

Model Building Experiences using Garp3: Problems, Patterns and Debugging

Jochem Liem and Floris E. Linnebank and Bert Bredeweg

Human-Computer Studies, Informatics Institute, University of Amsterdam, The Netherlands

Email: {J.Liem, F.E.Linnebank, B.Bredeweg}@UvA.nl

Abstract

Capturing conceptual knowledge in QR models is becoming of interest to a larger audience of domain experts. Consequently, we have been training several groups to effectively create QR models during the last few years. In this paper we describe our teaching experiences, the issues the modellers encountered and the solutions to solve them in the form of reusable patterns, and finally a structured way to debug models.

Introduction

Domain experts have been making more complex QR models the last few years. The models capture several processes and their interactions. However, different modellers seem to be reinventing modelling patterns to solve certain problems. This paper is meant to raise awareness in the model building community about the frequently encountered representational issues and possible solutions. The suggestions described in this paper are not prescriptive, but describe patterns that other modellers have found useful. As such, this paper is different than usual QR papers, since it is not focussed on algorithms, but instead offers modelling advice.

The paper introduces the groups we have trained the last few years, explains the representation used in the Garp3 QR modelling and simulation workbench, describes modelling and debugging issues and proposes solution patterns.

Modeller training

Three groups of modellers we have trained the last few years are particularly interesting.

- The first group are the (PhD-level) researchers who participated in the NaturNet-Redime EU project. These (non-computer scientist) domain-experts created models about ecology topics that they are actively researching.
- The second group are BSc. Future Planet Studies students who started their first semester of college education. As such, they had neither advanced computer-science knowledge, nor detailed knowledge about particular domains.
- The final group are PhD-students doing the SIKS research school Knowledge Modelling course¹. Most of them have backgrounds in fields close to computer-science.

The researchers have been working with the Garp3 software for about 2,5 years. Their training started at the second project meeting of the NaturNet-Redime project in Amsterdam, which consisted of 2 full days of hands-on practical sessions (including a 2 hour lecture), followed by a single day of working through the structured approach to building QR models (Bredeweg et al. 2008) (including a 2 hour lecture). Extra training was given during each following project meeting, which included a day of assignments and a day of debugging models in Sofia, a day of using the Sketch environment and the sharing and reuse functionality (Liem, Bouwer, and Bredeweg 2007) in Latvia, and a day of collaboratively improving the case study models in Germany. Additionally, the researchers were supported via bi-weekly Skype/Flashmeetings to discuss modelling issues. The results are complex models involving multiple interacting processes (Sánchez-Marrè et al. 2008).

The BSc. students doing an 8 week conceptual modelling course were divided in pairs. They spent the first 4 weeks learning to make Concept Maps (Cañas et al. 2004) and creating ontologies using Protégé (Knublauch et al. 2004). In the last 4 weeks the student pairs were learning to create QR models. The main goal for them was to create a small model about at least 2 processes relevant to the carbon cycle (and global warming). In addition to learning the QR technology, the students were asked to work towards this goal during these 4 weeks. Each week the students gave a 10-15 minute presentation about their current modelling progress towards the carbon cycle models. The students were supposed to spend about 8-10 hours a week on the course (including the weekly 3 hour practical session).

In first QR week, a 1 hour lecture was given contrasting QR models with concept maps and ontologies, explaining the general ideas of QR, and the applications of QR models. Following the lecture, the students worked on the Tree & Shade model (Bredeweg et al. 2006b). Nearly all students finished this exercise within the practical session. In the second QR lecture, the communicating vessels model (Bredeweg et al. 2006b) was used to explain the key representational aspects of QR models such as structure (entities and configurations), causality (proportionalities and influences), inequalities, correspondences and model fragments. In the rest of the session the students worked on recreating the population interaction model (Bredeweg et al. 2006b),

¹<http://hcs.science.uva.nl/SIKS/Siks2008>

which took them most of the assigned time. In the last two weeks, feedback was given on the models presented by students. In addition, we gave two extra 2 hour practical sessions to accommodate requests by students. Most students created excellent models focussing on two processes.²

The PhD students of the SIKS research school doing the Knowledge Modelling course³ got a 1 hour lecture on QR, followed by a two hour practical session in which they had to recreate the Tree & Shade model. The PhD students required on average about half an hour less time compared to the BSc. students to finish. This seems mostly to be due to their computer skills. However, those who also had some knowledge modelling skills were able to finish the modelling task the fastest (up to 30 minutes faster compared to the students without modelling skills).

QR Modelling and Simulation using Garp3

The introductory QR lectures are supposed to give the audience enough basic knowledge to allow some hands-on experience creating QR models using the Garp3 workbench⁴ (Bredeweg et al. 2006a). Garp3 allows modellers to capture their knowledge about the structure and the important processes in their system of interest as *model fragments*. These can be considered formalisations of the knowledge that applies in certain general situations. Model fragments can be considered rules indicating that if certain model ingredients are present (conditions), certain other model ingredients must also apply (consequences). The can be represented as: *conditions* \Rightarrow *consequences*.

Next to model fragments, different *scenarios* can be modelled. These represent specific start states of a system. Garp3 can run simulations of models based on a particular scenario. The result of such a simulation is a state graph, in which each state represents a particular possible situation of the system, and the transitions represent the possible ways a situation can change into another.

The state graph is generated by the simulation engine roughly as follows. The engine takes a scenario as input, and finds all the model fragments that apply to that scenario. The consequences of the matching model fragments are added to the scenario. The result is a state description including all consequences. Based on this augmented state description new knowledge can be inferred, such as the derivatives of quantities. Given a completed state description, the possible successor states are inferred. The complete state graph is generated by applying the reasoning to the new states.

The QR representation has a strict separation between structure and behaviour. The structure of a system is represented using *entities* (objects), *agents* and *configurations* (relations). For example, a lion hunting a zebra can be represented as two entities (lion and zebra) and a configuration (hunts). If the food web is considered to be the system, a hunter disturbing this system could be represented as an agent.

²This type of learning through modelling will be the main topic of research in the recently started DynaLearn EU project.

³<http://hcs.science.uva.nl/SIKS/Siks2008>

⁴<http://www.garp3.org>

One of the key behavioural model ingredients are *quantities*. Quantities represent the features of entities and agents that change during simulation. A quantity has a magnitude and a derivative, which represent its current value and trend. The magnitude and derivative are each defined by a quantity space that represent the possible values the magnitude and the derivative can have. Such quantity spaces are defined by a set of alternating *point* and *interval* values.

$M_v(Q_1)$ is used to refer to the current value of the magnitude of a quantity. $M_s(Q_1)$, the sign of the magnitude, indicates whether the magnitude is positive, zero or negative ($M_s(Q_1) \in \{+, 0, -\}$). $D_v(Q_1)$ is used to refer to the current value of the derivative of a quantity, which has to be a value from the predefined derivative quantity space ($D_v(Q_1) \in \{-, 0, +\}$). $D_s(Q_1)$ is used to refer to the current sign of a derivative.

As a shorthand to refer to the current magnitude and current derivative value of a quantity at the same time, we use the notation $Q[X, Y]$, where Q is the quantity, X is the current magnitude value and Y is the current derivative value. For example, $\text{Size}[+, -]$ indicates that the current size is positive and decreasing. This combination of the current magnitude and current derivative value is called the *quantity value*.

Causality

Important for QR models is the explicit notion of causality between different quantities. Garp3 represents the causal dependencies using direct and indirect influences (Forbus 1984). Direct influences, called *influences* for short, are represented as $Q_1 \overset{I\pm}{\rightarrow} Q_2$. Influences can be either positive (as above) or negative. The positive influence will increase $D_v(Q_2)$ if $M_s(Q_1) = +$, decrease it if $M_s(Q_1) = -$, and have no effect when $M_s(Q_1) = 0$. For a negative influence, it is the other way around.

The indirect influences, called *proportionalities*, are represented as $Q_1 \overset{P\pm}{\rightarrow} Q_2$. Similar to influences, proportionalities can be either positive or negative. The positive proportionality will increase $D_v(Q_2)$ if $D_s(Q_1) = +$, have no effect if it is stable, and decrease if it is below zero. For a negative proportionality, it is the other way around.

Other Behavioural Ingredients

Other behavioural ingredients essential for qualitative simulations in Garp3 are operators, inequalities, value assignments and correspondences. Operators (+ and -) are used to calculate the magnitude value of quantities (e.g. $Q_1 - Q_2 = Q_3$, to indicate $M_v(Q_1) - M_v(Q_2) = M_v(Q_3)$). Inequalities can be placed between different model ingredient types: (1) magnitudes ($M_v(Q_1) = M_v(Q_2)$)⁵, (2) derivatives ($D_v(Q_1) < D_v(Q_2)$), (3) values $Q_1(\text{point}(\text{Max})) = Q_2(\text{point}(\text{Max}))$ ⁶, (4) operator rela-

⁵Even if two quantities have the same qualitative value, they can still be quantitatively different (different points in the interval). An inequality can be used to indicate that they also have the same quantitative values.

⁶Values with the same name associated with different quantities do not necessarily have the same value. Points can represent to different quantitative values (e.g. the *maximum* heights of two

tions ($M_v(Q_1) - M_v(Q_2) < M_v(Q_3) - M_v(Q_4)$), (5) combinations of 1, 2, 3 and 4 (although magnitude and derivative items cannot be combined in a single expression). Value assignments simply indicate that a quantity has a certain qualitative value ($M_v(Q_1) = Q_1(Plus)$). Finally, correspondences are used to indicate that for certain values of one quantity, values of another quantity can be inferred. There are quantity correspondences ($Q_1 \overset{Q}{\leftrightarrow} Q_2$) and value correspondences ($Q_1(Plus) \overset{V}{\leftrightarrow} Q_2(Plus)$), which can both be either directed or undirected. The value correspondence indicates that if $M_v(Q_1) = Q_1(Plus)$ then $M_v(Q_2) = Q_2(Plus)$. If the value correspondence is bidirectional, the reverse inference is also possible. Quantity correspondences can be considered a set of value correspondences between each consecutive pair of the values of both quantities. There are also inverse quantity space correspondences ($Q_1 \overset{Q}{\leftarrow} Q_2$) that indicate that the first value in Q_1 corresponds to the last value in Q_2 , the second to the one before last, and so on.

Modelling Issues

Representing Structure

Entities or quantities? One of the main purposes of conceptual models is communication. QR models make an explicit distinction between structure and behaviour of a system to make models easier to understand. The quantities describing the behaviour of the system are attached to entities that describe the structure of the system. A balanced distribution between the number of quantities and the number of entities (i.e. only a few quantities per entity) improves the communicative value of a model.

The number of entities in a model should depend on the importance of those entities in the system. Otherwise they could be represented as quantities. For example, in the river restoration models (Sánchez-Marrè et al. 2008) we frequently noticed the use of quantities such as *oxygen concentration* and *Particulate Organic Matter (POM) concentration* as properties of an entity *river*. Since the POM and oxygen do not have important properties of their own for purposes of this model, they are modelled as quantities (the concentrations are properties of the river).

However, if we consider algae in the river, there is a modelling choice to be made. Algae can be modelled as an entity in the system (living in the river), or as a quantity of the river (Algae concentration). This choice depends on the importance of the Algae for the processes modelled in the system. For example, if the photosynthesis or biomass of the algae is important, Algae should become an entity with these features as quantities, since these quantities are features of the Algae and not of the river.

Configuration naming and direction In the investigation of the models created the last few years, it became apparent that in the modelling of the structure of a system, naming the configurations and choosing a direction is often experienced as being an issue. For example, when population A is preying on population B, is it better to

formalise this as Population A $\overset{\text{preys on}}{\longrightarrow}$ Population B, or as Population B $\overset{\text{is preyed on by}}{\longrightarrow}$ Population A?

This issue is analogous to writing in either active or passive voice. In our experience, the passive voice is frequently used. We propose that the active form should be consistently used for the naming of configurations. This shortens the configuration names, making the diagrams easier to read. Furthermore, if text is generated based on the contents of a model (e.g. a question generator or virtual character explaining the model), the quality of the text will be better.

Relationship reification There are relationships in systems that are difficult to formalise as configurations, since there are no verbs to describe them. For example, the Ants' Garden model (Salles, Bredeweg, and Bensusan 2006) describes the different interactions that populations can have with each other, such as commensalism, parasitism, and symbiosis. For parasitism a configuration *parasitises* could be defined, however no such verbs are available for commensalism and symbiosis. Using a configuration *lives in symbiosis with* seems suboptimal, since it has a long name and the direction seems arbitrary since the inverse is also true. Adding a second configuration to remedy this would only make the diagram more complex.

Another related issue is representing the speed of these processes, such as the parasitism rate (or other properties of the relationship). Assigning this rate (formalised as a quantity) to either of the populations participating in this relationship seems incorrect, as it is determined by the interaction of these populations, and not one particular population alone.

As a solution to these issues the relationship can be reified, i.e. represented as an entity. In the Ants' Garden model, the symbiosis relationship is described as an entity with *symbiont 1* and *symbiont 2* configuration relationships to the two populations. Although not in the Ants' garden model, the speeds at which these processes operate can be formalised as quantities attached to the reified relationships.

Representing causality

Choosing a proportionality or influence An important difficulty we encountered with all three groups is conveying the difference between influences and proportionalities. Moreover, even after having hands-on experience with creating models based on exercises, modellers still have trouble choosing whether to use an influence or a proportionality.

The key concept to understand is that only influences initiate change in a system and that proportionalities only propagate change. Specifically, the magnitude of the source quantity of an influence determines the derivative of the target quantity. As such, influences only cause change when the source quantity has a non-zero magnitude value. Proportionalities on the other hand determine the derivative of the target quantity based on the derivative of the source quantity, and thus only change the derivative of the target quantity when the source quantity is not stable.

We propose the following rule of thumb to decide whether an influence or a proportionality should be used when a modeller is sure that two quantities are causally linked. First, assume that the source quantity has a positive magnitude

container do not have to be equal).

value, but is stable (i.e. the derivative is zero). If the target quantity is supposed to change an influence should be used. Otherwise, a proportionality should be used. The reason this rule works is that a proportionality does not have an effect in this setting (since the derivative of the source quantity is stable), while an influence does since the magnitude of the source quantity is non-zero.

For example, consider water flowing from a tap into a bucket. The flow causes the amount of water in the bucket to increase. Should an influence or proportionality be used? Consider that the flow is positive but not changing. Since the amount of water should still increase, an influence should be used. The same rule of thumb can be used when considering the causal relation between the amount of water and the height of the water in the bucket. Consider that the amount of water is positive but stable. Since the height of the water should also remain stable, a proportionality should be used.

Causal chains Causal chains often start with an influence followed by several proportionalities that propagate the effect. Chains of proportionalities following each other occur quite often. These kind of causal chains are seen in many models. In contrast, it is unlikely that there is another influence in a causal chain. As such, causal chains with influences in them are more likely to be incorrect. Conceptually a causal chain should be seen as a process that affects several causally linked quantities. Other influences should therefore be part of other causal chains.

A special case of a causal chain is one that contains a loop of proportionalities. For example, $A \overset{P}{\rightarrow} B, B \overset{P}{\rightarrow} C, C \overset{P}{\rightarrow} A$. These loops of proportionalities should be avoided, as the value of the derivatives of these quantities can never be derived. The reason is that to derive the derivative of one of the quantities, the derivative of the quantity before it has to be known. However, to determine that derivative, the derivative of the quantity before that has to be known, etcetera.

Feedback loops A frequently asked question about QR models is whether feedback loops are supported by QR models. A feedback loop in a system is a situation in which the effect of a process will influence this same process. For example, the growth of a tree increases the size of the tree, but the size of the tree also increases the growth rate. Feedback loops frequently occur in QR models and are one of the most basic patterns that occur in most models. The mentioned tree example can be seen in the Tree & Shade model.

A feedback loop is represented in Garp3 by specifying an influence from a process quantity to a target quantity and a proportionality from the target quantity to the process quantity. For example, in the Tree & Shade model there is a positive influence from the growth rate of the tree to the size of the tree and a positive proportionality from the size of the tree to the growth rate. This pattern exactly captures the feedback loop in the system.

Such feedback loops do not have to be direct. There can be a causal chain from the process quantity through several quantities with the final quantity providing the feedback to the process quantity. One such example can be seen in the communicating vessels model. The flow in the pipe between

two containers has a negative influence on the volume of the liquid in the container (i.e. the flow reduces the amount of water). There are positive proportionalities from Volume to Height and from Height to Pressure to indicate that if volume changes, height will change in the same direction and if the height changes, the pressure will also change in the same direction. The feedback is represented in the form of a positive proportionality from the pressure to the flow. This proportionality indicates that the flow will increase if the pressure increases and decrease if the pressure decreases (as it will if water is flowing out of the container).

Causal Interactions As part of each introductory QR lecture we present the audience with a set of exercises in which two causal dependencies affect the same quantity. A member of the audience is asked what the resulting derivative value will be for the affected quantity. Each of the three groups of modellers had difficulty in deriving the correct derivative and explaining the result.

The exercises start with an exercise that tests whether the audience has understood the semantics of the causal dependencies. An example exercise is $Q1[-, 0] \overset{I}{\rightarrow} Q2[+, ?]$. The audience has to indicate that Q2 will decrease, since the magnitude of Q1 is negative and it affects Q2 through a negative influence. Several of the people in the audience are able to correctly derive the correct result and explain it to the rest of the audience.

In the following exercises the audience has to derive the derivative of the quantity that is affected by two causal dependencies. For example, $Q1[-, -] \overset{I}{\rightarrow} Q2[+, ?] \overset{I}{\leftarrow} Q3[-, -]$. The correct answer here is that the derivative of Q2 is ambiguous. The reasoning is as follows. Q1 has a negative magnitude which results in a negative effect on Q2 through the positive influence. Q3 has a negative magnitude value but influences Q2 through a negative influence. As a result the effect on Q2 is positive. Given that there is a positive and a negative result on Q2 the result is ambiguous.

Although not explained during the lecture, in more advanced modelling the ambiguity of this kinds of examples can be resolved by adding inequality knowledge. For example, knowing that $M_v(Q1) > M_v(Q3)$ allows us to derive a unique derivative value. Since the negative effect of the positive influence from Q1 is smaller (less negative) than the positive effect of the negative influence from Q3 (more negative), Q2 will increase.

Dealing with multiple competing causal dependencies

Many real-world problems involve multiple processes affecting single quantities. Although two competing influences of different types can be determined through a single inequality (see previous section), the more general case with multiple causal dependencies is more intricate. Consider two influences of the same type (both positive or both negative) affecting a single quantity, for example, the effects of release of CO_2 from the ocean (which can be negative to model the absorption of CO_2) and the burning of fossil fuels on the CO_2 concentration in the air. Given $release[-, +] \overset{I}{\rightarrow} concentration[+, ?] \overset{I}{\leftarrow} burning[+, +]$,

the derivative value of concentration is ambiguous.

Inequality knowledge between concentration and burning with not resolve the ambiguity. The knowledge that is needed is whether the absolute magnitude value of release is bigger or smaller than the absolute magnitude value of burning. In the former case, $D_v(\text{concentration}) = +$, while in the latter case $D_v(\text{concentration}) = -$. However, such representing absolute values and reasoning with them has not been solved in Garp3 yet.

A general pattern that can be used instead is specifying an inequality between the sum of all quantities with positive effect and the sum of all quantities with negative effects. In the example, the knowledge $\text{release} + \text{burning} < 0$ allow us to infer $D_v(\text{concentration}) = -$. If we also consider the effects of photosynthesis ($\text{photosynthesis} \stackrel{I}{\rightarrow} \text{concentration}$), we can again resolve the ambiguity by specifying the inequality $\text{release} + \text{burning} < \text{photosynthesis}$ ($D_v(\text{concentration}) = -$). This pattern allows modellers to specify what the result on the influenced quantity will be given a set of magnitude values (as conditions) of the processes.⁷

Correspondences as causality Correspondences are used to ensure that magnitude values always occur together. They are often paired with proportionalities. The correspondence assures that the magnitude values of the quantities always co-occur, while the proportionality ensures that the derivative values are equivalent (assuming there are no other causal dependencies on the quantity). In the case that there are multiple causal dependencies on the target quantity, a correspondence might be to strict.

Correspondences can be either directional or bidirectional. Directional correspondences are important when the magnitude value of one quantity can be inferred from another quantity, but not the other way around. For example, when the size of a population is zero, the birth rate should be zero, but the birth rate can be zero with positive population size. An example of a bidirectional correspondence is between the size of a population and their biomass.

Establishing Quantity Spaces

The selection of suitable quantity spaces for quantities is experienced as being a difficult task even by experienced modellers. A quantity space should contain just the right amount of distinctive values necessary to model the behaviour of that particular quantity in a qualitatively meaningful way. Inherent in this choice is the purpose of the model. For example, when modelling the effect of phytoplankton concentration on the amount of light it absorbs, modelling the critical concentration value when the other primary producers underneath them become significantly deprived of light in the quantity space is important. However, when modelling the effects of global warming on phytoplankton, this value is less important and could be left out of the model.

⁷When a quantity is affected by both a proportionality and an influence with opposing effects, the result is always ambiguous. Consequently, we argue that mixing different types of causal dependencies should be avoided. This rule is known as the *homogeneous influences adequacy constraint* (Rickel and Porter 1997).

When choosing a quantity space it is important to determine the qualitative distinct values a quantity can take that might cause a change in behaviour. This means thinking of particular value ranges in which a certain behaviour of the system takes place. These ranges are bounded by particular points that represent the thresholds between these ranges. This is the reason that quantity spaces in the QR representation consist of consecutively intervals (ranges) and points. In the above example about light, the concentration of phytoplankton could either be: no plankton, some plankton, a threshold representing the critical amount of plankton, and more than the critical concentration of plankton (e.g. {point(zero), positive, point(critical), hazardous}).

Note that the choice of the quantity space {zero, low, point(medium), high} for, for example, the size of a population is not ideal if medium is thought of as an interval (like low and high). Firstly, there seems to be no clear reason why this distinction is important from a behavioural point of view. Secondly, medium becomes a point value in this quantity space, and the behavioural properties of points are quite different than those of intervals. As a result certain behaviours of the system will not be simulated. The main reason for this is the epsilon ordering concept (de Kleer and Brown 1984), which indicates that changes from a point to an interval always have precedence over a change from an interval to a point. This means that a changing quantity can remain having the same interval magnitude value in consecutive states, but a changing quantity that is in a point value must change to the next magnitude value in the next state.

Consider two growing populations with size low. Given the quantity space discussed above there are only three possible behaviour paths. From the first state, there are three possible options, either the first population reaches medium first, the second population reaches medium first, or they reach medium simultaneously. Since medium is a point value, it is not possible for the population that is still low to reach medium before the other population has reached high, due to the epsilon ordering rule. As such, this possible behaviour is not captured in the model. Consequently, we argue that the choice for this quantity space should be avoided, and that in general modellers should make sure not to model intervals as points.

Actuator Patterns

Although conceptually changes in systems should either be caused by processes active in the system or by forces outside the system, there are several technical ways to initiate change within a QR model. For instance, it is possible to indicate that a certain quantity is always increasing. Several frequently occurring patterns can be used to initiate change in a QR model. We call these *actuator patterns*, since they put the system into action.

Process actuator Processes represent the causes of change within a system. Consider the Growth process ($\text{Growth} \stackrel{I}{\rightarrow} \text{Size}$ in a process model fragment *Growth*) represented in the Tree & Shade model. There are three variations of this actuator that are commonly used. In the simplest variation, the growth rate is simply assigned a positive mag-

nitude and a stable derivative through consequential value assignments. A drawback is that the growth of the tree cannot change and can never become zero. To resolve this issue, in the second variation a feedback is added between the size of the tree and its growth rate ($Size \overset{P\downarrow}{\rightarrow} Growth$). Consequently, no value assignments are needed in the model fragment, except a start value for the growth rate in the scenario. The third variation removes the need for the value assignment in the scenario. In the Growth model fragment a correspondence is added to indicate that a non-existing tree does not grow ($Size(zero) \overset{V}{\rightarrow} Growth(zero)$). Furthermore, a child model fragment is created that indicates that all trees grow ($M_v(Size) > zero \Rightarrow M_v(Growth) > zero$).

External Actuator Pattern and Exogenous Behaviour

The *external actuator* pattern models processes or effects of processes from outside the system. The patterns consists of an agent representing the source of the effect, and an associated quantity which represents an exogenous variable. "Human modelers treat a variable as exogenous only if it is approximately independent of the other variables in the model." (Rickel and Porter 1997). Garp3 allows exogenous behaviour to be specified for exogenous quantities (Bredeweg, Salles, and Nuttle 2007), which allows modellers to indicate that a quantity remains constant, is increasing, decreasing or steady, or has sinusoidal or random behaviour. Sinusoidal behaviour is used for cycles, such as day-night cycles, tides (monthly), and precipitation (yearly), while random behaviour is used for quantity behaviour that a modeller is unsure of and might unexpectedly change (e.g. rainfall over a shorter period of time).

There are two variants of the external actuator pattern. To model an external process (fully determined by forces outside the system) a quantity is combined with an influence. The influencing exogenous quantity tends to be set using a value assignment (as in the process actuator pattern), with either the derivative being set or determined by a feedback relationship. The second variant models the effects of external processes using an exogenous quantity and a proportionality. These external processes are often determined by giving the quantity an exogenous behaviour.

The choice between the two variants depends on what the exogenous quantity should do. For example, when the exogenous quantity fully determines a quantity in the system (e.g. with two corresponding large quantity spaces), this is modelled using a proportionality. For example, the nutrient run-off caused by farming fully determines the nutrient level in the Danube river and delta, and the average ambient temperature of the surrounding land determines the temperature in the river and delta (Sánchez-Marrè et al. 2008). In contrast, an exogenous process is used when an important process has to be modelled. For example, a fishery manager stocking young salmon in a river, or economical development activities increasing the number of anglers (Sánchez-Marrè et al. 2008).

Equilibrium Seeking Mechanisms The equilibrium seeking mechanism pattern models equalizing flows due to a potential difference. For example, energy exchange between two objects with different temperatures, or a

liquid flow equalizing the pressures in the communicating vessels system. Key in this pattern is the flow, which is determined by the difference between two state variables, e.g. of the temperatures of two objects ($Temperature1 - Temperature2 = Heat\ flow$). The heat flow reduces the heat from one object, and transfers it to the other ($Flow \overset{I\downarrow}{\rightarrow} Heat1, Flow \overset{I\uparrow}{\rightarrow} Heat2$). Finally, the two state variables determining the flow also determine the derivative of the flow. If the temperature of the first object increases, the flow will increase ($Temperature1 \overset{P\downarrow}{\rightarrow} Heat\ Flow$), while if the temperature of the second object would increase, the flow would decrease ($Temperature2 \overset{P\downarrow}{\rightarrow} Heat\ flow$). In the communicating vessels model, the pressure quantities determine the flow, while the flow changes the volumes of the water in the containers through influences.

Competing Processes The *competing processes* pattern consists of multiple interacting influences that model competing processes. There are at least two processes, such as the birth and death rate of a population, or more such as its immigration and emigration rates. The processes influence a single quantity, in this case the size of the population ($Birthrate \overset{I\downarrow}{\rightarrow} Size, Deathrate \overset{I\downarrow}{\rightarrow} Size$). There are also feedbacks: a larger population means a larger birth rate ($Size \overset{P\downarrow}{\rightarrow} Birthrate$) and a larger death rate ($Size \overset{P\downarrow}{\rightarrow} Deathrate$). More details on how to deal with these kind of interactions is explained in the Sections *Causal Interactions* and *Dealing with multiple competing causal dependencies*. The mentioned examples come from the 'Single population model with basic processes' model which is provided with the Ants' Garden model.

Issues when Running Simulations

Maximum simulation result Modellers often ask why their QR models generate so many states. One of the main reasons this question is asked is because modellers tend to underestimate the number of states a QR model can potentially generate. The maximum number of states that a model can generate is equal to the Cartesian product of all the quantity spaces of all the quantities in a model. So a model with 10 quantities with three possible magnitude values can generate at most $(3 \times 3)^{10} = 3486784401$ states, which is the number of magnitude values times the number of derivative values raised to the power of the number of quantities. As such, adding one quantity more to a model can potentially mean almost an order of magnitude more states (number of potential magnitude values times the number of derivative values). Note that this number includes only the possible different states due to different magnitude and derivative values and does not include different states due to different inequalities. So even more states are possible.

Successor states without correspondences A frequently seen reason for a large number of states is non-corresponding quantities. Consider that all changing quantities in a point value will change to an interval value in the next state due to the epsilon ordering rule (which states that changes from a point value to an interval are immediate).

Given a state in which quantities all have interval values, often a large number of successor states result if the quantities do not correspond in certain way. The reason is that for each quantity it is possible for it to either change or remain the same. Consequently, there is a successor state for each combination of changing or not changing quantities. The number of combinations for such binary variables is 2^n , however the combination in which no quantities change is not a successor but the state itself. As such, for a single state the number of successors s given a number on non-corresponding changing quantities q can be calculated though $s = 2^q - 1$.

Constraining behaviour Given that a model potentially results in an unusable large state graph, it is essential that its behaviour is constrained. Technically all behavioural relationships between quantities constraint behaviour, however correspondences and proportionalities are especially appropriate. Given two non-corresponding quantities, each combination of magnitude values is possible. Adding a correspondence assures that only each corresponding pair of values is possible. Also adding a proportionality removes the potential of the two quantities changing independently of each other (given that there is no other causal dependency on the targeted quantity). This combination of ingredients makes quantities behave equivalently, and thus allows them to be counted as a single quantity for purposes of determining the maximum number of states.

Inequality statements also help constrain the behaviour.⁸ For example, specifying that the birth and death rates are above zero when the population size is above zero removes behaviour. For purposes of simulation it might also be insightful to specify fixed values or ranges for quantities. These are modelled by adding new model fragments that indicate that if a specific assumption holds certain (in)equalities hold for quantities. For example, in the R-star model (Nuttall, Bredeweg, and Salles 2005), when the assumption 'Limited resource build-up' holds, the resources available to the plant population are smaller or equal to medium.

Inactive model fragments Modellers frequently ask why certain model fragments do not become active during their simulations. Modellers usually know that the reason is that certain conditions in their model fragment are not fulfilled by the state. However, their real question is how they can determine which conditions are not fulfilled. In many cases we encountered that there is a mismatch between the model fragment and the state (or scenario). For example, the direction of a configuration is reversed. Our advice is to rebuild the state as a scenario and try to run the simulation. Usually the inconsistency is detected in this process. In the other cases the scenario can be changed to determine what the inconsistency is.

⁸Modellers should take note that constraints should make sense for a domain perspective. For example, when a heater heats a pan, the heat of the heater cannot be set to stable, as this would make it impossible for the heat flow process to take heat from the heater. To make the stable heat possible, there should be at least another competing process that adds heat to the heater.

No states Beginning modellers often find it difficult to solve simulation results with no states. Having no states always means at least one model fragment was considered. Otherwise the simulation result would consist of at least one state which corresponds to the scenario. The inconsistency is caused either by the scenario and a matching model fragment, or by a combination of matching model fragments. The easiest way to find the inconsistency is by making all model fragments inactive. Consequently, there should be at least one state corresponding to the scenario. Model fragments can be activated one by one to detect the model fragment that causes the inconsistency.

Not all expected states Sometimes simulations do not generate all the expected states. Our advice to improve the model is creating a scenario that corresponds to the expected state. If the simulation results in no states there is an inconsistency that has to be resolved. Then, modellers should create a scenario that represents a predecessor state and determine if both these desired states are generated⁹. By working backwards in this way towards an already generated state allows the desired branch of behaviour to be simulated.

Inconsistencies Inconsistencies are caused by inconsistent inequalities. Determining which inequalities are inconsistent is a difficult issue and a topic on its own. The following is a list of sources of inequalities that should be checked when searching for the reason of an inconsistency.

- Magnitude (or derivative) value assignments in model fragments (or scenarios).
- Inequalities explicitly represented in model fragments or scenarios.
- Inequalities resulting from the calculations of operators (plus or minus). These calculations result in an inequality indicating that a quantity has a certain magnitude (or derivative) greater, smaller or equal to a specific value.
- Value assignments caused by correspondences. When a quantity A has a certain value, the corresponding quantity B also has to have that specific value.
- Value assignments resulting from influence resolution. The result of resolving of influences and proportionalities is a set of value assignments indicating whether quantities are increasing, stable or decreasing.
- Value assignments resulting from advanced (exogenous) quantity behaviour. It is possible to specify advanced quantity behaviour for specific quantities in scenarios. For example, a quantity can increase, change randomly, or move as a sinus. This behaviour sets the derivative of the quantity. It is also possible to generate all the possible magnitudes of a quantity.
- Another source of facts are engine rules. These rules indicate what is possible in a simulation results, and always apply. Engine rules impose these constraints by imposing inequalities. The most important rules to consider are the quantity constraints and the continuity constraints.

⁹Showing termination nodes (potential successor states) is helpful here.

- The quantity constraints simply specifies that each quantity space has to have a value within its quantity space. This is usually represented by two inequalities. The first indicates that it has a value greater or equal to its top value, and the second indicates that it has a value smaller or equal to its top value.
- The continuity constraints is a transition rule that indicates that a magnitude or derivative has to gradually change, e.g. a derivative cannot change from increasing to decreasing without passing through stable. For a derivative this would result in an inequality that indicates that the derivative is smaller or equal to zero when a quantity is decreasing. An example of when the continuity rule can cause conflicts is when one of a pair of opposing influences disappears.
- A special source of inequalities are the simulation preferences. These simulation preferences can be changed in the simulation preferences window. The most notable to consider are the two extreme values rules.
 - The 'Apply quantity space constraints on extreme values' rule indicates that the derivative of quantities has to be smaller or equal to zero (cannot increase) in their top magnitude value (if it is a point), and is greater or equal to zero (cannot decrease) in their bottom magnitude value (if it is a point). This rule applies to all extreme point values except zero.
 - The 'Apply quantity space constraints on zero as extreme value' applies the 'Apply quantity space constraints on extreme values' rule for zero as an extreme point value.

Conclusions and Future Work

This paper identifies frequently occurring model building issues, misconceptions and suboptimal modelling, and provides solutions, patterns and modelling advice. The issues and patterns originate from well-established models made by the groups we have trained over the last few years. We aim to contribute to the building of qualitative models raising awareness about the issues with model builders and providing them with the means to resolve them. In the coming years we will focus on providing better software support on resolving the presented issues and making frequently used patterns easier to represent.

Acknowledgements

The research presented in this paper is co-funded by the EC within FP6 (2002-2006) (project NaturNet-Redime, number 004074, www.NaturNet.org), and FP7 (2007-2013) (project DynaLearn, 231526, www.DynaLearn.eu). We thank our project partners, the students participating in our courses, and our reviewers for their insightful comments.

References

Bredeweg, B.; Bouwer, A.; Jellema, J.; Bertels, D.; Linnebank, F.; and Liem, J. 2006a. Garp3 - a new workbench for qualitative reasoning and modelling. In Bailey-Kellogg,

C., and Kuipers, B., eds., *20th International Workshop on Qualitative Reasoning (QR-06)*, 21–28.

Bredeweg, B.; Bouwer, A.; Liem, J.; and Salles, P. 2006b. Curriculum for learning about QR modelling. Technical Report Deliverable D6.9, STREP project FP6 no. 004074, NaturNet-Redime.

Bredeweg, B.; Salles, P.; Bouwer, A.; Liem, J.; Nuttle, T.; Cioaca, E.; Nakova, E.; Noble, R.; Caldas, A. L. R.; Uzunov, Y.; Varadinova, E.; and Zitek, A. 2008. Towards a structured approach to building qualitative reasoning models and simulations. *Ecological Informatics* 3(1):1–12.

Bredeweg, B.; Salles, P.; and Nuttle, T. 2007. Using exogenous quantities in qualitative models about environmental sustainability. *AI Communications* 20(1):49–58.

Cañas, A. J.; Hill, G.; Carff, R.; Suri, N.; Lott, J.; Gómez, G.; Eskridge, T. C.; Arroyo, M.; and Carvajal, R. 2004. Cmaptools: A knowledge modeling and sharing environment. In Cañas, A. J.; Novak, J. D.; and González, F. M., eds., *Concept maps: Theory, methodology, technology. Proceedings of the first international conference on concept mapping*, 125–133.

de Kleer, J., and Brown, J. S. 1984. A qualitative physics based on confluences. *Artificial Intelligence* 24(1-3):7–83.

Forbus, K. D. 1984. Qualitative process theory. *Artificial Intelligence* 24(1-3):85–168.

Knublauch, H.; Ferguson, R.; Noy, N.; and Musen, M. 2004. The protege owl plugin: An open development environment for semantic web applications. In McIlraith, S. A.; Plexousakis, D.; and van Harmelen, F., eds., *International Semantic Web Conference*, 229–243. Hiroshima, Japan: Springer.

Liem, J.; Bouwer, A.; and Bredeweg, B. 2007. Collaborative conceptual modeling: Share, search and reuse. In Price, C., ed., *21st International Workshop on Qualitative Reasoning (QR'07)*, 71–78.

Nuttle, T.; Bredeweg, B.; and Salles, P. 2005. R-star - a qualitative model of plant growth based on exploitation of resources. In Hofbaur, M.; Rinner, B.; and Wotawa, F., eds., *19th International Workshop on Qualitative Reasoning (QR'05)*, 47–53.

Rickel, J., and Porter, B. 1997. Automated modeling of complex systems to answer prediction questions. Technical report.

Salles, P.; Bredeweg, B.; and Bensusan, N. 2006. The ants garden: Qualitative models of complex interactions between populations. *Ecological Modelling* 194(1-3):90–101.

Sánchez-Marrè, M.; Béjar, J.; Comas, J.; Rizzoli, A.; ; and Guariso, G., eds. 2008. *Proceedings of iEMSs 2008: International Congress on Environmental Modelling and Software Integrating Sciences and Information Technology for Environmental Assessment and Decision Making: Workshop 5 - Qualitative models for sustainability*. p1971-2005.