Theoretical Steps Towards Modelling Resilience in Complex Systems

Cathy Hawes[1] and Chris Reed[2]

[1] Scottish Crop Research Institute, Invergowrie, Dundee, DD2 5DA
chawes@scri.ac.uk
[2] Department of Applied Computing, University of Dundee, Dundee, Scotland, DD1 4HN
chris.reed@computing.dundee.ac.uk

Abstract. This paper reports on theoretical work aimed at providing a harmonious set of tools for tackling the thorny problem of resilience in complex systems. Specifically, key features of resilience are laid out, and the ramifications on necessary theoretical and implementational machinery are analysed. These ramifications constitute a problem definition that, to the authors' knowledge, no extant system is sufficiently sophisticated to meet. It is, however, possible to identify existing components that can be combined to provide the necessary expressivity. In particular, theoretical ecology has individual based modelling approaches that are consonant with artificial intelligence techniques in multi-agent systems, and in philosophical logic, channel theory provides a mechanism for modelling both system energy and system information flow. The paper demonstrates that it is possible to integrate these components into a coherent theoretical framework, laying a foundation for implementation and testing.

1 Overview

Resilience is a property of complex systems that is widely used in describing such systems in both ecological theory and computational practice. In ecology, however, the term is used to refer to a wide variety of different phenomena, and in computer science it is often used either as a synonym of robustness or fault tolerance. In neither domain is resilience consistently or universally defined, and it is often employed simply to appeal to a reader's intuitions. The intuitions to which it appeals concern the ability of a large, complex system to recover from disturbance, to withstand external perturbation, to evolve, and to adapt. It seems, therefore, to be describing a set of properties that are essential measures of the health of a system, and therefore to be relevant not only in academic theory but also in public and political arenas.

Though there are many models of system change and resilience in ecology, and many applications of computational techniques to ecological systems, there are few that unite the two disciplines, placing ecological interactions at the heart of new computational algorithms.

The project of which this work forms a part aims to take ecological approaches to system function, and individual-based modelling in particular, as a starting point for development of a massively scaled multi-agent system that uses inter-agent communication to model the flow of energy through the system. The close analogy between energy flow and information flow is employed to build a bridge between ecological and computational theory, and to provide a basis for a domain-independent definition of resilience (generalising conclusions of, e.g. [1]). The system implementation and resilience analysis protocol will first be validated by comparison with existing ecological data, before then being applied to new problems of larger, more complex ecosystems, and thence to similar problems of large scale distributed and Grid computing. In this way, the we aim to develop a practical theory of resilience which can be reused in the design of artificial complex systems in eScience and e-commerce domains.

2 Motivation

There are a wide variety of platforms designed for modelling of complex systems using discrete events and, often, employing multi-agent systems as a enabling technology. One of the most mature is the Swarm platform [2] that aims to provide a generic platform for simulation, though other more recent developments extend and specialise these techniques to social simulation [3] and ecology [4]. These systems aim to provide an environment in which domain experts (e.g. sociologists or ecologists) can construct simulation models rapidly and easily. This inevitably has an impact upon the expressiveness of the languages they provide. In particular, previous approaches have omitted some or all of the following features:

- An explicit dynamic environment
- Agentive models (by which individuals might maintain partial, local, representations of their environment)
- Spatial extent (with individuals occupying different – and changing – volumes in the environment)
- Proactive behaviour (in which individuals' activities are not determined solely by Markovian reactive rules)
- Information flow (by which the transfer of energy, artefacts or data between individuals can be modelled)
- Dynamic topology (whereby the agents with which another can communicate change over time)

Bordini et al. [3] use ELMS to model the environment, but have no mechanism for capturing spatial extent; Swarm agents can model consumption, but act only reactively in a (primarily) static environment [4]; DSDEVS [5] can handle dynamic topology, but models only simple reactions, and so on. In order to explore the emergent character of resilience in both ecological and other complex systems, every one of these features is indispensable, and will be integrated into a single framework here. The resulting sophistication of the model as a whole is likely to limit its use to teams involving not only domain specialists but also computer scientists, but will allow those teams to probe fundamental questions of the properties of complex systems that previous approaches have simply not supported.

3 Foundation

The work builds upon three key notions: (i) using agents in a multi-agent system to model individuals in an ecosystem; (ii) modelling energy flow as information flow; (iii) supporting the primary ecological interactions (trophic, competitive, sexual and mutualistic) between system components in a common modelling framework.

Using agents for individual based modelling. Individual based modelling (IBM) is becoming established as a novel way of exploring ecosystem structure and function [6]. Multi agent systems (MAS) offer a computational platform upon which to construct decentralised distributed processing software [7]. Here, agents are used to embody the processing and characteristics corresponding to an individual. The paradigm of multi-agent system development offers a convenient computational metaphor for exploring individual behaviour and, in particular, the relationships between individuals.

Modelling energy flow as information flow. Energy flow is a powerful fundamental mechanism in describing ecosystem function [8]. Recent advances in multi-agent systems focus on semantically rich models of information flow to enhance robustness, fault tolerance and openness [9]. We can employ a mapping between these two to harness the developments in the latter to produce models of the former with greater explanatory power.

Complex interactions as inter-agent communication. Encounters between individuals in an ecosystem can be of four main types: (i) interactions involving energy flow (trophic); (ii) competition for resources (competitive); (iii) reproduction (sexual); and (iv) synergistic collaboration (mutualistic). Under different terminology, competitive and mutualistic interactions have formed the focus of much work in multi-agent systems, with the game theoretic approach often focusing on the former [10] and the distributed AI approach focusing on the latter [7]. Many recent systems have brought together the cooperative and competitive facets in fecund new research, with the Trading Agent Competition a now well-established example [11]. Building on these foundations, the two remaining classes of interaction are also directly implementable using the same principles. Trophic exchange can be modelled through a negotiation over energy flow; the semantics of that communication are then not denotational in the usual Tarskian sense, but rather directly operational in the fabric of the environment in which the agents operate. Sexual interaction is modelled similarly except that rather than negotiating over the flow of energy, agents instead negotiate of the exchange of genetic material.

4 Framework

The environment is structured as a set of contiguous, spatially explicit 3D cells characterised by a number of fixed and variable properties (including size, location and topography amongst the former and temperature and radiation amongst the latter). These cells can be managed by a class of “environment” agents to simplify implementation (though could be implemented as part of the structure of the environment itself given recent models of agent environments [3]).

Individual organisms within an environment are modelled by agents, but by itself, this is not sufficient for capturing spatial identity (i.e. location) and spatial extent (i.e. size). To model such spatial characteristics, agents are responsible for *projecting* the appropriate properties of the individual they represent into the physical, perceivable environment. Specifically, an agent that represents an individual that is located within an environment cell must communicate the appropriate properties of the individual to that cell, and similarly if the individual has spatial extent across multiple cells, then its properties must be communicated to all those cells. The properties that must be projected into the environment include chemical, physical and chromatic facets. So, for example, an agent representing an oil seed rape individual might project glucosinolates amongst its chemical properties, waxy surface amongst its physical properties and 80% reflected light in the 510nm band amongst its chromatic properties. An agent representing an aphid might similarly project honeydew and 60% reflected light in the 490nm band. Though the approach could be extended to handle a more detailed account of the three classes of properties, or indeed a richer range of properties, this level of abstraction suffices for exploring a range of interesting phenomena in high detail.

Agents interact with the environment and each other through a simple perceive-interpret-act cycle. Perception is straightforwardly characterisable as information flow, which in this model is implemented through communication, and in this case, communication with the environment (agents). Interpretation is the integration and interpretation of percepts with existing knowledge, a process entirely internal to individual agents, and one which concludes with the determination of appropriate action (including null actions). Finally, an agent's action is executed through the appropriate communication with either the environment (agents), other agents representing individuals, or both.

Agent communication is typically marshalled in traditional multi-agent systems through lookup services of one sort or another (JADE for example uses JINI [12]). Without further restrictions, any one agent in the system might communicate with any other. Here, agent communication is marshalled through a combination of the environment and an agent's own capabilities. So the possibility of communication between two individuals depends upon their spatial proximity and their innate communication capabilities.

The outcome of an initial, information-gathering communication between two agents will depend on the interpretation of that information and can take one of four forms: competitive, mutualistic, trophic or sexual.

Competition in ecosystems occurs through interference (direct, physical interaction) or exploitation (indirect, via utilisation of a shared resource). If a caterpillar consumes a leaf, there is no phloem for an aphid to consume; the caterpillar has exploited the resource to the detriment of the aphid, resulting in the indirect, exploitation type of competition. Such exploitation does not need to be modelled explicitly as this interaction operates solely through a change in resource availability handled by the extant modelling machinery. Interference competition, in contrast, involves a direct, aggressive encounter in which one individual threatens another (in territorial disputes, for example). This type of competition demands explicit handling in the model, through negotiation between individual agents based on perceived ownership and defence capabilities. Limitation of resources and subsequent dispute over those resources (whether explicit as interference, or implicit as exploitation) has been a fundamental part of the dynamics that drive scenarios tackled by agent systems. From early competitive tendering models such as CNET [13], through rich academic environments such as TAC [11], to commercial marketplace and auction protocol design [10], the ways in which resource limitation can drive sophisticated agent dialogue are well known. Rather less common are mechanisms by which agents can specifically employ the techniques of interference competition in making threats to disadvantage competitors (see, e.g. [14]). In those models the felicitous performance of a threat is dependent upon social relationships between agents, here they are dependent upon the ecological counterparts to those relationships.

Mutualism is handled in a similar way to interference competition, but rather than an aggressive encounter, negotiation is used to the benefit of both partners. Models of cooperative problem solving (such as [7]) explore many aspects of mutualism, but all are characterised by a strong focus on the reasoning capabilities of the individuals, and the communication structures by which they reach joint plans. This focus on information exchange and representation forms a cornerstone of the approach proposed here.

Trophic interactions concern the identification of one individual as food by another. In the model, they occur where the agent *perceives* the projection of another agent in the same environment cell, *interprets* this projection as an appropriately available food resource, and *acts* by requesting an amount of that resource from the perceived agent. An agent's food requirement and degree of specialism is defined in the set of internal properties of that agent.

Sexual interactions are not usually modelled (see [15] for a rare example) because they result in an open system in which agents can come into existence (and usually also terminate). Agent creation on-the-fly is also rarely tackled in academic work for the more prosaic reason that it is seldom an important or technically challenging part of a system. It is included here both for completeness and to support future extensions in adaptation.

Finally, all forms of interaction (competitive, mutualistic, trophic, sexual between individuals, and projective and perceptive between individuals and the environment) are mediated through *device constraints*. An individual without a device sensitive to light cannot detect chromatic percepts; an individual with a dermis-puncturing device may be able to consume sap, etc. Device constraints can be expressed as agent-relativised infomorphisms in Channel Theory [16]. In essence, an agent A might implement an infomorphism between a physical, environmental relation $\models_{Phys}$ that maps states of the world onto vectors of values representing the environment in that state, and its own perceptual relation $\models_{A\text{-}Percepts}$ corresponding to how the agent perceives that state of the world (and thereby its own internal representation of the state). This approach allows the model to capture the intuition that the information is "there", regardless of whether or not a particular agent is able to perceive and internalise it. A key technical novelty of this work is that consumption is modelled in exactly the same way, leading to parsimony and simplicity in implementation: the same agent A maintains a second infomorphism, $\models_{A\text{-}Consumes}$, from what exists in the environment to what it (decides to) consume.

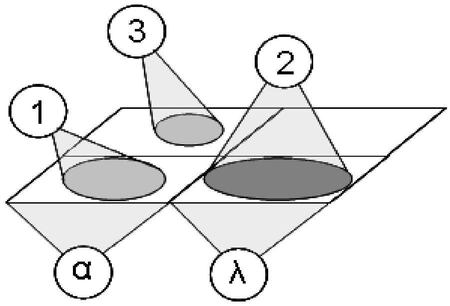

Fig. 1. Model schematic

To illustrate how these components fit together, a simple scenario involving perception and trophic interaction is presented (illustrated in Figure 1). At the first step, individual agent 1 (I_1) requests from environment agent α (E_α) a list of the properties associated with that part of the environment:

(1) $I_1 \rightarrow E_\alpha$: perception-request "*What properties are in this locale?*"
(2) $E_\alpha \rightarrow I_1$: $s_1 \models_{Phys} < (I_1 , \{...\}) > \;\; \Rightarrow_{f1}$
$s_1' \models_{I1\text{-}Percepts}$ < empty > "*None other than yours.*"
(3) I_1 Interpretation: *no resources available, internal food store sub-optimal, current plan is foraging. The appropriate action is to move.*
(4) $I_1 \rightarrow E_\alpha$: remove-projection "*Delete my properties from your locale.*"
(5) $I_1 \rightarrow E_\lambda$: add-projection($I_1$, {...}) "*Add my properties to your locale.*"
(6) $I_1 \rightarrow E_\lambda$: perception-request "*What properties are in this locale?*"
(7) $E_\lambda \rightarrow I_1$: $s_2 \models_{Phys} < (I_1 , \{...\}), (I_2 , \{...\}) > \;\; \Rightarrow_{f2}$
(8) $s_2' \models_{I1\text{-}Percepts} < (I_2 , \{glu, waxy, green\}) >$
"*There is an individual projecting glucosinolates, waxy surface, 80% reflection at 510nm*"
(9) I_1 Interpretation: *I_2 projection matches food profile. Appropriate action is to eat.*
(10) $I_1 \rightarrow I_2$: $s_3 \models_{I1\text{-}Energy} < (+I_2 , \{1ml\text{-}of\text{-}sap\}) > \Rightarrow_{f3}$
$s_3' \models_{I2\text{-}Energy} < (- I_1 , \{...\})$ *Transfer energy by consuming sap.*

The agent communication language is simple, comprising a few key primitives (such as *perception-request*), and the ability to exchange logical structures instantiating channel theoretic classifications, for example, that a given state is of a type with a given vector of variables. State types are described by a vector of individual / property-set pairs, in which property sets are arbitrarily complex descriptions of individuals (as trivial as a statement of the individual's volume, or as complex as a description of its chemical make-up). Communication is summarised by the application of an infomorphism (marked $\Rightarrow_f$) implemented in the recipient agent.

At step (1) I_1 initiates a perception event, the response to which at (2) is the description of the state of the environment managed by environment agent α. The physical classification of this state (by $\models_{Phys}$) is one which includes only I_1; in I_1's own classification (reached by the infomorphism *f1*), this is a state with no individuals present (that is, it ignores the part of the environment occupied by itself). At (3), its own local logic responds (in this case reactively, but including reference to the current intentional state) producing an appropriate action. After moving, its perception at (7) identifies another individual I_2 characterised by a set of percepts that at (8) it associates with food. At (9) it performs an energy transfer (that is it consumes some of I_2). It formulates its "request" according to its own local logic (that is, according to its own consumption device constraints), but for I_2 the request may have a different interpretation, built from the infomorphism *f3*.

In this way, the energy flow in the ecosystem is modelled by the exchange of information between the two agents, with the formal similarity between steps (7) and (9) emphasising the similar treatment of energy flow and perceptual data flow. Though the consumed agent (I_2) is compelled by the semantics of the interaction to release some energy package, the nature of that package is determined by three things: (i) the consumption device by which it has been requested; (ii) I_2's own status; and (iii) I_2's own predefined tailored response. In particular, this final component allows for modelling responses in which *permanent* defences are employed (such as a plant maintaining levels of chemical defences in its phloem, or waxes on it's leaf surface). *Induced* defences, in which a response is produced dynamically, can also be accounted for through I_2's internal processing in its immediate and subsequent interpretation phases.

The individuals and the infomorphisms that they embody represent packages of functional traits. System level properties emerge from the combination of these traits represented by the individuals and the interactions between them. Maintaining the functioning of complex systems has been shown to be dependent upon their diversity and productivity [17], and this gives a basis for exploring explicit definitions of system resilience.

5 Towards Resilience

Resilience is the ability of a system to return to a stable state following perturbation, and can viewed in two distinct ways: Engineering Resilience has been defined as the rate at which a system returns to a global equilibrium following perturbation [18]; Ecological Resilience is the amount of disturbance required to change from one steady or cyclic state to another [19]. The former focuses on maintaining efficient system function and predictability, and tends to assume a single, constant, global steady state. The latter focuses on maintaining the existence of function, and is more concerned with the factors that drive instabilities that cause the system to flip between multiple stable states [1]. A third approach, used in economic theory, is based on the engineering approach, but adapted to allow for multiple stable states through path dependency and other mechanisms [20].

The degree to which any complex system is resilient by these definitions depends to a large extent on the characteristics of the components of that system with respect to a given process or function. In ecological systems, resilience (frequently thought of in terms of the temporal stability of a process such as productivity) is thought to increase with functional diversity [21] and is affected by the strength of the interaction between different species in the system [22]. Ecological theory attempts to explain these patterns through a variety of models: "species richness-diversity" draws a linear relation between species richness and ecological function [23]; "idiosyncratic theory" proposes that stability depends on merely the identity of species present [24]; "functional redundancy" is based on the principle that compensating complementarity between species allows functions to be maintained in the face of species extinctions [25]; and finally the "keystone species hypothesis" extends the functional redundancy ideas to include differentiation in the strength of species functional roles in the community [26]. Theoretical frameworks have also been proposed that attempt to unify and simplify these models by combining their different elements or embedding them in gradients of the functional overlap between species or the degree of variation between them [27]. These theoretical models are based on the concept of species populations as the basic functional unit, and therefore fail to capture the complexity or importance of the interactions amongst individual organisms in the system. Here we attempt to overcome this subjective classification of organisms into potentially artificial groupings by modelling individual behaviour at the level that organisms interact with each other in real systems. Interaction strengths, overlap in ecological function and differentiation between individuals therefore emerge as a property of the simulated system rather than being imposed as a fixed or required feature. This allows the relative importance of functional overlap, interaction strength and diversity in determining system resilience to be tested directly. At the same time, it is possible to avoid oversimplifying the functional representation of individuals by eschewing traditional reactive models [4] in favour of autonomous proactive software entities (in the style of [7]) supporting goal-directed and coordinated behaviour.

Resilience can be measured in terms of change in a system level property and function following perturbation. Perturbation can be simulated as addition, removal or alteration of one or more system components (mirroring catastrophic ecological invasion, extinction or mutation processes). The system component is a class of individuals with similar trait packages that perform a particular function in the system (i.e. a set of individuals with the same resource requirement, method of acquiring that resource, and strength of interaction with other individuals in the system). The impact of addition, removal or modification of this system component on resilience will be measured through responses of system properties and processes. The system property is trait diversity (the number and frequency of each trait represented by the individuals in the system), and the change in trait diversity (over sufficient generations to allow the detection of stable or cyclically stable states). The system processes are energy capture efficiency (ECE) and change in ECE and, its related measure, total system productivity (SP) and change in SP. ECE is the proportion of energy entering the system that is captured by each individual and converted to new biomass. SP is the resulting cumulated biomass across all individuals in the system. The combination of traits represented by the individuals in the system and the interactions between them that result in enhanced system resilience can then be objectively evaluated. This framework therefore allows the exploration of the properties of a generic complex system whose emergent behaviour is dependent upon the diversity of functional traits of surviving or susccessful individual agents and the number, type and strength of the interactions between them.

6 Concluding Remarks

Our aim here has been twofold: first, to identify an omission in current research on properties of complex systems, and on resilience in particular; and second, to lay the groundwork upon which formal theories and then subsequently implemented systems might be developed. The preliminary framework sketched here has a number of important advantages not only for complex systems research, but also specifically for the two contributing fields harnessed here, namely, ecosystem theory and multi-agent systems. In particular, from an ecological modelling perspective, this framework provides the ability to integrate a wide range of interaction types and offers a generic approach that can handle individual behaviour with out the need to hard-wire specific behaviours or trophic structures. For multi-agent systems, it provides a mechanism by which to apply the techniques of information exchange which are increasing well understood in new modelling scenarios involving energy or artefact exchange. More broadly though, the approach offers avenues to developing domain-independent definitions of system resilience that hold the potential to provide new ways of exploring the dynamic, emergent character of complex systems.

References

1. Gunderson, L., Holling, C.S., Pritchard, L., Peterson, G.D.: Resilience. In: Mooney, H.A., Canadell, J.G. (eds.): SCOPE The Earth system: biological and ecological dimensions of global environmental change. Encyclopedia of Global Environmental Change (2002) 530-531
2. Daniels, M.: Integrating Simulation Technologies with Swarm. In: Working Notes of the Workshop on Agent Simulation: Applications, Models and Tools. University of Chicago (1999)
3. Bordini, R.H., da Rocha Costa, A.C., Hubner, J.F., Moreira, A.F., Okuyama, F.Y., Vieira, R.: MAS-SOC: A Social Simulation Platform based on Agent-Oriented Programming. Journal of Artificial Societies and Social Simulation 8 (3) (2005) http://jasss.soc.surrey.ac.uk/8/3/7.html
4. Langton, C.: The Swarm Simulation System and Individual Based Modelling. Santa Fe Institute Working Paper (1996)
5. Duboz, R., Cambier, C.: Small world properties in a DSDEVS model of ecosystem. In: Proceedings of the Open International Conference on Modeling and Simulation (OICMS-2005), (2005) 65-71

6. Breckling, B., Muller, F., Reuter, H., Holker, F., Franzle, O.: Emergent properties in individual-based ecological models – introducing case studies in an ecosystem research context. Ecological Modelling 186 (2005) 376-388
7. Jennings, N.R.: On Agent-Based Software Engineering. Artificial Intelligence 117 (2) (2000) 277-296
8. Schulze, E. D.: Flux control at the ecosystem level. TREE 10 (1995) 40-43
9. Foundation for Intelligent Physical Agents, ACL Spec (2002) www.fipa.org
10. Sandholm, T.: Agents in Electronic Commerce: Component Technologies for Automated Negotiation and Coalition Formation. Autonomous Agents and Multi Agent Systems 3 (1) (2000) 73-96
11. TAC: Trading Agent Competition (2005) home at www.sics.se./tac
12. JADE home at jade.tilab.com
13. Smith, R.G.: The contract net protocol: high level communication and control in a distributed problem solver. IEEE Transactions on Computers 29 (1980) 1104-1113
14. Sycara, K.: Persuasive Argumentation in Negotiation. Theory and Decision 28 (3) (1990) 203-242
15. Dumont, B., Hill, D.R.C.: Spatially explicit models of group foraging by herbivores: what can agent-based models offer? Animal Research 53 (2004) 419-428
16. Barwise, J., Seligman, J.: Information Flow: The Logic of Distributed Systems. CUP (1997)
17. Bolnick, D.I., Svanback, R., Fordyce, J.A., Yang, L.H., Davis, J.M., Hulsey, C.D., Forister, M.L.: The ecology of individuals: incidence and implications of individual specialization. Am. Nat 161 (2003) 1-28
18. Holling, C.S.: Engineering resilience versus ecological resilience. In: Schulze, E.D. (ed.): Engineering within ecological constraints. National Academy Press, Washington DC (1973) 31-43
19. Pimm, S.L.: The balance of nature. University of Chicago Press, Chicago (1984)
20. Clark, N., Juma, C.: Long-run economics: an evolutionary approach to economics growth. Pinter, London (1987)
21. Tilman, D., Wedin, D., Knops, J.: Productivity and sustainability influenced by biodiversity in grassland ecosystems. Nature 379 (1996) 718-720
22. Root, R.B.: Organization of a plant-arthropod association in simple and diverse habitats: the fauna of collards. Ecological Monographs 43 (1973) 95-117
23. MacArthur, R.H.: Fluctuations of animal populations and a measure of community stability. Ecology 36 (1955) 533-6
24. Lawton, J.H.: What do species do in ecosystems? Oikos 71 (1994) 367-74
25. Ehrlich, P.R., Ehrlich, A.H.: Extinction: the causes and consequences of the disappearance of species. New York: Random House (1981)
26. Walker, B.: Biological diversity and ecological redundancy. Conservation Biology 9 (1992) 747-752
27. Peterson, G., Allen, C.R., Holling, C.S.: Ecological resilience, biodiversity and scale. Ecosystems 1 (1998) 6-18