

Indexing the event calculus with Kd-trees to monitor diabetes

Citation for published version (APA):

Bromuri, S., Torre, A. B. D. L., Duboisson, F., & Schumacher, M. (2017). *Indexing the event calculus with Kd-trees to monitor diabetes*. arXiv <https://arxiv.org/abs/1710.01275>

Document status and date:

Published: 04/10/2017

Document Version:

Early version, also known as pre-print

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

<https://www.ou.nl/taverne-agreement>

Take down policy

If you believe that this document breaches copyright please contact us at:

pure-support@ou.nl

providing details and we will investigate your claim.

Downloaded from <https://research.ou.nl/> on date: 17 May. 2025

Open Universiteit
www.ou.nl



Indexing the Event Calculus with Kd-trees to Monitor Diabetes

Stefano Bromuri^{a,*}, Albert Brugues De la Torre^b, Fabien Duboisson^b,
Michael Schumacher^b

^a*Open University of the Netherlands,
Management, Science and Technology,
Valkenburgweg 177, 6419 AT, Heerlen, Netherlands*
^b*University of Applied Sciences Western Switzerland
Institute of Business Information Systems
TechnoArk 3,
CH-3960, Sierre, Switzerland*

Abstract

Personal Health Systems (PHS) are mobile solutions tailored to monitoring patients affected by chronic non communicable diseases. A patient affected by a chronic disease can generate large amounts of events. Type 1 Diabetic patients generate several glucose events per day, ranging from at least 6 events per day (under normal monitoring) to 288 per day when wearing a continuous glucose monitor (CGM) that samples the blood every 5 minutes for several days. This is a large number of events to monitor for medical doctors, in particular when considering that they may have to take decisions concerning adjusting the treatment, which may impact the life of the patients for a long time. Given the need to analyse such a large stream of data, doctors need a simple approach towards physiological time series that allows them to promptly transfer their knowledge into queries to identify interesting patterns in the data. Achieving this with current technology is not an easy task, as on one hand it cannot be expected that medical doctors have the technical knowledge to query databases and on the other hand these time series include thousands of events, which requires to re-think the way data is indexed. In order to tackle the knowledge representation and efficiency prob-

*Corresponding Author

Email address: stefano.bromuri@ou.nl (Stefano Bromuri)

lem, this contribution presents the kd-tree cached event calculus (CEC_{KD}) an event calculus extension for knowledge engineering of temporal rules capable to handle many thousands events produced by a diabetic patient. CEC_{KD} is built as a support to a graphical interface to represent monitoring rules for diabetes type 1. In addition, the paper evaluates the CEC_{KD} with respect to the cached event calculus (CEC) to show how indexing events using kd-trees improves scalability with respect to the current state of the art.

Keywords:

Diabetes type 1, Event Calculus, Kd-trees, Expert Systems, Rule management

1. Introduction

Chronic non communicable diseases are becoming a big challenge for the contemporary world, in particular because the people affected by such conditions are growing in numbers. This is a positive fact, as it implies that better cures are available to the public, but it also brings the consequence that healthcare costs rise considerably. In addition, doctors get flooded from information produced by the patients, who more and more often make use of personal health systems (PHS) [5] and wearable devices to monitor their own condition. PHS research has mostly focused on interconnection problems, namely on transferring the physiological data of the patient to the hospital infrastructure using interoperability standards. A second generation of PHS is though starting to borrow terminology from business intelligence. In particular, it is becoming of major importance to define tools that allow medical doctors to provide a description of the interesting events happening in the stream of information of the patient. Prescriptive medical reasoning, in the form of temporal rules applied to the stream of physiological values of the patient is becoming an interesting field of research. There are several issues that need to be solved to be able to define dynamic and personalizable knowledge driven approaches to monitor or query the physiological values of a patient. First of all, there is the need of a formalism, graphical or textual, that would allow medical doctors to specify patterns of interest coming from their own knowledge. Secondly, there is the problem associated with handling long streams of events happening in relatively shorts time spans, as for examples it happens with continuous glucose monitoring (CGM) that produces 288 events per day.

The Event Calculus (EC) is a formalism for reasoning about events and their effects in a computational logic framework. The original EC, pioneered by Kowalski and Sergot [17], has been extended and adapted to support multiple types of applications, ranging from AI planners [22], web-services [20], multi-agent environment platforms [6], and activity recognition systems [2].

A widely adopted version of the EC is the version presented in [21], which is often interpreted under the semantics of normal logic programs with negation as failure. What is appealing with this version is that the developer specifies which fluents are initiated and terminated by domain specific events, and it is then the domain independent axioms of the EC that cater for which fluents hold at different times. Despite the simplicity of this EC version, the domain independent axioms are computationally naive for large narratives, which can lead to impression that common EC specifications do not scale up.

In [9], Chittaro and Montanari study the computational complexity of EC within normal logic programs. Their proposal involves a Cached Event Calculus (CEC) that caches the maximum validity intervals (MVIs) for a fluent by moving some of the computational complexity from query to update time. The results achieved by [9] are important as they also clearly show the theoretical complexity of the CEC in health monitoring settings [10], showing that CEC is suitable for monitoring applications. One drawback of the current CEC, and of any EC dialect using Prolog, is that it relies on the indexing capabilities of a Prolog engine, which typically uses hash maps to index on the functor name and first argument of the facts stored in a logic based knowledge base. Such an indexing mechanism is inflexible as the application developer cannot change it, as in database systems.

The indexing problem is worsened by applications, such as patient monitoring, producing large event narratives . In this paper we propose CEC_{KD} , an integration of the CEC with an alternative indexing mechanism for events and MVIs using Kd-trees [4]. A Kd-tree is a space-partitioning data structure for organizing points in a k-dimensional space. Queries and updates in the Kd-tree operate on a hyper-plane region containing the multidimensional points and have tractable computational properties [13]. The contribution of the integration of CEC with Kd-trees in CEC_{KD} is twofold: (a) we study theoretically the computational aspects of CEC_{KD} knowledge representation, showing how this is an improvement with respect to CEC; and (b) we experimentally validate the complexity of CEC_{KD} specifications for monitoring applications that require long narratives, putting ourselves in comparison

with the state-of-the-art approach represented by CEC.

The research presented in this paper, takes place within the D1NAMO project, a Swiss national project aiming at monitoring patients affected by diabetes type 1 in order to define a non invasive PHS solution to monitor and understand hypoglycaemia events. The algorithms presented in this paper are therefore partially evaluated on D1NAMO data and partially evaluated through a simulation.

The reminder of this paper is structured as follow: Section 2 presents the graphical rule editor of D1NAMO and how this uses EC to calculate alerts concerning the health of the patient; Section 3 presents CEC_{KD} and its indexing mechanism; Section 4 evaluates CEC_{KD} performance with real events produced in the D1NAMO project; Section 5 discusses relevant related work; finally Section 6 concludes this work discussing about potential future work.

2. Rule Management in D1NAMO

Data analysis is becoming more and more an important feature for medical systems and thus this is important also from D1NAMO perspective. Most of data analysis systems include a component called datawarehouse. A datawarehouse is a component that, contrarily to a relational data base, takes the data and put it in a format that is convenient for applying data analysis algorithms. In this sense, D1NAMO datawarehouse serves as an event feeds for our rule management approach.

The Rule interface of D1NAMO is a component that interfaces with the datawarehouse of D1NAMO to collect the data of the patient. Effectively speaking we consider a Web service model in which the Rule interface is notified of changes by means of a REST interface receiving the CSV files in input. The CSV files are then saved in local by the rule interface for the data analysis. Thus, the CSV files received from D1NAMO datawarehouse are translated into events which then can be queried using D1NAMO rule interface.

D1NAMO rule interface makes use of a formalism called Event Calculus (EC) [17]. EC is based on a many-sorted first-order predicate calculus represented as logic programs that are executable in Prolog. The underlying time model is linear. The EC manipulates fluents. A fluent represents a property which can have different values over time.

Predicate	Meaning
initially(F=V)	The value of fluent F is V at time 0.
holds_at(F=V,T)	The value of fluent F is V at time T .
holds_for(F=V,[Tmin,Tmax])	The value of fluent F is V between time $Tmin$ and time $Tmax$.
initiates_at(F=V, T)	At time T the fluent F is initiated to have value V .
terminates_at(F=V, T)	At time T the fluent F is terminated from having the value V .
broken(F=V, [Tmin, Tmax])	The value of fluent F is either terminated at $Tmax$, or initiated to a different value than V between $Tmin$ and $Tmax$.
happens_at(E,T)	An event E takes place at time T updating the state of the fluents

Table 1: EC with multi-valued fluents: predicates.

The term $F = V$, denotes that fluent F has value V that has been initiated by an action at some earlier time-point and not terminated by another action in the meantime. Tab. 1 summarizes the main EC predicates we use in this contribution.

Predicates, function symbols and constants start with a lower-case letter while variables (starting with an upper-case letter) are universally quantified.

The specifications of the axioms of the EC are then represented below.

- (EC0) $\text{holds_at}(F=V, 0) \leftarrow$
 $\text{initially}(F=V).$
- (EC1) $\text{holds_at}(F=V, T) \leftarrow$
 $\text{initiates_at}(F=V, T_s),$
 $Tmin < T,$
 $\text{not broken}(F=V, [Tmin, T]).$
- (EC2) $\text{broken}(F=V_1, [Tmin,Tmax]) \leftarrow$
 $(\text{terminates_at}(F=V_1,T)$
 $Tmin < T, Tmax > T);$
 $(\text{initiates_at}(F=V_2,T),$
 $V_1 \neq V_2,$
 $Tmin < T, Tmax > T).$
- (EC3) $\text{initiates_at}(F=V, T) \leftarrow$
 $\text{happens_at}(Ev,T),$
 $\text{Conditions}[T].$
- (EC4) $\text{terminates_at}(F=V, T) \leftarrow$
 $\text{happens_at}(Ev,T),$
 $\text{Conditions}[T].$

(EC5) holds_for(F=V, [Tmin,Tmax])←
 initiates_at(F=V,Tmin)
 terminates_at(F=V,Tmax)
 End > Start, not broken(F=V, [Tmin,Tmax]).

(EC6) holds_for(F=V, [Tmin,infPlus])←
 initiates_at(F=V,Tmin),
 not broken(F=V,[Tmin,infPlus]).

(EC7) holds_for(F=V, [infMin,Tmax])←
 terminates_at (F=V,Tmax),
 not broken(F=V,[infMin,Tmax]).

Clause EC0 states that a property F holds at time 0 if an initially/1 predicate is true at time 0. Clause EC1 states that a property holds at a time T if it has been initiated at time Tmin and the holding of that property has not been broken between the starting time Tmin and the time of interest T. To decide when a property is broken, we use the clause EC2. This states that a property P is broken between time Tmin and Tmax, if it is terminated at a time T between Tmin and Tmax or initiated to a different value between Tmin and Tmax. The other clauses specify when a property is initiated

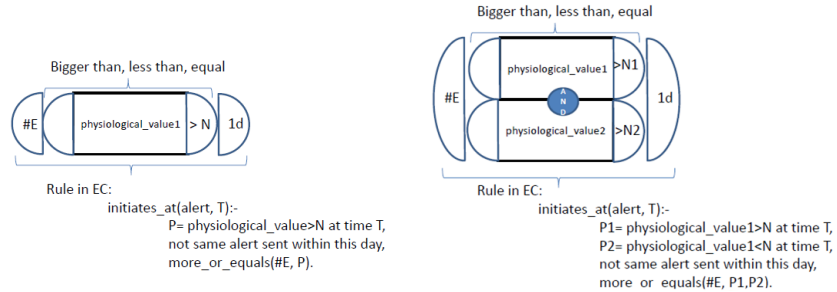



Figure 1: Simple Pattern Specification and Complex Pattern Specification.

(EC3) or terminated (EC4), in terms of the conditions holding in the current context, typically expressed in terms of the holds_at/2, holds_for/2 predicates, meaning that such clauses will change according to the particular domain being modeled with the EC. EC5-EC7 express the EC clauses that deals with validity intervals of fluents. In particular, EC5 specifies that a fluent F keeps a value V for an interval going from Tmin to Tmax if nothing happens in the middle that breaks such an interval. EC6-EC7 behave like EC5, but deal with open intervals.

The rule interface builds on top of the EC to specify the following types of patterns, ordered by complexity:

1. Simple Pattern: as shown in Fig. 1 a simple pattern simply consider a threshold value, like for example on glucose values to be violated a certain number of times within a time period.
2. Complex Pattern: as shown in Fig. 1 a complex pattern considers the occurrence of several threshold violations concerning multiple signals within a certain time period.
3. Sequential Pattern: as shown in Fig. 2 sequential patterns show the occurrence of a threshold violation followed by another threshold violation.
4. Complex Sequential Pattern: as shown in Fig. 2 complex sequential patterns are a combination of sequential patterns and complex patterns.




Rule in EC:

```

initiates_at(alert, T):-
    P1= physiological_value1>N happened between time Ts1, Te1
    P2= physiological_value1<N happened between time Ts2, Te2
    not same alert sent within this day,
    more_or_equals(#E, P1,P2), Ts1 < Ts2.

```



```

initiates_at(alert, T):-
    P1= physiological_value1>N happened between time Ts1, Te1
    P2= physiological_value2<N happened between time Ts2, Te2
    P3= physiological_value3>N happened between time Ts3, Te3
    P4= physiological_value4<N happened between time Ts4, Te4
    not same alert sent within this day,
    more_or_equals(#E, P1,P2),
    more_or_equals(#E, P1,P2),
    Ts1 < Ts3, Ts2 < Ts3, Ts1 < Ts4, Ts2 < Ts4.

```

Figure 2: Sequential Pattern Specification and Complex Sequential Pattern Specification.

The predicates shown in Figs 1-2 depend on a set of meta-predicates to evaluate the execution of the logical patterns and put it in relationship with the other patterns. In particular, we modelled More_or_equals/3:


```

more_or_equals_to(Frequency, Pattern, Time)←
  apply(Pattern,Time)
  findall(., Pattern,List), size(List,S), S=Frequency.

```

the pseudocode above takes the temporal pattern **Pattern** and it unifies it with the time variable **Time**, after this the **findall/3** predicate executes the pattern to see how many times it holds in the specified time. A complex pattern then uses the **More_or_equals/3** predicate on multiple patterns, using the declarativity of Prolog to combine the patterns. Sequential patterns use a **Constrained_more_or_equals/5** predicate, in particular this predicate takes into consideration the execution of the patterns between two periods of time that must be strictly happening in a sequential order.

```

constrained_more_or_equals_to(Frequency, Pattern, Pattern2, Period1, Period2)←
  apply(Pattern,Period1),
  apply(Pattern2,Period2),
  findall(., (Pattern,Pattern2,List), size(List,S),
  S=Frequency.

```

The definitions of **More_or_equals/3** and **Constrained_more_or_equals/5** are not recursive because we found that the semantics become unclear to a reader which is something we explicitly want to avoid given the medical settings of D1NAMO. Thus, to maintain understandability of the generated rules, we decided to allow only for a double level of nesting of the rules. Fig. 3 provides an example of a JavaScript implementation of the rule interface of D1NAMO.




Figure 3: Javascript Rule Interface in D1NAMO.

The rules created with the rule editor are compiled directly into **initiates_at/2** predicates. For example, if we wanted to specify a pattern to monitor continuous glucose events and heart rate at the same time, we could use a complex pattern as shown in Fig. 4.


```

not((query_kd(happens_at(sent_alert(generic_alert([doctor]), Th), [Tbefore,T]))).

initiates_at(generic_alert([doctor,rule1])=sent, T):-
happens_at(sent_alert(generic_alert([doctor,rule1]),T).

```

The pattern states to look for complex pattern in which a heart rate of 130 and a CGM value above 15 is then followed by a drop in CGM and heart rate within time frame of one day. Such a pattern can be used in D1NAMO to observe patients that present a variable diabetes pattern. The main issue with defining these patterns is that they depend on primitives of the EC such as the `holds_at/2` predicate. The main issue of EC predicates is that, being based on declarative programming and the **negation as failure** procedure of Prolog, they may be quite expensive to compute. For this reason, the next session presents the CEC_{KD} event calculus, an extension of the cached event calculus [8] (CEC) that indexes the events using kd-trees structures.

3. Dealing with Many Events in the Event Calculus

With respect to the rules that we model for D1NAMO, the EC formalism gets particularly slow when dealing with large amounts of events. Recent research has tried to overcome this issue in several ways [3, 8], but the amount of events that can be processed with such approaches is still very limited and not adequate to the issue of monitoring patients affected by a chronic illness, who notably produce many thousands of events even within one week. If we consider continuous glucose monitoring, for example, at least 288 values per day are produced. The normal versions of the EC currently available in research would not be able to cope with such a large number of events without becoming a bottleneck for the analysis to be performed. Within this deliverable we present and evaluate a solution for this problem. We described our extended EC as the CEC_{KD} event calculus, as it caches events like the cached event calculus, but it also indexes them using KD-trees. Such an event calculus has been defined with the idea of dealing with a large number of events. The necessity to define this extension resides in the fact that the medical domain typically implies heterogeneous types of events. For example, in D1NAMO we are faced with discrete events (point of care glucose samples) and continuous events (ECG, activity, continuous glucose monitors). As a consequence, to have a tool that can be truly useful, the temporal analysis must scale to thousands of events of different kind. The knowledge representation of the EC framework allows us an easy inclusion

of predicates to query different physiological values. The creation of a quick indexing scheme, as described in this deliverable, allows to deal with discrete and continuous domains.

3.1. K-dimensional Trees

Kd-trees [4] are binary trees optimised to deal with k-dimensional points. As reported in [13], given a set of k-dimensional points, we can generate a Kd-tree by splitting recursively the hyperplane containing the points at every level of the tree, alternating the coordinate that is split according to the depth of the tree. Fig. 6 shows how splits are performed on a 2-dimensional tree of depth 3, where at each level the value of the splitting coordinate is the median value, deciding if a new point should go to the left or to the right of an existing tree node.




Figure 6: Range Query on a Kd-tree [13]

Fig. 6 shows also the effect of searching a Kd-tree via a range query performed on it. The range query algorithm recursively searches for regions contained or intersected by the region specified in the range query. If the region found is contained in the region specified in the query, then the whole region is returned. If the region of the tree intersects the region specified by the query, then the points reported are only those ones included in the region of the query.

The Kd-tree data structure has a set of important properties when dealing with searches of multi-dimensional points: (a) a Kd-tree for a set P of n points uses $O(n)$ storage and can be constructed in $O(n \log(n))$ time; (b) the operations of adding or deleting a point have a complexity of $O(\log(n))$; (c) a rectangular range query on the Kd-tree takes $O(\sqrt{n} + k)$ time, where k is the number of reported points residing the rectangular area identified by the query.

These properties are fundamental to create a version of the EC that can scale up to be used in dynamic applications with large narratives.

Primitive Operations to Manage the Tree	Meaning	Theoretical Complexity
create_kdi(+L, Idx)/destroy_kdi(+L)	- Creates/Destroys a four-dimensional Kd-tree index <i>Idx</i> identified by label <i>L</i>	$O(1)$.
kdi(+L, -Idx)	Returns an existing four-dimensional Kd-tree index <i>Idx</i> identified by label <i>L</i>	$O(1)$.
insert_kdi(+Idx, (+Arg ₁ ,+Arg ₂ ,+Arg ₃ ,+Arg ₄),Value)/delete_kdi(+Idx, (+Arg ₁ ,+Arg ₂ ,+Arg ₃ ,+Arg ₄))	Inserts a four-dimensional key, whose coordinates are <i>Arg₁</i> , <i>Arg₂</i> , <i>Arg₃</i> , <i>Arg₄</i> , associated to a value <i>Value</i> in a Kd-tree index <i>Idx</i> . Such a predicate transforms the arguments in long integer values using an hashing function, whose range can be between $-\infty$ and $+\infty$. delete_kdi/2 deletes a value given a four dimensional key like for insert_kdi/2.	$O(\log(Events))$ for the events tree, $O(\log(MVI))$ for the MVI tree.
range_query_kdi(+Idx, (?R ₁ ,?R ₂ ,?R ₃ ,?R ₄), -Result)	Produces a four dimensional rectangular range query on the four dimensional kd-tree <i>Idx</i> , where the query is specified by the ranges <i>R₁</i> , <i>R₂</i> , <i>R₃</i> , <i>R₄</i> and the result is unified with the variable <i>Result</i> . The range arguments can be fixed values or specified as [StartValue, EndValue].	$O(\sqrt{ MVI } + K_{F,V})$ where <i>MVI</i> is the total number of MVIs and <i>K</i> is a constant related to the number of points reported depending on the value of <i>F</i> and <i>V</i> .
Cache Operations for Events and MVIs Indexing	Meaning	Theoretical Complexity
update(+Ev,+T)	Indexes an event happening at time <i>T</i> in a four dimensional Kd-tree index and caches its consequences in a MVI Kd-tree index.	$O(n(\sqrt{ MVI } + K_{F,V}))$. Where <i>n</i> is the number of queries to the context in the initiates.at/3/terminates.at/3 EC rules.
index(+Ev,+T)	Indexes an event happening at time <i>T</i> in a Kd-tree index.	like insert_kdi/5.
cache(+Ev,+T)	Caches an event happening at time <i>T</i> in a MVI Kd-tree index.	like insert_kdi/5.
close_interval(+Idx, (+F,+V,T _{end}))/open_interval(+Idx,(+F,+V,T _{start}))	close_interval/2 Indexes a closed MVI whose fluent is <i>F</i> , whose value is <i>V</i> , and whose ending time is <i>T_{end}</i> in <i>Idx</i> . open_interval/2 is like close_interval/2, but the time indexed is the starting time <i>T_{start}</i>	$O(\log(MVI))$.
intersect_query(+Idx, (?F,?V,-T _s , -T _e ,+WT _s ,+WT _e))	Uses range queries to find the MVIs that intersect <i>WT_s</i> , <i>WT_e</i> and unifies with <i>F</i> , <i>V</i> , <i>T_s</i> , <i>T_e</i> .	$O(\sqrt{ MVI } + K_{F,V})$.
cached_between(+WT _s ,+WT _{end} , mholds_for(?F=?V,[?T _s ,?T _{end}]))	Queries for the MVIs intersecting the time window between <i>WT_s</i> and <i>WT_e</i> .	$O(2 * \sqrt{ MVI } + K_{F,V})$.

Table 2: Predicates of CEC_{KD}. The symbols +, - and ? indicate respectively inputs, outputs and inputs/outputs.

3.2. The Cached Event Calculus with Kd-Trees

To obtain a version of the CEC that can scale up with respect to the number of events, we use two four-dimensional Kd-trees, one to index the events and one to index the MVIs in which a fluent holds. In addition to the CEC predicates, in CEC_{KD} we introduce new predicates to model our knowledge as stored in the Kd-trees. Table 2 summarises the predicates of CEC_{KD} and their theoretical complexities. These predicates take a query produced on an event or an MVI and translate the query to an insert point, delete point or range query on the event tree or on the MVI tree. In particular we do so maintaining a declarative approach, to keep the expressivity of CEC intact. We start with how we specify the addition of an event using the `update/2` predicate, which is specified as follows:

```

update(Ev,T)← index(Ev, T), cache(E, T).

index(Ev,T)← functor(Ev,Name,Arity), argument(Ev,1, Argument),
  kdi(happens_at, EvIdx), insert_kdi(EvIdx,(Name,Arity,Argument,T),Ev).

cache(Ev,T)← kdi(mholds_for, MVIdx)
  foreach(terminates_at(Ev,F1=V1,T),close_interval(MVIdx,(F1,V1,T))),
  foreach(initiates_at(Ev,F2=V2,T),open_interval(MVIdx,(F2,V2,T))).

```

To update the knowledge base with an event, we first index the event in a Kd-tree and then we add what is initiated and terminated. In CEC_{KD} the `index/2` predicate in `update/2` stores the produced event in a four dimensional Kd-tree, indexing its name, arity, first argument and time when it happened, using the `insert_kdi/3` predicate. Once the event is stored in the event Kd-tree, CEC_{KD} considers this as happened. The event is then used in the `cache/2` predicate to query the fluents whose values are initiated and terminated due to the event happening, which are then cached in the `MVIdx`. MVIs can be open to infinite or closed. We specify `close_interval/2` as follows:

```

close_interval(MVIdx,(F,V,T)←
  range_query_kdi(MVIdx, (F,V,[0,+∞],+∞), mholds_for(F,V,Ts,+∞)),
  delete_kdi(MVIdx,(F,V,Ts,+∞)),
  insert_kdi(MVIdx,(F,V,Ts,T),mholds_for(F=V,[Ts,T])).

```

The definition of `close_interval/2` takes into account the existence of an open interval, that has to be closed due to an event terminating the fluent F to have value V . The procedure to index open intervals is similar to the one for closed intervals with the difference that no MVI is retracted from the `mholds_for` Kd-tree. We then define the `holds_at/2` predicate, to query the value V of a fluent F as follows:

```
holds_at(F=V,T)←kdi(mholds_for,MVlidx),
  range_query_kdi(MVlidx, (F,V,[0,T],[T,+∞]), mholds_for(F=V,[T_s,T_e])).
```

This definition of `holds_at/2` looks for the interval of time that intersects T , speeding up the computation of `holds_at/2`. The `intersect_query/2` is implemented as a range query on a Kd-tree containing the MVIs, varying the starting time of the MVI from 0 to T , and the ending time from T to **positive infinite**. In the case that also F and V are variables, the range query is performed also on these coordinates, assuming a range for them between $-\infty$ and $+\infty$. If we want to query for the MVIs of a fluent F with value V we can redefine the `mholds_for/2` predicate as follows:

```
mholds_for(F=V, [T_s,T_end])←kdi(mholds_for,MVlidx),
  range_query(MVlidx,(F,V,T_s,T_end), mholds_for(F=V,[T_s,T_end])).
```

In the worst case, when both F and V are not defined, the predicate will backtrack through all the MVIs contained in the MVI tree, like the normal CEC would do. In the case that one between F and V is defined, the query will only select those intervals related to a particular fluent and return them. Furthermore, sometimes it is not worth to query the whole list of MVIs of a fluent, when a fluent is known to change quite often in time. To avoid performing expensive queries at update time, we define a further query on the MVI tree as follows:

```
cached_between(WT_s,WT_e, mholds_for(F=V, [T_s,T_e]))←
  kdi(mholds_for,MVlidx),intersect_query(MVlidx,(F,V,T_s,T_e,WT_s,WT_e)).

intersect_query(MVlidx,(F,V,T_s,T_e,WT_s,WT_e))←
  range_query(MVlidx,(F,V,[0,WT_s],[WT_s,+∞]),mholds_for(F=V,[T_s,T_e])).
intersect_query(MVlidx,(F,V,T_s,T_e,WT_s,WT_e))←
  range_query(MVlidx,(F,V,[WT_s,WT_e],[WT_e,+∞]),mholds_for(F=V,[T_s,T_e])).
```

The `cached_between/3` above uses a time window defined between WT_s and WT_e to query the MVI tree about the intervals in which a fluent F took value V . Such a query will return only those intervals which intersect the time window, leaving out the intervals happening before or after the time window.

With respect to CEC, the main improvement from the perspective of computational complexity can be seen in the `update/2` predicate shown in Table 2. If for CEC, in a context dependent theory, the update time complexity is exponential with respect of the number of context queries performed [9], for `CECKD` the complexity depends on the number of queries to the context

n multiplied by the square root of the total number of MVIs and number of MVIs reported by the query. This happens thanks to the fact that we do not rely on negation as failure that would check the whole temporal database for a solution, but we define time windows, which restrict the number of solutions returned, and we use an indexing based on kd-trees, that treat the time windows as range queries. Similar considerations apply for the query time.

4. Evaluation

We evaluated CEC_{KD} by comparing it with the update time and query time of CEC. In order to perform this comparison we adapted our CEC_{KD} theory to CEC ensuring the final behaviour of the two theories is the same. Firstly, the testing environment is a Intel Core i7, with 8GBs of RAM and




Figure 7: From Top to Bottom: Update Time of CEC vs CEC_{KD} , Results of a Ground holds_at/2 Query.

the tests that follows were performed by repeating them 50 times, averaging the results and producing their standard deviation to show their precision. CEC and CEC_{KD} were executed on 2Prolog version 3.0, using the Java interface of 2Prolog as the basis to connect the Kd-tree structure implemented in Java with Prolog predicates.

For the purpose of the evaluation, we used data collected in the D1NAMO project, this data included 7 patients affected by diabetes type 1, monitored for a total of 355 hours, with a total number of 6660 events.

The number of events that a DT1 patients may produce is variable, they general produce a large number of events within several weeks of observation. In D1NAMO, we did not have event histories longer than 1000 glucose events per patient, so we bootstrapped the available data in order to have event histories of about 10000 events including CGM, normal glucose readings and

other events such as weight and recording of meal consumption during the day and created 50 artificial patients. This is acceptable in this settings as we want to evaluate the scalability of the system and not the accuracy of the rules, that in any case are assumed to be dynamically defined by the medical doctors monitoring the patients.

We then simulated the production of such events and fed them to CEC_{KD} and CEC.




Figure 8: From Top to Bottom: holds_at/2 Query Completely Unbound and Ram Memory Consumption of CEC vs CEC_{KD} .

In what follows we compare the update time of CEC_{KD} and CEC in the DT1 scenario and the holds_at/2 predicate for querying the current state of the patient. The update time contains calls to all the predicates of CEC_{KD} and since holds_at/2 predicate is based on mholds_for/2 in CEC and on range_query_kdi/3 in CEC_{KD} , we have a complete picture of both the formalisms. The part on top of Fig. 7 shows the comparison between CEC and CEC_{KD} with respect to the update time while producing events in the Prolog database. On one hand, CEC demonstrates a linear dependency on the number of events produced on the EC database. In particular 2Prolog implements an indexing mechanism on the first term of the dynamic predicates asserted, to speed up the computation, but this is not enough to avoid a linear dependency on the events produced. On the other hand CEC_{KD} curve has an access time to the Kd-tree structure which does not make it start from 0 nanoseconds, but then it is almost flat throughout the whole simulation. This happens because we make large use of the Kd-tree to avoid computing the whole list of intervals, and by using as much as possible range queries. This results in a computation that depends, for range queries, on the square-root of the events or MVIs stored in the Kd-trees and, for deletion and insertion of events and MVIs, on the logarithm of the number of events

or MVIs currently stored. Furthermore the curve associated to CEC_{KD} looks almost flat because the cost of performing a range query on the Kd-tree is negligible if compared to the cost of performing a linear search on the MVIs stored in 2Prolog, that implies multiple expensive unifications.

The part on the bottom of Fig. 7 shows the curves resulting from a `holds_at/2` query where all the arguments are ground. The fact that all the arguments are ground in the query, improves the response of 2Prolog indexing but there is still a linear dependency on the events produced in the database which does not take place with CEC_{KD} . The ground `holds_at/2` query of CEC_{KD} is particularly fast thanks to the fact that we can perform a range query where we look for one interval intersecting the current time where the value and the fluent match the query.




Figure 9: Effect of Multi-Threading on the Update Time and RAM Memory Consumption of CEC.

On the top of Fig. 8 we show what happens when we perform a `holds_at/2` query where both the fluent and the value are variables and where we look for all the available solutions. This is probably the heaviest query for CEC as 2Prolog requires to perform a linear search on all the intervals to test whether these intervals intersect the time when the query is performed. For CEC_{KD} this query requires only to perform a `range_query_kdi/3` as shown in the previous section, which on one side is proportional to the square root of the number of MVIs stored in the MVI tree, and on the other side it avoids accessing the MVIs by unification.

The part on the bottom of Fig. 8 illustrates the memory consumption of CEC and CEC_{KD} in a monothread simulation. The result suggests that CEC_{KD} consumes twice as much RAM memory than CEC. This is explainable by the fact that we use two different trees for storing the events and the time periods and by the overhead introduced by the tree itself.




Figure 10: Effect of Multi-Threading on the Update Time and RAM Memory Consumption of CEC_{KD} .

Fig. 9 shows the effects of multi-threading on the computation of the update time in CEC. Despite running the tests on 2Prolog, which is optimised to run hundred of thousands of inferences, when adding multiple instances of the system loaded with CEC we have a situation where the instances compete for the CPU time, slowing themselves down, due to heavy use of backtracking and unification for each thread. Since the amount of patients that a PHS may monitor may be very large, using CEC may require to use a big amount of independent machines, which would imply a very expensive solution for the medical doctors running the PHS. The fact that the RAM memory consumption of CEC is quite contained does not help in the multi-threaded tests on the top of Fig. 10.

As far as CEC_{KD} is concerned, the use of Kd-trees allows our agents to avoid performing computationally demanding inferences using the 2Prolog engine and to query directly the Kd-trees containing the events and the MVIs of the fluents limiting the competition between the instances to obtain CPU time. In other words, the operation of accessing the nodes of the Kd-tree is less expensive than the unification procedure from the perspective of CPU usage, resulting in an efficient multi-threading behaviour in CEC_{KD} .

As shown on the bottom of Fig. 10 we can easily run 40 threads with 40 different CEC_{KD} instances and still having a quite acceptable update time. From the perspective of the memory consumption, the behaviour of CEC_{KD} is acceptable even when running 40 agents. As shown on the right of Fig. 10, 40 agents loaded with CEC_{KD} use around 500Mb of RAM, an amount available in most computers. Remembering that this is a crash test and that it is not the case that 40 medical doctors will use such an interface all at the same time, we conclude that CEC_{KD} is a good engine to support retrospective

analysis of time series in the chronic non communicable diseases settings.

5. Related Work

This paper is both related with works in the fields of complex event processing (CEP) [12] and in the field of efficient computing in logic programming. Concerning CEP the combination of a graphical formalism with EC rules finds a close contribution in [11], that uses the EC to model workflows. With respect to [11] the main contribution of this paper is that it models the logic formalism in blocks, using meta-programming techniques, where [11] rather focuses on execution problems concerning the workflow. Another prominent example is given by the ETALIS system [1]. ETALIS also defines how events evolve the state of a system, similarly to the EC. As a consequence of this similarity, it shares the same issues in terms of scalability. The main contribution of CEC_{KD} with respect to ETALIS is that CEC_{KD} has been developed to work on large streams of discrete and continuous events, whereas ETALIS focuses on discrete representations. Finally, another interesting contribution with respect to CEP with logic formalism is that provided by Teymourian et al in [23]. Such a contribution presents the combination of a logic formalism with issues such as subsumption and classes of events. Currently, the semantic representation of CEC_{KD} is flat, events do not have a semantic representation, in this sense producing an extension of CEC_{KD} that also consider semantics could be an interesting development.

There are a number of papers addressing efficient knowledge representation in the EC. We focus on those ones that keep the normal logic programming semantics as the original EC of [17]. For many practical applications we have found that the simple EC in [21] with recursive predicate definitions in the rules is sufficient. Formalisms that are more expressive [19, 16] often constrain some uses of recursion and are therefore beyond the scope of this work.

Before Chittaro and Montanari in [9, 10], Kowalski in [18] identified approaches to index events and validity intervals in EC influencing the research that took place subsequently. The Object Event Calculus (OEC) [15] can be seen as an extension of these ideas by relying on a simple version of the EC to model complex objects and their evolution in time. One of the attractive features of the OEC, is the ability to separate fluents representing single and multiple value attributes. Although the way these fluents were terminated by an event could be optimised, the overall knowledge representation required

more axioms for the initiation of attributes that cater for the object-based data model and their underlying inheritance structure.

Other attempts, such as the Reactive Event Calculus (REC), defined by Chesani et al. [7], builds on top of CEC, deriving the values of the maximum validity intervals by means of abductive reasoning on top of the SCIFF reasoner in order to achieve properties such as *irrevocability* and *soundness* of the EC. The axiomatisation of REC does not rely on `assert` and `retract` as in the other Prolog versions of the EC, but it relies on the constraint propagation mechanisms of SCIFF. This brings the advantage that the specification of REC is fully declarative, but, as reported in the developer notes of SCIFF¹, SCIFF indexes only on the functor of the produced events, which brings serious drawbacks on the computation time when dealing with large narratives.

Urovi et al. in [24] present a scalability evaluation of the Ambient Event Calculus (AEC) [6], a distributed and indexed version of the OEC. Urovi et al. evaluate the AEC with many thousands of events considering local and distributed settings. We do not consider distributed settings, we rather focus on scaling up CEC_{KD} for event recognition purposes, but, from the times reported in [24], where for 5000 events an update time and a query time of 2 seconds for the local tests, we can say that CEC_{KD} indexes better the events than AEC, which is based on the underlying indexing mechanisms of OEC.

In [2] Artikis et al. propose the LTAR-EC and they explain how the events are indexed by using a time window at compilation time. Our main differences with the Artikis et al. work is that we do not need to compile our theory with a fixed time window, we can specify, if needed, an arbitrarily wide time window at runtime. Similarly, we do not rely on the Prolog indexing capabilities as these may vary between different Prolog implementations, we rather use a Kd-tree indexing mechanism that allows us to represent our MVIs and events as multidimensional points of which we index several properties to speed up both the query and update time.

A recent contribution to the problem of indexing EC events is also discussed in [14]. Such a contribution defines as an activity recognition framework based on compiling the EC events happening within a certain time period, that comprises a month for the purpose of the application described

¹SCIFF developer Manual: <http://lia.deis.unibo.it/sciff>

in the paper. This certainly renders the EC more scalable, but it is dependent on the specific granularity of time selected. The most interesting aspect of CEC_{KD} is that it handles time as a continuous attribute, supported by the capabilities of Kd-trees of handling spatio-temporal multi-dimensional points. As a consequence, CEC_{KD} is time granularity invariant, although still dependent on the number of events happening, its indexing is not statically defined.

6. Conclusions and Future Work

This paper presented the strategy that D1NAMO uses to analyse events from CGM devices and point of care devices. Specifically, in this paper we have presented an efficient version of the Event Calculus that caches events and their effects using an indexing scheme based on Kd-trees. We have studied the benefits of such integration by showing how to revisit previous work to produce a new temporal reasoning system that we called CEC_{KD} . One of the main advantages of CEC_{KD} is that it allows us to support scenarios that produce long narratives on which intelligent agents can reason about how to monitor patients. We tested the formalism on events coming from D1NAMO dataset and compared the result with respect to the standard CEC, showing that our formalism outperforms the original CEC in terms of querying times, but implying a slight increase in the use of the RAM memory. Possible future work implies looking into the problem of creating an ontology of events, in order to allow description logic reasoning in terms of subsumption. Another possibility could be to change the indexing mechanism of CEC_{KD} and try different structures than KD-trees.

7. Acknowledgements

This work was partially funded by the FP7 Project COMMODITY12 Grant Agreement No. 287841 and by the Nano-tera.ch initiative through the D1NAMO project. We would like to acknowledge the precious comments of Professor Kostas Stathis concerning earlier versions of this paper.

References

- [1] Darko Anicic, Paul Fodor, Sebastian Rudolph, Roland Stühmer, Nenad Stojanovic, and Rudi Studer. Etalis: Rule-based reasoning in event

- processing. In Sven Helmer, Alex Poulouvasilis, and Fatos Xhafa, editors, *Reasoning in Event-Based Distributed Systems*, volume 347 of *Studies in Computational Intelligence*, pages 99–124. Springer, April 2011.
- [2] Alexander Artikis, Marek Sergot, and Georgios Paliouras. A Logic Programming Approach to Activity Recognition. In *Proceedings of the 2nd ACM international workshop on Events in multimedia*, EiMM '10, pages 3–8, New York, NY, USA, 2010. ACM.
 - [3] Alexander Artikis, Marek J. Sergot, and Georgios Paliouras. An event calculus for event recognition. *IEEE Trans. Knowl. Data Eng.*, 27(4):895–908, 2015.
 - [4] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517, 1975.
 - [5] Stefano Bromuri, Serban Puricel, René Schumann, Johannes Krampf, Juan Ruiz, and Michael Ignaz Schumacher. An expert personal health system to monitor patients affected by gestational diabetes mellitus: A feasibility study. *JAISE*, 8(2):219–237, 2016.
 - [6] Stefano Bromuri and Kostas Stathis. Distributed Agent Environments in the Ambient Event Calculus. In *DEBS '09: Proceedings of the third international conference on Distributed event-based systems*, New York, NY, USA, 2009. ACM.
 - [7] Federico Chesani, Paola Mello, Marco Montali, and Paolo Torroni. A Logic-Based, Reactive Calculus of Events. *Fundam. Inform.*, 105(1-2):135–161, 2010.
 - [8] L. Chittaro, A. Montanari, M. Dojat, and C. Gasparini. The event calculus at work: a case study in the medical domain. In *Intelligent Systems Engineering, 1994., Second International Conference on*, pages 195–200, sep 1994.
 - [9] Luca Chittaro and Angelo Montanari. Efficient Temporal Reasoning in the Cached Event Calculus. *Computational Intelligence*, 12:359–382, 1996.

- [10] Luca Chittaro, Marco Del Rosso, and Michel Dojat. Modeling Medical Reasoning with the Event Calculus: An Application to the Management of Mechanical Ventilation. In Pedro Barahona, Mario Stefanelli, and Jeremy C. Wyatt, editors, *AIME*, volume 934 of *Lecture Notes in Computer Science*, pages 79–90. Springer, 1995.
- [11] Nihan Kesim Cicekli and Ilyas Cicekli. Formalizing the specification and execution of workflows using the event calculus. *Inf. Sci.*, 176(15):2227–2267, August 2006.
- [12] Gianpaolo Cugola, Alessandro Margara, Matteo Matteucci, and Giordano Tamburrelli. Introducing uncertainty in complex event processing: model, implementation, and validation. *Computing*, 97(2):103–144, 2015.
- [13] Mark de Berg, Marc van Kreveld, Mark Overmars, and Otfried Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, third edition, 2008.
- [14] Ozgur Kafali, Alfonso Romero, and Kostas Stathis. Agent-oriented activity recognition in the event calculus: An application for diabetic patients. *Computational Intelligence*, 2017. To appear.
- [15] F. Nihan Kesim and Marek J. Sergot. A Logic Programming Framework for Modeling Temporal Objects. *IEEE Trans. Knowl. Data Eng.*, 8(5):724–741, 1996.
- [16] Tae-Won Kim, Joohyung Lee, and Ravi Palla. Circumscriptive Event Calculus as Answer Set Programming. In Craig Boutilier, editor, *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence, Pasadena, California, USA, July 11-17, 2009*, pages 823–829.
- [17] R Kowalski and M Sergot. A logic-based calculus of events. *New Gen. Comput.*, 4(1):67–95, 1986.
- [18] Robert A. Kowalski. Database Updates in the Event Calculus. *J. Log. Program.*, 12(1&2):121–146, 1992.
- [19] Erik T. Mueller. Event Calculus Reasoning Through Satisfiability. *J. Log. Comput.*, 14(5):703–730, 2004.

- [20] Esra Kirci Ozorhan, Esat Kaan Kuban, and Nihan Kesim Cicekli. Automated composition of web services with the abductive event calculus. *Inf. Sci.*, 180(19):3589–3613, 2010.
- [21] Murray Shanahan. The Event Calculus Explained. In *Artificial Intelligence Today*, pages 409–430. 1999.
- [22] Murray Shanahan. An abductive event calculus planner. *J. Log. Program.*, 44(1-3):207–240, 2000.
- [23] Kia Teymourian and Adrian Paschke. *Semantic Rule-Based Complex Event Processing*, pages 82–92. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
- [24] Visara Urovi, Stefano Bromuri, Kostas Stathis, and Alexander Artikis. Towards Runtime Support for Norm-Governed Multi-Agent Systems. In Fangzhen Lin, Ulrike Sattler, and Miroslaw Truszczyński, editors, *KR 2010*, Toronto, Ontario, Canada, May 9-13 2010. AAAI Press.