

Model-based Decision Support Systems - An Application to Water Treatment

P. Struss^{1,2}, W. Roque³, P. Salles⁴, U. Heller⁵

¹ Technische Universität München, Department of Computer Science,
Boltzmannstr. 3, D-85748 Garching, Germany
struss@in.tum.de

² OCC'M Software GmbH, Gleissentalstr. 22, D-82041 Deisenhofen, Germany
struss@occm.de

³ Instituto de Matemática, Universidade Federal do Rio Grande do Sul,
Porto Alegre, Brasil
Roque@mat.ufrgs.br

⁴ Instituto de Ciências Biológicas, Universidade de Brasília, Brasília, Brasil
Psalles@unb.br

⁵ ESG GmbH, Einsteinstr 174, Munich, Germany
Heller@in.tum.de

Abstract The paper presents some foundations for the application of model-based systems technology within a decision support system for drinking water treatment plants. This system aims at detecting deviations from normal plant operation, identifying their possible causes, and proposing adequate remedial interventions. Its basis is a library of model fragments that represent intended processes, disturbances, and possible interventions. Based on these fragments and the available observations, possible models of the disturbed plant behavior are generated automatically as a solution to the diagnostic task. An extension of such a model by models of interventions such that the result is consistent with remedial goals represents a possible therapy proposal. We discuss in more detail how the model and consistency-based problem solving can be exploited in the basic reasoning steps.

1 Introduction

Knowledge-based systems provide a means for making expert knowledge available to non-experts or specialists lacking comprehensive knowledge of a certain domain. This feature is of particular interest in the domain of ecology and environmental problems because every individual, institution, or company affects the environment (usually in a destructive way) but cannot be expected to have a proper understanding of this impact and of ways to prevent or counteract it. Environmental decision support systems have to capture domain knowledge of experts in the field of ecology and/or environmental issues and facilities for enabling the users to state information about

their particular problems, for generating answers and solutions based on this information and the domain knowledge, and for presenting them in a comprehensible manner.

We chose model-based systems to approach this ambitious goal. Some fundamental assumptions underlying our work are:

- The domain knowledge can be represented by a set of generic, independent, and, hence, re-usable model fragments (“processes”) that describe the relevant phenomena and are collected in a library (“domain theory”).
- The user faces two distinct tasks (**Fig. 1**): situation assessment (understanding “What goes on?”) and therapy proposal (“What can be done?”).
- A proper answer to situation assessment is given by a model that can be composed from the library and “explains” the partial information about a situation that is available to the user, and, similarly,
- an adequate therapy can be found by finding an extension of the situation model by feasible actions that satisfies a set of behavior goals.

The basis of our solution is a novel integration of logical theories and implementations of process-oriented modeling ([Forbus 84]) and consistency-base diagnosis ([de Kleer-Mackworth-Reiter 92], [Dressler-Struss 96]) in a generalized diagnosis engine, called G⁺DE. In this paper, we will not restate the formal theories and technical details which have been described in [Struss-Heller 01] and [Heller-Struss 02], but just summarize the concepts and focus on discussing how a decision support system can be based on them.

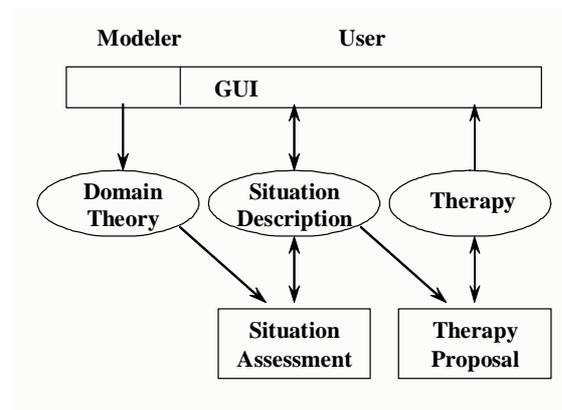


Fig. 1. Architecture of the decision support system

The application context for our work is a water treatment plant. While this is not exactly an environmental system, it seems to be a good starting point for a first practical evaluation of G⁺DE: On the one hand, it comprises physical, chemical, and biological processes with only partial and qualitative knowledge and information available, like an ecological system. On the other hand, it has a fixed structure and a limited set of relevant phenomena which makes it a good target for a first trial.

Furthermore, it is possible to validate and evaluate the decision support system for operators in real plants, and there is a smooth transition to real environmental and ecological systems by including phenomena that affect the natural water sources (e.g. algal bloom as in [Struss-Heller 01]).

In the next section, we briefly describe the application background, the water treatment process. In section 3, we summarize the modeling formalism. Based on this, we present the situation assessment and the therapy proposal components. Finally, we discuss some open issues we encountered regarding the use of deviation models for diagnosis and therapy.

2 The Treatment Process

The water to be treated is collected from a natural source like a river (e.g. Rio Guaíba) or a reservoir (e.g. Lomba do Sabão). After having been pumped to the water treatment plant, it passes through a sequence of process steps which we briefly describe below:

Arrival chamber

Here, the amount of algae or mollusks is reduced. The color (green or brown) indicates the presence of algae. The remedy is to add an oxidation agent (Copper Sulphate, Activated Carbon or Ozone).

Arrival canal (Pre-chlorination and pre-alkalinization)

If the alkalinity ≤ 19 and $\text{pH} \leq 6.5$, then actions are applied such that alkalinity ≥ 20 and $6.5 \leq \text{pH} \leq 7$. If necessary, first chlorine is added to the system to react with the organic materials and remove them.

Chicanery

The water goes through three processes: reaction with Aluminium Sulphate ($\text{Al}_2(\text{SO}_4)_3$), stirring of the water and coagulation. The objectives are the removal of dissolved substances (NH_3 , NO_2 , NO_3 etc.) and solid (undissolved) substances (algae, organic material, soil, salts etc). $\text{Al}_2(\text{SO}_4)_3$ dissociates in the water producing positive charges (Al^{+++}) and negative charges (SO_4^-) which attract opposite charges of dissolved and undissolved substances and form the flocks. The water is stirred in order to create bigger flocks (by agglutination of smaller ones). This is achieved by a slope and curves and obstacles in the tank. Thus, the probability of contact of molecules with different charges increases resulting in an increased size of the flocks. Separation between solid and liquid phases, that is, flocks and water. The water flows through the tank at a low speed. During this period, the flocks sink to the bottom of the tank, forming a growing layer. At the end of this step, color and turbidity are measured.

Decantation water canal

Chlorine is added to the water again (interchlorination) in order to guarantee the oxidation of any organic matter that has not been removed and to leave free chlorine in the water. This free chlorine will kill the bacteria that may still be found in the water.

Filters

During this stage, small flocks that have not been removed and, perhaps, algae and bacteria (micro organisms) are retained in layers of sand. The objective is to remove material and, thus, reduce color and turbidity. When the water leaves the filters, it has to be within the legal standards.

Filtered water canal

Final adjustments are made. Alkalinity, pH, color, turbidity, free chlorine concentration is measured at the exit of the filters.

3 Modeling

In the following, we summarize the modeling formalism that provides the foundation of the decision support system. According to the compositional modeling approach, this is split into two parts:

- The domain theory which represents the general knowledge in terms mathematical axioms and a library of behavior constituent types (i.e. generic processes).
- The situation description capturing the information about a specific system in a specific state stated in terms of perceivable objects, their interrelationships, and values of object variables (e.g. measurements) (see Table 1).

We follow the principles of structure-to-behavior reasoning and compositional modeling and provide a generalization of both component-based and process-based modeling paradigms. According to this view, the system model (or system description) consists of two parts: the domain theory and a situation description. The diagram in Table 1 provides an overview, and we briefly discuss each part.

3.1 Domain Theory

The domain theory captures what we know about the domain, i. e. all systems of a certain class (e. g. hydrological ecosystems or water treatment plants). We distinguish *structural elements* (objects and relations) from *behavior constituents* (which might be processes or other model fragments). The ontology consists of

Table 1. Structure of system models

Domain Theory
Ontology
Object Types (hierarchical)
Object Relations
Quantity Types
Behavior Constituent Types
Structural Conditions (objects and relations present)
Quantity Conditions (constraints on quantities)
Structural Effects (objects and relations created)
Quantity Effects (constraints and influences)
Basic Axioms
Situation Description
Objects
Relation Tuples
Quantity Value Assignments

- *object types* which occur in structural descriptions, for instance types of components in a device (resistor, broken wire), spatially distinguished entities (layers of a water body, pipes, tanks), etc. Object types can be structured hierarchically.
- *relations* for characterizing "configurations" of objects. Examples are spatial relationships (contained-in, below), connectivity of components, etc. Some important properties of relations (like uniqueness) can be specified.
- *quantities* as the basic elements for behavioral descriptions. Different quantity types (with different domains) can be defined and objects of a given type can be supplied with a number of associated quantities with given roles (e. g. the resistance of a resistor, the concentration of dissolved iron in a water tank etc.).

The domain theory also has to provide a vocabulary for behavior descriptions and the inferences that derive behavioral constituents from a structural description. It introduces

- *behavior constituent types*. These are physical phenomena which are considered to contribute to the behavior of the overall system. They can represent basic component laws (Ohm's Law, logical-or) or processes like in the Qualitative Process Theory, QPT ([Forbus, 1984]). Examples are alcalinization, water transport or algal blooms. They occur deterministically under certain conditions, and their occurrence generates particular effects.

Applying the distinction between structural and quantity aspects to both conditions and effects, we obtain

- *structural conditions*: assertions about the existence of relations and objects (e. g. of sedimental iron)
- *quantity conditions*: statements about values of quantities (e. g. a low pH in the reservoir)

- *structural effects*: creation or possibly even elimination of objects and relations (e. g. the "generation" of dissolved iron from the bound sedimental one)
- *quantity effects*: can be expressed as restrictions on variables (e. g. the dissolved iron concentration rises with the sedimental iron concentration and lower pH). Here, we also allow for partially specified effects in the form of influences as in QPT.

The abstract form of a behavior constituent type can then be written as

$$\text{StructuralConditions} \wedge \text{QuantityConditions} \Rightarrow \text{StructuralEffects} \wedge \text{QuantityEffects}$$

More precisely, we state that for each constellation of objects satisfying the structural and quantity conditions, an instance of the behavior constituent occurs and imposes the respective effects on the constellation. Fig. 2 shows a process in a graphical and a textual notation.

Additionally, we include a section for the fundamental laws that determine the mechanisms of model formation, how influences combine and prediction over time (continuity, integration etc.). These "basic axioms" cannot be specified arbitrarily by the modeler but rather represent the domain independent laws like the one of behavior constituent occurrence stated just above.

At this point, we make almost no commitment w. r. t. the quantity domains (symbolic, qualitative, real, ...), the formalism for specifying the quantity effects (constraints, differential equations, ...), and the expressiveness of structural conditions and effects (e. g. non-existence of certain objects as condition or destruction of objects as a structural effect).

3.2 Situation Description

A particular system under consideration is characterized by its *object structure*, i. e. instances of the object types and individual tuples of object relations (for instance the components and the connection structure of a device). In the following, we will refer to both objects and relation tuples as structural elements.

A particular situation of the system is characterized by *quantity value assignments*. Dependent on the task and context, they may represent actual measurements (e.g. an increased amount of iron in the drinking water), specification of goals (a certain amount of iron), mere hypotheses, etc.

3.3 Representing Actions

Actions may have certain physical preconditions for their applicability which can be stated in terms of structural conditions and quantity conditions. However, unlike ordinary processes, they do not become automatically active when these physical preconditions are satisfied. They have an additional precondition to become effective: some human intervention. Such an intervention can be an entire sequence of human activities (filling a container with some substance, connecting it to a treatment tank by

a pipe, and opening a valve) which has to be taken into account when actually planning the work or estimating its costs. However, from the point of view of reasoning about adequate therapies, it suffices to regard them as atomic entities.

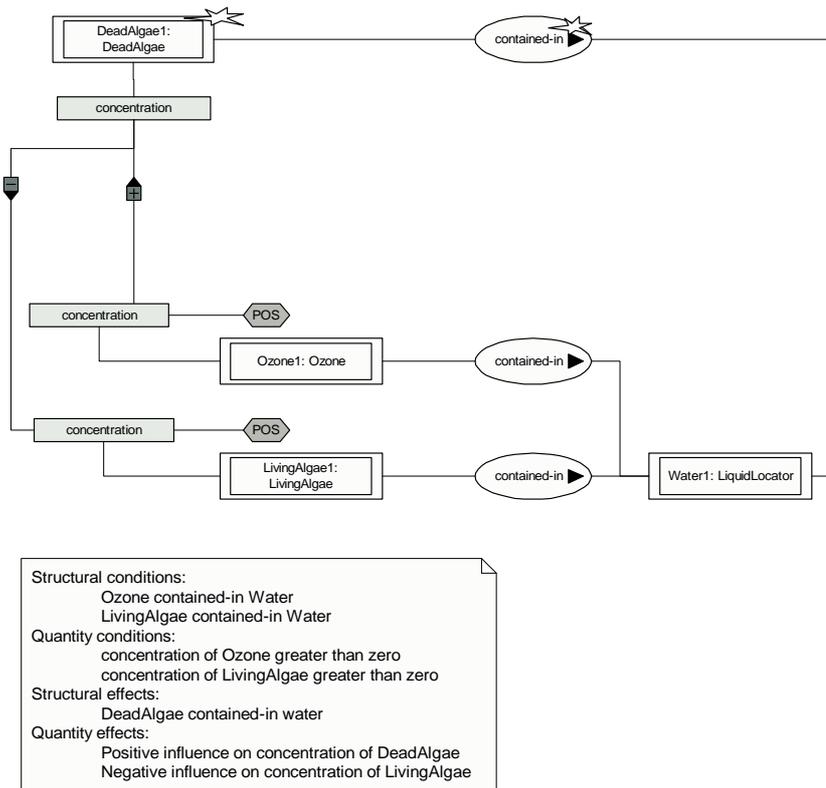


Fig. 2. Process „Ozone killing algae“

Hence, an easy way to integrate actions in the modeling formalism is to represent the human interventions as a special kind of objects, called **action triggers**, whose existence is a structural condition of the respective actions.

Since action triggers only depend on the decision and the respective activities by humans, they can never appear as structural effects of processes or other actions. Furthermore, we have to make sure that different action instances have different action trigger objects, even if they are instances of the same action type. Otherwise, several instances of an action type could be triggered by the same object. One way to achieve this within our modeling formalism without additional concepts is to guarantee that action triggers have a location that is unique to every instance of an action type. For instance, if a container has several connections to other tanks, the trigger of an *open_connection* action has to be specific to each opening rather than the container, because in the latter case, the opening of one connection would also trigger the opening of all others.

3.4 Deviation models

Often, it is not even relevant to situation assessment or therapy proposal to consider the absolute values of quantities. Rather, it can be sufficient to reason in terms of (qualitative) deviations from nominal values only. For example, they can be used to express that a chlorine supply which is higher than normal tends to make the pH lower than the specified range.

Descriptions of deviations can reflect the fact that it may be unnecessary or impossible to specify the normal behavior exactly and numerically. For each variable, such a deviation can be represented as

$$\Delta x := x_{\text{act}} - x_{\text{ref}}$$

The deviation models can be generated from the absolute models and propagate deviations from some nominal or reference behavior (which is possibly left unspecified). In the situation assessment phase, one may start from a deviation from the goal state which is indicated by measurements and try to find deviations of other quantities that cause the former.

4 Situation Assessment Component

As stated earlier, the first task to be solved is to determine hypotheses about the current situation based on available information about the system, i.e. observations. If the observations are true statements about all relevant objects involved in the system, situation assessment would just have to determine the active processes entailed by them. In general, observations are incomplete and may contain uncertain information. For instance, if the iron concentration of the incoming water is not measured, then a statement about iron is either missing or only a guess (e.g. a default value of the concentration). In this case, situation assessment has to complete the description provided by the user (e.g. hypothesizing the existence of iron) and/or revise uncertain information (e.g. the default concentration).

Therefore, we allow for user-defined assumptions to qualify quantity assignments. Assumptions can also be used for the existence of structural elements.

The completion of the situation description cannot be arbitrary. Some objects may simply be “introduced” without any further explanation, such as iron in the incoming water, whereas others are only accepted if they follow from the rest of the model, e.g. iron in the treated water. To this end, certain object types can be marked as being *introducible* to allow the addition of objects of the respective type to the system model. This provides the most important means for controlling the problem solving task, since a more restrictive set of introducibles requires to “deepen” the search for causes. Introducibles specify what cannot be expected to be explained and, hence, represent the model boundary. For example, one could accept that iron can appear in the incoming water without further justification, whereas iron in one of the treatment tanks and its concentration has to be derived from the activity of other processes (upstream), On the basis of defeasible assumptions as well as introducible elements,

one can now define the set of acceptable solutions of situation assessment as *the minimal consistent allowable structures in which a maximal set of user-defined assumptions holds*. A structure is defined to be *allowable*, if it contains at least the structural elements specified by the user as facts (without assumption) and all other structural elements are either introducible or a necessary consequence of an occurring behavior constituent (by ways of structural effects). Obviously, a structure missing necessary structural effects is not consistent with the laws of behavior constituent occurrence, therefore solutions always include all such effects.

Minimality is understood w. r. t. set inclusion. Note that we do not use the absolute cardinality of structural elements present in a solution as a criterion for preferring one solution over another one, we just exclude "unnecessary" or "superfluous" objects from being included in a solution.

Based on the four categories fact, assumption, introducible, and consequence, **Fig. 3** illustrates input and output of this step in the decision support system in an abstract way. Quantity assignments can concern the absolute values of variables, their derivatives, and deviations. The user input consists of facts and assumptions concerning structure and quantities. In particular, this includes (observed or assumed) deviations of variables from the goal behavior of the system), which in our application area, for instance, state that turbidity or iron concentration are too high in a certain stage of the treatment process. The situation assessment component starts from this and constructs a consistent model by adding introducibles and revising user assumptions respecting the minimality criterion described above.

Usually, the result is not unique, and there are several minimal solutions. In this case, more measurements may help to eliminate some of them, and techniques for measurement or test proposal as common in component-oriented diagnosis can be used. Otherwise (or additionally), the results have to be presented to the user and inspected by him or her. This raises the issue of explanation generation, especially because the result of the situation assessment contains elements that the user has not mentioned at all, but were introduced or inferred by the system.

5 Therapy Proposal Component

After having identified the current situation, the question arises whether this situation is compliant with the goal performance, or actions have to be taken to make it compliant. So, the input to therapy proposal is a result of the situation assessment and a description of the goals to be achieved by remedial actions (if any). In its most general form, a therapy would be a sequence of actions that ultimately leads to a system state that is consistent with the specified goals and requires planning. In this paper, we address only a more specific problem:

- We assume that the goals can be described by a set of quantity assignments (as opposed to complex constraints on several system variables), e.g. limiting iron concentration by a certain threshold.
- A therapy is defined as a set of feasible actions that move variables that deviate from their goal values in the right direction (without disturbing the proper ones) when applied in the current situation. For instance, if the iron

concentration is too high, then the intermediate is to reduce it and leave other goal variables unchanged.

Under these assumptions, we can be more specific about the input to therapy proposal:

- The intermediate goals: they can simply be expressed by the derivatives of goal variables taking a sign opposite to the deviation of the respective variable.

The current situation: this has to include everything as a fact that resulted from the situation assessment step (because we committed to it) with one exception: the derivatives of non-goal variables have to be turned to assumptions. The reason for this lies in the fact that we assume the actions and their effects are introduced instantaneously. Therefore, we have to allow that derivatives change their values (discontinuously) if influenced by the actions.

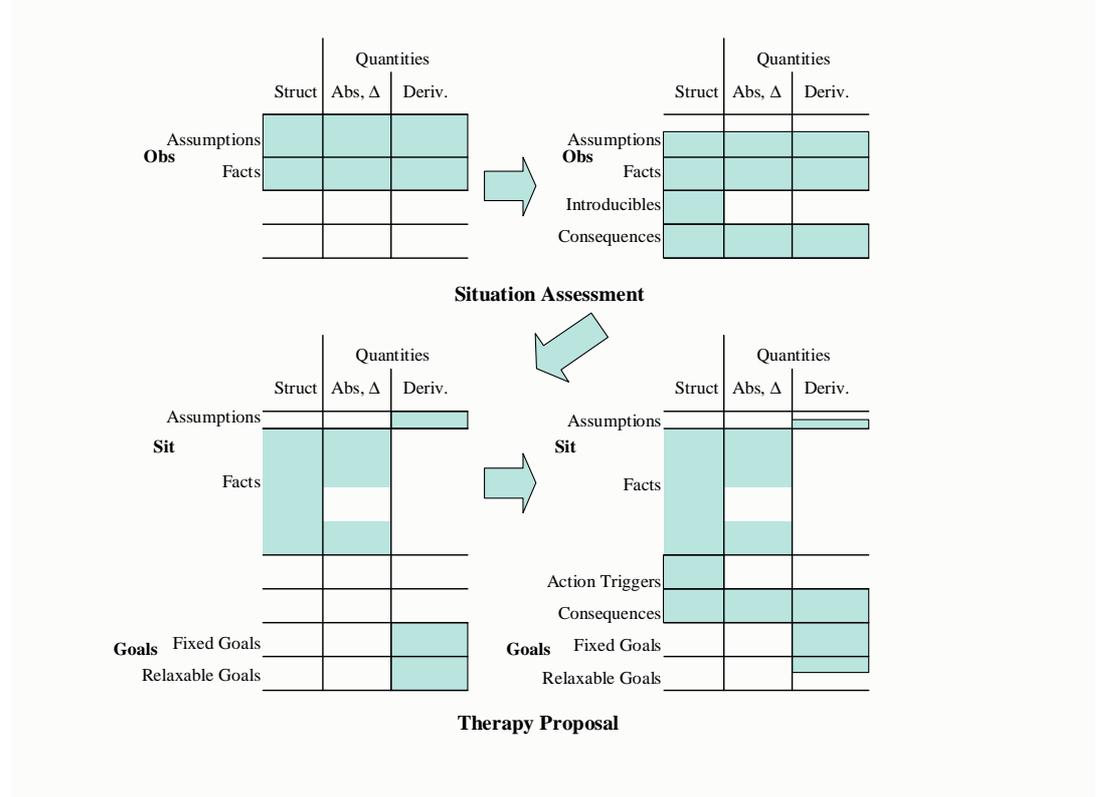


Fig. 3. Input and output of the two steps and their connection

The current situation: this has to include everything as a fact that resulted from the situation assessment step (because we committed to it) with one exception: the derivatives of non-goal variables have to be turned to assumptions. The reason for this lies in the fact that we assume the actions and their effects are introduced

instantaneously. Therefore, we have to allow that derivatives change their values (discontinuously) if influenced by the actions.

This is indicated in Fig. 3. Since it may be impossible to satisfy all intermediate goals, we also provide a means for dropping some of them. We can achieve this by stating the strict goals as facts and the defeasible ones as assumptions. Then a therapy is identified as a set of action triggers that, together with the current situation, yields a model involving *the minimal consistent allowable structures in which a maximal set of intermediate goals holds*. This is quite analogous to the situation assessment step (and uses the same diagnosis algorithms), but different in that the introducibles are the action triggers. Fig. 3 illustrates the therapy proposal step in a schematic way. Minimizing the set of actions will usually not suffice, and one might want to select the one which minimizes the cost.

6 Summary

The decision support system we outlined applies process-oriented modeling and consistency-based diagnosis to perform situation assessment and therapy proposal. It is based on a number of assumptions. The most important restrictions are due to the static perspective taken: in both steps, a solution is sought only by analyzing a snapshot of the system.

The first step determines only the current state rather than providing information about the development of a disturbance over time. While this seems appropriate as an input to therapy proposal which, indeed, has to be applied to the current state, one-step therapy may be unrealistic for many other applications which may require a sequence of interactions.

At the current state of the project, graphical editors for the domain theory and the situation description, the consistency-based problem solver, and a domain theory are available, and a first version of the decision support system is being implemented. This will be a highly interactive version which presents inconsistencies to the user and lets him guide the search for a consistent model. Based on this experiment, a version with automated generation of consistent solutions will be developed.

Acknowledgments

The AGUA project was supported by the German and the Brazilian government under the program of German-Brazilian Scientific Cooperation. Thanks to the other participants in this project, particularly the experts from DMAE, and to the members of the MQM group in Munich.

References

1. [de Kleer-Mackworth-Reiter 92] de Kleer, J.; Mackworth, A.; Reiter, R. 1992. Characterizing Diagnoses and Systems. *Artificial Intelligence*, 56.
2. [Dressler-Struss 96] Dressler, O. and Struss, P., The Consistency-based Approach to Automated Diagnosis of Technical Devices. In: Brewka, G. (ed.), *Principles of Knowledge Representation*, CSLI Publications, Stanford, pp. 267-311, 1996.
3. [Forbus 84] Forbus, K., Qualitative Process Theory. *Artificial Intelligence*, 24(1-3), 1984.
4. [Heller/Struss 02] U.Heller, P.Struss: Consistency-Based Problem Solving for Environmental Decision Support. In: *Computer-Aided Civil and Infrastructure Engineering* 17 (2002) 79-92. ISSN 1093-9687
5. [Struss-Heller 01] Struss P., Heller U., GDE⁺ - The Generalized Diagnosis Engine. In: 12th international workshop on principles of diagnosis, 2001.