

Analysing Temporal Relations – Beyond Windows, Frames and Predicates

Simon Brodt
Institute for Informatics
University of Munich
<http://www.pms.ifi.lmu.de/>
simon.brodt@ifi.lmu.de

François Bry
Institute for Informatics
University of Munich
<http://www.pms.ifi.lmu.de/>
bry@lmu.de

ABSTRACT

This article proposes an approach to rely on the standard operators of relational algebra (including grouping and aggregation) for processing complex event without requiring window specifications. In this way the approach can process complex event queries of the kind encountered in applications such as emergency management in metro networks. This article presents Temporal Stream Algebra (TSA) which combines the operators of relational algebra with an analysis of temporal relations at compile time. This analysis determines which relational algebra queries can be evaluated against data streams, i. e. the analysis is able to distinguish valid from invalid stream queries. Furthermore the analysis derives functions similar to the pass, propagation and keep invariants in Tucker’s et al. “Exploiting Punctuation Semantics in Continuous Data Streams”. These functions enable the incremental evaluation of TSA queries, the propagation of punctuations, and garbage collection. The evaluation of TSA queries combines bulk-wise and out-of-order processing which makes it tolerant to workload bursts as they typically occur in emergency management. The approach has been conceived for efficiently processing complex event queries on top of a relational database system. It has been deployed and tested on MonetDB.

Categories and Subject Descriptors

F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages—*Algebraic approaches to semantics, Operational semantics, Program analysis*; H.2.3 [Database Management]: Languages—*Query languages*

General Terms

Theory, Algorithms, Languages, Verification

Keywords

CEP, Event Processing, DSMS, Relational Algebra, Temporal Analysis, Punctuations, Frames, Predicate Windows

1. INTRODUCTION

Data stream management systems (DSMS) widely use tumbling, sliding, landmark and similar kinds of windows [4, 9, 20, 21] for handling “blocking operators” like negation or aggregation and “unbounded stateful operators” like join [25]. All these windows have in common, that they either have fixed bounds (tumbling/sliding windows) or last from a fixed bound to the present (landmark windows) and therefore are data independent. However a number of new applications, among others temperature, traffic and network monitoring [16, 23] require a data dependent way of detecting the “episodes of interest” [23] or “windows-of-interest” [16] within a data stream.

In emergency management applications like in metro networks, at airports or in power grids [30, 31] the relevant “episode” or part of a data stream is frequently defined by the start and the end of an emergency like a fire. Of course those queries responsible for detecting the start of an emergency need to be evaluated constantly and thus all or at least most of the data stream is relevant for those rules. However analytic queries like “the maximum smoke concentration within a room during a fire” should only be evaluated when actually needed i. e. in the presence of the corresponding emergency. Both the starting time as well as the temporal extension of the analysis is determined by events (the start and end of an emergency) and cannot¹ be defined in advance. Moreover, the precise knowledge of the start and end of some analysis period is sometimes interesting in itself, e. g. when monitoring an evacuation. Thus data independent window definitions are hardly suited for this kind of queries.

“Frames” [23] and “predicate-windows” [16] are two approaches towards detecting relevant parts of a data stream in an adaptive, data-dependent way. Interestingly both approaches introduce new operators although the standard operators of relational algebra² [2, 15] are actually able to express queries like the ones sketched for emergency management applications and (apparently) also those described in [23] and [16]. Instead of new operators we propose using the original operators of relational algebra for data stream processing. The challenge is that some relational algebra queries cannot be evaluated on data streams,³ i. e., are invalid. We address this challenge by proposing an in-depth compile-time analysis of temporal relations which first can distinguish valid from invalid stream queries and second en-

¹ Unfortunately not. Knowing the time and extent of an emergency in advance would be great.

² including grouping and aggregation

³ Or are not even well-defined like e. g. an infinite sum.

ables an incremental evaluation of relational algebra expressions on data streams.

The contributions of this article are as follows:

- We introduce a common data model for data streams and static relations generalizing the data model of relational algebra.
- We define Temporal Stream Algebra (TSA) which enhances the standard operators² of relational algebra by a mechanism for propagating constraints on temporal relations between attributes and on the correlation of attributes to the stream progress within the stream schema.
- We present an analysis for the propagated constraints that can determine the validity of TSA queries and derives functions similar to the invariants of [25] that define first when results can be passed on, second how information on stream progress is propagated and third which data needs to be kept.
- Based on these functions we finally describe an bulk-wise and out-of-order evaluation for TSA queries reducing the evaluation of TSA queries on data streams to a repeated evaluation of ordinary relational algebra expressions.

The rest of this article is organized as follows: Section 2 motivates our approach based on examples from emergency management. Section 3 presents our data model for data streams and introduces a schema for data streams based on constraints on temporal relations between attributes and on the correlation of attributes to stream progress. Section 4 describes the propagation of the constraints through the operators of relational algebra and gives a definition of formally “valid” queries. Section 4 also shows that “valid” queries can in fact be evaluated on data streams. Section 5 illustrates how different kinds of queries can be expressed in TSA. Section 6 describes the incremental evaluation for TSA queries. Section 7 discusses relevant prior work and Section 8 points to the presented achievements and future work.

2. EMERGENCY MANAGEMENT

Emergency management applications typically consist of three major kinds of rules or queries: First *detection rules* which detect the start and the end⁴ of an emergency. Second *analytic rules* which in case of an emergency gather and analyse further information on the emergency as to provide a comprehensive assessment of the emergency situation. Finally *reactive rules* are used to execute or propose⁵ proper reactions to an emergency [29, 30, 31]. Within this article we focus on the interplay of detection and analytic rules.

Event-controlled aggregation. Detection queries, like the detection of fire, need to be evaluated constantly. By contrast analytic queries only become important in the presence of the corresponding emergency. For example, there is hardly any use in computing the maximum smoke concentration in a room in absence of a fire. However in case of a fire, knowing the maximum smoke concentration of a room is very important. Smoke is actually the greatest threat for the life of people[29]. With the beginning of a fire it is therefore required to constantly evaluate the maximum smoke concentration of a room as to choose and adapt safe, i.e.

⁴ Usually a human operator needs to confirm the end of an emergency. However the system might propose that an emergency could be marked as being finished.

⁵ Some actions will need operator approval.

smoke free, evacuation paths and an appropriate ventilation regime. When the fire-fighters arrive they (manually) take over both the further evacuation and the ventilation regime in particular because they bring along a lot of own equipment that is not connected to the emergency management system.⁶ Thus the constant evaluation of the maximum smoke concentration becomes irrelevant. However the fire-fighter might be interested in a query like “the maximum smoke concentration within a room from the beginning of the fire to the arrival of the fire-brigade” giving a first impression of the quality of the so far evacuation management.

Obviously neither the beginning of a fire nor the precise time until the arrival of the fire-brigade are known in advance. Thus data-independent windows, e.g. tumbling or sliding windows, hardly seem to be suited for this kind of scenario. Instead the aggregation, i.e. the computation of the maximum smoke concentration, needs to be controlled by events indicating the start and the end of the aggregation, i.e. the detection of a fire and the arrival of the fire-brigade.

Access to Static Relations. Both detection and analytic queries frequently need to access static or hardly changing relations for correctly interpreting the incoming events. For example sensor messages usually carry sensor ids, but do not necessarily carry data indicating the locations of issuing sensors. The location is however essential for correlating messages from different sensors. So there should be an easy way to attach the location to the incoming sensor messages. More complex topological information is needed for handling an emergency in a power-grid but also in case of a fire in metro stations and airports[31] or other buildings. For example neighbouring rooms of a burning area, are directly threatened by the fire. Identifying those rooms needs to correlate the event of fire detection with the static neighbour relation of rooms.

User-defined Timestamps for Composite Events. In emergency management the “right” choice for the timestamp of an composite event is not always obvious and may differ from query to query. Consider for example the following two rules: Rule 1: A temperature sensor is considered to be malfunctioning if there is a message from that sensor which is not followed by another message within 30 seconds. Rule 2: A precautionary fire alarm is raised if there is a potential fire detection in an area and no report from the responsible warden is received within 2 minutes after the potential fire detection. Both queries have the same basic structure, a positive event that triggers the rule if it is not followed by another event within a certain time-frame. However for the first query, the timestamp for the “malfunction sensor” event should probably be the time of the first missing messages, e.g. 10 seconds after the last message from that sensor. By contrast the time for the precautionary alarm should be the time of the (potential) fire detection and not the time of the missing report. Therefore the user should be able to choose the most meaningful definition for the timestamps of a composite events on a query per query basis .

Multiple Time Lines. Emergency management requires queries referring to timestamps according to multiple different timelines. In a simple case these timelines are just application *and* system time. Reconsider the rule – A precautionary fire alarm is raised if there is a potential fire detec-

⁶ Actually we learned that the fire-fighters will not entrust their lives to any foreign emergency management system but prefer to do their own reconnaissance.

tion in an area and no report from the responsible warden is received within 2 minutes after the potential fire detection. The timespan of 2 minutes is defined between the detection of the potential fire, i. e. application time, and the reception of the report, i. e. system time. As to ensure that no more than two minutes will elapse between the potential detection and a reaction to that detection it is crucial that the 2 minutes period is defined in exactly that way, as transmission delays could otherwise block the emergency management system.

Further timelines are introduced by simulations used for predicting the physical evolution of an emergency, e. g. smoke propagation, which is very important for taking the right decisions, e. g. choosing safe evacuation routes. Beside application and system time, events returned by such simulations carry a timestamp denoting the time for which the prediction is meant to hold. Multiple timelines are also useful for first implementing slack [1] on a query per query bases and and in a way transparent for the user and second for skipping the derivation of composite events which have become irrelevant by lapse of time e. g. due to prioritization in favor of more important queries.

Supporting multiple timelines requires more than having multiple timestamp attributes as different timelines impose different orders on the events. Therefore in-order processing is only possible for at most one of the timelines. For all other timelines the evaluation is inevitably out-of-order. Thus, beside other benefits, out-of-order processing [21], seems to be almost mandatory for supporting multiple timelines.

Relational Algebra. The operators of relation algebra² turned out be sufficient for expressing queries like the ones described above.⁷ Thus it seemed natural to generalize the operators of relational algebra towards data streams. Using the operators of relational algebra has a number of advantages: The operators form a good basis for the implementation of a high-level declarative language like for example the event, state and action language Dura [28, 26, 27]. Data streams can be queried in the same way that is already familiar from database relations. The properties of the operators are well understood. For example the laws on operator permutations are very important for the optimization of relational algebra expressions. There exist reliable techniques for the evaluation and optimization of relational algebra expressions, like specialized join algorithms, heuristics for operator reordering or cost-models, that base on these laws and on other properties of the operators of relational algebra. Preserving the operators and their characteristics means that these techniques and algorithms can be reused or easily adopted for data streams.

3. STREAMS & TEMPORAL RELATIONS

A key observations on data streams made in [21] is that, though they might arrive out-of-order with respect to any of their attributes, data streams usually make “progress” with respect to some of their attributes. More precisely a data stream is making “progress” on an attribute a if it eventually exceeds any value v for that attribute. This lead to the definition of so-called “progressing streams” [21]. However the original definition of *progressing streams* is not fully suited

⁷ This does not mean that using the operators of relational algebra is a requirement for emergency management. It seems to be good choice, though.

for our purposes. First the cross-product of two progressing streams is not a progressing stream itself, second static relations are not covered, and third the definition implies a one-by-one arrival of tuples inducing some strict (though arbitrary) order on the tuples.⁸

Thus, we use the more general notion of *temporal streams*. In fact all progressing streams are temporal streams and a progressing stream its corresponding temporal stream have the same progressing attributes.⁹ Temporal streams differ from progressing streams as follows:

1. Prefixes of a data stream are defined with respect to time and not to tuple count
2. Data streams may be static, i. e. all tuples are available from the very beginning of the data stream
3. Progressing attributes are generalized to *progressing sets of attributes*

DEFINITION 1. (Temporal Streams)

1. A *data stream* R with attribute schema $\mathcal{A}(R) \subseteq ATTR$ ($ATTR$ is a set of attribute names) is a (possibly infinite) relation with finite prefix R^p at each point in time $p \in \mathbb{Q}$, i. e.

$$R = \bigcup_{p \in \mathbb{Q}} R^p \text{ with } p_1 \leq p_2 \Rightarrow R^{p_1} \subseteq R^{p_2}$$
 and $|R^p| < +\infty$ for all $p, p_1, p_2 \in \mathbb{Q}$
2. R is *static* iff $R^p = R$ for all $p \in \mathbb{Q}$
3. \emptyset is a *progressing set of attributes* for R iff R is static
4. $\{a_1, \dots, a_k\} \neq \emptyset$ is a *progressing set of attributes* for R iff $\{a_1, \dots, a_k\} \subseteq \mathcal{A}(R)$ and R eventually exceeds any upper bound $s_1, \dots, s_k \in \mathbb{Q}$ for a_1, \dots, a_k , i. e.

$$\exists p \in \mathbb{Q} : \{r \in R \mid r(a_1) \leq s_1 \wedge \dots \wedge r(a_k) \leq s_k\} \subseteq R^p$$
5. A data stream R is a *temporal stream* iff R has a progressing set of attributes

Progressing sets of attributes play quite a hidden role in the rest of the article although they are absolutely required. For input and output relations we are mostly interested in progressing sets of attributes with a single element, i. e. the so-called progressing attributes.⁹ Progressing sets of attributes are important particularly for intermediate results like the cross-product of two temporal streams and for the propagation of information on progressing attributes through the query expression.

The approach is centered around the following observation: If the progressing attributes of the temporal input streams of a relational algebra expression and the temporal relations specified in the expression imply at least one progressing attribute for the output stream then the expression is valid and can be incrementally evaluated on temporal streams. In fact there exists a (monotone) function that computes the minimum achievable progress of the output stream with respect to its progressing attributes based on the progress of the input streams with respect to their progressing attributes.

In the rest of this section we introduce the ingredients of a schema for temporal streams that can be used to describe, propagate and analyse information on progressing attributes and temporal relations.

⁸ Assuming a one-by-one arrival of events or tuples seems to be artificial in parallel or even distributed environments and the implicit postulation of an order (even an arbitrary one) seems to be inconvenient for a model of unordered streams.

⁹ An attribute a is called a *progressing attribute* of a temporal stream iff $\{a\}$ is a progressing set of attribute.

3.1 Temporal Relations

Temporal relations play a major role with respect to progressing attributes⁹ (or progressing sets of attributes). For example if an attribute a_1 is a progressing attribute for a temporal stream R and attribute a_2 is known to be greater than a_1 for all tuples of R , e.g. after a selection $\sigma[a_1 \leq a_2]$, then a_2 is also a progressing attribute for R . Even more if neither a_1 nor a_2 are progressing attributes for R , but $\{a_1, a_2\}$ is a progressing set of attributes for R , then a_2 is a progressing attribute for the derived stream $\sigma[a_1 \leq a_2](R)$.

Temporal relations are also introduced when new relative timestamps, e.g. for a composite event, are defined from existing ones. Consider the following definition of a new timestamp $a = \max(a_1, a_2)$. In that case it is known that a is greater than both a_1 and a_2 and therefore if $\{a_1, a_2\}$ is a progressing set of attributes for R then a is progressing attribute for $\iota[a = \max(a_1, a_2)](R)$ the stream with the newly introduced timestamp a .

In TSA new relative timestamps and explicit temporal relations are specified by means of *temporal terms* and *temporal relations formulas* (TRFs). Temporal terms can shift timestamps by a constant amount of time, they can express the maximum or minimum of a number of timestamps and they can derive less precise versions of a timestamp, e.g. $\text{floor}_{1\text{min}}(t)$ yields the last preceding full minute for timestamp t . The latter can be used to specify tumbling windows in a very elegant way (Section 5). Furthermore temporal terms can be nested. Beside attributes temporal terms may contain variables which are used when operators like projection discard attributes. In that case variables help to efficiently preserve transitive temporal relations. For example if $a_1 \leq a_2 \leq a_3$ and a_2 is discarded by a projection then $a_1 \leq v_{a_2} \leq a_3$ preserves the transitive relation $a_1 \leq a_3$ between a_1 and a_3 .¹⁰

DEFINITION 2. (Temporal Terms)

Temporal terms are defined inductively:

1. $v \in \text{ATTR} \cup \text{VAR}$ is an atomic temporal term
2. $t + c$ is a temporal term if t is a temporal term and $c \in \mathbb{Q}$
3. $\min(t_1, \dots, t_k)$, $\max(t_1, \dots, t_k)$ are temporal terms if t_1, \dots, t_k are temporal terms
4. $\text{floor}_c(t)$, $\text{ceil}_c(t)$ are temporal terms if t is a temporal term and $c \in \mathbb{Q}$

where VAR is a set of variables and $\text{VAR} \cap \text{ATTR} = \emptyset$.

TRFs consist of comparisons of timestamps defined by temporal terms which can be combined by conjunction, disjunction and negation. During the propagation of TRFs through the operators of TSA we sometimes need to express that a variable is not part of the namespace of a particular part of a formula. This is done by atoms $\text{ignore}(v)$.¹¹

¹⁰ Making the transitive relations explicit can cause a significant blowup of the formula.

¹¹ A full explanation on that point is out of the scope of this article. The basic reason for this is that the “variables” of temporal terms are used with two different semantics with respect to first order predicate logic, namely first as existentially quantified variables when the relations between attributes are concerned (compare Definition 9) or second as logic constants when the temporal distance (Definition 5) between *attributes* and *variables* is required. It is only when “variables” are used in the sense of logic constants that the formulas $\text{ignore}(v)$ have a meaning.

DEFINITION 3. (Temporal Relation Formulas)

Temporal relation formulas (TRFs) are defined inductively:

1. \top and \perp are atomic TRFs
2. $t_1 \text{ op } t_2$ is an atomic TRF for $\text{op} \in \{<, \leq, =, \geq, >, \neq\}$ if t_1 and t_2 are temporal terms
3. $\text{ignore}(v)$ is an atomic TRF if $v \in \text{VAR}$ is a variable
4. $G \wedge G'$, $G \vee G'$ and $\neg G$ are TRFs if G, G' are TRFs

Despite of the complex specifications that are possible using general TRFs they can be normalized to TRFs with a comparably simple structure. This is essential for the algorithmic analysis of the temporal relations. Note that the size of the normalized formula G^{norm} is only linear in the size¹² of the original formula G .

DEFINITION 4. (Normalized TRFs) A TRF G can be normalized to a TRF G^{norm} using the following equivalences and implications in left to right direction. The normalized TRF G^{norm} is equal to \top or \perp , or all atomic subformulas of G^{norm} have the form $v \leq w + c$ or $\text{ignore}(v)$ and negation occurs at most at literals of the form $\neg \text{ignore}(v)$.¹³

$$\text{Neg1} : \neg(G \wedge G') \Leftrightarrow (\neg G \vee \neg G')$$

$$\text{Neg2} : \neg(G \vee G') \Leftrightarrow (\neg G \wedge \neg G')$$

$$\text{Neg3} : \neg\neg G \Leftrightarrow G$$

$$\text{Neg4} : \neg(t_1 \text{ op } t_2 + c) \Leftrightarrow t_1 \text{ op}^{-1} t_2 + c$$

$$\text{Neg5} : \neg\top \Leftrightarrow \perp \text{ and } \neg\perp \Leftrightarrow \top$$

$$\text{Top} : G \wedge \top \Leftrightarrow G \text{ and } G \vee \top \Leftrightarrow \top$$

$$\text{Bot} : G \wedge \perp \Leftrightarrow \perp \text{ and } G \vee \perp \Leftrightarrow G$$

$$\text{Zero} : v \text{ op } w \Leftrightarrow v \text{ op } w + 0$$

$$\text{Eq} : t_1 = t_2 + c \Leftrightarrow (t_1 \leq t_2 + c) \wedge (t_1 \geq t_2 + c)$$

$$\text{Neq} : t_1 \neq t_2 + c \Leftrightarrow (t_1 < t_2 + c) \vee (t_1 > t_2 + c)$$

$$\text{Geq} : t_1 \geq t_2 + c \Leftrightarrow t_2 + c \leq t_1$$

$$\text{Gr} : t_1 > t_2 + c \Leftrightarrow t_2 + c < t_1$$

$$\text{Less}^{14} : t_1 < t_2 + c \Leftrightarrow t_1 \leq t_2 + (c - \varepsilon)$$

$$\text{Arith1} : t_1 + c \text{ op } t_2 \Leftrightarrow t_1 \text{ op } t_2 + (c)$$

$$\text{Arith2} : t_1 \text{ op } (t_2 + c) + d \Leftrightarrow t_1 \text{ op } t_2 + (c + d)$$

$$\text{Min1} : t \leq \min\{t_1, \dots, t_k\} + c \Leftrightarrow t \leq t_1 + c \wedge \dots \wedge t \leq t_k + c$$

$$\text{Min2} : \min(t_1, \dots, t_k) \leq t + c \Leftrightarrow t_1 \leq t + c \vee \dots \vee t_k \leq t + c$$

$$\text{Max1} : t \leq \max(t_1, \dots, t_k) + c \Leftrightarrow t \leq t_1 + c \vee \dots \vee t \leq t_k + c$$

$$\text{Max2} : \max(t_1, \dots, t_k) \leq t + c \Leftrightarrow t_1 \leq t + c \wedge \dots \wedge t_k \leq t + c$$

$$\text{Floor1} : \text{floor}_c(t_1) \leq t_2 \Rightarrow t_1 - c < t_2$$

$$\text{Floor2} : t_2 \leq \text{floor}_c(t_1) \Rightarrow t_2 \leq t_1$$

$$\text{Ceil1} : \text{ceil}_c(t_1) \leq t_2 \Rightarrow t_1 \leq t_2$$

$$\text{Ceil2} : t_2 \leq \text{ceil}_c(t_1) \Rightarrow t_2 < t_1 + c$$

for temporal terms t_1, \dots, t_k and $c, d \in \mathbb{Q}$ and TRFs G, G' and $v, w \in \text{ATTR} \cup \text{VAR}$ and $\text{op} \in \{<, \leq, =, >, \neq\}$ and $<^{-1} \mapsto >, \leq^{-1} \mapsto \geq, =^{-1} \mapsto \neq, \geq^{-1} \mapsto \leq, >^{-1} \mapsto <, \neq^{-1} \mapsto =$

A basic algorithm for the normalization of a TRF is very simple: Whenever the algorithm finds a syntactical match for the left side of one of the above equivalences and implications, then the matching subformula is replaced by the right side of the equivalence or implication. The algorithm stops if it does not find any further matches for the left side

¹² Size in number of terms not in number of atoms.

¹³ Literals of the form $\neg \text{ignore}(v)$ do not occur in practice.

¹⁴ Think of ε as an infinitely small value such that $c - \varepsilon < c$ but $d < c - \varepsilon$ if $d < c$ and $(c - \varepsilon) + (d - \varepsilon) = ((c + d) - \varepsilon)$.

of any of the equivalences or implications. The algorithm will always terminate with normalized form of the TRF. It is most efficient if it uses the equivalences and implications in top down precedence.

The temporal distance¹⁵ of two variables or attributes w and v with respect to a TRF is the maximum value that the difference $v - w$ can take, i.e. $v - w \leq \text{dist}(w, v)$ considering the restriction imposed by the TRF. The temporal distance is a core element of the compile-time analysis of temporal relations and is used to construct the functions (Definition 25) that propagate information on the progress of the inputs streams of a query to the information on the progress of the output stream of a query. Temporal distances can also be used for garbage collection (see end of Section 6) for which they have been proposed first [7]. The following definition significantly generalizes that of [7].

DEFINITION 5. (Temporal Distance) Let G be a temporal relation formula and R be a temporal stream with $\mathcal{A}(R) = A$. The *temporal distance* of two attributes or variables $v, w \in \text{ATTR} \cup \text{VAR}$ with respect to G is

$$\text{dist}_G(w, v) := \max_{C \in \text{dnf}(G^{\text{norm}})} \{ \text{dist}_C(w, v) \}$$

$$\text{dist}_C(w, v) := \begin{cases} -\infty & \text{if } C \text{ is inconsistent} \\ -\infty & \text{if } C \models \text{ignore}(v) \text{ or } C \models \text{ignore}(w) \\ \min\{c \mid \mathcal{T}\mathcal{D}, C \models_{\overline{\mathbb{Q}}} v \leq w + c\} & \text{else} \end{cases}$$

where $u, v, w \in \text{ATTR} \cup \text{VAR}$ and $c, d, +\infty, -\infty \in \overline{\mathbb{Q}}$ and G^{norm} is the normalized form of G and C is a conjunction of atomic TRFs of the form $v \leq w + c$ and $\mathcal{T}\mathcal{D}$ contains

$$\begin{aligned} \text{Ref} : & \quad v \leq v \\ \text{Trans} : & \quad u \leq v + c \wedge v \leq w + d \Rightarrow u \leq w + (c + d) \\ \text{Inf} : & \quad v \leq w + +\infty \end{aligned}$$

The algorithmic analysis of the temporal distances is closely related to the simple temporal problem (STP) [24] and the disjunctive temporal problem DTP [18]. Basically the (naive) analysis algorithm is as follows: The TRF is normalized and converted into disjunctive normal form. Each conjunction is an STP instance. The distance of all pairs of attributes and variables for this instance can be determined using any algorithm for the all-pair shortest path problem, e.g. the Floyd Warshall algorithm [14], or specialized algorithms for STP. The distance of two attributes or variables for the whole TRF is then the maximum distance of the two attributes or variables in any of the conjunctions.

3.2 Progressing Sets of Attributes

Stream bound formulas (SBFs) are used to describe the initial progressing sets of attributes of a temporal stream. The stream may have further derived progressing sets of attributes as illustrated in the beginning of Section 3.1. The formulas are called stream *bound* formulas because they actually tell which sets of attributes have the defining property of progressing sets of attributes, i.e. that *upper bounds* for the attributes yield a (finite) prefix of the stream.

DEFINITION 6. (Stream Bound Formulas) The set of *stream bound formulas* (SBFs) is defined inductively:

1. \top is an atomic SBF
2. $\text{bounded}(v, b)$ is an atomic SBF for $v \in \text{ATTR} \cup \text{VAR}$ where $b \in \text{BOUND}$ is a *stream bound identifier*.
3. $H_1 \wedge H_2$ and $H_1 \vee H_2$ are SBFs iff H_1 and H_2 are SBFs.

¹⁵Note that the temporal distance is usually asymmetrical.

For example the SBF for a temporal stream R with progressing attributes a_1 and a_2 has the form $\text{bounded}(a_1, b_{a_1}) \vee \text{bounded}(a_2, b_{a_2})$. The SBF for the cross-product of two temporal stream R and R with progressing attributes a and a' respectively, has the form $\text{bounded}(a, b_a) \vee \text{bounded}(a', b_{a'})$.

Given a TRF G and a SBF H we can determine all (initial and derived) progressing sets of attributes with respect to G and H in the following way:

DEFINITION 7. (Progressing Sets of Attributes)

$\{a_1, \dots, a_k\} \subseteq \text{ATTR}$ is a *progressing set of attributes* with respect to a TRF G and a SBF H iff

$\text{bounded}(a_1, b_1), \dots, \text{bounded}(a_k, b_k), G, \mathcal{T}\mathcal{D}, \text{SB} \models H$

for any $b_1, \dots, b_k \in \text{BOUND}$ ¹⁶ and SB contains

$$\begin{aligned} \text{BD} : & \quad v \leq w + c, \quad c < +\infty, \quad \text{bounded}(w, b_w) \Rightarrow \text{bounded}(v, b_v) \\ \text{IG} : & \quad \text{ignore}(v) \Rightarrow \text{bounded}(v, b_v) \end{aligned}$$

for $v, w \in \text{ATTR} \cup \text{VAR}$, $b_v, b_w \in \text{BOUND}$ and $c \in \overline{\mathbb{Q}}$

The algorithmic analysis of the stream bounds is similar to the one for temporal distances. The TRF G is normalized and converted into disjunctive normal form. For each conjunction C the following is done: First the distance between all attributes and variables in the conjunction is computed. Second for each atom in the SBF H , the atom is set to true if the distance from one of the attributes of the potential stream bound $\{p_1, \dots, p_k\}$ to the attribute or variable in the atom is finite. Otherwise the atom is set to false. If the H holds under this interpretation, then $\{p_1, \dots, p_k\}$ is a stream bound with respect to C . If $\{p_1, \dots, p_k\}$ is a stream bound with respect to all conjunctions, then $\{p_1, \dots, p_k\}$ is a stream bound with respect to G and H .

3.3 Temporal Stream Schema

The schema of a relation in relational algebra is just its set of attributes. Our schema for data streams complements the set of attributes with a TRF describing the temporal relations between the attributes and an SBF identifying the (initial) progressing sets of attributes. Using our stream schema we can propagate this information along the TSA operators and use it determining the validity of a TSA expression and for actually evaluating the expression. So the schema is key element of our approach. Note that our schema for temporal streams covers static relations. Thus we provide common schema for data streams and static relations.

DEFINITION 8. (Temporal Stream Schema)

1. A *temporal stream schema* is a triple $S = (A, G, H)$ where $A \subseteq \text{ATTR}$ is an attribute schema, G is a TRF with $\text{attr}(G) \subseteq A$ and H is a SBF with $\text{attr}(H) \subseteq A$
2. S is *static* iff H is equivalent to \top
3. $\{a_1, \dots, a_k\} \subseteq A$ is a *progressing set of attributes* for $S = (A, G, H)$ iff $\{a_1, \dots, a_k\}$ is a progressing set of attributes with respect to G and H
4. S is *valid* iff S has a progressing set of attributes

The definition of the matching data streams is straight forward. A data streams matches a temporal stream schema if it has the right set of attributes, the temporal relations are as specified in the TRF of the schema and all progressing sets of attributes for the schema are also progressing sets of attributes for the data stream.

¹⁶ Actually the stream bound identifiers of atomic SBFs do not play a role here.

DEFINITION 9. (**Matching Streams**) Let G be a TRF and H be a SBF and let R be a data stream.

1. G holds in R iff for all tuples $r \in R$ the instantiation $\sigma_r(G_{rel})$ of G is satisfiable in \mathbb{Q} :

$$\models_{\mathbb{Q}} \exists v_1, \dots, v_l : \sigma_r(G_{rel})$$

where $\mathcal{A}(R) = \{a_1, \dots, a_k\}$ and $\{v_1, \dots, v_l\} = vars(G)$ and $\sigma_r := \{a_1 \mapsto r(a_1), \dots, a_k \mapsto r(a_k)\}$ and G_{rel} results from G when replacing every atom $ignore(v)$ by \top if it occurs with positive polarity or by \perp if it occurs with negative polarity.¹⁷

2. H holds in R with respect to G , iff all progressing sets of attributes $\{a_1, \dots, a_l\} \subseteq \mathcal{A}(R)$ with respect to G and H are progressing sets of attributes for R .
3. R matches a temporal stream schema $S = (A, G, H)$ iff $\mathcal{A}(R) = A$ and both G and H hold in R .

PROPOSITION 10. If S a valid temporal stream schema then every data stream R matching S is a temporal stream.

PROOF. As S is valid there is a progressing set of attributes $A_{prog} \subseteq \mathcal{A}(S)$ for S . R matches S , therefore $\mathcal{A}(R) = \mathcal{A}(S)$ and A_{prog} is a progressing set of attributes for R . \square

Temporal stream definitions are used to name the different data streams (usually corresponding to event types) and to provide a formal schema definition for the streams. We impose some restrictions on the schema of temporal stream definitions. Most notably (b) implies that each progressing set of attributes contains at least one progressing attribute. This significantly simplifies the definitions of the functions (Definition 25) propagating information on stream progress from the input streams of a query to its output stream. The restrictions (a) and (c) are of technical nature.

DEFINITION 11. (**Temporal Stream Definition**)

A *temporal stream definition* is a pair $D = (n, S)$ where $n \in STREAM$ is a name for the temporal stream definition, $S = (A, G, H)$ is a valid temporal stream schema and

1. G and H do not contain variables
2. H is a disjunction of atomic SBFs, i. e.
 $H = bounded(a_1, b_1) \vee \dots \vee bounded(a_k, b_k)$
for some $a_1, \dots, a_k \in A$ and $b_1, \dots, b_k \in BOUND$
3. $sbid^D : a_i \mapsto b_i$ is an injection
with inverse $attr^D : b_i \mapsto a_i$

$STREAM$ is a set of names for temporal stream definitions. The schema S of $D = (n, S)$ is also denoted $\mathcal{S}(D)$.

4. OPERATORS RELOADED

TSA enhances the standard operators² of relational algebra with a mechanism for propagating constraints on temporal relations (TRFs) and progressing attributes (SBFs) to the schema of their output relation. So the interesting part of the following operator definitions is the definition of the schema for the result stream, not the definition of the result stream itself. Nevertheless we also give the definition for the output stream because first there hardly is *the one* representation of relational algebra, second we ourselves introduce two slight deviations from classical representations [2, 15] and third it seems convenient to have the definition of

¹⁷ An atom occurs with positive/negative polarity in C if it is found below an even/odd number of negations.

the result relations aside with the schema for these relations containing the propagated constraints. The two deviations from classical representations of relational algebra are as follows: First the basic TSA operators include a cross-product but no join operator. Second we separate the definition of new attribute from the projection operator and introduce the imbed operator ι for this purpose. Both deviations do not change the overall characteristics of the operators. In fact join, the classical projection and other frequently used operators like semi-join and anti-semi-join can be defined using the basic operators of TSA. However both deviations simplify the definition of constraint propagation as the operators become more orthogonal to each other, i. e., each performing a single task.

There are two operators that introduce new information on temporal relations namely selection σ and imbed ι . All other operators only combine and propagate the existing information in the schema of their input streams. The selection operator $\sigma[C]$ seeks to extract the maximum temporal information in its selection condition C into a TRF C_{temp} and adds C_{temp} to the TRF from the schema of its input stream. Basically C_{temp} results from C by replacing all non-temporal atoms in such a way that the condition C is fulfilled as much as possible.

DEFINITION 12. (**Selection σ**) Let E be a TSA expression with schema $\mathcal{S}(E) = (A, G, H)$ and let R be a temporal stream matching $\mathcal{S}(E)$ and C a condition with $attr(C) \subseteq A$.

$$\begin{aligned} \sigma[C](R) &:= \{r \in R \mid r \text{ satisfies } C\} \\ \mathcal{S}(\sigma[C](E)) &:= (A, (G \wedge C_{temp}), H) \end{aligned}$$

where C_{temp} is the TRF that results from C when replacing every non-temporal atom by \top if it occurs with positive polarity or by \perp if it occurs with negative polarity.¹⁷

The imbed operator ι adds the definitions of relative timestamps to the TRF of its input stream.

DEFINITION 13. (**Imbed ι**) Let E be a TSA expression with schema $\mathcal{S}(E) = (A, G, H)$ and let R be a temporal stream matching $\mathcal{S}(E)$. Let $a' \notin A$ be a new attribute, t the term defining a' and $a_1, \dots, a_k \in A$ the attributes occurring in t .

$$\begin{aligned} \iota[a' = t](R) &:= \{r' \in dom(A \cup \{a'\}) \mid \exists r \in R \text{ such that} \\ &\quad r'(a) = r(a) \text{ for } a \in A \text{ and} \\ &\quad r'(a') = f_t(r(a_1), \dots, r(a_k))\} \\ \mathcal{S}(\iota[a' = t](E)) &:= (A \cup \{a'\}, G', H) \\ G' &:= \begin{cases} G \wedge (a' = t) & \text{if } t \text{ is a temporal term} \\ G & \text{else} \end{cases} \end{aligned}$$

where $f_t : dom(a_1) \times \dots \times dom(a_k) \rightarrow dom(a')$ is the function of the values of $a_1, \dots, a_k \in A$ defined by t . If t is a temporal term, then t defines a *relative timestamp*.

Both the projection and the grouping operators propagate the TRF from the schema of the input stream by replacing the occurrences of the discarded attributes by variables.

DEFINITION 14. (**Projection π**) Let E be a TSA expression with schema $\mathcal{S}(E) = (A, G, H)$ and let R be a temporal stream which matches the schema of E and $A_1 \subseteq A$ the set of retained attributes.

$$\begin{aligned} \pi[A_1](R) &:= \{r' \in dom(A_1) \mid \exists r \in R \text{ such that} \\ &\quad r'(a) = r(a) \text{ for } a \in A_1\} \\ \mathcal{S}(\pi[A_1](E)) &:= (A_1, \xi(G), \xi(H)) \end{aligned}$$

where ξ is an injective substitution of the attributes in $A \setminus A_1$ by variables that do not occur in G or H .

The simple definition of grouping might be surprising, as grouping is one of the “blocking” operator which are usually considered “problematic”.

DEFINITION 15. (Grouping γ) Let E be a TSA expression with schema $\mathcal{S}(E) = (A, G, H)$ and let R be a temporal stream matching $\mathcal{S}(E)$ and $A_1 \subseteq A$ the set of grouping attributes. Let $a_1, \dots, a_k \in A \setminus A_1$ and $a'_1, \dots, a'_k \notin A_1$ and F_1, \dots, F_k aggregation functions like min, max, sum or avg.

$$\begin{aligned} \gamma[A_1][a'_1 = F_1(a_1), \dots, a'_k = F_k(a_k)](R) := & \\ & \{r' \in \text{dom}(A_1 \cup \{a_1, \dots, a_k\}) \mid \exists r \in R \\ & \text{such that } r'(a) = r(a) \text{ for } a \in A_1 \\ & \text{and } r'(a'_i) = F_i((r_\gamma(a_i))_{r_\gamma \in R_r}) \} \\ \mathcal{S}(\gamma[A_1][a'_1 = F_1(a_1), \dots, a'_k = F_k(a_k)](E)) := & \\ & (A_1, \xi(G), \xi(H)) \end{aligned}$$

where $R_r = \{r_\gamma \in R \mid r_\gamma(a) = r(a) \text{ for } a \in A_1\}$ and ξ is an injective substitution of the in attributes $A \setminus A_1$ by variables that do not occur in G or H .

The propagated TRF for the cross-product operator is basically the conjunction of the TRFs of the input streams.

DEFINITION 16. (Cross Product \times)¹⁸ Let E_1 and E_2 be TSA expressions with schema $\mathcal{S}(E_1) = (A_1, G_1, H_1)$ and $\mathcal{S}(E_2) = (A_2, G_2, H_2)$ that have disjoint attributes ($A_1 \cap A_2 = \emptyset$) and let R_1 and R_2 be temporal streams matching $\mathcal{S}(E_1)$ and $\mathcal{S}(E_2)$ respectively.

$$\begin{aligned} R_1 \times R_2 := \{r' \in \text{dom}(A_1 \cup A_2) \mid \exists r_1 \in R_1, r_2 \in R_2 \text{ with} \\ & r'(a) = r_1(a) \text{ for } a \in A_1 \\ & r'(a) = r_2(a) \text{ for } a \in A_2\} \\ \mathcal{S}(E_1 \times E_2) := (A_1 \cup A_2, \xi_1(G_1) \wedge \xi_2(G_2), \xi_1(H_1) \wedge \xi_2(H_2)) \end{aligned}$$

where ξ_1 and ξ_2 are injective substitutions, such that $\xi_1(G_1)$, $\xi_1(H_1)$ and $\xi_2(G_2)$, $\xi_2(H_2)$ use disjoint sets of variables.

The propagated TRF for the union operator is basically the disjunction of the TRFs of the input streams. Some details of the definition are of very technical nature,¹¹ particularly G'_1 and G''_2 .

DEFINITION 17. (Union \cup)¹⁸ Let E_1 and E_2 be TSA expressions with equal attributes, schema $\mathcal{S}(E_1) = (A, G_1, H_1)$ and $\mathcal{S}(E_2) = (A, G_2, H_2)$ and let R_1 and R_2 be temporal streams matching $\mathcal{S}(E_1)$ and $\mathcal{S}(E_2)$ respectively.

$$\begin{aligned} R_1 \cup R_2 := \{r \in \text{dom}(A) \mid \text{with } r \in R_1 \text{ or } r \in R_2\} \\ \mathcal{S}(E_1 \cup E_2) := (A, G', (\xi_1(H_1) \wedge \xi_2(H_2))) \\ G' := (G'_1 \wedge G''_1 \wedge \xi_1(G_1)) \vee (G'_2 \wedge G''_2 \wedge \xi_2(G_2)) \\ G'_1 := (\xi_1(p_1) = p_1) \wedge \dots \wedge (\xi_1(p_k) = p_k) \\ G''_1 := \bigwedge_{v \in \text{attr}(H_2) \cup \text{vars}(H_2)} \text{ignore}(\xi_2(v)) \\ G'_2 := (\xi_2(p_1) = p_1) \wedge \dots \wedge (\xi_2(p_k) = p_k) \\ G''_2 := \bigwedge_{v \in \text{attr}(H_1) \cup \text{vars}(H_1)} \text{ignore}(\xi_1(v)) \end{aligned}$$

where ξ_1 and ξ_2 are injective substitutions, that replace all attributes in A by variables and substitute variables by others such that $\xi_1(v) \neq \xi_2(w)$ for any two attributes or variables v and w occurring in A, G_1, G_2, H_1 or H_2 .

The propagated TRF for the set-difference operator is very technical nature.¹¹ The challenge is that the temporal relations between the *attributes* of the result stream are only determined by the temporal relations of the first input stream, i. e. by G_1 . Actually the propagated TRF G' is equivalent to G_1 with respect to the temporal relations between the *attributes* of the result stream. However G' also

¹⁸ The definition can easily be extended to a n-ary operator.

needs to propagate the information in the TRF G_2 of the second input. Otherwise the relations between the progressing sets of attributes for the result stream and the progressing sets of attributes for the second input stream get lost.

DEFINITION 18. (Set Difference \setminus) Let E_1 and E_2 be TSA expressions with equal attributes, schema $\mathcal{S}(E_1) = (A, G_1, H_1)$ and $\mathcal{S}(E_2) = (A, G_2, H_2)$ and let R_1 and R_2 be temporal streams matching $\mathcal{S}(E_1)$ and $\mathcal{S}(E_2)$ respectively.

$$\begin{aligned} R_1 \setminus R_2 := \{r_1 \in R_1 \mid \neg \exists r_2 \in R_2 \text{ with } r_1 = r_2\} \\ \mathcal{S}(E_1 \setminus E_2) := (A, G', (H_1 \wedge \xi_2(H_2))) \\ G' := (G_1 \wedge G'_1) \vee (G_1 \wedge G'_2 \wedge \xi_2(G_2)) \\ G'_1 := \bigwedge_{v \in \text{attr}(H_2) \cup \text{vars}(H_2)} \text{ignore}(\xi_2(v)) \\ G'_2 := (\xi_2(p_1) = p_1) \wedge \dots \wedge (\xi_2(p_k) = p_k) \end{aligned}$$

where ξ_2 is an injective substitution for attributes and variables, such that the substituted variants $\xi_2(G_2)$, $\xi_2(H_2)$ and G_1, H_1 use disjoint sets of variables.

DEFINITION 19. (TSA Query)

1. A TSA *query* is a pair $q = (D, E)$ such that E is a TSA expression, D is a temporal stream definition and the schema of E and D have the same set of attributes.
2. q is *valid* iff the schema $\mathcal{S}(E) = (A, G_E, H_E)$ of E matches the schema $\mathcal{S}(D) = (A, G_D, H_D)$ of D , i. e. G_E implies¹⁹ G_D and all progressing sets of attributes with respect to G_D and H_D are progressing sets of attributes with respect to G_E and H_E .

A TSA query assigns a TSA expressions to the temporal stream it contributes to. The most important property of *valid* TSA queries is, that they are non-blocking, i. e. can be evaluated incrementally.

PROPOSITION 20. (Non-Blocking Queries)

Let $q = (D, E)$ be a TSA query with schema $\mathcal{S}(D) = (A, G_D, H_D)$ and let D_1, \dots, D_k be the temporal stream definitions occurring in E and R_1, \dots, R_k temporal streams matching the schema $\mathcal{S}(D_1), \dots, \mathcal{S}(D_k)$ respectively.

If q is *valid* then q is *non-blocking*. This means and any limit $s \in \mathbb{Q}$ for a progressing attribute⁹ $a \in A$ of $\mathcal{S}(D)$, there is a point in time $p \in \mathbb{Q}$ such that for every point in time $p' \geq p$ the prefix of the result stream up to s is fully determined by the prefixes of the input streams at time p' .

$$\{r \in E(R_1, \dots, R_k) \mid r(a) \leq s\} = \{r \in E(R_1^{p'}, \dots, R_k^{p'}) \mid r(a) \leq s\}$$

PROOF (SKETCH). Let $\mathcal{S}(E) = (A, G_E, H_E)$. Without loss of generality one may assume that the temporal stream definitions D_1, \dots, D_k use different stream bound identifiers, as Proposition 20 and none of its indirectly referred Definitions depend on the actual names of stream bounds.

As q is a *valid* query, any progressing attribute⁹ $a \in A$ for $\mathcal{S}(D)$ is a progressing attribute⁹ for $\mathcal{S}(E)$.

Let G_E^{norm} be the normalized form of G_E (see Definition 4) and the DNF of G_E^{norm} be $\text{dnf}(G_E^{\text{norm}}) = C_1 \vee \dots \vee C_l$, where C_1, \dots, C_l are conjunctions of normalized atomic TRFs.

For each D_i , $1 \leq i \leq k$ and each C_j there must exist at least one²⁰ atomic SBF *bounded* $(v_{i,j}, b_{i,j})$ in H_E ²¹ where $v_{i,j} \leq a + \text{dist}_{C_j}(v_{i,j}, a)$ and $b_{i,j}$ belongs to D_i .

¹⁹Can be checked if right side does not contain variables. By definition G_D and H_D do not contain variables.

²⁰ If D_i describes a static relation this does not hold. But in that case there is nothing to show. One can choose $l_i = 0$.

²¹ Note that H_E is in CNF with exactly one clause per D_1, \dots, D_k (Proposition 21)

However $v_{i,j}$ is a renamed version of the progressing attribute⁹ $a_{i,j} := attr^{D_i}(b_{i,j})$ associated to $b_{i,j}$ in D_i . Thus if $a \leq s$ in the result stream then $a_{i,j} \leq s_{i,j} := s - dist_{C_j}(v_{i,j}, a)$ in the input stream D_i , at least for “case” C_j of G_H^{norm} .

Proposition 22 allows to apply the formal results from the definitions D_i to the actual relations R_i . Therefore

$$\{r \in E(R_1, \dots, R_k) \mid r(a) \leq s\} = \{r \in E(P_1, \dots, P_k) \mid r(a) \leq s\}$$

where

$$P_i := \bigcup_{1 \leq j \leq l} \{r_i \in R_i \mid r_i(a_{i,j}) \leq s_{i,j}\}$$

As $a_{i,j}$ is a progressing attribute of R_i there is a point in time $p_{i,j}$ such that $\{r_i \in R_i \mid r_i(a_{i,j}) \leq s_{i,j}\} \subseteq R_i^{p_{i,j}}$.

Therefore with $p := \max_{\substack{1 \leq i \leq k \\ 1 \leq j \leq l}} p_{i,j}$ and $R_i^p \subseteq R_i^{p'}$ iff $p \leq p'$

$$\{r \in E(R_1, \dots, R_k) \mid r(a) \leq s\} = \{r \in E(R_1^p, \dots, R_k^p) \mid r(a) \leq s\} \quad \square$$

The proof of Proposition 20. uses the following two auxiliary proposition on the structure of propagated SBFs and on the well-definedness of the propagation mechanism:

PROPOSITION 21. (Simple SBFs) For any TSA expression E with schema $\mathcal{S}(E) = (A, G_E, H_E)$ the SBF H_E is in conjunctive normal form (CNF).

PROOF. The proposition holds for temporal stream definitions (Definition 11). For each TSA operator the SBF of the composed expression is a conjunction of the SBFs of the subexpressions. \square

PROPOSITION 22. (Well-Definedness) Let E be a TSA expression with the *valid* schema $\mathcal{S}(E)$ where D_1, \dots, D_k are the temporal stream definitions occurring in E .

If R_1, \dots, R_k are temporal streams matching the schema of D_1, \dots, D_k respectively then the output $E(R_1, \dots, R_k)$ is a temporal stream and matches the schema $\mathcal{S}(E)$.

PROOF (SKETCH). The proof is straight-forward by induction over the structure of E . \square

PROPOSITION 23. (Relational Completeness) TSA is *relational complete* [10] for finite (non-stream) relations.²²

PROOF. On finite non-stream relations, TSA is equivalent to relational algebra. \square

5. TSA BY EXAMPLE

This section illustrates how frequent types of queries can be expressed using the operators of TSA. As to keep the TSA expressions more compact and readable, operators like rename δ , join \bowtie , semi-join \ltimes and anti-semi-join $\bar{\ltimes}$ are used in the examples. These operators can easily be composed from the basic TSA operators. All shown queries are valid TSA queries (Definition 19) and all but the query for frames without fragments can be garbage collected solely based on an analysis of temporal relations (see end of Section 6).

Stream and Static Data. The following examples shows a rule that takes raw events from temperature sensors $TRaw$ and supplements them with the location (*area*) of the sensors stored in static relation $TLoc$.

²² The term “relational complete” compares the expressive power of some formalism for querying *finite relations* to the expressive power of relational algebra. A generalization to temporal streams is, if at all possible, non-trivial.

$$\begin{aligned} \mathcal{S}(TRaw) &= (\{sid, time, temp\}, \top, bounded(TRaw:time, time)) \\ \mathcal{S}(TLoc) &= (\{sid, area\}, \top, \top) \\ \mathcal{S}(Temp) &= (\{sid, time, area, temp\}, \top, bounded(TRaw:time, time)) \\ Temp &\leftarrow \\ &\delta[t.* \rightarrow *, l.area \rightarrow area](\\ &\quad \pi[t.sid, t.time, t.temp, l.area](\\ &\quad \quad \sigma[t.sid = l.sid] \\ &\quad \quad \delta[* \rightarrow t.*](TRaw) \times \delta[* \rightarrow l.*](TLoc) \\ &\quad))) \end{aligned}$$

Tumbling Windows. The following rule is a typical example for tumbling windows. The rule computes the average speed of air movement measured by a certain airflow sensor within the last 20 seconds. This information is required as input for the simulation of smoke propagation.

$$\begin{aligned} \mathcal{S}(AFRaw) &= (\{sid, time, speed\}, \top, bounded(AFRaw:time, time)) \\ \mathcal{S}(Airflow) &= (\{sid, time, speed\}, \top, bounded(Airflow:time, time)) \\ Airflow &\leftarrow \\ &\delta[time20sec \rightarrow time, avgSpeed \rightarrow speed](\\ &\quad \gamma[sid, time20sec][avgSpeed = avg(speed)](\\ &\quad \quad \iota[time20sec = ceil_{20sec}(time)](\\ &\quad \quad \quad AFRaw \\ &\quad \quad))) \end{aligned}$$

Lazy Aligned Sliding Windows. The following example shows a sliding window, which is only evaluated when actually needed (lazy), and which is aligned to the event initiating the evaluation. The implemented query reads as follows: From the start to one hour after the end of an fire compute the average smoke concentration in room in the last 10 seconds with updates every second. The implementation of the windows uses the safe, i.e., non-blocking form of recursion provided by TSA.²³

$$\begin{aligned} \mathcal{S}(Smoke) &= (\{time, area, konz\}, \top, \\ &\quad bounded(Smoke:time, time)) \\ \mathcal{S}(FireStart) &= (\{fid, time\}, \top, \\ &\quad bounded(FireStart:time, time)) \\ \mathcal{S}(FireEnd) &= (\{fid, time\}, \top, \\ &\quad bounded(FireEnd:time, time)) \\ \mathcal{S}(Windows) &= (\{start, end\}, start = end - 10sec, \\ &\quad bounded(Windows:start, start) \vee \\ &\quad bounded(Windows:end, end)) \\ \mathcal{S}(AvgSmoke) &= (\{time, area, avgKonz\}, \top, \\ &\quad bounded(AvgSmoke:time, time)) \\ AvgSmoke &\leftarrow \\ &\delta[s.area \rightarrow area, w.end \rightarrow time](\\ &\quad \gamma[s.area, w.end][avgKonz = avg(konz)](\\ &\quad \quad \sigma[w.start \leq s.time \leq w.end](\\ &\quad \quad \quad \delta[* \rightarrow s.*](Smoke) \times \delta[* \rightarrow w.*](Windows) \\ &\quad \quad))) \\ Windows &\leftarrow \\ &\pi[start, end](\\ &\quad \iota[start = end - 10sec](\\ &\quad \quad \delta[time \rightarrow end](\\ &\quad \quad \quad FireStart \\ &\quad \quad))) \\ Windows &\leftarrow \\ &\delta[nstart \rightarrow start, nend \rightarrow end](\\ &\quad \pi[nstart, nend](\\ &\quad \quad \iota[nstart = w.start + 1sec, nend = w.end + 1sec](\\ &\quad \quad \quad \bar{\ltimes} [w.start \leq f.time + 1h < w.end](\\ &\quad \quad \quad \quad \delta[* \rightarrow w.*](Windows), \\ &\quad \quad \quad \quad \delta[* \rightarrow f.*](FireEnd) \\ &\quad \quad \quad) \\ &\quad \quad))) \end{aligned}$$

Frames. The following query implements a frame [23] starting at the beginning of a fire, lasting till the end of that fire and computing the maximum smoke concentration in every area during that period when the duration of the fire is at least one minute.

²³ Details on recursion are out of the scope of this article.

$$\mathcal{S}(\text{MaxSmoke}) = (\{start, end, area, maxConz\}, \\ end - start \geq 1min, \\ bounded(\text{MaxSmoke}:end, end))$$

$$\text{MaxSmoke} \leftarrow \\ \delta[fs.time \rightarrow start, fe.time \rightarrow end, \\ fs.fid \rightarrow fid, s.area \rightarrow area](\\ \gamma[fs.time, fe.time, s.area][maxConz = max(conz)](\\ \sigma[fs.time \leq s.time \leq fe.time, \\ fe.time - fs.time \geq 1min, fs.fid = fe.fid](\\ \delta[* \rightarrow fs.*](\text{FireStart}) \times \\ \delta[* \rightarrow s.*](\text{Smoke}) \times \\ \delta[* \rightarrow fe.*](\text{FireEnd}) \\))))$$

Frame Fragments. The following example implements frame fragments [23] for the preceding query. The first frame fragment is produced after one minute when it is known that a frame exists. Each minute another frame fragment is produced until the end of the frame, i. e., the end of the fire is detected. Then the final result for the frame is produced. Frame fragments increase efficiency, as smoke events only need to be stored for and averages only need to be computed over a period of one minute. The implementation uses the safe, i. e., non-blocking form of recursion provided by TSA.²³

$$\mathcal{S}(\text{MaxSmokeFrag}) = (\{start, end, area, maxConz\}, \\ end - start \geq 1min, \\ bounded(\text{MaxSmokeFrag}:end, end))$$

$$\text{MaxSmokeFrag} \leftarrow \\ \iota[end = start + 1min](\\ \delta[fs.time \rightarrow start, s.area \rightarrow area](\\ \gamma[fs.time, s.area][maxConz = max(conz)](\\ \times [fs.time \leq fe.time < fs.time + 1min](\\ \sigma[fs.time \leq s.time < fs.time + 1min](\\ \delta[* \rightarrow fs.*](\text{FireStart}) \times \\ \delta[* \rightarrow s.*](\text{Smoke}) \\) \\ \delta[* \rightarrow fe.*](\text{FireEnd}) \\))))$$

$$\text{MaxSmokeFrag} \leftarrow \\ \pi[start, end, area, maxConz] \\ \delta[fs.start \rightarrow start, fs.area \rightarrow area](\\ \iota[end = fs.end + 1min, \\ maxConz = max(fs.maxConz, NMC)](\\ \gamma[fs.start, fs.end, fs.area][NMC = max(conz)](\\ \times [fs.end \leq fe.time < fs.end + 1min](\\ \sigma[fs.end \leq s.time < fs.end + 1min, \\ fs.area = s.area](\\ \delta[* \rightarrow fs.*](\text{MaxSmokeFrag}) \times \\ \delta[* \rightarrow s.*](\text{Smoke}) \\) \\ \delta[* \rightarrow fe.*](\text{FireEnd}) \\))))$$

$$\text{MaxSmoke} \leftarrow \\ \pi[start, end, area, maxConz] \\ \delta[fs.start \rightarrow start, fs.area \rightarrow area](\\ \iota[end = fe.time \\ maxConz = max(fs.maxConz, NMC)](\\ \gamma[fs.start, fe.time, s.area][NMC = max(conz)](\\ \sigma[fs.end \leq s.time < fe.time, \\ fs.area = fs.area \\ fs.end \leq fe.time < fs.end + 1min](\\ \delta[* \rightarrow fs.*](\text{MaxSmokeFrag}) \times \\ \delta[* \rightarrow s.*](\text{Smoke}) \times \\ \delta[* \rightarrow fe.*](\text{FireEnd}) \\))))$$

6. INCREMENTAL EVALUATION

The incremental evaluation is obviously crucial for TSA. It allows to derive results continuously as the data arrives on the stream. The incremental evaluation bases on the important property of valid TSA queries shown in Proposition 20:

Any prefix of the result stream of a query q depends only on finite prefixes of the input streams of q .

The major challenge for the incremental evaluation of TSA queries is to pass only those results to the output that are complete or stable, i. e. don't change on the arrival of further input data. This is realized by *pass conditions* (Definition 27) which accept only those result tuples that can safely be output based on the current progress of the input streams. Pass conditions are very similar to the "pass invariants" of [25]. *Pass formulas* (Definition 24) are parametrized versions of the pass conditions which can be derived at compile-time. The actual pass condition is obtained from the pass formula by replacing the parameters $b \in BOUND$ (formally called "stream bound identifiers") by their so-called *progress values* (Definition 26). Progress values basically tell about the current progress of a data stream with respect to its different progressing attributes.⁹ Progress values are similar to linear punctuations [21] for the output stream of a query with the main difference that progress values are not imbedded into the output stream.

DEFINITION 24. (Pass Formulas) Let H be a SBF. The *pass formula* ΔH to H is defined recursively:

$$\begin{aligned} \Delta \top &= \top \\ \Delta bounded(p, b) &= p \leq b \\ \Delta (H_1 \wedge \dots \wedge H_k) &= \Delta H_1 \wedge \dots \wedge \Delta H_k \\ \Delta (H_1 \vee \dots \vee H_k) &= \Delta H_1 \vee \dots \vee \Delta H_k \end{aligned}$$

where $p \in ATTR \cup VAR$ and $b \in BOUND$. ΔH results from H by replacing each atom $bounded(p, b)$ in H by $p \leq b$.

When instantiating a pass formula to obtain the pass condition for the current increment computation of a query those progress values are used that will hold after the current increment computation has finished. The progress for the output stream that is achievable with the current increment computation depends on the progress of the input streams at the beginning of the increment computation, more precisely the progress values for the output stream at the *end* of the increment computation depend on the progress values of the input streams at the *beginning of the increment computation*. The actual function that computes the progress values for the output stream based on the progress values of the input streams is called *propagation function* (Definition 25) and is derived at compile-time. Propagation functions are closely related to the propagation invariants of [25].

DEFINITION 25. (Propagation Functions)²⁴

The *propagation function* for a TSA query $q = (D, E)$ with $\mathcal{S}(D) = (A, G_D, H_D)$ and $\mathcal{S}(E) = (A, G_E, H_E)$, and a stream bound identifier $b \in BOUND$ with corresponding attribute $a = attr^D(b_D) \in A$ (Definition 11) for which the atom $bounded(a, b)$ is contained in H_D , is

$$\begin{aligned} f_{b,q} &= f_{b_D, G_E, \Delta H_E} \\ f_{b, G_E, \Delta H_E} &= \min_{C_E \in dnf(G_E^{norm})} f_{b, C_E, \Delta H_E} \\ f_{b, C_E, \top} &= +\infty \end{aligned}$$

²⁴ Without loss of generality the definition assumes that all temporal stream definitions of a TSA program use different stream bound identifiers $b \in BOUND$. This can easily be realized if $BOUND = STREAM \times ATTR$ and for temporal stream definition $D = (n, S)$ and progressing attribute⁹ p of S the stream bound identifier $b = b^D(p)$ corresponding to p in D has the form $b = n : p$.

$$f_{b,C_E,(v \leq b_E)} = \begin{cases} +\infty & \text{if } \text{dist}_{C_E}(a, v) = -\infty \\ -\infty & \text{if } \text{dist}_{C_E}(a, v) = +\infty \\ b \leq b_E - \text{dist}_{C_E}(a, v) & \text{else} \end{cases}$$

$$f_{b,C_E,(\Delta H_{E_1} \wedge \dots \wedge \Delta H_{E_k})} = \min(f_{b,C_E,\Delta H_{E_1}}, \dots, f_{b,C,\Delta H_{E_k}})$$

$$f_{b,C_E,(\Delta H_{E_1} \vee \dots \vee \Delta H_{E_k})} = \max(f_{b,C_E,\Delta H_{E_1}}, \dots, f_{b,C,\Delta H_{E_k}})$$

where $v \in \text{ATTR} \cup \text{VAR}$, $b_E \in \text{BOUND}$ and C_E is a conjunction of atoms of the form $v \leq b_E$ and $\text{ignore}(v)$.

The idea is, to enforce $v \leq b_E$ in the input relation using a and b_D from the output relation and the inequation

$$v \leq a + \text{dist}_{C_E}(a, v) \leq b + \text{dist}_{C_E}(a, v) \stackrel{!}{\leq} (b_E - \text{dist}_{C_E}(a, v)) + \text{dist}_{C_E}(a, v) \leq b_E$$

The minimum in the first case results from the fact that H_E must hold in each case of G_E (compare with Definition 6).

The propagation function $f_{b,q}$ can be normalized to be $+\infty$ or $-\infty$ or not to contain $+\infty$ or $-\infty$ at all. If $f_{b,q} = -\infty$, then a is not a progressing attribute⁹ with respect to G_E and H_E . This will raise a compile error, as it actually means that the query is invalid (Definition 19). If $f_{b,q} = +\infty$ then G_E and H_E belong either to a query that refers only to static relations or to a query that has an empty output stream as it imposes inconsistent temporal relations.

As already mentioned progress values tell about the current progress of a data stream with respect to its different progressing attributes.⁹ The progress values of the output stream of a query are computed by applying the corresponding propagation functions to the progress values of the input streams of the query. The progress value for a derived data stream is the minimum progress value of the output streams of all queries contributing to the derived data stream.

DEFINITION 26. (Progress Values)²⁴

Let P be a set of TSA Queries. For a TSA Query $q = (D, E)$ with $\mathcal{S}(D) = (A, G_D, H_D)$ let $s_q^i \in \mathbb{Q}$ and $e_q^i \in \mathbb{Q}$ denote the start and end time of the i th ($i \geq 1$) increment computation for q . The following is defined simultaneously:

1. Progress value of b for query q at the i th increment computation for q and $i \geq 0$

$$\text{prop}_q^0(b) = -\infty \quad ^{25}$$

$$\text{prop}_q^i(b) = f_{b,q}(\text{prop}_{s_q^i}(b_1), \dots, \text{prop}_{s_q^i}(b_l))$$

where b occurs in H_D and $f_{b,q}$ is the propagation function corresponding to q and b (Definition 25) and b_1, \dots, b_l are the stream bound identifiers occurring in $f_{b,q}$.

2. Progress value of b for query q at time $p \in \mathbb{Q}$

$$\text{prop}_q^p(b) = \begin{cases} \text{prop}_q^0(b) & \text{for } p < e_q^1 \\ \text{prop}_q^i(b) & \text{for } e_q^i \leq p < e_q^{i+1} \end{cases}$$

3. Value of b for temporal stream definition D at time p

$$\text{prop}_D^p(b) = \min_{\substack{q' \in P \\ \text{out}(q')=D}} \{\text{prop}_{q'}^p(b)\}$$

The progress value $\text{prop}_D^p(b)$ for input streams must be provided from outside.²⁶

4. Progress value of b at time p

$$\text{prop}^p(b) = \text{prop}_{D_b}^p(b)$$

²⁵ For recursive programs the initial values have to be chosen more carefully. This is out of the scope of this article.

²⁶ See [25, 20] for a discussion on how to obtain punctuations for input streams.

where D_b is the temporal stream definition corresponding to b .²⁴

The definition is well defined as $s_q^i < e_q^i$ for $i \geq 1$.

For hierarchical programs the incremental evaluation also works with conservative approximations of the progress values.²⁷ This is particularly useful for a parallel or even distributed evaluation of a TSA program, as the progress values do not need to be perfectly synchronized. Furthermore conservative approximations for the propagation functions $f_{b,q}$ could be used. The propagation functions $f_{b,q}$ of Definition 25 are optimal with respect to the achievable progress within one step, however approximate, i. e. simpler versions of the functions, may help to reduce the computational overhead that is potentially introduced by the computation of the progress values.

The i th increment expression, i. e. the expression that is actually executed for performing the i th increment computation of a query, consists of the pass condition (and the negated pass condition of the last increment computation) placed on top of the original query expression.²⁸

DEFINITION 27. (Increment Expression)

Let $q = (D, E)$ be a *valid* TSA query with output schema $\mathcal{S}(D) = (A, G_D, H_D)$ and $\Delta H_D = a_1 \leq b_1 \vee \dots \vee a_k \leq b_k$.

1. The *pass condition* for $i \geq 0$ is:
 $\Delta H_D^i = a_1 \leq \text{prop}_q^i(b_1) \vee \dots \vee a_k \leq \text{prop}_q^i(b_k)$
2. The *increment expression* for $i \geq 1$ is:
 $\Delta E^i = \sigma[\Delta H_D^i \wedge \neg \Delta H_D^{i-1}](E)$

The increment expression for a TSA query can be seen as parametrized relational algebra expression as it is only evaluated against finite relations (the currently buffered part of the input streams) and the only changes from one increment computation to the other are the progress values within the pass conditions. This is very useful as it avoids to compile each of the increment expressions at the run-time of a TSA program. Instead the parametrized version of the increment expression can be compiled only once. Furthermore an efficient implementation of relational algebra on finite standing relations can easily be enhanced to an implementation of TSA, i. e. relational algebra on data streams.

PROPOSITION 28. (Correctness) The consecutive execution of the increment expressions $\Delta E_1, \Delta E_2, \Delta E_3, \dots$ for a TSA query $q = (D, E)$ yields the same result as if E had been applied to the whole stream at once.

PROOF (SKETCH). The proof is similar to the proof of Proposition 20. The proof of Proposition 20 just uses the fact that there must exist some “progress values”. Thus the proof of Proposition 20 basically has to show that the “progress values” required by the proof of Proposition 20 are actually just those progress values computed by the propagation functions of Definition 25. This is the case as Definition 25 is well aligned with Definition 7 of progressing sets of attributes and Definition 5 of temporal distance. \square

²⁷ In the case of a recursive program this may not hold for the progress values of queries and temporal stream definitions in a strongly connected component of the program. It still holds for the hierarchical parts of the recursive program.

²⁸ The query expression can be enhanced with relevance filters that accept only those tuples from the input streams that could contribute to new tuples in the result stream. This keeps intermediate results small and avoids recomputations.

Relevance & Garbage Collection

The incremental evaluation of TSA uses the fact that upper bounds to the result stream of a valid TSA query impose upper bounds to required parts of the input streams. For many TSA queries the same also holds for lower bounds: Lower bounds to the result stream impose lower bounds to the required parts of the input streams. This observation can be used to determine the relevance of tuples in the input streams of a query for further results of the query. More precisely, if the result stream of a query progresses beyond a certain point, indicated by the corresponding progress values (Definition 26), then for each input stream there exist so-called *keep values* such that only tuples with timestamps greater than these keep values can contribute to further results. Analog to the propagation functions (Definition 25) for the incremental evaluation, we can derive so-called *keep functions* from the temporal relations specified in the query. The keep functions are able to compute the current keep values for the input streams of a query based on the current progress values of the result stream of a query. The details are out of the scope of this article.

7. RELATED WORK

Windows. Windows appear in the form of tumbling, sliding or landmark windows [4, 9, 1], generalized window operators like in the WID [20], band predicates [21, 11] and validity intervals [19]. As shown in [23, 16] and in this article, new applications require queries that cannot be conveniently expressed using those data independent windows.

Pattern Matching. The (evaluation) semantics of pattern matching approaches like CEDR[6], SASE+[3], AFAs[8] and CAYUGA [17] usually bases on some automaton model. Inherently their semantics strongly relies on some order on the processed tuples or events. The approach is very well suited for applications like in financial markets where queries are formulated against a single timeline, queries with a “triggering” event at the beginning of the queried period²⁹ and extremely low response times are required. However it is unclear how automaton-based approaches could be extended to queries referring to multiple timelines, as in the presence of multiple timelines the order of the events is usually unclear.

Frames & Frame Segments. The motivation of frames [23] is similar to the idea of “event-controlled aggregation”. Generally the borders of a frames are determined by means of changes in the arriving data. In the case of T+D frames the value of some attribute must be above some threshold for all tuples within the borders of a frame. The “Threshold-Frame” and “FillFrame” operators proposed for T+D frames can be expressed by a composition of basic TSA operators. As frames may have arbitrary duration [23] introduces frame fragments for first detecting frames early, i.e. when their existence but not their end is known, and second enabling precomputations of subsequent operators as to reduce memory requirements and avoid workload bursts at the end of a frame. The intentions of frame fragments can partially be realized using the safe form recursion supported by TSA.²³

Predicate Windows. Predicate windows are an approach towards data dependent windows. Predicate windows are

motivated by the perspective that a window defines a set of currently valid tuples where tuples are entering into and expiring from this set based on a predicate assigned to the window. The approach apparently consists of two components: First predicate windows are used to interpret a data stream as a sequence of updates and deletes to a set of currently valid tuples and second a streamed version, i.e. no materialization of result relations, of a view maintenance algorithm propagates changes through subsequent operator.

Punctuations & Out-of-order Processing. Punctuations have been proposed first in [25] as an alternative to windows for unblocking operators and for limiting the state that an operators has to maintain. The WID approach in [20] showed that punctuation can be used to handle disordered streams and introduced operators that hardly need to reorder tuples for processing. The architecture for out-of-order processing (OOP) proposed in [21] bases on the WID approach and uses linear punctuations for unblocking operators and propagating stream progress. [21] shows that out-of-order processing can clearly out-perform in-order processing. The linear punctuations of OOP are almost equivalent to the progress values used in the out-of-order incremental evaluation of TSA expressions. DataCell [22] also performs out-of-order processing however using (primarily tuple-based) windows instead of punctuations.

CERA. The Complex Event Relational Algebra (CERA) [12, 13] serves as operational semantics for the high-level complex event processing language XChange^{EQ}. CERA uses an analysis of temporal relations for garbage collection but not for incremental evaluation like TSA. CERA has some capabilities for the definition of relative timestamps but the timestamps of composite events are mostly predefined. Basic temporal relations can be combined using conjunctions, disjunctions and negation are missing, though. To some extend CERA can express windows with variable size. CERA allows simultaneous events, i.e. does not need a strict order on events, however the evaluation CERA needs to processes the sets of simultaneous events in temporal order. Consequently CERA only support a single timeline.

8. CONCLUSION & FUTURE WORK

We introduced a common data model for data streams and static relations and define Temporal Stream Algebra (TSA) enhancing the standard operators² of relational algebra by an mechanism for propagating constraints on temporal relations between attributes and on the correlation of attributes to the stream progress within the stream schema. We presented an analysis for the propagated constraints that can determine the validity of TSA queries and derives functions that define first when results can be passed on and second how information on stream progress is propagated and third which data needs to be kept. Based on these functions we describe a bulk-wise and out-of-order evaluation for TSA queries. In this way TSA can do without windows with fixed, data independent, sizes.³⁰ Outstanding features of TSA are the support of event-controlled aggregation, user-defined timestamps for composite events and multiple timelines and the smooth integration of static relations.

TSA serves as operational semantics of the event, state and action language Dura [28, 26, 27]. A prototype implementation of TSA and Dura on top of MonetDB has been

²⁹ Queries with negation or aggregation to the past, like “Return event B if there was no event A in the past 5 minutes”, where the “trigger” for the query is at the end of the queried period, are problematic for automaton-based approaches.

³⁰ Nevertheless such windows can be expressed when needed.

completed and used in prototypes of emergency management applications for metro networks, airports and power-grids[31, 27]. The prototype of TSA implements garbage collection based on the analysis of temporal relations. Currently we work on further improvements on the garbage collection and on the optimization of the incremental evaluation using so-called relevance filters.

The prototype currently supports a limited but safe, i.e. non-blocking, form of “recursion to the future”.³¹ We would like to examine other forms of recursion for example how the Flying Fixed-Point operator [23] could be generalized to strongly connected components in a recursive TSA program.

The incremental evaluation of TSA performs a bulk-wise out-of-order processing of events. We expect that this contributes to a high throughput, though it probably does not achieve minimum response times for single events, and helps to cope with peaks in the event load. We currently implement the Linear Road Benchmark [5] as to verify these claims. Moreover incremental evaluation enables an asynchronous query processing. We are very interested whether this allows an easy distributed processing.

The combination of asynchronous and bulk-wise processing opens a number of interesting questions with regards to scheduling. Though any fair execution sequence for the increment expressions of the queries will yield the correct results, scheduling affects the efficiency and the response-time of the evaluation. Determining a good or even optimal sequence is another of our current research issues. The same holds for a situation dependent query prioritization.

9. ACKNOWLEDGMENTS

The work presented in this article is part of the research project EMILI (<http://www.emili-project.eu/>) funded by the European Commission under grant agreement number 242438. Special thanks go to Fabian Groffen and Martin Kersten from the MonetDB team at CWI for their support in the prototype implementation of TSA and to Steffen Hausmann for the excellent cooperation in the design of Dura and TSA and for numerous fruitful discussions.

10. REFERENCES

- [1] D. J. Abadi, D. Carney, U. Çetintemel, M. Cherniack, C. Convey, S. Lee, M. Stonebraker, N. Tatbul, and S. B. Zdonik. Aurora: A new model and architecture for data stream management. VLDB 2003.
- [2] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [3] J. Agrawal, Y. Diao, D. Gyllstrom, and N. Immerman. Efficient pattern matching over event streams. SIGMOD 2008.
- [4] A. Arasu, S. Babu, and J. Widom. The CQL continuous query language: Semantic foundations and query execution. VLDB 2006.
- [5] A. Arasu, M. Cherniack, E. F. Galvez, D. Maier, A. Maskey, E. Ryvkina, M. Stonebraker, and R. Tibbetts. Linear road: A stream data management benchmark. VLDB 2004.
- [6] R. S. Barga and H. Caituiro-Monge. Event correlation and pattern detection in CEDR. In *Proc. Int. Workshop Reactivity on the Web*. Springer, 2006.
- [7] F. Bry and M. Eckert. On static determination of temporal relevance for incremental evaluation of complex event queries. DEBS 2008.
- [8] B. Chandramouli, J. Goldstein, and D. Maier. High-performance dynamic pattern matching over disordered streams. VLDB 2010.
- [9] S. Chandrasekaran, O. Cooper, A. Deshpande, M. J. Franklin, J. M. Hellerstein, W. Hong, S. Krishnamurthy, S. Madden, V. Raman, F. Reiss, and M. A. Shah. Telegraphcq: Continuous dataflow processing for an uncertain world. In *CIDR*, 2003.
- [10] E. F. Codd. Relational completeness of data base sublanguages. In: *Database Systems*, Prentice Hall, 1972.
- [11] C. Cranor, T. Johnson, and O. Spatschek. Gigascope: a stream database for network applications. SIGMOD 2003.
- [12] M. Eckert. *Complex Event Processing with XChange^{EQ}: Language Design, Formal Semantics and Incremental Evaluation for Querying Events*. PhD thesis, Institute for Informatics, University of Munich, 2008.
- [13] M. Eckert, F. Bry, S. Brodt, O. Poppe, and S. Hausmann. Two Semantics for CEP, no Double Talk: Complex Event Relational Algebra (CERA) and its Application to XChange^{EQ}. In *Reasoning in Event-based Distributed Systems*. Springer, 2011.
- [14] R. W. Floyd. Algorithm 97: Shortest path. *Commun. ACM*, 5(6):345–, June 1962.
- [15] H. Garcia-Molina, J. Ullman, and J. Widom. *Database Systems: The Complete Book*. Prentice Hall, 2001.
- [16] T. M. Ghanem, W. G. Aref, and A. K. Elmagarmid. Exploiting predicate-window semantics over data streams. SIGMOD 2006.
- [17] M. Hong. *Expressive and Scaleable Event Stream Processing*. PhD thesis, Cornell University, 2009.
- [18] M. K. Kostas Stergiou. Backtracking algorithms for disjunctions of temporal constraints. *Artif. Intell.*, 2000.
- [19] J. Krämer and B. Seeger. Semantics and implementation of continuous sliding window queries over data streams. *ACM Trans. Database Syst.*, 34(1), April 2009.
- [20] J. Li, D. Maier, K. Tufte, V. Papadimos, and P. A. Tucker. Semantics and evaluation techniques for window aggregates in data streams. SIGMOD 2005.
- [21] J. Li, K. Tufte, V. Shkapenyuk, V. Papadimos, T. Johnson, D. Maier. Out-of-order processing: a new architecture for high-performance stream systems. VLDB 2008.
- [22] E. Liarou, R. Goncalves, and S. Idreos. Exploiting the power of relational databases for efficient stream processing. EDBT 2009.
- [23] D. Maier, M. Grossniklaus, S. Moorthy, and K. Tufte. Capturing episodes: may the frame be with you. DEBS 2012.
- [24] J. P. Rina Dechter, Itay Meiri. Temporal constraint networks. *Artif. Intell.*, 49(1-3):61–95, May 1991.
- [25] P. A. Tucker, D. Maier, T. Sheard, and L. Fegaras. Exploiting punctuation semantics in continuous data streams. *IEEE Trans. on Knowl. and Data Eng.*, 2003.
- [26] S. Brodt, S. Hausmann, and F. Bry. D4.6: Modularization Mechanisms. www.emili-project.eu, 2012.
- [27] S. Brodt, S. Hausmann, and F. Bry. D4.7: Refinement of the implementation of event processing and ECA rules for SITE, 2012.
- [28] S. Hausmann, S. Brodt, and F. Bry. D4.3: Dura - concepts and examples. <http://www.emili-project.eu/>, 2011.
- [29] N. Seifert and M. Bettelini. D3.1: Use Cases Requirements Analysis and Specification (Main Report). www.emili-project.eu, 2010.
- [30] N. Seifert, M. Bettelini, and S. Rigert. D3.2: Concrete Use Case Models. www.emili-project.eu, 2011.
- [31] N. Seifert, et al. D3.3: Use case modelling for implementation in SITE. www.emili-project.eu, 2012.

³¹ This means a current event of some event type may contribute to future events of the same event type, but neither to present or past events.