ELECTRONIC NEWSLETTER ON REASONING ABOUT ACTIONS AND CHANGE

Issue 98047 Editor: Erik Sandewall 24.5.1998

Today

Debates

In the review debate about the recent submitted article by Cervesato, Franceschet, and Montanari, we have today the answer of Montanari to the questions asked by Liberatore (invited discussant). A number of fine points in the article are covered.

We also have contributions by Brandano and by myself in the discussion about ontologies of time. Recently, one part of that discussion has concerned some examples of uses of a less-than relation on intervals. It seems to me (in a moderator's role) that this problem does not require the attention of the full ENRAC readership. I have therefore asked those immediately concerned to resolve the disagreement in the way proposed by Leibniz (that is, by checking out carefully what are the consequences of the chosen axioms), and to report jointly when the question has been resolved.

That issue aside, the debate still seems to address several questions in parallel, including:

- What are examples of concrete problems that cannot (conveniently) be represented using a "standard" view of the time axis as e.g. timepoints = the real numbers; time-intervals = sets of points each defined by a lower bound a and an upper bound b.
- Given that both timepoints and time-intervals are needed to have sufficient expressivity, what is the best way of defining those concepts and their relation?

ENRAC Facilities

Postscript versions of *News Journals* are now available from the ETAI/ ENRAC web site for all months up to and including March. The November, December, and March issues will be printed two weeks from now - please check for correctness if you have contributed.

These editions of the News Journals are the result of fairly extensive postediting of the original contributions, including a brushup of formulas (to make use of the expressiveness offered by latex compared to plain text) and of references. Incomplete references to published articles have been replaced by correct references drawing on the ACRES database. This has facilitated uniformly correct and complete citations. For the future it also paves the way for making direct electronic links available in an incremental fashion in the HTML version of the News Journals, within the restrictions, of course, that may be imposed by the respective copyright holders.

In your future contributions to our debates, please note that references may be cited simply as e.g. `c-ecai-94-266`, meaning conference - ECAI - 1994 - article starting on page 266. The full corresponding bibliographic information will be fetched automatically from the database and added to your contribution, at least for the main conferences and journals in our area. Journals are coded like e.g. `j-jlc-4-26`, meaning Journal of Logic and Computation, volume 4, page 26 ff. Conventions for additional cases (tech reports, etc) are available on request.

Don't forget that postscript versions of *Newsletters* (the daily issues) continue to be produced at the same time as the plaintext and html versions. If you prefer to read the present Newsletter off-line, you may wish to try the postscript variant. It's available from the web site like everything else.

ETAI Publications

Discussion about received articles

The Complexity of Model Checking in Modal Event Calculi with Quantifiers

by Iliano Cervesato, Massimo Franceschet, and Angelo Montanari

Comment from: **Angelo Montanari**

To:**Paolo Liberatore**

Dear Paolo,

thank you very much for your comments and questions. In the following, we (try to) answer them.

> *Page 2, col 2, second-last para: "We represent an MVI for p as $p(e_i, e_j)$ ...". I found this sentence a little misleading. Indeed, the notation $p(e_i, e_j)$ refers to* any *3-tuple ⟨property, event, event⟩. The aim of queries is to decide whether such a 3-tuple is a MVI. From the sentence written in the paper, it seems that $p(e_i, e_j)$ is used to* state *that e_i, e_j is an MVI for p. Perhaps this sentence should be rewritten in some way.*

In the following, we quote a bigger paper excerpt, including the sentence: "We represent an MVI for p as $p(e_i, e_j)$ ...". We slightly changed the notation to make it easier to read.

> Given an EC-structure H = (`E, P, [.>, <.], ].,.[`) and a knowledge state w, the Event Calculus (EC for short) permits inferring the maximal validity intervals (MVIs) over which a property p holds uninterruptedly. We represent an MVI for p as $p(e_i, e_t)$, where e_i and e_t are the events that respectively initiate and terminate the interval over which p holds maximally. Consequently, we adopt as the query language of EC the set
>
> $$\{p(e_1, e_2) \mid p \text{ belongs to } P \wedge e_1, e_2 \text{ belong to } E\}$$
>
> of all such property-labeled intervals over H. We interpret the elements of the language as propositional letters and the task performed by EC reduces to deciding which of these formulas are MVIs in the current knowledge state w and which are not.

We believe that reading the sentence in its context should allow one to conclude that every MVI is 3-tuple $\langle property, event, event\rangle$, but not every 3-tuple $\langle property, event, event\rangle$ is an MVI.

The definition of MVI states that between e_1 and e_2 there is no event that initiates p. I would like you to spend a few words for explaining why.

We discussed this point in some detail in previous publications (cf. reference [4] as well as Chittaro L., Montanari A., Efficient temporal reasoning in the Cached Event Calculus, Computational Intelligence, 12(3):359-382, 1996). Your comment indicates that it is probably useful to add a brief explanation immediately after Definition 2.2.

In the following, we report part of the discussion provided in [4].

Definition 2.2 states that an event e interrupts the validity of a property if it initiates or terminates p itself or a property q which is incompatible with p. This rule adopts the so-called strong interpretation of the initiate and terminate relations: a given property p ceases to hold over a time interval if an event e which initiates or terminates p, or a property incompatible with p, occurs during this interval. The strong interpretation is needed when dealing with incomplete sequences of events or, as in our case, incomplete information about their ordering. For example, consider a switch that can take two different positions: on and off. Its behavior can be described by means of two types of event: one that changes the position from off to on (turn-on), the other from on to off (turn-off). While two turn-on events cannot occur consecutively in the real world, it may happen that an incomplete sequence consisting of two consecutive turn-on events, followed by a turn-off event, is recorded in the database. The strong interpretation of the initiate relation allows EC to recognize that a missing turn-off event must have occurred between the two turn-on events. However, since it is not possible to temporally locate such an event, EC only concludes that the switch is on between the second turn-on event and the turn-off event, and it considers the first turn-on event as a pending initiating event. It is worth mentioning that an alternative interpretation of the initiate and terminate relations, called weak interpretation, is also possible. According to such an interpretation, a property p is initiated by an initiating event unless it has been already initiated and not yet terminated (and dually for terminating events). The weak interpretation is needed to aggregate homogeneous states.

Further details about the strong/weak distinction, including formalization and other examples, can be found in [4].

Page 9, col 1, para 3: the bound $\mathcal{O}(n^2)$ for reachability is due to the fact that you consider the number of events as the size of the input. Indeed, reachability in a graph is linear in the total size of the graph, that is, is $\mathcal{O}(n+m)$, where n is the number of vertices and m is number of edges. This implies for example that the cost of EC is $\mathcal{O}(s^2)$ if s is the total size of the input and not only the number of events. Is there any reason for choosing the number of events to represent the size of H?

An instance of the model checking problem we analyzed in the paper is a triple (H, w, φ), where H is an EC-structure, w is a (strict) partial order on the events in E, and φ is a well-formed formula.

Given an EC-structure H, E is the set of events (we can assume that they are instances of a finite set of domain-specific event types), which can be arbitrarily large, while P is the set of relevant properties, which are fixed once and for all and characterize the considered domain. For instance, in the light example, we may have two events e_1 and e_2, both of type *turn-on*, which initiate the property *on*. Since the cardinality of P does not change from one problem instance to another one (unless we change the application domain), while the cardinality of E may grow arbitrarily, we chose the cardinality of E, that is, the number n of events, as the size of H (it is usually called data complexity).

The cardinality of w, here denoted by m, is bounded by n (more precisely, it is less than or equal to $n(n-1)/2$). We did not choose m as a complexity parameter, because we are interested in explicitly characterizing the amount of available ordering information with respect to the given number of events. As an example, we are interested in distinguishing between situations in which $m = O(n)$ (sparse graphs) from situations in which $m = O(n^2)$ (dense graphs).

Finally, the size k of φ in not bounded by n, and it can arbitrarily grow, independently from n (clearly, not in the case of basic EC). Hence, we took k as another relevant complexity parameter (which is usually called query complexity).

> *Page 10, col 2, comment after Theorem 5.5: if k is just a bit high (for instance, 4), the cost of solving the problem becomes something in the order of $\mathcal{O}(n^7)$, which is usually considered not very useful. In the conclusions, you mention that the cost of the implemented version of QCMEC is very close to the theoretical value. Do you have empirical evidence that the efficiency of EQCMEC is also similar to the upper bound you proved? If so, the problem should be considered not actually feasible.*

This last question is concerned with the following sentence, at the end of the complexity section:

> The directness of the implementation allows checking easily that the complexity of the implemented algorithms coincide with the bounds proved in this section for the problems they implement."

Well, you are right, this sentence is probably a little misleading. Its intended meaning is that, given the directness of the proposed implementation, we expect that its worst-case complexity coincides with the proved bounds. We can possibly reformulate the above sentence as follows:

"The directness of the implementation allows checking easily that the complexity of the implemented algorithms coincide, in the worst case, with the bounds proved in this section for the problems they implement."

However, we are convinced that the actual costs of the implemented versions of EC, ECMEC and EQCMEC are NOT close to the theoretical upper bounds, even though we have not empirical evidence of this fact.

Take, for example, the case of EQCMEC. Here, the critical factor is the level of nesting of quantifiers. The worst case holds when we have an EQCMEC-formula involving q nested quantifiers at its top-level. Let $Q.F$ be such a formula, where Q is the quantification part, and let d be the dimension of F, q the number of quantifiers in Q, and n the number of events. Exploiting the unfolding strategy, we are led to check a formula of size $d \cdot n^{\{q\}}$. Since both d and q are less than or equal to the size k of $Q.F$, we get the upper bound $k \cdot n^{\{k\}}$ for the size of the unfolded formula (without loss of generality, we can interpret the size k of $Q.F$ as the number of propositional connectives, quantifiers and modal operators occurring in $Q.F$ plus one).

Furthermore, as mentioned in the paper, practical applications using modal event calculi with quantifiers are expected to model situations involving a large number of events, while the size of the queries, and in particular the nesting of quantifiers, will in general be limited. The hardware fault diagnosis example falls into this category (the nesting level of quantifiers is 2).

Finally, it is possible to show that the upper bound of $O(n^3)$ for EC can actually be lowered of one degree by exploiting the notions of transitive reduction and closure of an ordering relations. Preliminary results in this direction can be found in [2].

From the above elements, it follows that the hardware fault diagnosis problem can actually be solved in $O(n^4)$.

Debates

Ontologies for time

From: Sergio Brandano

In reply to Pat Hayes [ENRAC 17.5]

> *Sergio's point seems to be that one can describe the line in terms of conventional open and closed intervals in such a way that no 'gaps' appear, so that the 'interval' between the open intervals (a, b) and (b, c) is the closed interval $[b]$ containing a single point.*

In my example, the temporal domain S consists *exclusively* of intervals from the real-line and no total order relation was imposed, so that the interpretation for which one can describe the line with that domain is surely not correct. If one restricts S to the case of intervals from the positive real-line, for example, then S is a tree with a continuum of branches.

> *The question is whether this standard mathematical view of the line is the most suitable for capturing linguistic intuition or for action reasoning. For example, we want to be able to assert that a light is off before time t and on after time t*

In your approach the switching action takes place at time t, while the process of turning on the light does not necessarily require a single time-point. Anyway, according to the classical approach (linear point-based temporal-domain) one may easily capture the proper intuition with a simple function of time:

```
 \[ to\_switch(s,\tau,t) =
\cases{ {\rm off}              & if $\tau = s$       \cr
```

```
                {\rm on} \vee {\rm off} & if $\tau \in (s,t)$ \cr
                {\rm on}                & if $\tau = t$}
\]
```

where the interval $[s, t]$ is the switching interval, which may also be point-like, of course. The switching interval is properly that time-interval where the action of switching on the light is performed. Please note that s and t may also be variables.

Now, assume the light is initially "off". Then it will remain "off" until s, because of the assumption of inertia. From s to t the value will change according to the definition of the switching action, while at t the light will be surely "on". The light will be "on" at t and at every timepoint after t until another action will occur to change state. As simple as this...

In any case, many users of temporal ontologies do not wish to assume continuity or even density, for reasons of their own, and so a general-purpose temporal ontology should therefore not make such unnecessarily strong assumptions as Sergio's 'completeness' axiom. Temporal database technology usually assumes times are discrete, for example.

(a) When implementing whatever formalism on a resource-bounded machine, compromises are never enough, otherwise we shall bound ourselves when designing them. On the other hand, I like to observe that humans formulated the notion of continuity in their mind, although humans too are resource-bounded. It is also true that not every human knows about continuity, but some do. In any case the notion of continuity does exist, at least on paper.

(b) The problem that comes when aiming at generality (as in this debate?), is one shall include all possible cases, instead of excluding those which are of no interest for someone. Personally, I like to consider problems within Newtonian physics as being an important part of Reasoning about Actions and Change, where continuity plays an fundamental role when idealizing the physical world. Furthermore a continuous temporal domain subsumes the case of integer time, as well as the case of rational time, so that it is general enough to embrace all possible cases.

Therefore, in my opinion, the axiom of completeness is a good assumption, at least within the classical point-based approach. The purpose of my example with S as an ad-hoc example to show the axiom of completeness adequate as well for an interval-based temporal domain. As an immediate consequence of that result, the dividing instant problem is then solved for those interval-based formalisms where this axiom will be properly adopted. Still remains the question whether at least one problem exists that needs to introduce intervals into the temporal domain. Personally, I think this problem does not really exist.

Finally, I would like to resume Jixin's original statements in ENRAC 13.3:

(1) For general treatment, both intervals and points are needed.

(2) To overcome the so-called Dividing Instant Problem, that is the problem in specifying whether intervals are "open" or "closed" at their ending-points, both intervals and points should be treated as primitive on the same footing.

My reply to (1) was: [ENRAC 8.5]

> Does there exist at least one problem (within R.A.C) that needs to introduce intervals into the temporal domain?

My reply to (2) was the example with S, where just intervals are needed.
Best Regards Sergio

From: Erik Sandewall

Pat,
You wrote (ENRAC 17.5]:

> *The question is whether this standard mathematical view of the line is the most suitable for capturing linguistic intuition or for action reasoning. For example, we want to be able to assert that a light is off before time t and on after time t without having to commit ourselves to its being either on or off* **at** *that time, but also without sacrificing the assumption that it is always either on or off.*

The most natural way of dealing with such a situation, it seems to me, is to admit that one's axioms allow two kinds of models: those where the light is on at time t, and those where it is off. The assumption that light is either on or off holds in each of the models, but the axioms don't imply one or the other.

The alternative that has occurred in the present discussion, and proposed as a way of dealing with the "dividing instant problem" is the use of a *punctuated timeline* where the domain of timepoints is chosen e.g. as Re - B where Re is the real numbers and B is a finite set of "breakpoints", typically chosen as the times where actions start or end. If we ignore questions of how this domain is specified in terms of axioms or how intervals are formally defined, this seems to be the essence of how one would deal with the DIP without going to *very* unconventional concepts of time.

Given that the standard view already exists, it seems worthwhile to understand what the concrete reasons would be for replacing it with the punctuated timeline ontology. Apart from the purely subjective reasons (that is, some people *prefer* to do it that way) I wonder what *results* have been obtained using punctuated or other nonstandard ontologies, and which are not trivially translatable back to the standard view. The following would seem to be interesting *results* for this purpose:

- Expressiveness in a concrete scenario. This ought to be a scenario that can't be properly expressed in the standard ontology, but which can be expressed in the nonstandard one.
- Algorithms or other implementation techniques. Is there some algorithm that works with a punctuated timeline and that outperforms those that don't?
- Validation and range of applicability results. These are results stating that or when a certain entailment method is correct, that is, it obtains exactly the intended models viz conclusions. Have some such results been obtained using nonstandard ontologies?

Please feel free to add more categories to the list. (In the methodology discussion at the NRAC workshop at IJCAI last year we got to more than ten such categories).

I realize of course the purely philosophical and/or logical interest in analyzing different possible concepts of time, but from an AI point of view there is no point in pursuing an approach if it doesn't deliver any results. I haven't been able to see any indications of such concrete results from the present discussion or from the articles that have been referenced in it. (That most work in temporal databases uses discrete time doesn't say anything about the choice between a standard or a punctuated real timeline, does it?)

Although I'd prefer to see algorithmic results or range of applicability results, a few words about expressiveness since after all that's also a relevant type of achievement. One nice scenario for hybrid change (that is, continuous and discrete) is the *impact problem* that Persson and Staflin used in their ECAI 1990 paper [c-ecai-90-497]. It goes like this: two solid spheres B and C are lying side by side on a horizontal surface, B to the left of C. A third sphere, A, comes rolling from the left and hits B. As we know from physics, the result of the impact is that A stops, B stays, and C starts moving towards the right. The representation of this particular exercise has been solved in 1990, but is there now some harder one that makes essential use of punctuated time ontology?

Erik

References:

[c-ecai-90-497] Tommy Persson and Lennart Staflin. *A Causation Theory for a Logic of Continuous Change.* Proc. European Conference on Artificial Intelligence, 1990, pp. 497-502.

This newsletter is issued whenever there is new news, and is sent by automatic E-mail and without charge to a list of subscribers, in plaintext format. The present postscript issue is a bona fide rendering of the same information in alternative typography. Instruction for contributors and other additional information is found at the following URL:

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