Electronic News Journal on Reasoning about Actions and Change

Editor: Erik Sandewall

Volume 1, Number 5

Contributions during December, 1997

ISSN, print: 1403 - 2031
on-line: 1403 - 204X

Linköping University Electronic Press
http://www.ep.liu.se/ej/enrac/1997/05/

Publication Trail and Update Rules

The Electronic News Journals (ENJ) are a medium for exchange of scientific information and debate. In particular, they serve as the forum where articles received by the ETAI (Electronic Transactions on Artificial Intelligence) are discussed publicly for review, in separate News Journals for each ETAI area.

The Electronic News Journal on Reasoning about Actions and Change (ENRAC), in particular, uses a publication trail where contributed information is first distributed on a daily basis as *Newsletters* using electronic mail and web pages in HTML. Then, contributions during the same month are compiled into an HTML version of the monthly *News Journal*. Finally, the same information is converted via Latex to a postscript version that is suitable for printing on paper and reading off-line. It is formally published and archived by the Linköping University Electronic Press in both its electronic and its paper form, as described on page (iv).

Offprints of individual segments. Each News Journal issue consists of a sequence of *segments* relating to specific topics. "Offprints" of these can be obtained from the URL:s indicated at the head of the first page of each segment. It is intended to maintain these URL:s and their contents for the foreseeable future. Misprints of minor significance are corrected in these offprints (but not in the E-Press version since it is considered archival).

Corresponding HTML edition with dynamic links. The HTML edition contains considerable numbers of links to other pages and structures on the net: articles that are available on-line, home pages of conferences and of individual researchers, links to other part of the ETAI structure, and so on. Due to the natural limitations of the paper medium, only some of the WWW links have been retained here as footnotes. In order to retrieve articles and other information that are referenced in the present issue, it is recommended to look up the corresponding issue in HTML and to use its link. The HTML issues of the News Journals on Reasoning about Actions and Change can be found at the following URL:

`http://www.ida.liu.se/ext/etai/actions/njl/`

It is intended that the HTML issues will by updated continously to the largest extent possible, for example by replacing URL links to the home pages of authors that have changed to another site.

Table of Contents

Dates of Publication

Since the date of publication may be understood either as the date of first public appearance, or as the day of reproduction on paper in many copies, and since both of these definitions may be difficult to apply in the case of electronic publication, we make the following clarifying statement.

The contents of the present issue were put on-line in their original, HTML version during the month of December, 1997. Then the contents were edited and formatted, resulting in the present, formatted version which was published on June 10, 1998, in two concurrent editions: an on-line edition and a paper edition. The on-line edition was timestamped electronically and put on-line by Linköping University Electronic Press at the URL specified on page (i). The paper edition was obtained by printing the on-line edition on a standard computer printer. It was reproduced in 200 copies, legally archived, and made available for distribution.

Electronic News Journal on Reasoning about Actions and Change 159
Vol. 1: *159-161*, 31.12.1997
http://www.ida.liu.se/ext/etai/actions/nj/9712-2/

The Month of December

Miscellaneous Newsletter Items

Reflections at the end of the year

From the Editor

[31.12.1997] The development of the Electronic Transactions on Artificial Intelligence turns out to be an exploratory activity, even more than what we had anticipated when it started. The original idea was to create a new publication medium for scientific articles which would make the best use of the Internet availability. We foresaw a number of changes to the conventional practices that have been developed for paper-based journals, and in particular, we emphasized the importance of separating publishing (in the sense of "making public" and "making available") from the scientific quality control, that is, from the reviewing and refereering process. We also emphasized the importance of opening one part of the reviewing process, so that the traditional, confidential peer review would be replaced by a combination: an open discussion part where reviewers appear without the shroud of anonymity, and after it a rapid pass/fail decision by anonymous referees.

Finally, and maybe the most important of all, we emphasized the importance of having a publication medium where the authors retain the copyright of their articles, unlike what is the case in conventional journal publishing. The insight that this is a very important question has spread rapidly in our community during the past year.

These concepts were presented in the spring of 1997, and the ETAI was formally announced in May. In comparison with what we expected, there has been a small minus and a big plus. The minus is that the start was slower that we had thought: it took a while before papers started coming. However, I believe that the number of contributions will increase when the paper version of the ETAI begins to spread. The ETAI is "electronic" in the sense that it is stored and transmitted electronically, but noone expects you to read long, technical papers directly from the screen, and the paper printouts of the ETAI issues do look like any other journal. Professional-looking issues containing professional-quality articles will be our best advertisement.

The plus is that a whole new layer of communication concepts have evolved during the second half of 1997. Three layers can be identified here:

- The *publication layer*, consisting of the First Publication Archives whereby articles can be published while allowing the author to retain the copyright.

- The *quality control layer* involving reviewing and refereeing. The original ETAI concept foresaw the use of these two layers.
- The *communication layer*, where the E-mail-based Newsletter is the decisive component. As researchers begin to make use of this communication medium for exchange and confrontation of ideas, we see the emergence of a completely new "functionality", to use our technical jargon. The kind of information that used to be communicated very locally, in personal mail or during the coffee breaks at conferences, is now being exchanged on a global level. The concept of a "global village" is being instantiated in very concrete ways in our "electronic colloquium".

It is my firm conviction that we are just at the beginning of a very important development. Certainly we will see a lot more of it during the coming year. The future is bright!

Other Publications

Research articles

[30.12.1997] The following three articles from the Doherty group have been published recently, and describe how they deal with (in turn) concurrency, qualification, and ramification. In particular, these articles provide the detailed answers to Vladimir Lifschitz's questions to Sandewall and Doherty, in the ENRAC Newsletter of 27.11.

The first two articles have been submitted for reviewing in another journal (not ETAI), and can therefore not be included among the articles presently being reviewed by ETAI. We have had some previous cases where the present Newsletter included references and links to current articles that are in the publication channel for elsewhere, for example the papers by Judea Pearl and his group (ENRAC 97002, 22.9.1997). We would like to encourage readers to make use of this possibility of making their current work known to the community.

The third article mentioned below is a documentation of the details of the approach: precise definition of the logic, solutions for test examples, and so forth. This is the kind of material that is not traditionally published by our journals or conferences, but which is important for any detailed analysis of an approach, and for comparisons between approaches. It is therefore reference material in the strong sense of the word: appropriate to use as a reference in conventional articles, and in order to document the details of the method being proposed. Again, we welcome similar documentations from all our readers.

[f-cis.linep.se-97-014] Lars Karlsson and Joakim Gustafsson. *Reasoning about actions in a multi-agent environment.* Linköping University Electronic Press, 1997, Nr. 14.

[f-cis.linep.se-97-016] Patrick Doherty and Jonas Kvarnström. *Tackling the Qualification Problem using Fluent Dependency Constraints: Preliminary Report.* Linköping University Electronic Press, 1997, Nr. 16.

[f-cis.linep.se-97-020] Patrick Doherty. *PMON+: A Fluent Logic for Action and Change: Formal Specification, Version 1.0.* Linköping University

Electronic Press, 1997, Nr. 20.

These reports are permanently available on-line at
`http://www.ep.liu.se/ea/cis/1997/*/` where * = `014`, `016`, and `020`.

Calendar

Forthcoming conferences and workshops

LAP-98: Language Action Perspective on Communication Modelling.
Stockholm, 25.6-26.6, 1998. Papers due: 10.3 1998.
CFP: http://www.ida.liu.se/labs/vits/lap98/
MSG: http://www.ida.liu.se/ext/brs/confi/cfp/LAP-98.txt

FCS-98: Formalization of Commonsense Reasoning.
London, U.K., 7.1-9.1, 1998. Papers due: 10.10 1997.
INFO: http://www.dcs.qmw.ac.uk/conferences/CS98/index.html

Workshop on Action and Causality at KR-98.
Trento, Italy, 30.5 - 1.6, 1998. Papers due: 23.2 1998.
CFP: http://www.cs.utexas.edu/users/vl/nmr98.html
MSG: http://www.ida.liu.se/ext/brs/confi/cfp/002-98.txt

Belief Revision Workshop at KR-98.
Trento, Italy, 30.5 - 1.6, 1998. Papers due: 23.2 1998.
CFP: http://infosystems.newcastle.edu.au/belief.revision/br@kr.html
MSG: http://www.ida.liu.se/ext/brs/confi/cfp/003-98.txt

ESSLI-98 Workshop on Reasoning about Actions: Foundations and Applications.
Saarbrücken, Germany, 17-21.8, 1998. Papers due: 15.2 1998.
CFP: http://www.dis.uniroma1.it/esslli98-actions
MSG: http://www.ida.liu.se/ext/brs/confi/cfp/004-98.txt

A Note on Refereeing Procedure

The following note was made because the area editor submitted an article himself.

[19.12.1997] In any journal, a special procedure is needed when the editor himself or herself submits an article. In the case of ETAI's procedure for discussion and refereeing, no special procedure seems to be required for the discussion phase, since the entire discussion is done in public anyway. When we get to the refereeing phase, I will ask the area editor for one of the adjacent ETAI areas to be in charge of the refereeing.

If some reader should feel a need to communicate a message to the present editor without revealing his identity, then please relay through the area editor of one of the other ETAI areas, or through the ETAI policy committee.

Electronic News Journal on Reasoning about Actions and Change 162
Vol. 1: *162-162*, 31.12.1997
http://www.ida.liu.se/ext/etai/received/actions/005/

José Júlio Alferes, João Alexandre Leite, Luís Moniz Pereira, Halina Przymusinska, and Teodor Przymusinski:

Dynamic Logic Programming

Abstract of the article

The original version of the full article has been published by Linköping University Electronic Press, and is permanently available at
`http://www.ep.liu.se/ea/cis/1997/018/`

In this paper we investigate updates of knowledge bases represented by logic programs. In order to represent negative information, we use generalized logic programs which allow default negation not only in their bodies but also in their heads.

We start by introducing the notion of an update $P \oplus U$ of a logic program P by another logic program U. Subsequently, we provide a precise semantic characterization of $P \oplus U$, and study some basic properties of program updates. In particular, we show that our update programs generalize the notion of interpretation update.

We then extend this notion to sequences of logic programs updates $P_1 \oplus P_2 \oplus \ldots$, defining dynamic program updates, thereby introducing the paradigm of dynamic logic programming. This paradigm significantly facilitates modularization of logic programming, and thus modularization of non-monotonic reasoning as a whole.

Suppose that we are given a set of logic program modules, each describing a different state of our knowledge of the world. Different states may represent different time points or different sets of priorities or perhaps even different viewpoints. Consequently, program modules may contain mutually contradictory as well as overlapping information. The role of the dynamic program update is to use the mutual relationships existing between different states to precisely determine, at any given state, the declarative and the procedural semantics of the combined program, resulting from all these modules.

Electronic News Journal on Reasoning about Actions and Change 163
Vol. 1: *163-164*, 31.12.1997
http://www.ida.liu.se/ext/etai/received/actions/005/

Protocol of on-line discussion during December, 1997
about the following research article:

José Júlio Alferes, João Alexandre Leite, Luís Moniz Pereira, Halina Przymusinska, and Teodor Przymusinski

Dynamic Logic Programming

Q1. **Erik Sandewall** (12.12)

Your paper addresses update of knowledge bases represented by logic programs - a topic which is known and understood by only a limited part of the reasoning about actions community. Because of the similarity of research goals and the difference of background, I think that a discussion between you as authors and the readers in our area is particularly important and valuable; it can hopefully facilitate very much the understanding of this specific work as well as the "school" that you represent. Towards the end of the paper, you explicitly mention "reasoning about actions" as one of the intended applications, which of course is of particular interest to us.

Let me start out with a question about the prehistory of the approach you have chosen. You refer to Marianne Winslett's 1988 article [c-aaai-88-89] as an early reference for an "interpretation update" approach. Based on an example where it does not seem to give the intended result, you propose that the principle of inertia should be applied to the rules of the initial program rather than to the individual literals in a model. However, Winslett's article was written in response to an earlier article by Ginsberg and Smith [s-Brown-87-233] where they proposed exactly this: to define update on a set of logic formulae. Winslett pointed out some examples where the approach of Ginsberg and Smith did not work as intended, which is what led her to propose interpretation update. (An even earlier reference would of course be to the work by Lewis on counterfactuals [mb-Lewis-73].

My first question, therefore, is to what extent is there a difference: does your approach avoid the problems observed by Winslett, and if it does, what is the key to this improvement?

My second question is with respect to updates in the presence of observations and action laws. One of your results is that if the initial program is just a set of facts, then program updates and model updates coincide. However, in the case of reasoning about actions, one typically deals both with facts about the world at various points in time ("observations") and with rules characterizing some of the effects of actions ("action laws", "effect laws"). If update methods are used for characterizing ramification, which is what Winslett's article was all about, then presumably one wishes to prefer changes of "facts" (that is, sign reversal of literals) over changes of the action laws, at least as a first approximation. Only in the presence of accumulated evidence is it reasonable to revise a well established action law. How would you foresee representing such cases: will action laws be written out explicitly as logic-program rules, and what updates will then be obtained on the current state? On the other hand, if action laws are not represented as rules, how are they represented and how are the results in your article to be used?

References:

[c-aaai-88-89] Marianne Winslett. *Reasoning about actions using a possible models approach.* Proc. AAAI National Conference on Artificial Intelligence, 1988, pp. 89-93.

[mb-Lewis-73] D. Lewis. *Counterfactuals.* Harvard University Press, 1973.

[s-Brown-87-233] Matthew L. Ginsberg and David E. Smith. *Reasoning About Action I: A Possible Worlds Approach.* In: Brown (ed): The Frame Problem in Artificial Intelligence, pages 233-258. Morgan Kaufmann Publishers, Inc., 1987.

Electronic News Journal on Reasoning about Actions and Change 165
Vol. 1: *165-167*, 31.12.1997
http://www.ida.liu.se/ext/etai/received/actions/006/

Erik Sandewall:

Logic-Based Modelling of Goal-Directed Behavior

Summary of the article

The original version of the full article has been published by Linköping University Electronic Press, and is permanently available at
`http://www.ep.liu.se/ea/cis/1997/019/`

We address the problem of characterizing goal-directed robotic behavior using a logic of actions and change. The logicist formalization is required to be such that the set of models is exactly the set of acceptable goal-directed behaviors in the application at hand, each model being a representation of a possible history of the world. It is not required from the logic that it should *represent* psychologically related concepts, such as goals or intentions, but merely that the resulting behavior obtained from the formalization should have the characteristics of being goal-directed.

Frame of reference and motivation

We view intelligent robotic behavior as something which has both a reactive level and a deliberative level. The deliberative level is in charge of such things as receiving instructions from the commander, pursuing goals, predicting the possible outcomes of intended actions, avoiding dangers, diagnosing faults, planning means of acquiring knowledge, and so on. The underlying, reactive level may be characterized by stimulus-response behaviors and by program-like composition of simpler behaviors into more complex ones.

Logic is clearly relevant for the design and analysis of the deliberative level of such a design. However, the logic being used then must be able to characterize its different functionalities, including its goal-directedness. It must also be able to characterize the underlying, reactive level, not in every detail, but with sufficient precision for the needs of the deliberative layer.

Procedural vs goal-directed actions

The present article focuses on the characterization of goal-directed behavior as one aspect of an intelligent robot. However, we do not introduce "goals" as a separate logical construct. Instead, our approach is to distinguish two kinds of actions: *procedural* actions which are defined using a subroutine, and *goal-directed* actions which are performed through a process involving tries, possibly failures, and corrective action and new tries until the goal has been reached. Procedural actions may therefore be used for representing reactive-level behaviors to the deliberative layer; goal-directed actions characterize processes within the deliberative layer itself.

Success vs failure of actions

For both kinds of actions (goal-directed and procedural ones), the logic expresses explicitly whether the action *succeeds* or *fails*. Each execution of a goal-directed action is also characterized by a number of *breakpoints* where some sub-action has been failed, and another sub-action is selected for one more attempt to arrive at the desired goal. The logic is used for characterizing the selection of sub-actions at breakpoints, and the success or failure of the goal-directed action in terms of the success or failure of the sub-actions.

The term "goal" is used in the sense used in AI planning, that is, as the concrete goal (goal state) that is to be achieved by a plan = a sequence of actions. We do not use the word in the sense of "general goal" (goal in life).

Formalism for actions and change

We use a logic of time and action, that is, a multi-sorted first-order logic where points in time is one of the sorts. The following predicates are used:

$\mathsf{H}(t, p)$: the "propositional fluent" p holds at time t.

$\mathsf{X}(t, f)$: the fluent f is occluded at time t

$\mathsf{G}(s, a)$: the action a is invoked ("go") at time s

$\mathsf{A}(s, a)$: the action a is applicable at time s

$\mathsf{D_s}([s, t], a)$: the action a is executed successfully over the time interval $[s, t]$: it starts at time s and terminates successfully at time t.

$\mathsf{D_f}([s, t], a)$: the action a is executed but fails over the time interval $[s, t]$: it starts at time s and terminates with failure at time t.

In fact, these predicates are viewed as abbreviations for underlying, more elementary constructs indicating the applicability and the failure of actions.

Actions are composed using the conventional operators such as sequential composition (;), conditionals, etc.

The article proposes a set of axioms that characterize possible structures using these and related predicates.

Formalism for goal-directed behavior

A few specialized relations are introduced for characterizing the goal-directed behavior. The key relation is:

$Option(s', g, s, a)$: while pursuing the goal set by invoking the goal-directed action g at time s', one is now at a timepoint s (a breakpoint), and one considers performing the action a. This action may be a composite one, which means that it has the character of a plan. Several follow-up actions may be considered at the same point in time.

The article proposes a set of axioms that characterize goal-directed behavior in terms of this predicate and the ones mentioned under the previous heading.

Entailment methods

A logic characterizing goal-directed use of actions must deal with multiple knowledge sources, including action laws characterizing the effects of actions, observations of the state of the world, plans for achieving specific goals, and so forth. Each of these aspects of the system may call for some kind of nonmonotonicity in order to be sufficiently selective. It is not trivial how to combine the different types of nonmonotonicity that are required. The present article proposes and uses a general approach for dealing with this problem.

Electronic News Journal on Reasoning about Actions and Change 168
Vol. 1: *168-171*, 31.12.1997
http://www.ida.liu.se/ext/etai/received/actions/002/

Protocol of on-line discussion during December, 1997
about the following research article:

Michael Thielscher

A Theory of Dynamic Diagnosis

Q2. Marie-Odile Cordier (23.12)

I have three main comments on this paper :

1) The example on page 4-6 is quite interesting and highlights very clearly what happens when dealing with interactive faults ($ab(r1)$ causes $ab(re1)$ when $s1$ is closed). But, I did not really agree with the conclusion the author is drawing from it.

What it clearly highlighted, in my opinion, is that "minimizing abnormality" cannot be used when dealing with interactive faults. Most research papers on diagnostics suppose implicitly that faults are independent and equiprobable, and in these cases, "minimizing abnormalities" is a good way of selecting the most probable diagnoses. However, as soon as you are dealing with interactive faults the preferred diagnoses have no good reason (no probabilistic foundation) to be the minimal ones. In the example, the probability of $ab(re2)$, knowing $ab(r2)$ and $closed(s2)$ is equal to 1 whatever the probability of $re2$ is of being faulty from its own. Then, $d1$ has to preferred rather than $d2$ and $d3$.

The key point is that "minimization" (minimizing abnormalities) is not a good preference strategy in case of interactive faults. This seems to me to be the very reason why one doesn't get the expected results in this example.

Another point concerns when this "preferring" step has to be done. The author argues that it has to be done at the starting point and uses the example as a justification. I don't contest this fact (see below), but I contest that it follows from the example.

From the following example, it can be seen that "minimizing abnormalities" is not a correct solution even if it is done at the starting state. Let us suppose that it is known that $closed(s1)$, $closed(s2)$, $closed(s3)$ and $off(light)$ are true in the initial state. Whatever the action might be, for example "open s1" or the empty action, you are going to prefer a state where $ab(re1)$ is true rather than the one where $ab(r1)$ and $ab(re1)$ are true, which is not at all justified from a probabilistic point of view. It is even problematic from a diagnostic point of view, for replacing $re1$ by an unfaulty relay (instead of replacing $r1$) will lead to breaking $re1$ again as soon as you will close $s1$. In fact, example 1 (page 20) exhibits similar results which are not satisfying by forgetting $ab(r1)$ and $ab(r2)$ as possible faults.

2) The term "dynamic diagnosis" is used throughout the paper to denote diagnosis on systems on which you are performing actions (tests). As far as I understand, the systems are supposed to be static ones; they are not supposed to evolve by their own; they don't have any proper dynamic behaviour. The only way to make them change is to perform actions. This is the reason why you can predict the resulting state by looking only to the effects of the action. Unpredictable events are not taken into account, for exemple faults (or more simply evolutions of the system) that occur during the sequence of actions.

The term "dynamic diagnosis" is then misleading, at least for the diagnostic community for which dynamic diagnosis usually means diagnosing systems evolving in time by themselves, without explicit exogeneous events making them change.

In this context, the problem which is proposed is very similar to that of postdiction : knowing some observed facts resulting from an action (or a sequence of actions), you want to infer the actual state of the system. Faults cannot happen during the sequence of actions, and then dealing with an action or a sequence of actions makes no difference.

Consequently, it is quite justified to apply the "preference step" on the initial point. You have to determine the most probable sequences of steps (or histories, scenarii, trajectories?) starting from initial states, leading to some final states in which observations are true, and corresponding to a given sequence of actions. There are no unknown events; there is no uncertainty in the actions; no uncertainty wrt their effects; the only uncertainties concern the initial states. The preference between sequences depends directly on preferences on initial states, which explains why the "preference step" concerns the initial states.

This scheme is a restricted case of a most general scheme in which you take into account the possible occurences of events (as faults) interleaved with the actions, the probability of such events, and the probability of an action to produce some effects. Selecting some of these scenarii according to preference criteria corresponds to what we called "event-based diagnosis" in [Cordier 94]. It is also close to McIlraith's approach; see [McIlraith 94], [McIlraith 97]. The main difference between these approaches is that we tried to define diagnostics independently of the mechanism used for modeling actions and changes, whereas Sheila's proposal is clearly dependent on the formalism used to model actions (situation calculus).

3) The last point concerns the ramification problem and the use of a causal model to predict the effects of an action. I realize that this point is not the main subject of this paper, since it is devoted to diagnostics. However, an important point related to this paper is to examine whether it can be applied when dealing with dynamic systems.

This theory of action based on causal relationships is very attractive as long as you are looking for the effects of an action or a sequence of actions, and as long as the concerned system has no proper evolution. Fluents which are not affected by an action are then supposed not to change, by virtue of a minimal change principle. But as soon as you are concerned with dynamical systems, (which is not really the case in the paper), such a causal model would probably not be sufficient and you will need a "transition model", describing the way things evolve along time. This happens for example if you want to model the dynamics of a system, or the possible events as faults that occur as you are monitoring a system, or the natural ageing of components.

An answer could be that there cannot be changes without causes, but most of the time you don't want to model these causal chains or you are not even able to model the primary causes of such evolutions (for example, the ageing of components or the sudden occurence of a fault), but you want nevertheless to take them into account as far as possible.

The basic idea of the proposal we made in [Cordier 95] was that a transition model (that is, a set of possible (partially ordered) transitions) is needed in order to decide what is the most plausible state after an update, or equivalently, an action. A causal model is certainly quite adequate when considering static systems reacting to actions. More than that, in my opinion, a causal model is a very nice formalism allowing to acquire the partial orderings that exist between transitions, in a natural way. There is probably a strong correspondence between your "influences" and our "partially ordered transitions" which would be worth studying more deeply. However, transitions seem to have a broader scope in that they allow to represent any changes from one world to a next one, whereas causal relations or influences are restricted to represent "causative changes" (changes for which one can exhibit the causes). This is a problem when dealing with dynamical systems.

References:

[Cordier 94] Marie-Odile Cordier and Sylvie Thiébaux. *Event-based diagnosis of evolutive systems.* Proceedings of the 1994 DX Conference.

[Cordier 95] Marie-Odile Cordier and Pierre Siégel. *Prioritized transitions for updates.* Proceedings of the 1995 ECSQARU Conference, pp. 142-151.

[McIlraith 94] Sheila A. McIlraith. *Towards a theory of diagnosis, testing and repair.* Proceedings of the 1994 DX Conference, pp. 185-192.

[McIlraith 97] Sheila A. McIlraith. *Explanatory diagnosis: conjecturing actions to explain observations.* Proceedings of the 1997 DX Conference, pp. 69-77.

Q3. **Wolfgang Nejdl** (27.12)

One basic remark about the title and intention of your paper:

When I first read your paper, I had some difficulties in connecting your work to the usual diagnosis literature, as you basically refer only to papers about reasoning about action and change (which is ok, considering the content, but should be changed, considering the title). Also, the current title is somewhat misleading, as the term "dynamic diagnosis" in the diagnosis community is usually reserved for systems which monitor and diagnose continuous and/or time-varying systems. I would have suggested something like "diagnosis and actions" or similar within the title.

Anyway, here are some more specific questions, which came into my mind, while I was trying to comprehend your approach.

1. It seems to me, that one main problem (chapter 2) you are considering are dependent failures like "$ab(c1)$ implies $ab(c2)$", which are usually neglected in many papers. Could you elaborate more on the advantages of your formalism when these dependent failures are not present? In such a case,

what exactly do you gain by including explicit causal relationships (considering that most diagnosis systems use just ordinary state constraints)?

2. A second thread which seems to emerge in chapter 4 is the integration of test actions. Have you thought about which test actions one should take, or is this only a side issue in this chapter?

3. Also, could you comment some more about the relationship of your approach to the one of Sheila McIlraith?

Best regards,

Wolfgang Nejdl