

Point-Based Temporal Extensions of SQL and Their Efficient Implementation

David Toman*

Department of Computer Science, University of Toronto
Toronto, Ontario M5S 1A4, Canada
`david@cs.toronto.edu`

Abstract. This chapter introduces a new approach to temporal extensions of SQL. In contrast with most of the current proposals we use *single time points* as references to time, while still achieving efficient query evaluation. The proposed language, SQL/TP, naturally extends the syntax of SQL/92 : it adds a single new data type that represents a linearly ordered universe of time instants. The semantics of the new language extends the standard SQL in the expected way: the new data type behaves identically to the existing types. We also eliminate or fix many of the problems connected with defining a precise semantics in interval-based languages. In addition we provide an efficient query evaluation procedure based on a compilation technique that translates SQL/TP queries to SQL/92. Therefore existing off-shelf database systems can be used as back-ends for managing temporal data.

1 Why yet another temporal extension of SQL?

After more than a decade of research in the area temporal databases there is still no universal consensus on how temporal features should be added to the standard relational model. Instead, there are dozens of mutually incompatible models and proposals. The more practical of these are often based on (often ad-hoc) extensions of existing relational languages, e.g., TQUEL [17] and various temporal extensions of SQL: TSQL2 [18], ATSQL2 [5], and SQL/Temporal [19], the current proposal of temporal extension of SQL3 to the ISO and ANSI standardization committees.

The goal of this chapter is twofold: First, we point out severe problems common to the majority of current proposals, namely representation dependent handling of temporal values and limited temporal dimensionality of the underlying temporal models. We argue that these problems are inherently tied to the use of interval-valued temporal attributes and cause major problems when a precise semantics of the temporal query languages is to be defined. Second, we propose

* The research was supported by a NATO/NSERC PDF fellowship, 1996-98. Extended abstract of this paper appeared in Proc. DOOD'97, Montreux, Switzerland, 1997.

an alternative solution that avoids these problems. Our proposal is based on a clean separation of the syntax and semantics (what is stored in the database and how do we query it) from the underlying compact representation of the temporal information (how is it stored in the database). While our technique can be applied to a wide range of relational query languages, we have chosen standard SQL with its duplicate semantics and aggregation operators as the starting point of our proposal. This choice demonstrates that we are indeed interested in real-life query languages rather than in toy examples.

The rest of the chapter is organized as follows: the remainder of this section exposes problems inherent to the current proposals. Section 2 formally introduces the temporal data model: abstract and concrete (interval-based) temporal databases, following the terminology introduced in [9]. Section 3 defines the syntax and semantics of SQL/TP and gives examples of temporal queries. We also include a brief discussion of compatibility issues and migration of SQL queries to SQL/TP (Section 3.5). Section 4 provides foundations for the proposed compilation technique. The chapter is concluded with several open questions and directions of future research. Appendix A summarizes the BNF of the core SQL/TP language and Appendix B briefly compares SQL/TP with the SQL/Temporal proposal to the ISO/ANSI SQL3 committee.


1.1 Current Proposals

Most of the current proposals recognize that storing temporal data as ordinary tuples leads directly to enormous space requirements: a tuple has to be repeated for every time instant at which the represented fact holds in the modeled reality. Instead, tuples are associated with compact encodings of a set of time instants (often called period of validity). The sets of time instants are commonly represented by (finite sets of) intervals [17, 19], bitemporal elements [4, 14], or other fixed-dimensional products of intervals (hyper-rectangles). The syntax of the chosen encoding then provides a domain of values for temporal attributes, e.g., pairs of interval endpoints. Indeed, ATSQL [5] and SQL/Temporal [19] use the **BEGIN** and **END** keywords to extract the endpoints of intervals, **PERIOD** keyword to construct new interval timestamps, and Allen's interval algebra [2] operators to compare the timestamps. However, this is only possible if the user knows that the timestamps are encoded using intervals. Moreover, this approach leads to a tension between the syntax of the query languages and their intended semantics: the data model and the semantics of the languages are point-based¹ [4, 9], while temporal attributes refer to the actual encoding for sets of time instants (e.g., interval endpoints). This conflict leads to several unpleasant surprises when precise semantics needs to be defined. Most importantly, it is easy to show examples of queries whose answers depend on the choice of the particular encoding rather than on the underlying meaning of the data; cf. Example 1 below. Moreover, it is extremely hard to avoid this behavior in an elegant way. In many cases

¹ The truth is associated with individual time instants rather than with sets of instants (intervals).

uniqueness of answers can only be guaranteed by operational means, e.g., by prescribing a particular evaluation order. This in turn leads to a very complicated and cumbersome semantics (if one exists at all). The problems become even more apparent and critical when a temporal extension for query languages with duplicate semantics has to be defined. Consider the following situation:

Example 1 Let D be a temporal relation (or an answer to a temporal query) that represents the region in the figures below.



It is important to understand that all the three figures represent the same relation. However, it is also clear that we can distinguish (2) and (3) using a first-order query in, e.g., SQL/Temporal. We call such queries *representation dependent*. These queries cannot be given meaningful semantics in the point-based temporal models. Moreover, even very simple queries, e.g., counting the number of regions along the axes, give different results depending on the particular representation.

The common temporal query languages, hoping to avoid representation dependencies (among other reasons), try to mimic a point-based semantics using set operations on the encoded timestamps and a *normalization* procedure on temporal relations. In a single-dimensional case, the representation dependency problem can indeed be successfully evaded using *coalescing* [6]. TSQL2's informal semantics (including many examples of queries in TSQL2) is implicitly based on this assumption [3, 18].

In the rest of this section we argue that the situation in Example 1 naturally arises during temporal query evaluation and cannot be avoided in general.

First, we show that a single temporal dimension is not sufficient to formulate general temporal queries. Consider the query “are there two distinct time instants when a given relation contains exactly the same tuples?” [1] and [21] have independently shown that this query cannot be formulated in first-order temporal logic. A direct corollary of this result is that this query cannot be expressed in any single-dimensional temporal relational algebra². Moreover, [21] shows that to express all first-order queries the number of temporal dimensions cannot be bounded by any constant. Therefore, unbounded number of temporal dimensions cannot be avoided during query evaluation even if the final result is a single-dimensional temporal relation or boolean. This fact, combined with the use of explicit interval-valued temporal attributes, leads directly to situations similar to Example 1.

There are other reasons for including multiple temporal dimensions in a temporal database system, e.g., the need for representing both valid and transaction

² A relational algebra over the universe of single-dimensional temporal relations.

time [14]. However, we would like to emphasize that the need for unbounded number of temporal dimensions originates from the inherent properties of first-order queries alone, even if the temporal database and the results of the queries are single-dimensional.

Second, there is no unique normal form based on coalescing in the case the number of temporal attributes (the temporal dimension) is greater than one.

Now it is easy to see why the coalescing-based approaches fail to guarantee representation independence: To guarantee fixed size of tuples in temporal relations, region (1) in Example 1 has to be represented by a finite union of rectangles, e.g., using the representation (2) or (3) above. While both (2) and (3) are coalesced, they can still be distinguished by a first-order query with interval-valued temporal attributes.

In addition, in many cases the user has no control over the representation of intermediate results since the coalescing is performed implicitly by the system. With two or more temporal dimensions in queries coalescing leads to serious problems: the user has no knowledge if region (1) in Example 1 is actually represented as (2) or (3), but the results of queries often depend on this information.

1.2 The Point-based Proposal: SQL/TP

The above problems are inherent all temporal query languages with temporal attributes ranging over intervals. Therefore our proposal follows a different path to avoid all of the above problems: we let temporal attributes in our language range over the domain of single time instants.

Our approach is based on several recent results in the area of temporal and constraint query languages [1, 15, 20, 21]. In addition we define a meaningful approach to duplicate semantics and aggregation, independent of the chosen encoding (extending [10, 11]).

In addition to simple and elegant syntax and semantics we propose an *efficient* query evaluation procedure for SQL/TP over compactly encoded temporal databases. While we mostly concentrate on efficient evaluation of temporal queries over an interval-based encoding of time, the design of SQL/TP allows the use of additional encodings for sets of time instants, e.g., the *linear repeating points* [23] for periodic events, without the need for new syntax and semantics. There are several other features of the proposal:

- SQL/TP statements can be compiled to standard SQL/92³ [12]; the translated queries can be evaluated using an off-shelf database system. This way we can build a SQL/TP front-end to an existing RDBMS and provide temporal capabilities without modifying the underlying database system itself.
- SQL/TP can express all representation independent SQL/Temporal queries. Moreover, SQL/TP is first-order complete (in the sense of [8]). The results in [1, 21] show that this is not the case for any of the temporal query languages

³ Other relational languages can be used as well, provided they have sufficient expressive power.

based on a fixed-dimensional temporal relational algebra, e.g., [7]; this issue is not clear for TSQL2-derived languages [5, 18, 19] due to the presence of explicit coercion operators that convert encoded temporal attributes to data attributes.

- The proposal can easily be extended to support migration requirements [19] for upward temporal compatibility with SQL. While SQL/TP itself does not literally follow these requirements, the compatibility can be easily achieved using a very simple syntactic manipulation of the source queries and adding tags to distinguish the particular compatibility modes, cf. Section 3.5.

Before we start the technical part of the chapter, we would like to reiterate (to avoid any misunderstanding) that we are interested in intervals to store sets on time instants. This is the main difference between our approach and approaches taken by various interval logics [2], where intervals represent points in a two-dimensional (half-)space. However, due to the natural multidimensional character of SQL/TP, we can represent the *true* intervals using pairs temporal attributes.

1.3 Summary of Contributions

The three main technical contributions of our proposal are: (1) the definition of a *representation-independent* temporal extension of SQL: we decouple the syntax and semantics of the language from the underlying data representation. We support both the set- and the duplicate-based semantics of SQL (including aggregation), (2) a *query compilation technique* for this extension that allows SQL/TP queries to be efficiently evaluated using a *standard* RDBMS, and (3) the definition of nouveau *normalization* technique that facilitates evaluation of temporal queries over an interval-based encoding of timestamps. We would also like to note that a naive direct translation of time instants to singleton intervals [16] *fails* as an efficient query evaluation technique: it causes an exponential blowup in complexity.

2 The Data Model for Temporal Databases

We start with the definition of the underlying data model: the domain of time is viewed as a discrete⁴ countably infinite linearly ordered set without endpoints (e.g., the integers). The individual elements of the set represent the actual time instants while the linear order represents the progression of time. The actual granularity of time is not important in our proposal⁵. Besides the data type for time instants we also use all the other data types defined in standard SQL: strings, integers, floats, etc. As usual, these data types do not have an a-priori assigned interpretation. We summarily refer to those data types as *the uninterpreted constants*.

⁴ A dense linearly ordered time can be used with only a minor adjustment.

⁵ For our purposes any fixed granularity will do.

The relationships between time instants and uninterpreted constants are captured in a finite set of *temporal relations* stored in the database. Following the terminology of [9] we distinguish the *abstract* temporal databases from the *concrete* temporal databases:

Definition 2 (Abstract Temporal Database) *The signature of a predicate symbol R is the tuple $(a_1 : t_1, \dots, a_k : t_k)$ where a_i are distinct attribute names, t_i the corresponding attribute types, and k the arity of R . Attributes of type **time** are the temporal attributes, the remaining attributes are the data attributes. A database schema is a finite set of relational symbols $R_1, \dots, R_k$ paired with their signatures. An abstract temporal database is a set of tables defined by a database schema.*

In general we do not restrict the cardinality of abstract temporal tables: we allow infinite tables in general. However, in order to define meaningful operations on the tables, we require that the number of occurrences (duplicates) is finite for every distinct tuple.

Example 3 In the rest of the chapter we use an abstract temporal database with the schema $\{\text{indep}(\text{Name}, \text{Year})\}$ as a running example. The particular instance of the *indep* relation we use in our examples captures independence of countries in Central Europe:

indep		...		...	
Name	Year	Czech Kingdom	1198	Czech Republic	1995
Poland	1025	...	...	...	...
...	...	Czech Kingdom	1620	...	...
Poland	1794	Czechoslovakia	1918	Slovakia	1940
Poland	1918	...	...	...	...
...	...	Czechoslovakia	1938	Slovakia	1944
Poland	1938	Czechoslovakia	1945	Slovakia	1993
Poland	1945	...	...	...	...
...	...	Czechoslovakia	1992	...	...

We do not impose any restrictions on the number of temporal attributes in relations (unlike, e.g., TSQL2 [18]). Indeed, in general we may want to record relationships between different time instants as well. While we may want to restrict the users from creating such multi-dimensional tables, we need this feature in the later sections to translate SQL/TP queries to SQL/92.

The abstract temporal databases provide a natural data model for representation and querying temporal data. However, it would be impractical (and often impossible) to store the temporal databases as plain bags of their tuples: a particular tuple is often related to a large and possibly infinite set of time instants. Rather than storing all these tuples one by one, we use a *compact encoding* of sets of time instants. The choice of a particular encoding—in our case the *interval*-based encoding—defines the class of *concrete temporal databases*:

Definition 4 (Concrete Temporal Database) Let R be a relational symbol with signature E . A concrete signature corresponding to E is defined as a tuple of the attributes that contains (1) a for every data attribute $a \in E$ and (2) $t_{\min}$ and $t_{\max}$ for every temporal attribute $t \in E$. The attributes $t_{\min}$ and $t_{\max}$ denote endpoints of intervals. We denote the concrete signature of R by $\bar{E}$. A concrete temporal database schema corresponding to a given abstract database schema is a set of relation symbols and their concrete signatures derived from the signatures in the abstract database schema⁶. A concrete temporal database is a set of finite relations defined by a concrete database schema.

To capture the relationship between an abstract and a concrete temporal database we define a *semantic mapping* that maps a concrete temporal database to its meaning—an abstract temporal database. The meaning of a single concrete tuple $\mathbf{x} = (t_{\min}, t_{\max}, a_1, \dots, a_k)$ is a set of tuples $\|\mathbf{x}\| = \{(t, a_1, \dots, a_k) : t_{\min} \leq t \leq t_{\max}\}$; analogously for tuples with multiple temporal attributes. The meaning $\|\mathbf{R}\|$ of a concrete relation $\mathbf{R}$ is the duplicate preserving union of $\|\mathbf{x}\|$ for all concrete $\mathbf{x} \in \mathbf{R}$. We say that $\mathbf{R}$ *encodes* $\|\mathbf{R}\|$. We extend the $\|\cdot\|$ to concrete temporal databases in the standard fashion.

The *encodes* function also defines a subset of the abstract temporal databases that can be encoded using concrete temporal databases. We call this subset the *finitary temporal databases*. Note that the encoding is not unique and thus two or more distinct concrete temporal databases can encode the same abstract temporal database (cf. Example 1). We call such pairs of concrete temporal databases $(\|\cdot\|, \|\cdot\|)$ -equivalent.

Example 5 The database instance from Example 3 is infinite. However, it is finitary and can be encoded using the following concrete temporal database:

indep			...	
Name	Yearmin... Yearmax			
Czech Kingdom	1198	... 1620	Slovakia	1940... 1944
Czechoslovakia	1918	... 1938	Slovakia	1993... ∞
Czechoslovakia	1945	... 1992	Poland	1025... 1794
Czech Republic	1993	... ∞	Poland	1918... 1938
			Poland	1945... ∞

All queries in the rest of the chapter are evaluated over this database while preserving answers with respect to the original relation in Example 3.

3 The Language SQL/TP

In this section we define the syntax and semantics of SQL/TP. This includes the data definition, data query, and data manipulation parts of the language. In all three cases we show that SQL/TP is a natural syntactic extension of SQL over

⁶ We use the same names for both the abstract and concrete relations. The actual meaning of the symbol is always clear from the context.

the abstract temporal databases. Moreover, the proposed semantics of SQL/TP is essentially identical to the semantics of SQL (safely) extended to abstract temporal relations.

3.1 Data Definition Language

We start with the *Data Definition Language*: it is essentially identical to standard SQL/92:

```
create table <rid> ( <signat> )
create view  <rid> ( <query> )
```

where <rid> is a table identifier and <signat> is a signature of the new table. For views the signature is derived from the signature of the <query> expression (cf. Section 3.2). The only difference is that the temporal attributes are declared using a new data type *time* that supports *modifiers* that determine how the sets of time instants are stored in a concrete temporal table:

using points: The time instants are stored as atomic values similarly to all other data types. This choice is suitable for representing single atomic events that happen at a specific time, e.g., when a particular tuple was inserted in the database.

using [bounded | unbounded] intervals: Continuous sets of time instants associated with a particular data tuple are encoded using intervals. This encoding is suitable for representing durations of time. The **bounded** and **unbounded** keywords specify if the $-\infty$ and ∞ may be used as endpoints of intervals. This choice affects, what aggregate operations are allowed for that particular attribute; cf. Section 3.2.

It is important to understand that these modifiers affect only the way the table is stored, not the semantics of the queries (similarly to specifying, e.g., a sort order or a key for the table). In the future this list may grow to accommodate different ways of encoding sets of time instants. The modifiers are the only place in SQL/TP where the syntax reflects the chosen encoding. The default modifier **unbounded time** is assumed for all temporal attributes unless explicitly stated otherwise.

The following table defines how the modifiers interact with standard relational operators (“p” is a shorthand for **points**, “b” for **bounded intervals**, and “u” for **unbounded intervals**):

op	p op p	p op b	p op u	b op b	b op u	u op u	p op p	p op b	p op u	b op b	b op u	u op u
$\cap$	p	p	p	p	b	b	p	b	u			
—	p	p	p	b	b	b	u	u	u			
$\cup$	p	b	u	b	b	u	u	u	u			

The interaction of temporal attributes in joins and selections is captured by their behavior in the $\cap$ (intersection) operation, projection does not affect the remainders of the tuples.

Example 6 The table `indep` in Example 3 is created as follows:

```
create table indep (name char(20),
                   year time using unbounded intervals)
```

In the rest of the chapter we discuss only the interval-based encoding; encoding time instants using points does not introduce any problems over the traditional data types. In addition we assume the time instants can be represented by integers (using a fixed granularity) and we allow integer-like operations on the new data type to reduce the amount of superfluous syntax.

3.2 The Query Language

For sake of simplicity we discuss only a syntactic subset of full SQL/TP. This choice does not affect the generality of our proposal: it is an easy exercise to show that the proposed fragment forms a (first-order) complete query language [8]. Moreover, all representation independent SQL/Temporal queries, including queries with aggregation and universal subqueries, can be equivalently formulated in this fragment.

Syntax. The syntactic subset of SQL/TP uses two basic syntactic constructs:

Select block. Similarly to the standard SQL the select block is the main building block of the query language. It has the usual form

```
select <slist> [from <flist>] [where <cond>] [group by <glist>]
```

where

<slist> is a list of attribute identifiers, constants, and (aggregate) expressions with the possibility of renaming the output column using <sexp> as <id>. Columns defined by expressions or aggregation have to be given a name in this way,

<flist> is a sequence of relation identifiers or subqueries, again with the usual possibility of assigning correlation names,

<cond> is a selection condition built from atomic conditions using boolean connectives. The atomic conditions depend on the data types of the involved attributes: in the case of temporal attributes we allow conditions of the form <id> *op* <id> + *C* where *op* $\in \{<, \leq, =, \geq, >\}$, and *C* a constant denoting a *length* of a time period, and

<glist> is a list of attribute identifiers that specifies how the result of the select block is grouped. The usual SQL rules that govern the grouping operations apply here as well.

We extend the definition of signature to SQL/TP expressions: The *signature* of an expression is tuple of attribute names in the result paired with the corresponding data types (including modifiers for the temporal types).

Set Operations. Besides nesting queries in the **from** clause of the select block we can combine the individual select blocks using *set operations* as follows:

(<exp>) <setop> (<exp>)

where <setop> is one of the **union** (set union with duplicate elimination), **union all** (additive union), **except** (set difference with duplicate elimination), **except all** (monus), **intersect** (set intersection with duplicate elimination), and **intersect all** (duplicate preserving intersection). We require the signatures of both the expressions to match⁷. The resulting signature is the common signature of the expressions involved in the operation.

The proposed syntax omits two common SQL constructs: subqueries nested in the **where** clause and the **having** clause. Both these constructs can be expressed in the presented fragment using nesting in the **from** clause of the select block and therefore can be considered to be only syntactic sugar.

To achieve signature compatibility for temporal attributes we allow the use of a special constant pseudo-relation **true**(**t**: **time**) true for all elements of the temporal domain. This relation allows us to pad the attribute lists involved in the set operations (cf. Section 3.3) and to formulate queries that involve, e.g., the complementation over the temporal domain.

Semantics. SQL/TP is essentially SQL/92 extended with an additional data type **time**. The main feature of this proposal is that we can use the standard SQL-like semantics over the class of the *abstract temporal databases*. This way we completely avoid all problems connected with representation dependencies while maintaining compatibility with SQL. Moreover, changes in the encoding (the physical representation) do not affect the syntax and semantics of queries.

However, we have to be careful when extending relational operations to infinite tables: we have to ensure that we never produce tables with infinite duplicates of a single tuple. It is easy to see that all the relational operations, with the exception of duplicate preserving projection, meet this requirement. However, the duplicate-preserving projection can produce such tables, e.g.:

$$\{(\text{"Poland"}, [1945, \infty))\} \xrightarrow{\parallel \cdot \parallel} \{(\text{"Poland"}, n) : n \geq 1945\} \\ \xrightarrow{\pi_1} \{(\text{"Poland"}), \dots, (\text{"Poland"}), \dots\}$$

The result of the projection contains infinite duplication of the tuple ("Poland"). This cannot be allowed as other relational operators, e.g., the bag difference, are not well defined over such tables. Therefore we restrict the use of duplicate-preserving projection to attributes of bounded types, i.e., bounded intervals, time points, and data types.

⁷ SQL only requires the types to match. However, we require both the names of the attributes and their types to match. This is not a restriction as the renaming can be conveniently done within the **select** clauses.

Closure over Interval-based Concrete Databases. While the above restriction guarantees a well defined semantics, it is too weak to guarantee closure of SQL/TP queries over the chosen class of concrete temporal databases. The main source of problems are the order dependencies among temporal attributes. Consider the following example:

Example 7 It is easy to find SQL/TP expressions that *do not* preserve closure over the class of finitary abstract temporal databases. Consider the expression:

```
Q: select  r1.name as name, r1.year as t1, r2.year as t2
      from    indep r1, indep r2
      where   r1.name = r2.name and r1.year < r2.year
```

The attributes t_1 and t_2 are *correlated* by an inequality $t_1 < t_2$ in the result of the query:

$$\{ ("Poland", 1945, 1946), ("Poland", 1945, 1947), \dots, ("Poland", 1945, 1950), \dots \\ ("Poland", 1946, 1947), \dots, ("Poland", 1946, 1950), \dots \\ \vdots \\ ("Poland", 1949, 1950), \dots \}$$

Obviously the triangle-like result can not be described by a product of intervals.

To avoid this problem we use the notion of *attribute independence*. Rather than a semantic definition of attribute independence [11] we use a syntactic inference system to guarantee attribute independence in a SQL/TP expression:

Definition 8 (Attribute Independence) Let t_1 and t_2 be two temporal attributes in the signature of a SQL/TP expression *exp*. We say that t_1 and t_2 are independent in *exp* if

1. *exp* is a base relation,
2. *exp* is a select block, t_1 and t_2 are names of t'_1 and t'_2 assigned in the select clause, t'_1 and t'_2 are independent in all expressions in the **from** list, and an order relationship between t'_1 and t'_2 is not implied by the **where** clause.
3. *exp* is (e_1) setop (e_2) and t_1 and t_2 are independent in both e_1 and e_2 .

In addition all the data attributes (and point temporal attributes) are mutually independent.

The inference system is sufficient to infer *independence* of attributes. We could also analyze the compositions of the selection conditions on temporal attributes and check for tautologies. We have chosen not to pursue this direction in the current version of the proposal for sake of simplicity. However, the theory of linear order is decidable and thus such an extension is feasible; note that these tests are performed at *compile time* and thus do not affect the data complexity of the queries.

For similar reasons we restrict the use of aggregate operations: we require the aggregated attribute to be independent of the **group by** attributes [10]. Moreover, the aggregation has to obey the restrictions in Figure 1.

We also restrict the use of duplicate-preserving projection on all temporal attributes encoded by intervals. We have already seen that duplicate-preserving projection is not possible for unbounded data types. On the other hand, for bounded data types we could implement the duplicate preserving projection by creating the appropriate number of copies of the remainder of a tuple. However, such an operation would make the query evaluation very inefficient and almost certainly unusable in practice. Consider the following example:

Example 9 Let $R(x, t) = \{(a, [0, 2^n - 1])\}$ be a concrete temporal relation where n is a large integer. Clearly the size of R (in bits) is $|a| + n$. However, the size of $\pi_x(R)$ is $2^n \cdot |a|$ as the result of duplicate preserving projection has to contain 2^n tuples (a) .

Allowing such projections would cause an exponential blowup in the (space) complexity of query evaluation. Note that the duplicate preserving projection is used in SQL for two main reasons: (1) to avoid duplicate elimination or (2) to facilitate correct aggregates. The first use does not apply to SQL/TP—we deal with redundant duplicate elimination in the optimization phase of our compilation procedure. The aggregates are handled using a rewriting technique that allows us to avoid the duplicate-preserving projections. We can evaluate a vast majority of representation-independent aggregate queries even under the above restriction: note that all other relational operations preserve duplicates (cf. Section 3.3). Therefore we exclude the duplicate-preserving projections of all temporal attributes encoded by intervals in order to maintain the polynomial complexity bound.

We define the SQL/TP queries to be the subset of SQL/TP expressions obeying the following rules:

Definition 10 Let Q be a SQL/TP expression that obeys the following rules:

1. temporal attributes encoded by intervals cannot be projected out in a **select all** clause.
2. all attributes in the (top-level) signature of the expression are pairwise independent.
3. the attributes in the **group by** clause are independent of the remaining attributes, and
4. the attributes used in aggregation operators follow the rules in Figure 1

Then we say that Q is a SQL/TP query over the class of concrete interval-based temporal databases.

It is easy to verify that all SQL/TP queries preserve closure over the class of finitary temporal databases:

Theorem 11 Let D be a finitary database and Q a SQL/TP query. Then $Q(D)$ is finitary.

The requirement of attribute independence seems like a rather severe restriction. However, the independence is required only for the temporal attributes present in

Type of the group by attr's	Type of the aggregated attribute		
	data	bounded intervals	unbounded intervals
data	min,max,sum,count	min,max,count	min,max
bounded int	min,max,sum,count	min,max,count	min,max
unbounded int	min,max	min,max	min,max

When the **group by** clause contains multiple attributes we take the intersection of the allowed aggregate operations. We have excluded the **sum** aggregate on the temporal attributes: while it is definable in our framework, it makes little sense from the semantics point of view.

Fig. 1. Allowed aggregates.

the signature of the top-level query, not for all temporal attributes that appear in the query. Therefore all the representation-independent TSQL2 queries, and all first order queries with a single temporal attribute in their signature in general, can be expressed as SQL/TP queries.

Theorem 12 *The first-order fragment of SQL/TP is expressively equivalent to range restricted two-sorted first order logic (temporal relational calculus).*

We can also express queries shown not to be expressible in TRA [7], e.g., the query “is there a pair of distinct time instants, when exactly the same countries were independent?” [1, 21]. This is not possible in any temporal query language that assumes a fixed number of temporal dimensions in its data model.

3.3 Examples of Queries

In this section we provide illustrative examples of SQL/TP queries. The examples are chosen to highlight the ease of formulating queries in SQL/TP. In addition some of the examples, e.g., example 3, can not be easily (and correctly) be formulated in TSQL2 derivatives.

1. The first example is a simple PSJ query “List all countries that were independent while Czech Kingdom was independent”.

```
select  r1.name
from    indep r1, indep r2
where   r2.name = 'Czech Kingdom' and r1.year = r2.year
```

Note also that the result is a standard non-temporal relation. Over the database from Example 3 we get:

```
name
-----
Czech Kingdom
Poland
```

2. Formulating more complicated queries in SQL/TP, e.g., the query “List all years when no country was independent”, is also very natural:

```
(select t as year from true) except (select year from indep)
```

Note the use of the `true` pseudo-relation to achieve signature compatibility. The result of the query is

yearmin	yearmax
-infinity	1024
1795	1917
1939	1939

While the output—a concrete table containing all the periods when no country was independent—has two columns, it is essential to understand that it is only a compact representation of an abstract table with a single column `Year`.

3. In addition to first-order queries, the aggregate operations in SQL/TP also naturally interact with the rest of the language, e.g., in the query “List all countries that became independent before Slovakia”:

```
select name
from   indep, ( select min(year) as y0
                from indep where name = 'Slovakia' )
where  year < y0
```

The result is:

name
Czech Kingdom
Czechoslovakia
Poland

4. SQL/TP also supports a natural way of aggregating over the temporal attributes: “For every country (that has been independent during the 20th century) list the number of years of independence within the 20th century”:

```
select name, count(year) as years
from   indep
where  1900 <= year < 2000
group by name
```

The aggregation is made possible by the `where` clause: it restricts the otherwise unbounded attribute `year`. The result is:

name	years
Czechoslovakia	67
Czech Republic	7
Poland	75
Slovakia	11

Note that in query languages with interval-valued temporal attributes we would have to use a special syntactic construction to *measure* the size of the intervals explicitly.

5. Moreover, SQL/TP supports *grouping* by temporal attributes: “For every year list the number of independent countries (if any)”:

```
select year, count(name) as numofc from indep group by year
```

The result is:

yearmin	yearmax	numofc
1025	1197	1
1198	1620	2
1621	1794	1
1918	1938	2
1940	1944	1
1945	1992	2
1993	infinity	3

This query is quite hard to ask in temporal query languages that use coalescing implicitly: the input table is coalesced, and re-coalescing after the *name* column is eliminated loses the duplicate information we want to compute.

6. Another query hard to express in the current proposals is “List all pairs of years, when exactly the same countries were independent”:

```
select  y1, y2 from (
  ( select r1.t as y1, r2.t as y2 from true r1, true r2 )
except
  ( select y1, y2 from
    ( ( ( select name, year as y1, t as y2 from indep, true )
      except
        ( select name, t as y1, year as y2 from indep, true ) )
    union
      ( ( select name, t as y1, year as y2 from indep, true )
        except
          ( select name, year as y1, t as y2 from indep, true ) )
    ) ) )
```

The answer to this query is (the output shows only part of the answer):

y1min	y1max	y2min	y2max
-infinity	1024	-infinity	1024
-infinity	1024	1795	1917
-infinity	1024	1938	1938
...	...	...	...
1918	1938	1918	1938
1918	1938	1945	1992
...	...	...	...

Again the output relation is just a representation of a binary relation with two columns *y1* and *y2*. The results in [1,21] show that the two dimensionality is

inherent to this query (this remains true even when we consider the existential closure of this query).

3.4 Database Updates

Besides considering the query language, in a truly practical approach we also need to address updates of temporal relations. We propose two constructs:

```
insert [all] into R ( <query> )
delete [all] from R ( <query> )
```

The updates preserve semantics with respect to the abstract temporal databases while manipulating only the concrete representation in a similar way queries do. The `delete` construction is more powerful than the SQL/92 version (as it handles duplication correctly).

3.5 Compatibility with SQL/92

To allow easy migration from SQL/92 to SQL/TP we follow [19] and introduce two *compatibility modes* that allow standard SQL queries to be evaluated over a temporal database:

Temporal Upward Compatibility (TUC). The first level of compatibility treats standard SQL queries as queries operating with respect to the current time instant. This goal can be easily achieved by a simple query transformation that replaces every reference to a base table R by the expression

$$Q[R_i \mapsto \pi_V \sigma_{t_i=N}(R_i)]$$

where V is the list of R_i 's data attributes, t_i is R_i 's temporal attribute, and N is a constant representing the current time instant (e.g., a `CURRENT DATE` in DB2 [13]). This transformation removes all temporal attributes and thus the original SQL query is evaluated on the current snapshot of the database.

Sequenced Queries. The second level evaluates the standard SQL queries over the temporal database with respect to every time instant returning a temporal relation as a result⁸. The application of a sequenced query is defined by a transformation

$$\{Q[R_i \mapsto \pi_V \sigma_{t_i=t}(R_i)] \times \{t\}\}.$$

The resulting query performs exactly the required operation.

⁸ The query evaluation algorithm defined in Section 4 executes such queries efficiently with respect to the interval encoding, rather than by separate evaluation for every time instant.

We do not specify explicit syntax for the compatibility modes: all queries that come from legacy applications use one of the compatibility modes (usually the TUC mode) by default, and thus we assume that they are preprocessed using the above definitions *before* they are submitted to the SQL/TP query processor.

For all practical purposes, such an arrangement is sufficient to provide the upward compatibility for SQL. However, it is easy to see that we could add tags to queries that distinguish the compatibility modes on the level of syntax (similarly to [19]). The tagging mechanism would guarantee a *syntactic* temporal upward compatibility, should it become necessary.

In all the cases the SQL/TP queries generated by the compatibility modes satisfy assumptions of Theorem 20 and thus can be efficiently translated to SQL/92 by our algorithm.

4 Evaluation of SQL/TP Statements

Starting with this section we focus on the second and third results of the chapter: the compilation technique that translates point-based SQL/TP queries to equivalent SQL queries over interval-based concrete temporal databases. The subtle point here is that the resulting queries are *efficient*. This is not completely trivial as the semantics of the original queries is defined with respect to *abstract temporal databases*. A naive query evaluation procedure [16] would indeed refer to all points in the active domain of the corresponding *abstract temporal database*—an immediate exponential blowup in the data complexity of the query evaluation.

While most approaches to query evaluation in temporal databases take the path of adding specialized temporal operations to a standard relational system, we take an alternative approach: we define a *translation procedure* that allows us to *compile* SQL/TP statements to standard SQL/92 statements. This way:

- we guarantee that the query evaluation uses only the active domain of the concrete temporal database (the result is an ordinary SQL/92 query),
- we can implement a SQL/TP interface as a preprocessor on top of an ordinary relational system, and
- we isolate a single intrinsic operation needed for temporal query evaluation over the interval-based encoding.

While the result of the translation can be defined using a SQL/92 query, we may decide to implement parts of it, e.g., the normalization operation; cf. Definition 15, natively to improve the efficiency of query evaluation. The translation also utilizes the *quantifier elimination procedure* for linear order [22] to replace references to individual time instants in the queries with references to interval endpoints. In the rest of this section we give a sketch of the SQL/TP to SQL/92 translation and is based on an extension of results in [20] to duplicate semantics and uses a nouveau normalization technique.

4.1 Data Definition Language


The translation of the data definition language statements is fairly simple: we merely convert abstract signatures to their concrete counterparts. The SQL/TP statement in Example 3 is translated to:

```
create table indep (name char(20), yearmin Time, yearmax Time)
```

where **Time** is a user defined type (UDT) for an integer like time. The data type **Time** is equivalent to **INTEGER**⁹ extended with two special elements, $-\infty$ and ∞ . We define the successor and order for this new type by lifting the operations from the **INTEGER** type.

4.2 The SQL/TP Queries

The crux of our approach lies in the translation of SQL/TP queries. The natural correctness criterion is the preservation of query semantics. This requirement is captured by the following diagram:



We show that the proposed translation algorithm guarantees commutativity for the left part of the diagram. The commutativity of the right part is backed up by the reliability of the underlying relational system. The rest of this section gives a proof to the following theorem:

Theorem 13 *Let D be a concrete temporal database and Q a SQL/TP query. Then $Q(\|D\|) = \|compile(Q)(D)\|$.*

Note that $compile(Q)$ is executed at query compilation time—before the actual execution over the temporal database begins. Thus it does not affect the data complexity of the query evaluation algorithm. Before we present the steps of the translation itself, we introduce three auxiliary definitions.

The definition of SQL/TP queries requires the temporal attributes in the signature of a query to be independent. However, it does not prevent us from writing queries whose subqueries do not share this property. To deal with such attribute dependencies during the translation we introduce the notion of a *conditional query*:

Definition 14 (Conditional Query) *Let Q be a SQL/92 query and φ a quantifier-free formula in the language of linear order such that t is a free variable in φ if and only if t_{min} and t_{max} are temporal attributes in the signature of Q . We call $Q\{\varphi\}$ a conditional query.*

⁹ Often we can take advantage of a built-in data type provided by the RDBMS, e.g., data type **DATE** in DB/2.

Attribute	$\max(a)$	$\min(a)$	$\text{count}(a)$	$\text{sum}(a)$
data	$\max(a)$	$\min(a)$	$\text{sum}(\text{CNT}_G^{\mathcal{O}}(\mathbf{x}))$	$\text{sum}(a \cdot \text{CNT}_G^{\mathcal{O}}(\mathbf{x}))$
temporal	$\max(\text{amax})$	$\min(\text{amin})$	$\text{sum}(\text{CNT}_G^{\mathcal{O}}(\mathbf{x}))$	N/A

Fig. 2. Translation of Aggregate operations.

While the translation algorithm uses the conditional queries to translate subqueries of a SQL/TP query, the attribute independence of the top-level attributes guarantees that no such dependencies remain in the result of the translation.

The second challenge lies in the definition of relational operators that preserve semantics over the interval-based encoding. For this purpose we introduce a nouveau *normalization* technique. The idea behind the technique is quite simple:

Definition 15 Let $\{Q_1, \dots, Q_k\}$ be a set of SQL/92 queries with compatible signatures, X a subset of their data attributes, and t a temporal attribute. $Q_1, \dots, Q_k$ are t -compatible on X if for all concrete temporal databases D and all $0 < i \leq j \leq k$ whenever two concrete tuples $\mathbf{a} \in Q_i(D)$ and $\mathbf{b} \in Q_j(D)$ such that $\pi_X(\mathbf{a}) = \pi_X(\mathbf{b})$ then the sets $\pi_{\{t\}}(\|\mathbf{a}\|)$ and $\pi_{\{t\}}(\|\mathbf{b}\|)$ are identical or disjoint.

$Q_1, \dots, Q_k$ are time-compatible on X if $Q_1, \dots, Q_k$ are t -compatible on X for all temporal attributes t in the common signature.

The definition of a time-compatible set of queries essentially says that if the data portion of a tuple is related to an interval in $Q_i(D)$ and to another interval in $Q_j(D)$, then it is always the case that these two intervals coincide or are disjoint. This way we can guarantee the intervals behave like points with respect to set/bag operations (cf. Figures 3 and 4). This definition is non-trivial even for singleton sets of queries as the answers to queries are *bags* of tuples.

It is also easy to see that we can define a *normalization operation* that transforms a set arbitrary queries to a t -compatible set of $(\|\cdot\|)$ -equivalent queries. Moreover, this operation can be defined using a first-order query¹⁰:

Lemma 16 Let $\{Q_1, \dots, Q_k\}$ be a set of SQL/92 queries with compatible signatures such that X a subset of their data attributes and t a temporal attribute. Then there are first-order queries $N_X^t[Q_i; Q_1, \dots, Q_k]$ such that

1. $\|Q_i(D)\| = \|N_X^t[Q_i; Q_1, \dots, Q_k](D)\|$ for all concrete databases D .
2. $\{N_X^t[Q_i; Q_1, \dots, Q_k] : 0 < i \leq k\}$ are t -compatible on X .

To define a time-compatible set of queries the above lemma is used for all temporal attributes in the common signature. It is easy to see that the normalization operation can be performed in $O(n \log n)$ where n is the combined size of the results of $Q_i(D)$.

The last obstacle is the translation of aggregate operations. To translate the aggregation operators correctly we need to know how many tuples are encoded

¹⁰ Similarly to coalescing; a native implementation of the normalization can often be made more efficient.

by every single concrete tuple in the relation we aggregate over: we define a function CNT_G^φ for this purpose: it tells us how many duplicates would be in the $\|\cdot\|$ -image of the result of projecting a concrete tuple on G after applying the selection condition σ_φ . More formally:

Definition 17 *Let E be an abstract signature, $G \subset E$, φ a quantifier free formula in the language of linear order over temporal variables in $E - G$, and $\mathbf{x}$ a concrete tuple in the signature $\bar{E}$. Then $\text{CNT}_G^\varphi(\mathbf{x}) = \text{card}(\sigma_\varphi\|\pi_{E-G}(\mathbf{x})\|)$.*

Note that CNT_G^φ maps concrete tuples to natural numbers. However, if we used a dense model of time then CNT would be a measure on the sets of time instants and could return non-integral counts, e.g., 1.5 years. For details on aggregation and measures see [10].

Lemma 18 *Given fixed E , $G \subset E$, and φ , the function CNT_G^φ can be defined using an integer expression over the value of $\mathbf{x}$.*

The CNT function operates on single tuples and thus contributes only a constant to the overall data complexity of queries.

Now we are ready to proceed with the translation itself: every SQL/TP expression Q is translated to a set of conditional queries $Q_1\{\varphi_1\}, \dots, Q_n\{\varphi_n\}$ while maintaining the invariant $Q(\|D\|) = \sigma_{\varphi_1}\|Q_1(D)\| \cup \dots \cup \sigma_{\varphi_n}\|Q_n(D)\|$. The translation itself is defined inductively on the structure of the SQL/TP query.

Translation of the Select Block. Consider the SQL/TP statement:

select [all] X **from** $E_1, \dots, E_n$ **where** φ **group by** G

where X is the set of attributes in the answer, E_i subqueries or base table references, φ the selection condition, and G the set of grouping attributes. In addition, let A be the set of all aggregate expressions in X and E be the set of all attributes in $E_1, \dots, E_n$. We assume that we have already translated the subqueries to $Q_1\{\varphi_1\} \in \text{compile}(E_1), \dots, Q_n\{\varphi_n\} \in \text{compile}(E_n)$ ¹¹. We compose the partial results to get a set of conditional queries equivalent to the original select block as follows (we proceed by translating every clause of the original select block one by one):


from $E_1, \dots, E_n$: For every $Q_1\{\varphi_1\}, \dots, Q_n\{\varphi_n\}$ the **from** clause gives us a SQL/92 query

select $\bar{E}$ **from** $Q_1, \dots, Q_n\{\psi\}$

where $\psi = \varphi_1 \wedge \dots \wedge \varphi_n$.

where φ : To apply the selection condition φ we need to determine the relationships between the endpoints of intervals, the corresponding point-valued attributes, and the selection formula. We use the quantifier elimination procedure for linear order to achieve this goal. Let $Q(E)\{\psi\}$ be the result of the

¹¹ for the base tables we use a trivial condition *true* in place of φ_i .



S has data attribute x and a temporal attribute t . The boxes in the figure represent the t - x graphs of the involved tables. Similar technique is used for aggregation: it is easy to see that we could easily count duplicates on the normalized relation (cf. Section 3.3, query 6).

Fig. 3. Projection with Duplicate Elimination

previous step. We define

$$\begin{aligned}\psi_1 &:= \text{QE}(\exists T(\psi \wedge \varphi \wedge \bigwedge_{t \in T} (tmin \leq t \leq tmax))) \\ \psi_2 &:= \text{QE}(\exists \bar{T}(\psi \wedge \varphi \wedge \bigwedge_{t \in \bar{T}} (tmin \leq t \leq tmax)))\end{aligned}$$

where T is the set of all temporal attributes in E (encoded by intervals), $\bar{T}$ is the set of all attributes except those in T and constants, and QE is the quantifier elimination procedure for linear order. Now we define

$$\text{select } \bar{E} \text{ from } Q \text{ where } \psi_1\{\psi_2\}$$

to be the result of applying the original **where** clause on $Q\{\psi\}$.

group by G : To apply grouping we first normalize the result of the previous step with respect to the attributes in G . This allows us to use the standard SQL/92 grouping construction (cf. Figure 3). Let $Q(E)\{\psi\}$ be the result from the previous step. As the attributes in G and $E - G$ are independent we can split ψ to ψ_1 involving only attributes in G and ψ_2 involving only attributes in $E - G$. We generate

$$\text{select } \bar{G}, \bar{A} \text{ from } N_G[Q; Q] \text{ group by } \bar{G}\{\psi_1\}$$

Note that the aggregates in A have to be transformed using the rules in Figure 1 and Lemma 18 applied on E , G , and ψ_2 .

select X : The translation of the final projection depends on the use of duplicate preserving vs. duplicate-eliminating projection. Let $Q(E)\{\psi\}$ be the result from the previous step.

- for **select distinct** queries we get **select distinct** $\bar{X}$ from $N_X[Q; Q]\{\psi\}$ and
- for **select all** queries we get **select** $\bar{X}$ from $Q\{\psi\}$

Note that only queries that use aggregation or duplicate elimination use the **N** operation.

For each of these steps we can easily verify that the transformation invariant is preserved. Moreover, in the actual implementation the subqueries generated by the above four steps are merged into as few nested blocks as possible, e.g., the first two steps can be always merged into a single **select** block, etc.




Fig. 4. Set Difference using Normalization.

Translation of the Set Operations. The translation of the set operations follows similar path as the translation of the duplicate elimination: we need to find conditions under which the set operations on the encoding are equivalent to set operations on the abstract relations. Clearly a direct use of SQL/92 set operations does not preserve the semantics of SQL/TP queries.

Lemma 19 Let Q_1, Q_2 be time-compatible SQL/92 queries with common signature. Then $\|Q_1\| \text{ op } \|Q_2\| = \|Q_1 \text{ op } Q_2\|$ where op is one of the union [all], except [all], or intersect [all].

The above lemma extends to conditional queries: Let $Q_i\{\varphi_i\} \in \text{compile}(Q)$ for $i \in I$, $R_j\{\psi_j\} \in \text{compile}(R)$ for $j \in J$, and X the set of data attributes in the common signature. Then

$$\begin{aligned} Q \text{ union } R &\mapsto \{Q_i\{\varphi_i \wedge \neg \bigvee_{j \in J} \psi_j\}, R_j\{\psi_j \wedge \neg \bigvee_{i \in I} \varphi_i\}, Q'_i \text{ union } R'_j\{\varphi_i \wedge \psi_j\}\} \\ Q \text{ except } R &\mapsto \{Q_i\{\varphi_i \wedge \neg \bigvee_{j \in J} \psi_j\}, Q'_i \text{ except } R'_j\{\varphi_i \wedge \psi_j\}\} \\ Q \text{ intersect } R &\mapsto \{Q'_i \text{ intersect } R'_j\{\varphi_i \wedge \psi_j\}\} \end{aligned}$$

where $Q'_i = N_X[Q_i; Q_i, R_j]$ and $R'_j = N_X[R_j; Q_i, R_j]$. We can omit the normalization for the union operation. The duplicate-preserving operations are transformed analogously.

The result of the translation can unfortunately be exponential in the depth of nesting of the original query. Note that this does *not* affect the data complexity of the query evaluation as the translation is performed at compile time. Moreover, for large class of queries we can show:

Theorem 20 Let Q be a query composed of attribute independent subqueries with size at most k for a fixed constant $k \in N$. Then $|\text{compile}(Q)| \in O(|Q|)$.

Thus *small* subqueries can violate the attribute independence requirement while still maintaining the bound on the size of the translated query. Using this result we can show that, e.g., the first-order temporal logic queries can be efficiently translated through SQL/TP: all temporal operators are translated into small fixed-size subqueries and views with exactly one temporal attribute [3]. Similar result holds for all TRA [7] based languages.

View Definitions in SQL/TP. A side effect of this observation is a natural definition of *views* in the SQL/TP. Again, we use the familiar SQL/92 syntax:

```
create view <rid> ( <query> )
```

where <rid> is a view name and <query> an SQL/TP query. According to Theorem 20 queries that can be expressed using views can be translated with only linear increase in size.

4.3 Optimization of the Translation

It is easy to see that the normalization operation N is idempotent in a similar way as projection:

$$Q_i(D) = N_X^t[Q_i; Q_1, \dots, Q_k](D)$$

for any queries $Q_1, \dots, Q_k$ that are t -compatible on X . Note that in this case we are comparing the concrete representations rather than the images under $\|\cdot\|$. The optimization are based on the following two lemmas:

Lemma 21 Let S be a set of SQL/92 queries t -compatible on X . Then all queries in the closure of S under π_V , σ_φ , and op are also t -compatible on X where $X \cup \{t\} \subset V$, op is a set operation, and φ is a selection condition.

Lemma 22 If $Q_1, \dots, Q_k$ are t -compatible on X then they are t -compatible on Y for all $Y \supseteq X$.

Thus these rules can be used to push and eliminate superfluous normalization steps. In addition, during the optimization phase we flatten nested simple select blocks created during the translation (using trivial attribute renaming), merge their **where** conditions (and check for satisfiability), and perform a “general” clean-up of the resulting query.

4.4 Updates

A naive implementation of the update operations is straightforward. However, an efficient execution model of database updates over the concrete databases is a topic for further research (cf. Section 5). A major step towards efficient updates *in place* would be development of an in place normalization technique for temporal relations. This way we could solve the update problem by normalizing (cf. Section 4) the updated relation with respect to the update query and then process the actual update the same way SQL does. However, the normalization operation is costly; the hope is that the cost would amortize over time as the normalization is idempotent; cf. Section 4.3.

5 Conclusion

We have shown that a high-level point-based approach to temporal extension of SQL has many advantages over the common approaches that use interval-based attributes: simple syntax and semantics, meaningful aggregation, and possibilities of advanced query optimization. All this is achieved while maintaining efficient query evaluation over temporal databases based on interval encoding of timestamps. We have also shown that all representation independent TSQL2 queries are expressible in SQL/TP (follows from [20]).

Future Work. This proposal is only the first step towards the implementation of SQL/TP. There are still many open questions:

- Can we use more complex temporal domains? In our proposal we used a discrete linear order with a limited way of counting. Is it possible to use richer temporal domains while maintaining the properties of the proposed language? What are the tradeoffs?
- We have chosen an interval-based encoding of sets of time instants in the concrete data model. While this encoding can compactly describe periods of time, it fails, e.g., for periodic events. Is it possible to extend the encoding scheme and this way to enlarge the class of finitary temporal databases? Note that we have to be careful here: we still need to perform complementation and projection in a data-independent way during the translation of queries. Choosing too rich encoding can prevent us from achieving this goal.
- What optimization techniques can be used in conjunction with our query translation procedure? We have provided only a high-level description of the compilation technique that allows SQL/TP queries to be evaluated over an interval-based temporal database. We also sketched several optimization techniques specific to this translation. However, we need to analyze how the resulting queries interact with the existing database optimizer and the underlying representation. This leads to questions of proper index selection, the use of spatial operations (spatial joins), etc.
- How do we perform updates efficiently? The area of updates presents a completely new set of problems, the main problem being the in-place updates of the encoded temporal relations. This problem goes hand in hand with defining various normal forms [6, 20] of temporal relations and enforcing them over updates¹².
- Can the standard indices built-in relational systems aid the query evaluation based on the proposed compilation technique? What are the tradeoffs comparing to specialized indices?

There are few answers to these questions even for the established temporal query languages like TSQL2. Our technique allows us to reuse most of the efforts aimed towards boosting performance of temporal DBMS, e.g., the development of efficient temporal and spatial joins, and sophisticated access methods.

¹² In this chapter we did not assume any normal form for the temporal relations.

References

1. Abiteboul, S., Herr, L., Van den Bussche, J. Temporal versus First-Order Logic to Query Temporal Databases. *Proc. ACM PODS 1996*, 49–57, 1996.
2. Allen, J. F. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843, 1983.
3. Böhlen, M. H., Chomicki, J., Snodgrass, R. T., Toman, D. Querying TSQL2 Databases with Temporal Logic. In *Proc. EDBT'96*, LNCS 1057, 325–341, 1996.
4. Böhlen, M. H., Jensen, C. S. Seamless Integration of Time into SQL. University of Aalborg, <http://www.cs.auc.dk/~boehlen/Software/Tiger/atsql.ps.gz>, 1996.
5. Böhlen, M. H., Jensen, C. S., Snodgrass, R. T. Evaluating and Enhancing the Completeness of TSQL2. Technical Report TR 95-5, Computer Science Department, University of Arizona, June 1995.
6. Böhlen, M. H., Snodgrass, R. T., Soo, M. D. Coalescing in Temporal Databases. *Proc. 22nd Int. Conf. on Very Large Databases*, 180–191, 1996.
7. Clifford J., Croker A., Tuzhilin A. On Completeness of Historical Relational Query Languages. *ACM Transactions on Database Systems*, Vol. 19, No. 1, 64–116, 1994.
8. Codd, E. F. Relational completeness of database sublanguages. In Rustin, R.(ed.) *Courant Computer Science Symposium 6: Data Base Systems*, 65–98, Prentice-Hall, 1972.
9. Chomicki J. Temporal Query Languages: a Survey. *Proc. International Conference on Temporal Logic*, July 1994, Germany, Springer-Verlag (LNAI 827), 506–534.
10. Chomicki, J., Kuper, G. M. Measuring Infinite Relations. *Proc. ACM PODS 1995*, 78–85, 1995.
11. Chomicki, J., Goldin, D. Q., Kuper, G. M. Variable Independence and Aggregation Closure. *Proc. ACM PODS 1996*, 40–48, 1996.
12. Date, C. J., Drawen, H. A Guide to the SQL Standard (3rd ed.), Addison-Wesley, 1993.
13. IBM Database 2, SQL Reference for common servers. IBM Corp., 1995.
14. Jensen, C. S., Snodgrass, R. T., Soo, M. J. Unification of Temporal Data Models. *Proc. 9th Int. Conf. on Data Engineering*, 262–271, 1993.
15. Kanellakis, P. C., Kuper, G. M., Revesz, P.Z. Constraint Query Languages. *Journal of Computer and System Sciences* 51(1):26–52, 1995.
16. Ladkin, P. The Logic of Time Representation. PhD Dissertation, University of California, Berkeley, 1987.
17. Snodgrass R. T. The Temporal Query Language TQuel. *ACM Transactions on Database Systems*, 12(2):247–298, June 1987.
18. Snodgrass R.T. (editor). *The TSQL2 Temporal Query Language*. Kluwer Academic Publishers, 674+xxiv pages, 1995.
19. Snodgrass, R. T., Böhlen, M. H., Jensen C. S., Steiner, A. Adding Valid Time to SQL/Temporal, ISO/IEC JTC1/SC21/WG3 DBL MAD-146r2 21/11/96, (change proposal).
20. Toman, D. Point-based vs. Interval-based Temporal Query Languages *Proc. ACM PODS 1996*, 58–67, 1996.
21. Toman, D., Niwinski, D. First-Order Temporal Queries Inexpressible in Temporal Logic *Proc. EDBT'96*, Arpes, Bouzeghoub (eds.), LNCS 1057, 307–324, 1996.
22. Williams, H. P. Fourier–Motzkin Elimination Extension to Integer Programming Problems. In *Journal of Combinatorial Theory (A)* 21, 118–123, 1976.
23. Kabanza, F., Stevenne, J.-M., Wolper, P. Handling Infinite Temporal Data. *JCSS* 51(1): 3–17, 1995.

A Syntax of SQL/TP Core Language

We summarize the syntax of SQL/TP by the following BNF grammar:

```

<exp>      ::= <ddl_exp> | <dql_exp> | <dml_exp>

<ddl_exp> ::= create table <rid> ( <signat> ) |
              create view  <rid> ( <dql_exp> )

<dql_exp> ::= select [all] <s_list> [from <f_list>]
                  [where <cond>] [group by <g_list>] |
                  ( <dql_exp> ) union [all] ( <dql_exp> ) |
                  ( <dql_exp> ) except [all] ( <dql_exp> ) |
                  ( <dql_exp> ) intersect [all] ( <dql_exp> )

<dml_exp> ::= insert [all] into <rid> ( <dql_exp> ) |
              delete [all] from <rid> ( <dql_exp> )

<signat>   ::= <id> <type> [, <signat>]

<s_list>   ::= <expr> [as <id>] [, <s_list>]

<f_list>   ::= { <rid> [<rid>] | ( <dql_exp> ) <rid> } [, <f_list>]

<cond>     ::= [<rid>.]<id> <rel> [<rid>.]<id> [+ <const>] |
              [<rid>.]<id> <rel> <const> |
              <cond> and <cond> | ...

<g_list>   ::= [<rid>.]<id> [, <g_list>]

<type>     ::= time [using <enc>] | integer | real | char(N) | ...

<expr>     ::= [<rid>.]<id> | <aggr> ( [<rid>.]<id> ) | ...

<enc>      ::= points | bounded intervals | unbounded intervals

<aggr>     ::= min | max | count | sum

```

where <rid> is a valid relation identifier, <id> a valid attribute identifier, and <rel> a valid comparison type-compatible with the involved attributes.

B Comparison with SQL/Temporal

Now we are ready to compare SQL/TP with the most current proposal of temporal extensions to the SQL-3 ISO/ANSI committees—SQL/Temporal [19]. The

comparison is slightly complicated as SQL/Temporal comes in three separate flavors (with distinct semantics). These variants follow the migration requirements stated informally in Section 3.5 (for full description see [19]).

TUC. The temporal upward compatibility semantics is based on mere restriction of all temporal relations to the current state. This way the query evaluation is essentially performed on a standard relational database. From this point of view, this mode doesn't involve any manipulation of the temporal attributes (and we include it only for completeness here).

Sequenced semantics. This is, in fact the only mode in which the temporal dimension is truly managed by the database system. Unfortunately, the class of queries to which this semantics applies is very restricted: only plain SQL queries applied point-wise to the whole database are allowed. Moreover, this part of the semantics is based on a *valid-time relational algebra* (VTA) and the results in [1, 21] show that no single-dimensional temporal relational algebra can express all first-order queries (e.g., the existential closure of query 3 in Section 3.3) and thus even extending VTA with additional temporal operators, e.g., *since* and *until*, can not achieve first-order completeness (in contrary with the popular belief, e.g., [7]).

Nonsequenced semantics. In order to formulate *general temporal queries* (as none of the above semantics allows that), SQL/Temporal introduces the *non-sequenced semantics*. This semantics is indeed first-order complete. However, the completeness comes with a high price tag: the semantics is essentially based on a first-order logic over *interval-based temporal domain* [20], rather than the individual time instants and thus is plagued by all the problems discussed in Section 1.

The separate semantic layers may also cause quite a bit of confusion to the users: many operations, e.g., the set difference or aggregation, behave completely differently in the sequenced semantics and in the nonsequenced semantics: the sequenced semantics is defined with respect to the individual time instants (similarly to SQL/TP), while the nonsequenced semantics is defined with respect to the encoding of sets of time instants—intervals—and the user is asked to manipulate the the representation of timestamps on her own (this can be clearly seen by formulating queries 2 and 3 from Section 3.3 in SQL/Temporal).

From our point of view the second (sequenced) semantics is the most promising part of SQL/Temporal: the temporal information is handled correctly with respect to the underlying *point-based* semantics. We consider SQL/TP a fully expressive temporal extension of the sequenced semantics, rather than the nonsequenced semantics. Section 3.5 supports this claim by showing how the *sequenced semantics* can be naturally embedded into SQL/TP.

However, the encoding-dependent part of SQL/Temporal is solely connected with the nonsequenced semantics. We propose to replace this part of the language with SQL/TP: this change would lead to a natural extension of the sequenced semantics to all temporal queries.