AI Communications 18 (2005) 305–317
IOS Press

The Ants' Garden: Complex interactions between populations and the scalability of qualitative models

Bert Bredeweg [a,*] and Paulo Salles [b]
[a] *University of Amsterdam, Faculty of Science, Informatics Institute, Human Computer Studies, Kruislaan 419 (Matrix I), 1098 VA Amsterdam, The Netherlands*
E-mail: bredeweg@science.uva.nl
[b] *University of Brasília, Biological Sciences Institute, Campus Darcy Ribeiro, Brasília – DF, 70.904-970, Brazil*
E-mail: psalles@unb.br

Abstract. Ecological theories often explain the behaviour of communities in terms of the underlying interactions that take place between the species that are part of the community. This closely relates to the idea of compositionality within Qualitative Reasoning, in which the model that simulates the behaviour of a larger system is automatically assembled from previously defined elementary units. A recurring question however concerns the viability of this approach, particularly the scalability of such small partial descriptions. In this article we present a fully implemented model and accompanying simulations of the Ants' Garden, a complex system consisting of multiple species with multiple interactions. The Qualitative Reasoning engine automatically assembles these models by reusing a previously created library with partial models of basic processes that govern population behaviour and interactions. The simulations show the typical behaviour of the Ants' Garden as currently known and described by experts, and support the idea that our previously developed library is adequate and scalable to simulate complex system behaviour.
Keywords: Population ecology, qualitative reasoning and simulation, compositional modelling and scalability

1. Introduction

Conceptual models aid understanding and are important means for research, management, and education. Qualitative Reasoning (QR) engines can be used to automate and support reasoning with conceptual knowledge. A characteristic of the QR approach is compositionality [5]: the development of partial models that represent the behaviour of elementary units that can be assembled to reason about the behaviour of larger systems. Scalability is a typical problem, because each additional unit (in a bigger system) increases the number of possible interpretations often due to ambiguity.

In this article we describe qualitative models and simulations of the *Ants' Garden*, a complex system consisting of four species: ants, their cultivated fungi, virulent parasitic fungi that may attack the garden and bacteria that produce antibiotics against the parasitic fungi [3,4]. The purpose of our research is twofold. To support the ecological understanding of such systems and to show the reusability and scalability of previously defined qualitative models for population ecology (e.g., [24]). Interactions between populations have been a hot topic in ecological theory and practice. Competition, for instance, is still seen as a driving force for shaping biological communities. Notice that such interactions closely match the idea of reusable elementary units in QR. Food webs, for instance, are complex systems built from elementary interactions.

Traditional modelling approaches, mostly based in differential and difference equations, are limited in capturing complex interactions between species. In general they are difficult to build, very hard to calibrate and their results are almost impossible to understand for non-experts. The required "good quality" numerical data is often not available. Even when such models can be created they are "black boxes", because they have no clear representation of the system's

*Corresponding author.

structure and do not capture an explicit explanation of how the system works [7]. Applied to more complex systems the predictive capacity of numerical population models may be jeopardized by other problems. It has been shown that even simple mathematical models may produce complex trajectories, with stable points, cycles and chaotic behaviour (cf. [12,13]). For example, modelling interactions among populations with ordinary differential equations, Gilpin [8] demonstrated that chaos could be observed when there are at least three populations. Cyclic behaviour has been observed in many populations, but the existence of chaotic behaviour in natural populations is still an open question (e.g., [12,14]).

QR may play an important role in this discussion. Qualitative models are conceptual models and implement an explanation of the system's behaviour. Using QR we describe models and simulations that capture the behaviour of the Ants' Garden. Our models are constructed by reusing a previously developed library of partial models implementing single population dynamics, community succession [21,23], and interactions between two populations [24]. These models are "simpler" than the complex mathematics used to represent population and community dynamics but powerful enough to support useful conclusions about the system's behaviour.

The content of this article is as follows. Below we first present the Ants' Garden problem. We then discuss previous work that will be reused for modelling the Ants' Garden. The fourth section presents and discusses the models and simulations of the Ants' Garden. After that there is a short section on related work and a section with our main conclusions.

2. The Ants' Garden

Because few organisms cultivate their own food, fungi gardening by ants is considered to be a major breakthrough in evolution. It is a symbiosis in which organisms of two different species (ants from the family *Formicidae* and fungi mostly from the family *Lepiotaceae*) benefit from each other and create a system that can successfully survive in a number of different environments, being the dominant herbivores in the Neotropics. Recent studies [3,4] show that the Ants' Garden is far more complex than initially understood. A third species, the specialized garden parasite fungi of the genus *Escovopsis* is often involved and may destroy the system, by attacking the cultivated fungi. However, it almost never happens because ants carry on their body colonies of bacteria (genus *Streptomyces*) that produce antibiotics that suppress the growth of *Escovopsis*. Therefore the system consists of four species and a complex balance of interactions in which eventually the Ants' Garden survives. The basic idea of the Ant's Garden is shown in Fig. 1. Further studies indicate that the bacteria produce metabolites (vitamins, amino acids) that may enhance the growth of the cultivated fungi [3]. This fifth interaction is not shown in Fig. 1.

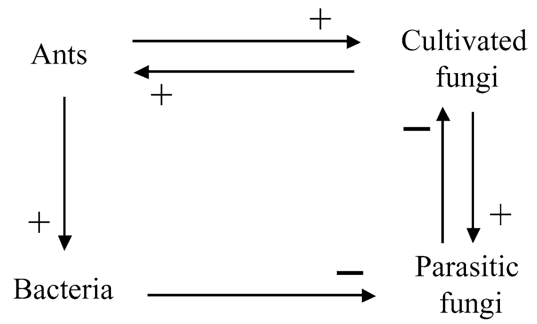

Fig. 1. The Ants' Garden: positive and negative interactions between populations.

How does the Ants' Garden work? How can the behaviour of the overall system be explained in terms of the behaviour of the constituents? We present qualitative models that simulate the Ants' Garden phenomenon. Such conceptual models are urgently required to enhance the understanding of complex ecological systems. Qualitative models are interesting candidates for this purpose.

3. Previous work on QR and population ecology

Our models are built in GARP [2], a domain *independent* reasoning engine that implements a compositional modelling approach [5] to qualitative simulation. The engine requires as input a set of scenarios and a library with model-fragments. Scenarios specify initial situations for the simulator to start a behaviour prediction. Model-fragments capture knowledge about the structure and behaviour of partial systems and are used to assemble states of behaviour. Assumptions may be used to further detail the applicability of a model-fragment. After selecting a scenario the reasoning engine proceeds with the prediction task by recursively consulting the library of model-fragments for applicable fragments. This search is exhaustive and each consistent subset of applicable model-fragments

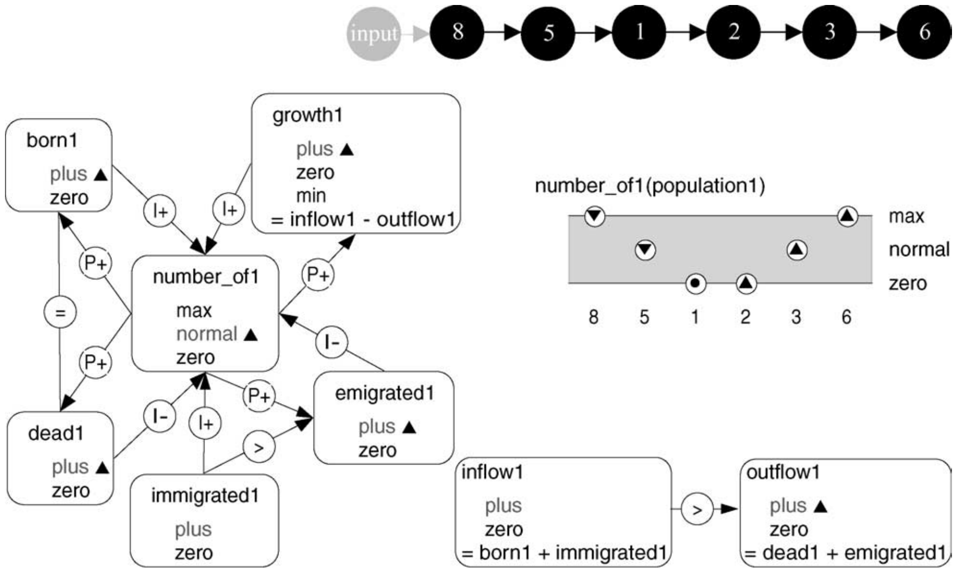

Fig. 2. Basic processes influencing a population [23]. Names used in the drawings relate to the text as follows: number_of (*Nof*), dead (D), born (B), emigrated (E), and immigrated (*Im*). The numbers 1, 2, etc. are used to distinguish between different populations. Thus number_of1 refers to the *Nof* individuals of the first population and number_of2 refers to the *Nof* individuals of the second population (if there is a second one). The same procedure is used for the other quantities.

represents a behaviour interpretation that matches the selected scenario. How many interpretations will be found depends on the kind of scenario, particularly on the amount of detail and constraints that have been specified in it.

We have adopted the process-oriented ontology [6]. Two modelling primitives are used to represent the effects of processes: *influences* (I+, I−) and *qualitative proportionalities* (P+, P−). Every process is associated to a *rate* and this quantity has a *direct* effect on another quantity (state variable). Changes in directly influenced quantities may propagate to other quantities via qualitative proportionalities. Both influences and qualitative proportionalities have thus a *causal* meaning and a mathematical *rationale*. Direct influences represent the derivative ($\mathrm{d}X/\mathrm{d}t$) of the quantity and cause change in the state variable (X). Proportionalities represent monotonic functions involving the quantities, so that when one of them is increasing or decreasing it causes changes on the other quantity (that may also increase or decrease). Quantities represent behaviour. Quantity values are characterized by <*magnitude*, *derivative*>, representing the current value and the direction of change, respectively. The possible values a quantity can take on are represented in its Quantity Space (QS). The derivative can take on three values {*min*, *zero*, *plus*}, referring to decreasing, steady and increasing, respectively.

The basic idea of our "single population" models and simulations is shown in Fig. 2. This figure is generated with VisiGarp [1] a graphical user interface that can be used to run and inspect simulations with GARP. The figure shows three typical outputs when running a qualitative simulation: the dependencies in the *causal* model (Left Hand Side (LHS) and bottom Right Hand Side (RHS)), the possible behaviours in the state-graph (top RHS) and enumeration of the values for certain quantities in the value-history (middle RHS). In this simulation the state variable is the population size "Number_of" (*Nof*). It has the $QS = \{zero, normal, max\}$, meaning that it can be non-existing (zero), have an average size (normal), or be at its maximum size (max). The black triangle (pointing upwards) signifies that *Nof* is increasing (a circle means steady and a triangle pointing downwards means decreasing). Most of the other quantities represent rates from currently active processes: birth rate (B), death rate (D), immigration rate (*Im*) and emigration rate (E). They all use the $QS = \{zero, plus\}$ meaning that they are active or inactive. They all have the value plus and increase, except immigration whose derivative is unknown (no triangle nor a black circle). Notice that each process rate has a direct influence on *Nof*, namely: $I+(Nof, B)$; $I+(Nof, Im)$; $I-(Nof, D)$ and $I-(Nof, E)$. It represents a qualitative version of the "growth" equation that is typically found in textbooks on Ecology:

$$Nof(t+1) = Nof(t) + (B + Im) - (D + E).$$

We used proportionalities to model the feedback loop involving the state variable (*Nof*) and the rates capturing the idea that when *Nof* is increasing or decreasing, the rate is changing in the same direction. For example: $P{+}(B, Nof)$. This relation does not apply to the immigration process because the size of the population seldom influences the number of immigrating individuals. We represented also an aggregated process (*Population growth*) to combine all the basic population processes, and its rate (*growth rate*) uses the $QS = \{min, zero, plus\}$. These values capture situations in which $(B + Im)$ is smaller than, equal to, or bigger than $(D + E)$, respectively. Finally, there are inequalities, such as $B = D$, and two additional quantities representing the total *inflow* $(B + Im)$ and *outflow* $(D + E)$.

The state-graph shows only one behaviour-path, namely: $\{8 \rightarrow 5 \rightarrow 1 \rightarrow 2 \rightarrow 3 \rightarrow 6\}$ (notice that the state numbers are just identifiers and that they do not necessarily reflect the order of the states). The value-history enumerates the quantity values for *Nof* in each of those states. It tells us that *Nof* starts in state 8 with magnitude *max* and that it decreases (*min*). State 8 is followed by state 5: *Nof* = <*normal*, *min*>, which is followed by state 1: *Nof* = <*zero*, *zero*>. After becoming extinct, the colonization process establishes a new population in state 2. After that the population grows and via state 3 ends in state 6. In state 6 it has again its highest value and is still increasing: *Nof* = <*max*, *plus*>. Notice that the causal model discussed before is part of the facts that hold in state 3. *Each* state has its *own* causal model, and causal models can thus differ among states. For instance, in state 1 there are no processes active, hence none of the rates (B, Im, D, E) influences the state variable *Nof*.

Starting from the basic processes that govern the behaviour of a single population we have constructed models that simulate the behaviour of two interacting populations. The interactions covered by these models are amensalism, commensalism, predation (similar to parasitism and herbivory), symbiosis, and competition (following [19]). A general mechanism was defined as the basis for these interactions (Fig. 3). Interactions are classified as combinations of the signs $\{+, 0, -\}$ according to the effects they cause on the growth of the other population (positive, neutral or negative, respectively). The mechanism assumes that the effect one population can have on another must be represented by a specific intermediate quantity (*Effect*). This way,

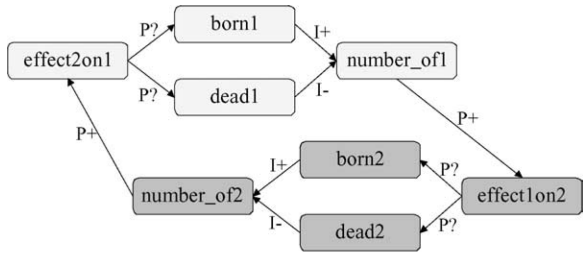

Fig. 3. The base model for pairs of interacting populations [24].

there are two quantities that represent the effects of population1 on population2 and vice-versa: *Effect1on2* and *Effect2on1* which can be seen as *Benefit* in the case of positive interactions or *Harm* in the case of negative interactions. Depending on the type of interaction the $P?$ in Fig. 3 becomes a negative (P−) or a positive (P+) proportionality. Simulations with these models show the typical qualitatively distinct behaviours that experts describe for each of these interaction types. For details see Salles et al. [24].

Qualitative models have a number of features that support their use to approach complex systems such as the ants and their gardens. They express concepts using a large vocabulary, create explicit representations of causality that may support explanations, and use a compositional modelling approach so that a library of situation-independent model-fragments can *automatically* be combined to form a set of distinct models and simulations. Compositionality involves re-using "simpler" partial models to scale up to more complex models. In the work presented in this article, models about single population behaviour and interactions between populations are re-used to model the Ants' Garden. We take the approach that the Ants' Garden may be represented as a combination of interactions between pairs of populations. This way, models of symbiosis, predation (parasitism), amensalism and commensalism are re-used to describe the interactions between ants, their cultivated fungi, parasitic fungi, and bacteria. Given that model-fragments implemented in our previous work are problem-independent, the knowledge already encoded in GARP should be sufficient and adequate to model the Ants' Garden.

4. Qualitative models of the Ants' Garden

As discussed above, the goal is to implement a qualitative model of the Ant's Garden by reusing previously defined model-fragments. Table 1 presents the relationships between the four species involved in the Ants'

Garden system, described by Currie and her collaborators [3,4].

The table should be read as the effects that the population in the row has on the populations in the column. In the table, (+) refers to a population positively influenced by another; (−) refers to a population negatively influenced by another; (?) means that there is no information available; (**) refers to "self-interaction". For example, the second row reads as follows: Cultivated fungi have a positive effect (+) on Ants and Parasitic fungi, and an unknown effect (?) on Bacteria (see also Fig. 1). Using the mechanism depicted in Fig. 3, the following minimum set of interactions is required to model the Ants' Garden:

- Ants/Cultivated fungi = symbiosis (+, +);
- Parasitic fungi/Cultivated fungi = parasitism (+, −);
- Ants/Bacteria = commensalism (0, +);
- Bacteria/Parasitic fungi = amensalism (0, −).

4.1. Single population models revisited

To address the modelling problem, let us go back to the models defined for the behaviour of a single population. Discussing these models with domain experts[1] have resulted in a few modifications. A summary of that is shown in Fig. 4. Experts argued that the aggregated population growth process (representing the whole set of basic processes) was unnecessary, and should not be included in one model at the same level of detail as the basic processes. Moreover, the derivative of the *Nof* already represents growth. In the new model the growth process is therefore left out, which means that the total *inflow* and *outflow* quantities become superfluous and can also be left out. Notice that this modification has no effect on the simulation outcome of the models. It only simplifies the representation.

A second issue concerns the question whether a quantity may increase while being in its highest point. Consider state 5 in the state-graph and value-history shown in Fig. 4 on the RHS, for which holds: *Nof* = <*max*, *plus*>. In the simulation shown on the LHS this state of behaviour does not happen. The simulation on the LHS only has state 4 in which: *Nof* = <*max*, *zero*> (which also appears in the simulation shown on the RHS). Is state 5 in Fig. 4 on the RHS appropriate or not? There is no single answer to this issue, because it depends on what the modeller wanted to represent with the upper limit in the quantity space. Is it a real limit beyond which the population cannot grow? Or is it a virtual limit: the population can in principle grow beyond this landmark but the values above the landmark are not of interest? Not being able to resolve this discussion with experts has inspired us to augment the qualitative reasoning engine with user definable features controlling the problem solving behaviour of the engine. In total 18 of such features have been identified. Some of the most typical ones are listed in Table 2. For the simulations discussed below the idea is that *Nof* cannot increase beyond its maximum value. Hence, the engine is set to not allow increase in the

[1] QSER (http://www.qrser.de/), QR03 (http://www.unb.br/ib/necbio/QR03/), and NNR (http://www.naturnet.org/).

Table 1

Qualitative relations between the four species involved in the Ants' Garden

	Ants	Cultivated fungi	Parasitic fungi	Bacteria
Ants	**	+	?	+
Cultivated fungi	+	**	+	?
Parasitic fungi	?	−	**	?
Bacteria	?	+	−	**

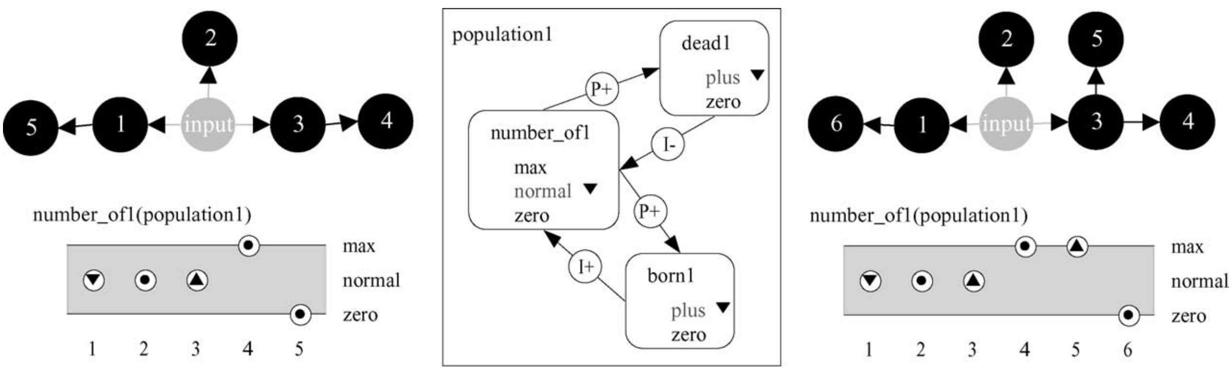

Fig. 4. Single population behaviour revisited.

Table 2

Some of the problem solving features in GARP that users can change

Type	Description	Default
Closed-world assumption	Applying the closed-world assumption during influence resolution results in assuming that non-influenced quantities are steady.	off
Inequality terminations assumption	Inequalities such as $q1 > q2$ only terminate when their derivatives are unequal. When this information is unknown, all possibilities are tried.	off
Extra termination rules	Additional termination rules can be defined and included, e.g., domain specific ones.	off
Epsilon ordering rule	Changes from a point (and equality) happen instantly and therefore precede changes that take at least some small amount of time.	on
No derivative constraints on extreme point values	Allow the derivative of a quantity to keep increasing (or decreasing) in the highest (or lowest) point of the quantity space.	off
Inactive quantities removal	When processes stop related quantities (mainly rates) may become superfluous. These quantities are removed.	off

highest and decrease in the lowest point values of a quantity space.

Explanations on interactions between pairs of populations usually do not require migration. We use the idea of operating assumptions [5] to develop models that include migration or exclude migration. In our previous work "no migration" was implemented as: $Im = E = 0$ and $\partial Im = \partial E = 0$ [24]. In our current model we have taken a slightly different approach. We still use they idea of operating assumptions to set the range of model applicability. However, the superfluous quantities are now fully removed from the model. Notice that this does not affect the simulation results of interacting populations. As before, it only simplifies the representation.

4.2. Interaction models revisited

Interactions between population pairs are based on the base model discussed above (Fig. 3). For predation this is implemented as follows. The prey benefits the predator by influencing D negatively and B positively. The predator harms the prey by positively influencing D. These influences depend on the size (*Nof*) of the populations and are mediated via the (*Effect*) quantities *supply* and *consumption*. The simulation shown in Fig. 5 starts with a scenario in which both prey (population2) and predator (population1) start at *normal* with an unknown direction of change: *Nof* = <*normal*, ?>. The state-graph has 13 states. States 1, 2, 3, and 4 are possible interpretations of the scenario. From those the following main behaviours are possible.

- **Balanced co-existence.** In state 2 the two populations have a natural balance; they co-exist without further changes.
- **Populations grow to maximum.** State 4 leads to 5 and shows the case in which both populations grow to their maximum size. Alternatively 4 may lead to 6 in which case only the prey grows to its maximum value.
- **Populations get extinct.** State 1 leads to 10, or to 12 via 11, showing the cases in which both populations get extinct.
- **Predator gets extinct.** State 3 leads to 7, to 8, or to 13 via 9, showing that the predator may get extinct while the prey survives. State 3 may also lead to 6, a situation in which the predator stabilizes at normal. Finally, state 1 may lead to 8 in which case the predator gets extinct but the prey does not reach its highest value.

Two issues should be pointed out with this simulation. First, one may wonder why there is apparently little ambiguity in the simulation results. Consider for instance D for the predator. It has multiple influences and could show ambiguity when for instance *supply* increases while *Nof* for the predator increases. D would then be ambiguous resulting in the simulator generating all possibilities for its derivative. If we consider state 3 then there should have been at least two other states with the same values for both sizes, but with different derivatives for D, namely decreasing, steady, and increasing. In fact this source of ambiguity would lead to a state-graph consisting of 48 states, that is, 35 additional states compared to the graph shown in Fig. 5 due to ambiguity in the B and D rates of both populations. Why does this not happen in our simulation? Although the additional 35 states represent correct predation behaviour, they do not add much to the insights explaining the main behaviours. This triggered the idea that the variations in the derivates of the rates should be made dependent on changes of the *Nof*. Put differ-

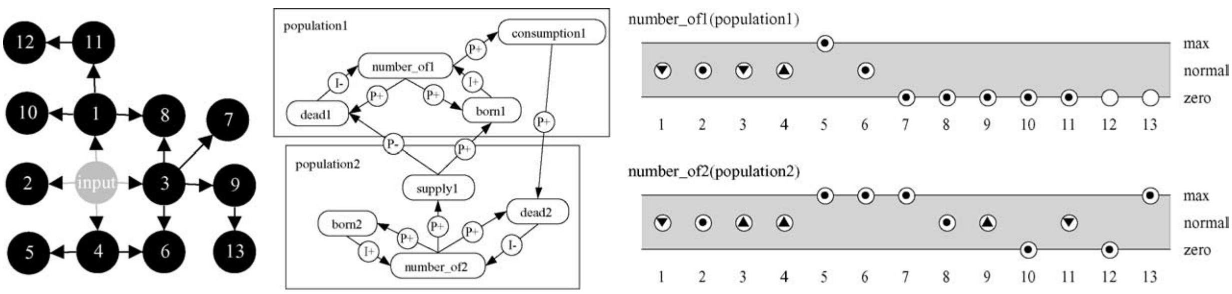

Fig. 5. Simulating predation based on processes determining single population behaviour.

ently, it is assumed that changes in the rates always follow the change in the population size. To implement this idea we introduce a new primitive: *derivative correspondence* (of two types, directed or undirected). It can be used to have the change in one quantity following the change of another quantity. By using this dependency between *Nof* and its D and B the superfluous ambiguity on rates is removed, leaving an easy to interpret simulation result.

In the simulation results shown in Fig. 5 (RHS) state 12 seems similar to state 10, and state 13 similar to state 7. The pairs differ because in state 12 and 13 the derivate of the *Nof* of the predator is unknown (designated by the empty white circles), while in states 10 and 7 this derivative is zero (steady). This reflects the difference between *known* and *unknown* information. Let us focus on state 13. Its predecessor is state 9 (see state-graph, LHS). In state 9 the predator already became extinct *Nof* $=$ <*zero*, *zero*>, while the prey still increases *Nof* $=$ <*normal*, *plus*>. Because of that, in state 9 there are no processes, and thus no influences, acting upon *Nof*. When moving from state 9 to 13 the prey changes values, but for the predator nothing changes. The predator's *Nof* magnitude is *zero* and stays *zero*, because its derivative was *zero* in state 9. But, how can we determine the derivative for *Nof* in state 13? There are no influences acting upon *Nof*, so strictly speaking the derivative is unknown, which is different from being able to calculate that it is *zero*. If migration was included, state 13 might be a state in which colonization could start. But as mentioned before, migration is excluded, and for simulating predation the distinction between unknown and being zero is not really important. In fact, to simplify matters it would be better if the simulator considered state 12 "equal" to 10, and 13 "equal" to 7. This can be realised using the idea of the "closed-world assumption" [6]. It states that for influence resolution we may assume that all relevant information is known and that unknown information can be considered to have no influence. Specifically this leads to the following rules for doing influence resolution. If Q_1 directly influences Q_2, thus I $+$ / $-$ (Q_2, Q_1), and the magnitude of Q_1 is unknown, then assume that it is zero and hence the influence of Q_1 on Q_2 becomes zero (the magnitude of Q_2 does not change). Similarly, if Q_1 indirectly influences Q_2, thus P$+$/$-$$(Q_2, Q_1)$, and the derivative of Q_1 is unknown, then assume that it is zero (steady) and hence the indirect influence of Q_1 on Q_2 is zero (and Q_2 remains steady). In GARP the user can decide to turn the closed-world assumption on or off (Table 2). Turning it on reduces the simulation for predation to 11 states, merging state 12 with 10, and 13 with 7. For the simulations discussed below, the closed-world assumption is set to being on.

4.3. The Ants' Garden – four interactions

Within the Ants' Garden there are at least four interactions between the species (Fig. 1). They map onto the basic interactions described above as parasitism, symbiosis, commensalism and amensalism (Table 1). Table 3 enumerates the causal dependencies that implement these interaction types and can be summarized as follows. For parasitism the parasite increases D for the host (*Consumption*), while the host increases B and reduces D for the parasite (*Supply*). For symbiosis both populations benefit (*Benefit1* and *Benefit2*), each other reducing D and increasing B. For amensalism, one population harms the other (*Pollution*), decreasing its B and increasing its D, while the influencing population itself remains unaffected. For commensalism, one population benefits the other (*Benefit*), increasing its B and reducing its D, while the influencing population itself remains unaffected.

Figure 6 shows the main simulation results for these models, except for parasitism because in the Ants' Garden model discussed here it is implemented in the same

Table 3

Causal dependencies implementing the interactions (RHS quantity influences LHS quantity)

Parasitism	Symbiosis
P + (Consumption, Number_of1)	P + (Benefit1, Number_of1)
P + (Supply, Number_of2)	P + (Benefit2, Number_of2)
P + (Dead2, Consumption)	P − (Dead2, Benefit1)
P − (Dead1, Supply)	P + (Born2, Benefit1)
P + (Born1, Supply)	P − (Dead1, Benefit2)
	P + (Born1, Benefit2)
Amensalism	Commensalism
P + (Pollution, Number_of1)	P + (Benefit, Number_of1)
P + (Dead2, Pollution)	P − (Dead2, Benefit)
P − (Born2, Pollution)	P + (Born2, Benefit)

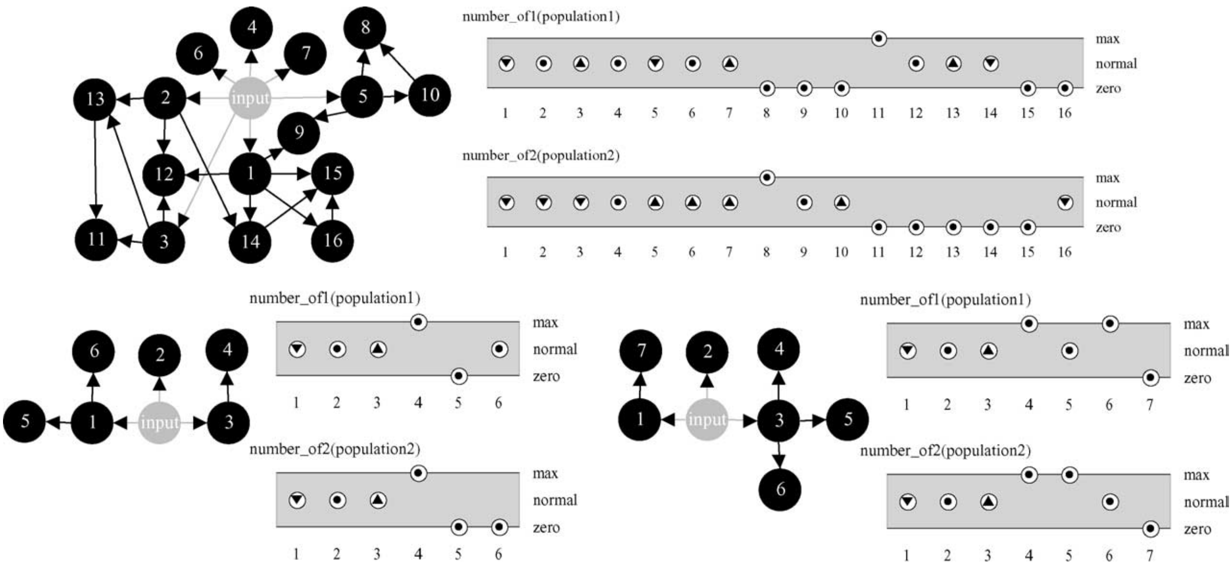

Fig. 6. Simulation results (state-graph and value-history) for the qualitative models of amensalism (top), commensalism (bottom, LHS), and symbiosis (bottom, RHS).

way as predation.[2] Notice that the main simulation results are the same as those presented in [24], but details differ because of the modifications discussed in the previous sections. Most notable difference is that in the models discussed here the *Nof* may not increase in its highest point value.

At this point all the ingredients required to simulate a scenario of the Ants' Garden are established. The only thing that needs to be done, is to actually construct a scenario and give it as input to the engine. Based on that, GARP will consult its library and find model-fragments detailing single population behaviour and interactions between populations. Based on the information found in the applicable model-fragments it will generate a state-graph with possible behaviours (if any exist). The backbone of GARP is a description of the entities involved and definitions of how they are related. This is achieved by using a subtype hierarchy representing the entities and a set of labels that specify structural connections between them. From these definitions the structural specifications in scenarios and model-fragments can be built. A scenario uses it to specify the structural organisation of the system that is subject of the simulation. Within model-fragments it is used to define the structural conditions to which the fragment applies. Figure 7 depicts the structural details for the scenario of the Ants' Garden. It reads as

[2]Another option (also implemented) is to have the *harm* affect both D and B(and not only D as done for predation).

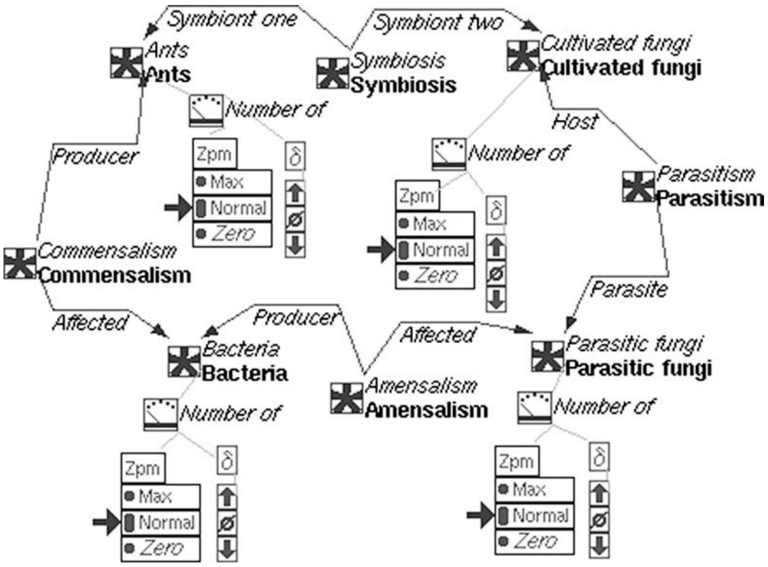

Fig. 7. Scenario for the Ants' Garden consisting of four populations and four interactions. All the populations start at magnitude *normal* and have an unknown derivative: *Nof* = <*normal*, ?>.

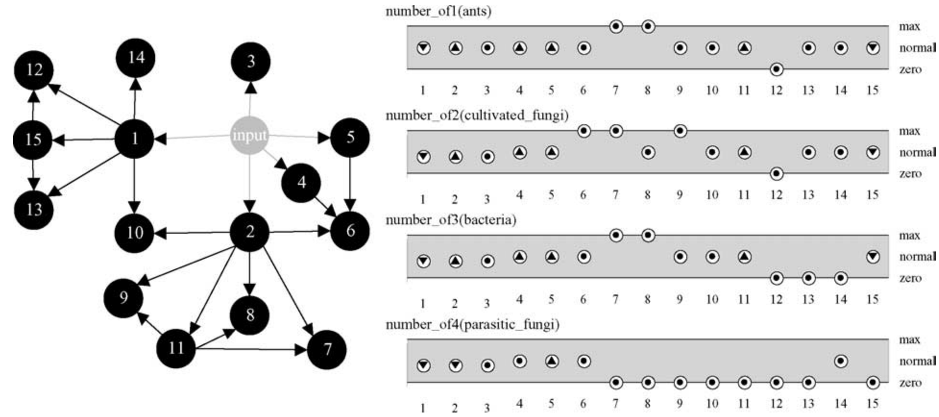

Fig. 8. Simulation results (state-graph and value-history) for the Ants' Garden (4 interactions).

follows. There are four populations[3]: ants, cultivated fungi, bacteria, and parasitic fungi. Each population starts with *Nof* = <*normal*, ?>. There are four interactions: symbiosis, commensalism, amensalism, and parasitism. For commensalism, ants produce the benefit and bacteria are affected. For amensalim, bacteria produce the harm and parasitic fungi are affected. For parasitism, cultivated fungi are the host and parasitic fungi the parasite. Finally, for symbiosis ants are symbiont one and cultivated fungi are symbiont two.

Feeding the scenario as shown in Fig. 7 to GARP results in the simulation shown in Fig. 8. The state-graph (LHS) has 15 states, from which 9 are end-sates (that is, states that do not produce any follow up): [3,6–10,12–14]. The values for the main quantities (*Nof*) are enumerated in the value-history (RHS). The initial scenario leads to four states [1–4]. Thus, according to the simulator there are four possible interpretations of this initial situation. In each state the magnitude of *Nof* is *normal* and the states differ on the derivatives calculated for the population sizes. In state 1, all populations decrease. In state 2, all populations increase except the parasitic fungi decreases. In state 3, all populations are steady. In state 4, all populations in-

[3]The fact that the ants and the other organisms are populations can be found in the subtype hierarchy. This hierarchy is not shown in Fig. 7.

crease except the parasitic fungi is steady. Following these initial states there are 17 possible behaviours: [3], [1 → 14], [1 → 12], [1 → 15 → 12], etc. The following main behaviours can be found in the state-graph:

- Different ways of coexistence, e.g., [3].
- Complete extinction of the garden, e.g., [1 → 15 → 12].
- Ants, bacteria, and cultivated fungi reaching their maximum size, e.g., [2 → 11 → 7].
- Elimination of parasitic fungi, e.g., [1 → 10].

Each state has approximately 35 model-fragments that specify behavioural details captured by the state. Consider for example state 1. For each single population 5 model-fragments are found: "population" (defines the structural details of a population, and is a condition for other population-related model-fragments), "assume_closed_population" (introduces B and D, and the related indirect causal dependencies, while ignoring migration), "existing_population" (*Nof* > *zero*, distinguishes it from an extinct population, as for instance in state 12), "mortality" (introduces the direct negative influence from D on *Nof*), and "natality" (introducing the direct positive influence from B on *Nof*). For each interaction type there are at least 3 model-fragments, e.g., for symbiosis: "symbiosis" (defines the structural details for the interaction to become active, introduces the main quantities, and is a condition for the other model-fragments detailing the interaction to apply), "symbiosis_interaction" (introduces the causal dependencies that implement the interaction), and "symbiosis_assumptions" (specifies interaction specific assumptions if needed). With respect to the latter, consider for instance commensalism. To set the strength of the benefit on the affected population we use two versions of commensalism: medium and high impact. Medium impact means that the benefit is partially responsible for changes in the other population. High impact means that changes in the other population are fully determined by the benefit. The "assumptions" model-fragment is used to represent such details.

The causal model that is assembled by these applicable model-fragments for state 1 is shown in Fig. 9. It shows the four populations, each being influenced by its basic processes (B and D), and the interactions between them (*Benefit*, *Supply*, *Consumption*, and *Pollution*).

4.4. The Ants' Garden – five interactions

The literature on the Ants Garden discusses the possibility of a fifth interaction in addition to the four already discussed. This interaction concerns bacteria

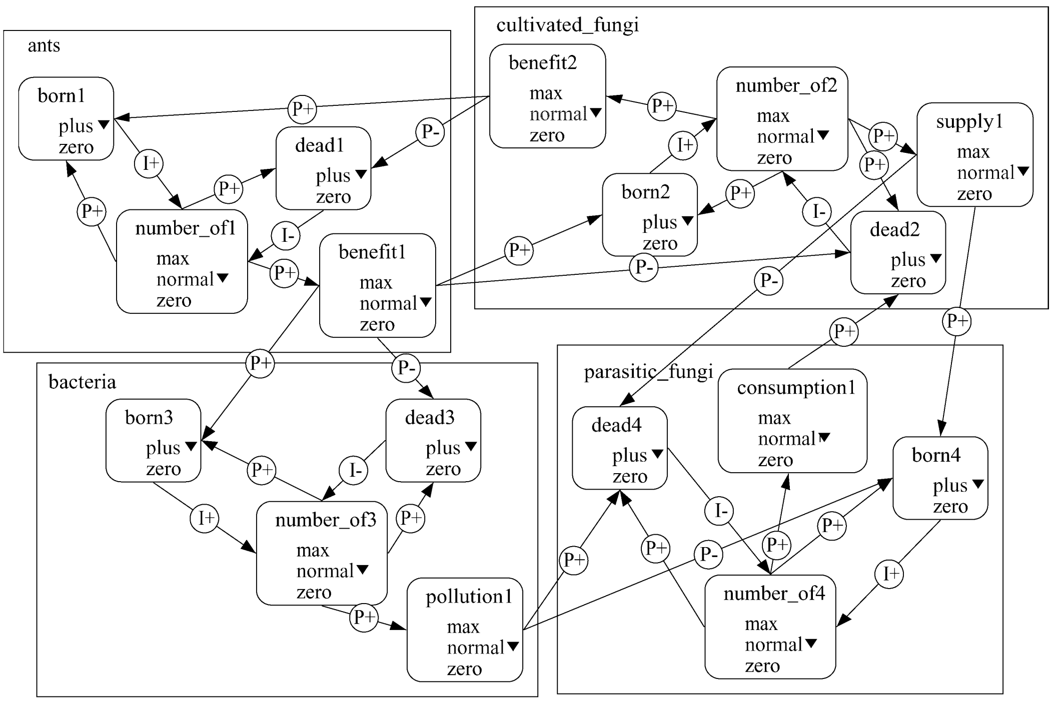

Fig. 9. Causal model of the Ants' Garden in state 1 (4 interactions).

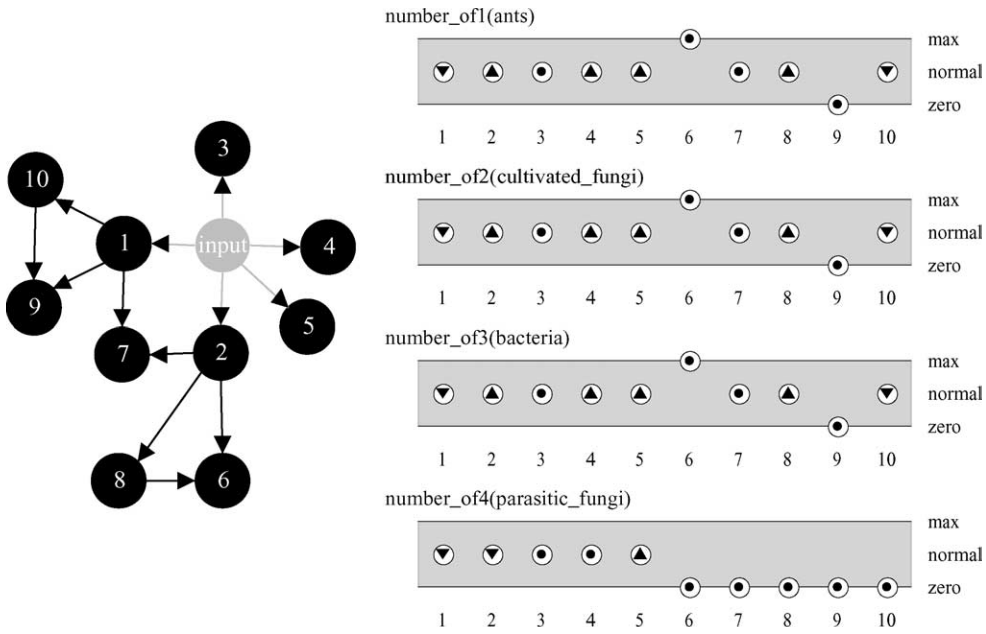

Fig. 10. Simulation results (state-graph and value-history) for the Ants' Garden (5 interactions).

producing metabolites that may enhance the growth of cultivated fungi [3]. To simulate this hypothesis we have augmented the scenario shown in Fig. 7 with a commensalism between bacteria and cultivated fungi. The simulation results for this scenario are shown in Fig. 10. The state-graph has 10 states of which 5 states are end-states: [3–6,9], and shows 9 behaviours. Again we can observe the typical behaviour of the Ants' Garden: different ways of coexistence (e.g., [3]), complete extinction of the garden (e.g., [1 $\rightarrow$ 10 $\rightarrow$ 9]), ants, bacteria, and cultivated fungi reaching their maximum size (e.g., [2 $\rightarrow$ 8 $\rightarrow$ 6]), and elimination of the parasitic fungi (e.g., [2 $\rightarrow$ 7]). The causal model for this simulation closely matches the one with the four interactions. The main difference being the additional benefit from bacteria to cultivated fungi when those two populations exist (as for instance in state 1), which introduces the following dependencies: P+(*benefit3*, *number_of3*) & P+(*born2*, *benefit3*) & P−(*dead2*, *benefit3*).

Compared to the state-graph resulting from four interactions the graph based on the five interactions is smaller. Apparently less behaviours are possible. This is in line with the understanding that ecologists have of communities, namely that there is a limit to the number of species that can in total interact and depend on each other. Finally, both models and their simulation results support the idea that our qualitative models of population dynamics and interactions can be reused and scale-up to model more complex situations.

5. Related work

Noble and Slatyer [18] and Moore and Noble [16,17] proposed an approach for building qualitative models about the dynamics of communities subject to recurrent disturbance (such as fire). This approach is based on a small number of attributes of the plant's life history (vital attributes) that can be used to characterise the potentially dominant species in a particular community, under different types and frequencies of disturbance. Simulations typically produce a replacement sequence that depicts the major shifts in composition and dominance of species that occur following a disturbance.

Guerrin and Dumas [9,10] describe models representing empirical knowledge of freshwater ecologists on the functioning of salmon spawning areas and its mortality in early stages, aiming at predicting and explaining the survival rate of fish under various scenarios. Their approach represents processes that occur at different time-scales (fast and slow) and introduces a real time dating and duration in a purely qualitative model.

State-transition modelling is often used to describe community dynamics. For example, McIntosh [15] presents a rule-based modelling language to describe succession in communities stressed by fire and grazing using this paradigm. Pivello and Coutinho [20] also describe a state-transition model about changes in the

Brazilian Cerrado vegetation under influences of fire, wood-cutting and grazing. Although successful in certain ways, these approaches do not explain *why* things happen, because there is no representation of the underlying mechanisms that cause changes in the system.

One of the most important studies on complexity and stability of biological communities was a qualitative analysis of the results produced by differential equation models about interactions between populations [11]. May's research question was to investigate what can be said if only the topological structure of the trophic web is known, i.e., knowing only the signs $\{-, 0, +\}$ of the interaction between the species and reasoning with changes over time, that is, with the derivatives of the quantities. May showed that the "common-sense wisdom" that more complexity means increased stability may not be true. In his simulations, a less complex community met the conditions for stability, while the more complex one was not stable. It is pointed out that this can be a useful approach for modelling quite complex food webs and to capture the general tendencies of the system, bypassing long and complicated steps required by numerical models. However, larger populations violate some of the criteria required for stability analysis and, in these cases, the signs of the interactions alone are not enough, and the interaction magnitudes should be taken into account.

6. Conclusions

Qualitative Reasoning techniques can be used to develop *conceptual* models of complex systems such as the Ants' Garden. These techniques provide a rich vocabulary to describe objects, quantities, relations, causality, situations, mechanisms of change, and conditions for changes to start and finish, and are thus important means to capture and simulate, and further develop, insights and explanations that experts have of systems and their behaviour. Compositional modelling is an important principle of qualitative reasoning. It refers to the idea of developing partial models (model-fragments) that capture the workings of elementary units so that the behaviour of larger and more complex systems can be generated by reusing these model-fragments. The models and simulations of the Ants' Garden presented in this article support this idea. The simulator automatically generates them by reusing a previously developed library of model-fragments on basic processes determining population growth and interactions. Although some details in this library have been improved, based on discussions with domain experts, the basic approach has not been changed. This is an important result both for research on Artificial Intelligence and on Ecology. It also supports the idea that our previously developed library of partial models is adequate and scalable to simulate the behaviour of complex system. For additional discussion on the impact of this work for Ecology, see Salles et al. [22].

Further research may focus on trying to model alternative interactions within the Ant's Garden to answer some of open questions that can be found in the literature. It would also be worthwhile to explore the behaviour of other communities, reusing and further improving the library develop thus far.

Acknowledgements

We would like to thank Waldenor Barbosa da Cruz for drawing our attention to this beautiful ecological problem and Nurit Bensusan for her support on the work presented here.

References

[1] A. Bouwer and B. Bredeweg, VisiGarp: Graphical representation of qualitative simulation models, in: *Artificial Intelligence in Education: AI-ED in the Wired and Wireless Future*, J.D. Moore, G. Luckhardt Redfield and J.L. Johnson, eds, IOS Press/Ohmsha, Osaka, Japan, 2001, pp. 294–305.

[2] B. Bredeweg, Expertise in qualitative prediction of behaviour, PhD thesis, University of Amsterdam, Amsterdam, The Netherlands, 1992.

[3] R.C. Currie, J.A. Scott, R.C. Summerbell and D. Malloch, Fungus-growing ants use antibiotic-producing bacteria to control garden parasites, *Nature* **398** (1999), 701–704.

[4] R.C. Currie, U.G. Mueller and D. Malloch, The agricultural pathology of ant fungus gardens, *Proc. Natl. Acad. Sci. USA* **96**(14) (1999), 7998–8002.

[5] B. Falkenhainer and K. Forbus, Compositional modeling: Finding the right model for the job, *Artificial Intelligence* **51**(1–3) (1991), 95–143.

[6] K.D. Forbus, Qualitative process theory, *Artificial Intelligence* **24** (1984), 85–168.

[7] M. Gillman and R. Hails, *An Introduction to Ecological Modelling: Putting Practice into Theory*, Blackwell Science, Oxford, 1997.

[8] M.E. Gilpin, Spinal chaos in a predator-prey model, *Am. Nat.* **113** (1979), 306–308.

[9] F. Guerrin and J. Dumas, Knowledge representation and qualitative simulation of salmon reed functioning, Part I: qualitative modelling and simulation, *BioSystems* **59** (2001), 75–84.

[10] F. Guerrin and J. Dumas, Knowledge representation and qualitative simulation of salmon reed functioning, Part II: qualitative model of redds, *BioSystems* **59** (2001), 85–108.

[11] R.M. May, Qualitative stability in model ecosystems, *Ecology* **54**(3) (1973), 638–641.

[12] R. May, Biological populations with non-overlapping generations: stable points, stable cycles and chaos, *Science* **186** (1974), 645–647.

[13] R. May, Simple mathematical models with very complicated dynamics, *Nature* **261** (1976), 459–467.

[14] R. May and G.F. Oster, Bifurcations and dynamic complexity in simple ecological models, *Am. Nat.* **110** (1976), 573–599.

[15] B.S. McIntosh, Qualitative modelling with imprecise ecological knowledge: a framework for simulation, *Environmental Modelling and Software* **18** (2003), 295–307.

[16] A.D. Moore and I.R. Noble, An individualistic model of vegetation stand dynamics, *Journal of Environment Management* **31** (1990), 61–81.

[17] A.D. Moore and I.R. Noble, Automatic model simplification: the generation of replacement sequences and their use in vegetation modelling, *Ecological Modelling* **70** (1993), 137–157.

[18] I.R. Noble and R.O. Slatyer, The use of vital attributes to predict successional changes in plant communities subject to recurrent disturbances, *Vegetatio* **43** (1980), 5–21.

[19] E.P. Odum, *Ecologia*, Rio de Janeiro, Discos CBS, 1985 (translation of *Basic Ecology*, 1983).

[20] V.R. Pivello and L.M. Coutinho, A qualitative successional model to assist on the management of Brazilian Cerrados, *Forest Ecology and Management* **87** (1996), 127–138.

[21] P. Salles and B. Bredeweg, Building qualitative models in ecology, in: *Proceedings of the 11th International Workshop on Qualitative Reasoning*, L. Ironi, ed., Instituto di Analisi Numerica C.N.R., Pubblicazioni no. 1036, Pavia, Italy, 1997, pp. 155–164.

[22] P. Salles, B. Bredeweg and N. Bensusan, The Ants' Garden: Qualitative models of complex interactions between populations, *Ecological Modelling* (in press).

[23] P. Salles and B. Bredeweg, Qualitative reasoning about population and community ecology, *AI Magazine* **24**(4) (2003), 77–90.

[24] P. Salles, B. Bredeweg, S. Araújo and W. Neto, Qualitative models of interactions between two populations, *AI Communications* **16**(4) (2003), 291–308.