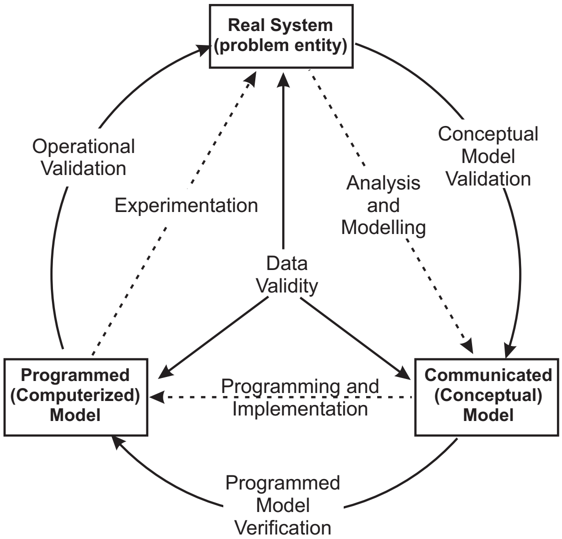

Fig. 2. Overview of model verification and validation process, taken from [33]

- The expressive power of the representation technique. Can the model be fully captured with the particular technique?
- The technical knowledge of the people to whom this model will be communicated. Other people, involved in the study, should be able to understand the notation used to describe the model.
- The application of formal analysis and verification of the model in the particular form.
- Automation of the transformation from communicated to programmed model. How can the conceptual model be transformed (implemented) into a form which can be animated and used in a simulation study?

Given the variety of representation techniques, there is no generic solution to all of these problems. The success of a particular representation schema depends on the system under study as well as the environment and people involved in the study. Therefore we avoid proposing a particular technique as part of the framework. Nevertheless, we strongly believe that a formal representation should be used in order to develop the model of the individual agent. By using formal methods the investigator can verify that the model of the individual agent is correct and can therefore focus its attention on the validation of the communication and interaction between the agents. However, not all formal methods are suitable for developing models of agents. In general, in order for a formal method to be useful for modelling agents, it should satisfy the following criteria [15]:

- to model both the data and the internal changes of an agent,
- to model separately the behaviours of an agent and the ways in which the behaviours interact with each other,
- to be intuitive, practical and effective towards the implementation of an agent,
- to facilitate the development of correct agents.

While using formal techniques may be applicable to an individual agent, it is practically infeasible to formally verify a complex multi-agent system with stochastic interaction [13]. Consequently the only viable solution is to validate the interacting multi-agent model by means of simulation. However, in order to do so the conceptual model needs to be implemented (programmed) in a form which can be animated. The transition from the conceptual to the programmed model is the second step in the validation and verification process presented in Figure 2. The primary concern in this step is to ensure correctness and correspondence of the implementation to the conceptual model design. In the best case this transformation could be done using an automated tool (which has been extensively tested) in order to ensure correctness. However, depending on the problem and the representation technique used, such a tool might not be available. In this case the transformation needs to be performed manually. In this context a variety of verification techniques which are used in software engineering could be applied [26]. Detailed analysis of various testing techniques is beyond the scope of this discussion; for a comprehensive overview see [36].

At a point when the transition from the conceptual to the programmed model is complete, the investigator has the means to validate the complete multi-agent model. This is the final verification and validation step in Figure 2, labelled operational validity, which aims at validating the correspondence between the behaviour exhibited by a programmed model and the real system. The final behaviour of the programmed model must have a reasonable accuracy in respect to the real system and exhibit the required emergent phenomenon. According to Sargent [33] a general division of the operational validation process suggests two main types:

Objective approaches, which usually rely on statistical or mathematical proof of the correspondence between the model. This approach is much more demanding in terms of time and effort, but it has a higher credibility (compared to the second type). Nevertheless, it is not always applicable, especially in the case of complex systems.

Subjective approaches can rely on different techniques to observe and to some extent evaluate the model. However, the final decision whether the model is appropriate representation of the real system is made by the development team in the way they see fit. These kinds of approaches are usually used to ensure operational validity of complex systems, since in most cases it is practically impossible to perform this process in a formal way.

The idea that we propagate through the framework is to use mathematical validation where applicable; however, since this is rarely the case with complex systems, we suggest an iterative process which will enable the investigator to gradually build confidence in the behaviour of the model.

5.4 Two-Way Experimental Approach

The experimental phase follows the development and verification phase and commences once the investigator has assessed that the model as correct (valid for the purpose of the study). As previously discussed, the main aim in this stage is to explore the causal relations between the micro- and the macro levels through an iterative experimental approach. In each experiment the role of the investigator is to formulate a hypothesis and then test it by executing an appropriate simulation. In order to do so, the simulation conditions or even the model itself can be modified. This process in essence resembles the so called "general scientific method", where by testing a hypothesis the knowledge about the system is gradually increased.

The examination of the simulation execution in the proposed framework is addressed in a manner similar to the two-way approach for investigation of emergence as proposed by Conte and Castelfranchi in [9] as well as the experimental method proposed by Edmonds and Bryson in [14]. Consequently the process encompasses both bottom-up and top-down processes in order to analyse micro-macro connections.

The **bottom-up process** facilitates attaining collective behaviour from the individual agents. It offers insight on how their behaviour is combined and aggregated. The role of the investigator during this process is to identify interactions that have an immediately visible result at the system level. In addition, through observation of the model animation the investigator has the opportunity to gain insight into the behaviour and stability of the emergent phenomena. This knowledge can be very useful during the analysis and evaluation.

The **top-down process** is concentrated on analysis of the data gathered during the simulation. The analysis process should address several issues:

Behaviours of micro entities. Deduce the behaviour of the individual element (agent) from the global behaviour of the system. This includes identification of how the micro elements behave at a given time instance, what behaviour should be visible at the micro level given the behaviour of the overall system, as well as definition of how the behaviour of the system imposes restrictions on the behaviour of an agent (top-down feedback).

Behaviour of emergent phenomenon. Define a set of global variables which indicate different aspects of the observed phenomenon. Where applicable avoid specifying binary variables. If possible devise a metric for each of the variables. Define values for the variables for each time instance. Compare changes in the variable values (if more than one variable). Correlate the changes in a variable with micro level events.

Associate roles and states. Identify the possible roles and role transitions for the micro level entities (agents) in terms of responsibilities, permissions and activities. Identify the possible states and state transitions for the system at the macro level. Determine the possible roles of the micro level components for a particular state at the macro level. Identify how state changes at the macro levels influence the role changes at the micro level. Associate the role changes at the micro level with the state transitions at the macro level.

Communication and conflicts. Identify the micro level communication (interaction) and coordination mechanisms and if possible determine tolerable conflicts and inconsistencies. This includes identification of the type of communication (interaction) exhibited by the agents as well as the reason for the initiation of communication. It is also important to determine if there are repetitive interaction patterns and how a particular interaction pattern in the micro level yields an observable system behaviour at the macro level.

It has to be noted that both processes (bottom-up and top down) are complementary to each other and share a common goal. Therefore the findings of the bottom-up observation and top-down analysis need to complement each other. Any contradictory findings need to be further examined either by re-examining the data or by repeating the simulation in an iterative manner until the conflict is resolved. All the findings need to be consistent in order to determine the micro level factors which have observable influence on the emergents at the macro level.

5.5 Summary

The experimental process proposed in the methodology can only yield beneficial results when it is applied on a "correct" model of the system. Therefore ensuring the correctness of the model is of paramount importance. The best way to deal with this issue is to use formal techniques to verify the correctness of the components (agents) and simulation in order to verify the correctness of the entire multi-agent model. The verification through the model animation should follow an iterative process which should gradually build the confidence of the model performance. The experimental phase, on the other hand, is essentially a two-way (bottom-up and top-down) iterative simulation study which aims at defining micro-macro relations in order to build a coarse cause and effect taxonomy of the specific emergent behaviour.

6 CONCLUSIONS

In recent years the phenomenon of emergence, as one of the fundamental properties of complex systems, managed to capture significant attention from the scientific community. There are several reasons behind this development. First of all, emergence seems to be everywhere in nature. It appears in different forms and shapes in a variety of systems from simple to the most complex. It is responsible for a variety of fascinating properties and behaviours. The ability to engineer emergent phenomena can be very beneficial in many areas of science and technology. On the other hand, emergence can also be viewed as negative phenomenon, it can significantly infringe the functional performance of engineered systems. This prospect is especially disturbing since there is a growing trend in engineering open distributed systems with high complexity. Nevertheless, to the best of our knowledge, at the moment there are no studies which deal with the practicalities of constructing a framework (comprised of a well defined process supported by a set of practices and tools) which will guide the analysis of existing emergent phenomena.

We address this issue by proposing a framework for empirical exploration of emergent formations. The core idea is to offer a structured approach which utilizes iterative multi-agent simulation as means for experimental examination of emergent manifestations. The goal of the process is to gradually increase the understanding of the causal relations between the individual (micro) and emergent (macro) levels of the system under study. It is essentially an analytical process aimed at detection of causal relations in the model, through evaluation of hypotheses about the expected behaviour of the model under certain conditions. Nevertheless, such a study can only yield beneficial results when it is applied on a "correct" model of the system. Therefore ensuring the correctness of the model is of a paramount importance. However, formal validation of the complete multi-agent model, on the other hand, is often impossible or too expensive in terms in time and effort. Consequently we propose a validation of the model through iterative simulation and refinement.

The future work will be concentrated on the evaluation of the proposed process and practices as well as assessment of the practical applicability of the proposed framework. Towards this end we have committed to examination of the herd formation as an emergent phenomena in the herd dynamics case study. Although the work is ongoing, for brief overview of the simulation model and the mechanism for automated detection of herd formations please refer to [28]. In addition we are exploring the several tools (like for example the automatic detection of invariants) which will help us facilitate the different activities. Once this work is completed, the study will move to the experimental examination in order to explore the causal relations responsible for formation of herds.

REFERENCES

[1] Anderson, P. W.: More Is Different. Science New Series, Vol. 177 (4047), 1972, pp. 393–396.

[2] Balci, O.: Guidelines for Successful Simulation Studies. In Proceedings of the Winter Simulation Conference, 1990, pp. 25–32.

[3] Bar-Yam, Y.: Dynamics of Complex Systems. Perseus Books, 1997.

[4] Bedau, M.: Weak Emergence. Philosophical Perspectives: Mind, Causation, and World, Vol. 11, 1997, pp. 375–399.

[5] Bonabeau, E.—Dessalles, J.-L.: Detection and Emergence. Intellectica, Vol. 2, 1997, No. 25, pp. 8594.

[6] Bullock, S.—Cliff, D.: Complexity and Emergent Behaviour in ICT Systems. Hewlett Packard, Technical Report HPL-2004-187, 2004.

[7] Cariani, P.: Adaptivity and Emergence in Organisms and Devices. World Futures: Journal of General Evolution, Vol. 32, 1992, pp. 111–132.

[8] Chalmers, D. J.: Varieties of Emergence. Technical report/preprint, Department of Philosophy, University of Arizona, USA, 2002.

[9] Conte, R.—Castelfranchi, C.: Simulating Multi-Agent Interdependencies: A Two-Way Approach to the Micro-Macro Link. In Proceedings of Social Science Microsimulation, Troitzsch K. G., Mueller U., Gilbert N. and Doran J. (Eds.), Springer-Verlag Berlin, 1996, pp. 394–415.

[10] Csuhaj-Varjú, E.—Dassow, J.—Kelemen, J.—Păun, Gh.: Grammar Systems: A Grammatical Approach to Distribution and Cooperation. Gordon and Breach, London, 1994.

[11] Darley, V.: Emergent Phenomena and Complexity. Artificial Life, Vol. 4, 1994, pp. 411–416.

[12] de Wolf, T.—Holvoet, T.: Towards a Methodology for Engineering Selforganising Emergent Systems. In Self-Organization and Autonomic Informatics (I), H. Czap, R. Unland, C. Branki, H. Tianfield (Eds.), Frontiers in Artificial Intelligence and Applications, Vol. 135, 2005, pp. 18–34.

[13] Edmonds, B.—Bryson, J.: The Insufficiency of Formal Design Methods – the Necessity of an Experimental Approach for the Understanding and Control of Complex MAS. In Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems, 2004, pp. 938–945.

[14] Edmonds, B.: Using the Experimental Method to Produce Reliable Self-Organised Systems. In Proceedings of the 2nd International Workshop on Engineering Self-Organising Applications (ESOA 2004), 2004, pp. 84–99.

[15] Eleftherakis, G.—Kefalas, P.—Sotiriadou, A.—Kehris, E.: Modelling Biology Inspired Reactive Agents Using X-Machines. In Proceedings of the International Conference on Computational Intelligence (ICCI 04), 2004, pp. 93–96.

[16] Epstein, J. M.: Agent-Based Computational Models and Generative Social Science. Complexity, Vol. 4, 1999, No. 5, pp. 41–60.

[17] Freund, R.—Kelemen, J.—P˘aun, Gh.: A Note on Emergence in Multi-Agent String Processing Systems. Computing and Informatics, Vol. 22, 2003, No. 6, pp. 623–637.

[18] Fromm, J.: Types and Forms of Emergence. Complexity Digest, Vol. 25, 2005, No. 3.

[19] Fromm, J.: Ten Questions About Emergence. Complexity Digest, Vol. 40, 2005, No. 15.

[20] Goldstein, J.: Emergence as a Construct: History and Issues. Emergence: Complexity and Organization, Vol. 1, 1999, pp. 49–72.

[21] Gueron, S.—Levin, A.—Rbenstein, D. I.: The Dynamics of Herds: From Individuals to Aggregations. Journal Theoretical Biology, Vol. 182, 1996, pp. 85–98.

[22] Havel, I. M.: Causal Domains and Emergent Rationality. In Proceedings of the 23rd International Wittgenstein Symposium, 2000, pp. 129–151.

[23] Holland, J. H.: Emergence: From Chaos to Order. Reading, MA: Addison-Wesley, 1998.

[24] Jennings, N. R.—Wooldridge, M.: Intelligent Agents: Theory and Practice. Knowledge Engineering Review, Vol. 10, 1995, No. 2, pp. 115–152.

[25] Jones, S.: Organizing Relations and Emergence. In Proceedings of the Artificial Life VIII: The 8th International Conference on the Simulation and Synthesis of Living Systems, 2002, pp. 418–422.

[26] Kehris, E.—Eleftherakis, G.—Kefalas, P.: Using X-Machines to Model and Test Discrete Event Simulation Programs. In Proceedings of Systems and Control: Theory and Applications, N. Mastorakis (Ed.), World Scientific and Engineering Society Press, 2000, pp. 163–168.

[27] Kubik, A.: Toward a Formalization of Emergence. Artificial Life, Vol. 9, 2003, p. 4166.

[28] Paunovski, O.—Eleftherakis, G.—Cowling, T.: Fuzzy Approach to Detection of Emergent Herd Formations in Multi-Agent Simualtion. Proceedings of the 9th Annual International Workshop Engineering Societies in the Agents World (ESAW 08), Spetember 24–26, 2008, Saint Etienne, France.

[29] Mogul, J. C.: Emergent (Mis)Behavior vs. Complex Software Systems. HP Laboratories, Palo Alto, Tech Rep. HPL-2006-2, 2006.

[30] Reynolds, C. W.: Flocks, Herds, and Schools: A Distributed Behavioral Model. In Computer Graphics, 1987, pp. 25–34.

[31] Ronald, E. M.—Sipper, M.—Capcarre, M. S.: Design, Observation, Surprise! A Test of Emergence. Artificial Life, Vol. 5, 1999, No. 3, pp. 225–239.

[32] Rosen, R.: Anticipatory Systems. Pergamon, Oxford, 1985.

[33] Sargent, R. G.: Verification, Validation, and Accreditation of Simulation Models. In Proceedings of the 2000 Winter Simulation Conference, 2000, pp. 50–59.

[34] Simon, H. A.: The Architecture of Complexity: Hierarchic Systems. Proceedings of the American Philosophical Society, Vol. 106, 1962, pp. 467–482.

[35] Wegner, P.: Why Interaction Is More Powerful Than Algorithms. In Communications of ACM, Vol. 40, 1997, No. 5, pp. 80–91.

[36] Whitner, R. B.—Balci, O.: Guidelines for Selecting and Using Simulation Models Verification Techniques. Department of Computer Science, Virginia Polytechnic Institute and State University, Virginia, Tech. Report TR-89-17, 1989.

[37] Wolfram, S.: Universality and Complexity in Cellular Automata. Physica D Nonlinear Phenomena, Vol. 10, 1984, No. 1–2, pp. 1–35.

Ognen Paunovski is a Ph. D. candidate at the South-East European Research Centre (SEERC), an institution jointly established by the University of Sheffield and the CITY College of Thessaloniki, Greece. He holds a B. Sc. in computer science and an MSc degree in software engineering and telecommunications. He is a member of the Intelligence, Modelling & Computation research group at the CITY College and the Verification and Testing group at the University of Sheffield. His primary research interests, focused on different aspects of emergence, include definition and classification of emergent manifestations, modelling of systems exhibiting emergence, analysis of emergent behaviours and other topics relevant to the exploration of emergence. He is also an active researcher in the area of distributed systems and has published several papers dealing with distributed discovery in decentralized bio-inspired agent systems.

George Eleftherakis is a Senior Lecturer, Research Director and Undergraduate Project Coordinator of the CS department at CITY College of Thessaloniki, an affiliated institution of the University of Sheffield, and the leader of the ICT research track of the South Eastern European Research Centre (SEERC). He holds a B. Sc. in physics and an MSc and a Ph. D. in computer science. He is the leader of the Intelligence, Modelling & Computation and the founder and leader of the Information Security, Networks and Mobility research groups at CITY. His main research work is in the area of formal methods, biologically inspired computing and information security. He is the founder and leader of Thessaloniki's Java User Group. He gave more than 10 invited talks to universities, conferences and companies around the world. He has published more than 50 papers and edited 4 books, most of them in the area of formal methods. He is a member of BCS, ACM (Sponsoring an ACM student chapter) and the Greek Computer Society (serving as a member of the administration board). Finally he organised four scientific conferences, in formal methods, in membrane computing, in mobile communications and embedded systems, and in ICT solutions for justice, as well as seven international student conferences.

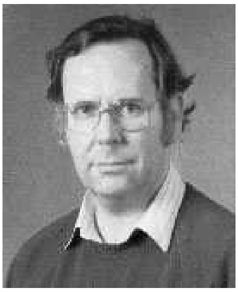

Tony Cowling is a senior lecturer in the Department of Computer Science at the University of Sheffield, and a member of the Department's Verification and Testing research group. He has a broad range of interests within software engineering, imcluding in particular software engineering education, the integration of system development methods, software testing and the contributions that the formal methods can make to this, and empirical software engineering, as an approach for establishing whether theoretical methods actually work as intended. Within the work on methods integration he is interested in a range of possible approaches to solving the practical problems of engineering complex software systems, but has focused particularly on the roles of two formal methods: one using X-machines, communicating X-machine systems, and related machine models, and the other using a data flow algebra model which he is developing.