

# 23

## Evaluating the Completeness of TSQL2

**Michael H. Böhlen, Christian S. Jensen, and  
Richard T. Snodgrass**

---

The question of what is a well-designed temporal data model and query language is a difficult, but also an important one. The consensus temporal query language TSQL2 attempts to take advantage of the accumulated knowledge gained from designing and studying many of the earlier models and languages. In this sense, TSQL2 represents a constructive answer to this question. Others have provided analytical answers by developing criteria, formulated as completeness properties, for what is a good model and language.

This paper applies important existing completeness notions to TSQL2 in order to evaluate the design of TSQL2. It is shown that TSQL2 satisfies only a subset of these completeness notions.

---

## 1 Introduction

The temporal database community has been prolific in its production of temporal data models and query languages. Over the past fifteen years, more than two dozen temporal relational data models have been proposed, each with one or more associated query languages [24]. This has left the community with a wide, confusing—but also challenging—variety of alternatives.

As one response to this state of affairs, a committee of eighteen temporal database researchers has recently released the TSQL2 Language Specification [25], which defines a temporal extension to the SQL-92 standard [22]. TSQL2 was created partly in an attempt to consolidate, in a single consensual model and language, the insights and experiences gained from the development of the previous data models and languages.

As a quite different approach, other efforts (e.g., [7, 21, 6, 5]) have put focus on the properties of temporal data models and query languages, as well as on the design alternatives available when developing these. This has led to precise definitions of model and language properties that can be used to characterize and evaluate the many models and languages. In the spirit of Codd's original definition of relational completeness [11], some of these properties have been stated as different kinds of completeness.

It then seems appropriate to use the body of work on model and language properties to study the design of TSQL2—this paper does exactly that. It is a fundamental assumption of the paper that when evaluating a data model and query language, both the functionality and the syntax for expressing a certain functionality are important. The completeness notions that we adopt in the investigation thus include both functionality-related and syntactical criteria.

Specifically, we formalize the notion of a data model being *upwards compatible* with another data model and show that TSQL2 is upwards compatible with SQL-92. Briefly, this means that a smooth transition from SQL-92 to TSQL2 is possible.

One of the most widely cited distinctions among temporal data models is that between first normal form and non-first normal form models. This distinction has been formally captured by the concepts of *temporally ungrouped* and *grouped* data models [7]. We show that TSQL2 is temporally ungrouped and not temporally grouped. As this property is inherent in the model, we do not propose to change it. Rather, we put focus on the implications of a model being ungrouped or grouped.

The last two completeness notions considered in this paper are *temporal semi-completeness* and *temporal completeness* [5]. The former notion essentially states that a temporal relational data model must contain temporal generalizations of *all* snapshot relations and queries. Further, temporally generalized queries must be *syntactically similar* to the snapshot queries they generalize. Temporal completeness

adds further functional and syntactical requirements, accounting for query language aspects not covered by temporal semi-completeness. It is shown that TSQL2 does not fully satisfy these completeness notions.

Related work on completeness has been primarily in the context of non-temporal databases. It is possible to distinguish two basic approaches. The first one takes a particular calculus (usually first order relational calculus) as a metric. Any language having at least the expressive power of the calculus is said to be complete. Original work along these lines was done by Codd for relational databases [11]. There have been generalizations for entity-relationship databases [1] and for temporal databases [30]. One inherent problem with these approaches is the degree of appropriateness of the calculus that is used as a metric. There is no guarantee that the calculus captures *all* reasonable queries. For example, it has been shown [3] that first order relational calculus cannot express the transitive closure of binary relations.

The second approach is to define an appropriately large set of queries and require query languages to express all queries in this set. This kind of completeness was investigated by Bancilhon [4] and Chandra and Harel [8]. The definitions of temporal semi-completeness and temporal completeness are in this spirit. They (in particular temporal semi-completeness) take the set of queries that are expressible by a non-temporal language as a reference and ensure that temporal generalizations of the non-temporal language can express all these queries. Additionally, they establish syntactic restrictions a temporal language must obey, which we believe is also important.

Briefly, the contributions of the paper are twofold. First, the paper further formalizes some existing definitions of completeness of relevance for temporal data models and query languages, namely the notions of upwards compatibility, temporal semi-completeness and temporal completeness. Second, the paper explores the design of TSQL2 by applying these completeness notions and the notion of temporal (un)groupedness to TSQL2. It is shown that TSQL2 satisfies some of these notions, but does not satisfy all of them.

The paper is structured as follows. Each of Sections 2–5 first defines a particular type of completeness. They then evaluate the completeness of TSQL2 in the context of each completeness notion. During this investigation, some deficiencies of TSQL2 are uncovered.<sup>1</sup> Section 6 summarizes the paper and points to directions for future research.

---

<sup>1</sup>To answer these deficiencies the design of Applied TSQL2 (ATSQL2) was initiated. An extended version of this paper (<ftp://ftp.cs.arizona.edu/reports/1995/TR95-11.ps.gz>) sketches syntax and semantics of ATSQL2. A more comprehensive discussion as well as an implementation for public consumption are under way and should be completed by Fall 1995.

## 2 Upwards Compatibility

Completeness is generally a relative property of a data model or a query language. Thus, a model or language satisfies some notion of “completeness” if it is related to another model or query language in a certain way. In this section, we introduce the first of the three types of completeness. Specifically, we formalize the notion that a data model is upwards compatible with respect to another data model. We subsequently consider the upwards compatibility of TSQL2 with respect to SQL-92.

### 2.1 Definitions

When a new database management system, with an associated data model, is introduced into an organization, often that system replaces an existing system, also with an associated data model. For software engineering reasons, to be discussed in more detail below, it is an important property that the existing data model be upwards (or, forwards) compatible with the new data model. Put differently, the new data model should be a strict superset of the existing data model.

We will adopt the convention that a data model consists of two components, namely a set of data structures and a language for querying the data structures. For example, the central data structure of the relational model is the relation, and the central, user-level query language is SQL. Notationally,  $M = (DS, QL)$  then denotes a data model,  $M$ , consisting of a data structure component,  $DS$ , and a query language component,  $QL$ . Thus,  $DS$  is the set of all databases expressible by  $M$ , and  $QL$  is the set of all queries in  $M$  that may be formulated on some database in  $M$ . We will use  $db$  to denote a database and  $q$  to denote a query.

**Definition 1 (upwards compatibility)** Let  $M_2 = (DS_2, QL_2)$  and  $M_1 = (DS_1, QL_1)$  be two data models. Model  $M_2$  is *upwards compatible with* model  $M_1$  if and only if

- $DS_2 \supseteq DS_1$ , and
- for each instance  $db$  in  $DS_2$  and for each query expression  $q$  in  $QL_1$ ,  $q$  is also a legal query expression in  $QL_2$ , and the results of evaluating  $q$  on  $db$  is the same in  $M_1$  and  $M_2$ . □

This concept captures the conditions that need to be satisfied in order to allow a smooth transition from a current system, with data model  $M_1$ , to a new system, with data model  $M_2$ . The first condition implies that all existing databases in the old system are also legal databases in the new system and thus need not be modified when the new system is adopted. The second condition guarantees that existing queries will remain legal and will compute the same results in the new system as

in the old system. Thus, the bulk of legacy application code is not affected by the transition to a new system.

The definition of upwards compatibility is related to the traditional notion of Codd completeness (Codd originally used the term relational completeness) [11], as formulated in the context of the standard relational model. To see the similarity and differences, we review that completeness notion.

Essentially, a relational or extended relational data model is Codd complete if all queries that can be formulated on arbitrary conventional relations expressible in the model are a superset of all relational algebra queries that can be formulated.

**Definition 2 (Codd completeness)** Let  $M = (DS, QL)$  be some data model, and let  $(SR, RA)$  be the relational model with the relational algebra as its query language. Model  $M$  is *Codd complete* if and only if  $DS$  contains  $SR$  and each query in  $RA$  has an equivalent counterpart in  $QL$  when all  $db$  in  $DS$  are considered.  $\square$

Two query expressions are equivalent if they always yield mutually identical results when supplied identical arguments. The relational algebra comes in numerous versions<sup>2</sup>, and while the definition is dependent on the particular version chosen, the choice is not important in the remainder of this paper.

The similarity between upwards compatibility and Codd completeness is apparent, but there are also important differences. First, Codd completeness is restricted to use the relational algebra as a yardstick for measuring the expressive power of other query languages. Thus, the relevance of Codd completeness is dependent on how “natural” or well-chosen the relational algebra is. On the other hand, upwards compatibility is not tied to any particular data model.

Second, Codd completeness strictly concerns functionality while upwards compatibility concerns both functionality and the syntax for expressing the functionality. Specifically, Codd completeness is defined in terms of the existence of *equivalent*, but possibly different, query expressions. Upwards compatibility requires query expressions that yield identical results to also be syntactically *identical*. Thus, a model being upwards compatible with the relational model/algebra is a stronger criterion than the model being Codd complete.

## 2.2 Upwards Compatibility among SQL-92 and TSQL2

Clearly, it is an important property for a new data model, such as TSQL2, to be a strict superset of the data model it is intended to supersede, i.e., SQL-92. We now consider this issue.

---

<sup>2</sup>The relational algebra used in conjunction with the original definition of Codd completeness [11] included “cartesian [sic] product,” “union,” “intersection,” “difference,” “projection,” “ $\theta$ -join,” “division,” and “restriction” (a special case of selection).

In TSQL2, there are six kinds of relations<sup>3</sup>: snapshot relation, valid-time event relation, valid-time state relation, transaction-time relation, bitemporal event relation, and bitemporal state relation. The first is the kind of relation found in the relational model; the remaining five are temporal relations. As all the schema specification statements of SQL-92 are included in TSQL2, it follows that the data structures of TSQL2 include those in SQL-92.

TSQL2 is also a strict superset of SQL-92 in its query facilities. In particular, if an SQL-92 select statement does not incorporate any of the constructs added in TSQL2 (e.g., the valid clause, the `VALID( )` and `TRANSACTION( )` expressions, and extensions to the from and group by clauses), and mentions only snapshot relations in its from clause(s), then the language specification states explicitly that the semantics of this statement is identical to its SQL-92 semantics.

It should be noted that the preliminary TSQL2 language specification released in March, 1994 [25] did not have that property. In particular, SQL-92 `INTERVALS` were termed `SPANS` in the preliminary TSQL2 specification, and TSQL2 `INTERVALS` were not present at all in SQL-92. The final TSQL2 language specification [25] retained SQL-92 `INTERVALS` and added the `PERIOD` data type, which was previously called `INTERVAL` in preliminary TSQL2 (confusing, isn't it?). Additional changes to the datetime literals were also made to ensure that TSQL2 was a strict superset of SQL-92.

Hence, both conditions are satisfied, demonstrating that TSQL2 is upwards compatible with SQL-92.

As discussed previously, this directly implies that TSQL2 is Codd complete.

Finally note that, while upwards compatibility is a highly desirable property, it says absolutely nothing about constructs added to a data model or query language to support time. This notion of completeness is thus quite limited in scope, as seen from a temporal data-model perspective.

### 3 Temporal Groupedness

In this section, we first review the previously proposed notions of temporally ungrouped and grouped data models. We then investigate the temporal groupedness of TSQL2. In contrast to upwards compatibility, temporal groupedness speaks directly to the support of time-varying information in the temporal data model.

#### 3.1 Definition

In temporal data modeling, an informal division among temporal relational data models into first normal form (1NF) and non-first normal form (NFNF) models has

---

<sup>3</sup>In this paper, we use the terminology Codd introduced [10]: relation, tuple, and attribute, rather than the more prosaic terminology used in SQL-92 and subsequently in TSQL2: table, row, and column.

developed over the years, and each type of model has attracted its followers<sup>4</sup>.

With one objective being to clarify this distinction, Clifford et al. [7] have recently given a formal definition of two types of relation structures, termed temporally ungrouped and temporally grouped. While it is debatable whether the data model of TSQL2 is strictly a 1NF model in the generic sense<sup>5</sup>, we will show that the model is temporally ungrouped. To set the stage, we review the definition of a temporally ungrouped data model.

A data model is temporally ungrouped if its data structure component is isomorphic to a particular canonical temporally ungrouped data structure, i.e., an onto and 1–1 mapping must exist between the canonical structure and the structure of the model to be proved temporally ungrouped. The canonical structure is defined next.

**Definition 3 (canonical temporally ungrouped relation structure)** [7, pp. 69–70] Let  $U_D = \{D_1, D_2, \dots, D_{n_d}\}$  be a set of non-empty value domains, and let  $\mathbf{D} = \cup_{i=1}^{n_d} D_i$  be the set of all values. Let  $\mathbf{T} = \{t_0, t_1, \dots, t_i, \dots\}$  be a non-empty, finite or countably infinite set of times with “<” as the total order relation. Finally, let  $U_A = \{A_1, A_2, \dots, A_n\}$  be a set of attributes, and let *TIME* be a distinguished time attribute.

A canonical temporally ungrouped (TU) relation schema is defined as a triple  $\langle \mathbf{A}, \mathbf{K}, \mathbf{DOM} \rangle$  where

- (1)  $\mathbf{A} \cup \{\textit{TIME}\}$  ( $\mathbf{A} \subseteq U_A$ ) is the set of attributes of the schema.
- (2)  $\mathbf{K} \cup \{\textit{TIME}\}$  ( $\mathbf{K} \subseteq \mathbf{A}$ ) is the key of the schema, i.e.,  $\mathbf{K} \cup \{\textit{TIME}\} \rightarrow \mathbf{A}$ .
- (3)  $\mathbf{DOM}$  is a function from  $\mathbf{A} \cup \{\textit{TIME}\}$  to  $U_D \cup \{\mathbf{T}\}$  that assigns domains in  $U_D$  to attributes in  $\mathbf{A}$  and *TIME* to  $\mathbf{T}$ .

A TU database schema is a finite set of TU relation schemas. A TU tuple  $\mathbf{t}$  on schema  $\langle \mathbf{A}, \mathbf{K}, \mathbf{DOM} \rangle$  is a function from  $\mathbf{A} \cup \{\textit{TIME}\}$  to  $\mathbf{D} \cup \mathbf{T}$  that assigns a value in  $\mathbf{DOM}(A_i)$  to each attribute  $A_i$  in  $\mathbf{A}$  and a value in *TIME* to  $\mathbf{T}$ . A TU relation is then a finite set of TU tuples that satisfy the key constraint in (2) above. A TU database is a finite set of TU relations.  $\square$

**Example 1** The following is a sample TU database with one relation.

<sup>4</sup>In this particular context, the NFNFness is due to *how time is added* to the relational model, so most NFNF temporal data models do not support general NFNF relations, and the distinction is different from the distinction between the 1NF and the various general NFNF relational data models (e.g., [17]).

<sup>5</sup>First normal form (1NF) states that each attribute value is *atomic* [10]. This certainly holds for TSQL2’s explicit attributes, which can have as types any of the SQL–92 data types or the new type `PERIOD`. Hence, considering only values of explicit attributes, TSQL2 is a 1NF model. However, the timestamp associated with each tuple in TSQL2 is a *temporal element*, a finite union of periods [15]. While the timestamp is not an explicit attribute, it can be referenced within a query. We thus feel that timestamps should also satisfy the property. Since the partitioning construct in the from clause of TSQL2 (designated “(`PERIOD`)”) effectively iterates over the maximal periods of a temporal element, timestamps are not treated as atomic. Thus, TSQL2 is *not* a 1NF model in this strict sense.

A	B	TIME
$a_1$	$b_1$	1
$a_3$	$b_2$	1
$a_2$	$b_1$	2
$a_3$	$b_3$	2

The relation schema is the structure  $\langle \{A, B\}, \{A, B\}, f \rangle$ , where  $f$  assigns domains  $\{a_1, a_2, a_3\}$  and  $\{b_1, b_2, b_3\}$  to  $A$  and  $B$ , respectively, and the natural numbers to  $TIME$ .  $\square$

A data model cannot be temporally ungrouped and temporally grouped (see below), and as we will prove that the TSQL2 data model is isomorphic to TU, we will not give a formal definition of a the canonical temporally grouped relation structure, TG. Rather, we give an example and point to what makes TG grouped.

**Example 2** The schema of a temporally grouped relation consists of the same three components as that of an ungrouped relation, with the exception that the component **DOM** assigns a domain of functions to each attribute in **A**. These functions map times to some value domain. A tuple, then, consists of some specific function for each attribute in **A**. In addition, a tuple has an associated lifespan, a set of times. The functions of a tuple map each time in the tuple's lifespan to some value.

For example, a TG relation schema may have attributes  $A$  and  $B$ . The key may be the combination of these attributes, and **DOM** may assign functions to  $A$  and  $B$  that map from the natural numbers to  $\{a_1, a_2, a_3\}$  and  $\{b_1, b_2, b_3\}$ , respectively. A sample tuple may have lifespan  $\{1, 2\}$  and may have the mappings  $[1 \rightarrow a_1, 2 \rightarrow a_2]$  as its  $A$ -value and  $[1 \rightarrow b_1, 2 \rightarrow b_1]$  as its  $B$ -value. A relation instance with this and one more tuple is given next.

A	B	<i>lifespan</i>
$1 \rightarrow a_1$	$1 \rightarrow b_1$	
$2 \rightarrow a_2$	$2 \rightarrow b_1$	$\{1, 2\}$
$1 \rightarrow a_3$	$1 \rightarrow b_2$	
$2 \rightarrow a_3$	$2 \rightarrow b_3$	$\{1, 2\}$

In comparison with the TU instance given before, this instance adds a grouping of the temporal information. As before, there are four rows. However, these rows may now be combined in several ways to form tuples. Other legal TG instances with the same rows are the following ones.

A	B	<i>lifespan</i>
$1 \rightarrow a_1$	$1 \rightarrow b_1$	$\{1\}$
$2 \rightarrow a_2$	$2 \rightarrow b_1$	$\{2\}$
$1 \rightarrow a_3$	$1 \rightarrow b_2$	
$2 \rightarrow a_3$	$2 \rightarrow b_3$	$\{1, 2\}$

A	B	<i>lifespan</i>
$1 \rightarrow a_1$	$1 \rightarrow b_1$	$\{1\}$
$2 \rightarrow a_2$	$2 \rightarrow b_1$	$\{2\}$
$1 \rightarrow a_3$	$1 \rightarrow b_2$	$\{1\}$
$2 \rightarrow a_3$	$2 \rightarrow b_3$	$\{2\}$

Put informally, a TG structure contains more elements than a TU structure. This



indicates why it is not possible for a relation structure to be both temporally ungrouped and grouped.  $\square$

### 3.2 TSQL2 and Temporal Groupedness/Ungroupedness

The canonical ungrouped relation structure TU is a quite simple one. The relation structure of TSQL2 is more elaborate. TSQL2 relations come in several variations. First, relations may support valid time, transaction time, or both. Second, valid-time support may be for either state or event relations. While each of the resulting six types of relations are important in practice, it is advantageous in this context to consider only valid-time state relations. This permits a focus on the important concepts and is consistent with existing work [7]<sup>6</sup>. With this restriction, the relation structures of TSQL2 and TU are quite similar.

The central difference is that in TU, tuples are stamped with a single *TIME* value from domain **T** while in TSQL2, tuples are stamped with sets of times, valid-time elements, from domain **T**. As we shall now demonstrate, this difference is not essential when groupedness is considered.

To show that TSQL2 is temporally ungrouped, we devise an isomorphic mapping between TSQL2 and TU. This mapping takes as argument an arbitrary TSQL2 relation with schema  $(A_1, A_2, \dots, A_n \parallel \mathbf{T})$  where the  $A_i$  are explicit value attributes and **T** is the implicit, set-valued time attribute (the vertical double-bar is used to emphasize that the  $A_i$ 's are explicit attributes and that **T** is a distinguished, implicit attribute). It maps each TSQL2 tuple in turn. A TSQL2 tuple

$$(a_1, a_2, \dots, a_n \parallel \{t_{i_1}, t_{i_2}, \dots, t_{i_k}\})$$

is mapped to the set

$$\{(a_1, a_2, \dots, a_n \parallel t_{i_1}), (a_1, a_2, \dots, a_n \parallel t_{i_2}), \dots, (a_1, a_2, \dots, a_n \parallel t_{i_k})\}$$

of TU tuples. Note that one TU tuple is generated for each time in the timestamp of the argument TSQL2 tuple. No duplicates are introduced as TSQL2 timestamps are sets of times. Note also that no duplicate tuples are introduced between the sets of tuples generated from individual TSQL2 tuples. This is so because TSQL2 relations do not contain value-equivalent tuples [18] (tuples are *value-equivalent* if they agree on all explicit attribute values [15]).

It should be clear that this mapping is defined for all TSQL2 relations. Next, for any TU relation instance, there exists an TSQL2 instance that maps to it, i.e., the mapping is *onto* the set of all TU relations. To see this, pick an arbitrary TU relation. For each set of value-equivalent tuples, form a single TSQL2 tuple with the

---

<sup>6</sup>For simplicity, we also assume that the attribute domains of TU and TSQL2 are the same and that all the domains, including the totally ordered time domain, are finite.

same explicit values and with a timestamp that is the union of the timestamps of the value-equivalent tuples. The result is a legal TSQL2 relation, and that relation maps to the initial TU relation. Finally, there is exactly one TSQL2 relation that maps to any TU relation, i.e., the mapping is 1–1. To see this, observe that two different TSQL2 relations map to different TU relations. In conclusion, the mapping is an isomorphism.

It is worth noting that TU and TSQL2 agree regarding duplicates. A TU relation is defined as a *set* of tuples and thus excludes duplicates. TSQL2 relations do not contain value-equivalent tuples, and a timestamp is a *set* of times. A version of TSQL2 changed to allow value-equivalent tuples with overlapping timestamps would contain more instances than the original TSQL2 and would thus not be temporally ungrouped.

It may also be shown that if  $\{A_{j_1}, A_{j_2}, \dots, A_{j_i}\}$  is a temporal key [19] of a TSQL2 relation then  $\{A_{j_1}, A_{j_2}, \dots, A_{j_i} \parallel \text{TIME}\}$  is a key of the corresponding TU relation.

We have now seen that TSQL2 is temporally ungrouped and thus not temporally grouped. It is then a natural question to ask whether it would be useful to make TSQL2 grouped. It is our contention that it would not be useful to pursue temporal groupedness in the context of TSQL2. One reason for this is that making TSQL2 relations temporally grouped would change the data structures and thus the query language fundamentally. Another reason is that groupedness does not mix well with the restructuring capabilities of the TSQL2 query language. These capabilities, inspired by, e.g., Gadia [13], have proven very useful [16, 26] when formulating queries, but are meaningless in a grouped model. The following example illustrates in what sense temporal groupedness and restructuring do not mix well.

**Example 3** In TSQL2 queries, it is possible to declare and use “tuple” variables that range over groups of tuples. To illustrate this, consider the almost identical TSQL2 version, termed `relationAB` (see also below), of the TU instance in Example 1. The TSQL2 from clause

```
FROM relationAB ( A ) AS relationA
```

declares a “tuple” variable `relationA` that ranges over groups of `relationAB` tuples with the same `A` value. In the where clause, it is then possible to reference the `A` value (using `relationA.A`) and the timestamp (using `VALID(relationA)`), but not the `B` value, of this variable.

Observe that the two instances (in Figure 1) essentially correspond to the last two temporally grouped instances depicted in Example 2. However, restructuring as in TSQL2 (and, e.g., TempSQL [14]) cannot produce the first instance depicted in Example 2. This indicates that temporally groupedness and grouping as in TSQL2 do not coexist harmoniously.  $\square$

A	T
$a_1$	{1}
$a_2$	{2}
$a_3$	{1, 2}

relationA

A	B	T
$a_1$	$b_1$	{1}
$a_3$	$b_2$	{1}
$a_2$	$b_1$	{2}
$a_3$	$b_3$	{2}

relationAB

Figure 1: Grouping Tuples in TSQL2

More generally, it is our contention that temporal ungroupedness versus groupedness are fundamental design decisions to be made when designing a data model.

Ungroupedness and groupedness have relative advantages and weaknesses. Clifford et al. [7] lucidly describe the advantages of groupedness. Others have pointed to potential advantages of ungroupedness. To exemplify, assume that attribute A records employee names and attribute B records department names. With grouping, one must decide at schema definition time whether tuples in the relation are intended to represent departments (i.e., one tuple records all employee names for a department which may change its name over time) or employees (i.e., one tuple records all department names for an employee which may change name over time). Without grouping and assuming that departments and employees do not change names over time, this decision may be deferred to when a query is formulated, which is more flexible.

## 4 Temporal Semi-Completeness

This section first gives refined definitions of *temporal semi-completeness* and *temporal completeness* [5]. The definitions presented here add additional syntactical requirements that were intended in the original definitions, but were not stated explicitly.

These notions reflect a belief that both functionality and syntactical requirements are important when evaluating a data model. Both types of requirements are relative to some chosen non-temporal data model. While the definitions are applicable to any pair of a temporal and a non-temporal data model, they are intended to be applied to pairs of a temporal relational data model and the particular version of the snapshot relational model that the temporal model extends.

The section ends with an evaluation of TSQL2 according to each definition. We emphasize that the contribution is not in the definitions per se, but in their application to TSQL2, yielding new insights into this language.

#### 4.1 Definition

To define temporal semi-completeness, we first introduce the auxiliary notion of a snapshot reducible query. We will use  $r$  and  $r^v$  for denoting a snapshot and a valid-time relation instance, respectively. Similarly,  $db$  and  $db^v$  are sets of snapshot and valid-time relation instances, respectively.

The definition uses a valid-timeslice operator  $\tau_c^{M^v, M}$  (e.g., [27, 5]) which takes as arguments a valid-time relation  $r^v$  (in the data model  $M^v$ ) and a valid-time instant  $c$  and returns a snapshot relation  $r$  (in the data model  $M$ ) containing all tuples valid at time  $c$ . In other words,  $r$  consists of all tuples of  $r^v$  whose valid time includes the time instant  $c$ , but without the valid time. We assume that the valid timeslice preserves duplicates, i.e., if  $r^v$  contains value-equivalent tuples that are valid at time  $c$  then  $\tau_c^{M^v, M}(r^v)$  will contain duplicates. This becomes important later, when we consider SQL-92 relations with duplicates.

**Definition 4 (snapshot reducibility)** [28] Let  $M = (DS, QL)$  be a snapshot relational data model, and let  $M^v = (DS^v, QL^v)$  be a valid-time data model. Also, let  $db^v$  be a database instance in  $DS^v$ . A valid-time query  $q^v$  in  $QL^v$  is *snapshot reducible with respect to* a snapshot query  $q$  in  $QL$  if and only if

$$\forall db^v \forall c (\tau_c^{M^v, M}(q^v(db^v)) = q(\tau_c^{M^v, M}(db^v))). \quad \square$$

Graphically, snapshot reducibility implies that for all  $db^v$  and for all  $c$ , the commutativity diagram shown in Figure 2 must hold.

$$\begin{array}{ccc}
 db^v & \xrightarrow{q^v} & q^v(db^v) \\
 \text{timeslices at } c \downarrow & & \downarrow \text{timeslice at } c \\
 \tau_c^{M^v, M}(db^v) & \xrightarrow{q} & q(\tau_c^{M^v, M}(db^v)) = \tau_c^{M^v, M}(q^v(db^v))
 \end{array}$$

Figure 2: Snapshot Reducibility of Query  $q^v$  With Respect To Query  $q$

Temporal semi-completeness of a temporal data model with respect to a snapshot data model requires first that all relation instances in the snapshot data model can be produced by taking timeslices of some relation instance in the temporal data model. Further, it is required that each query  $q$  in the snapshot model has a counterpart  $q^v$  in the temporal model that is snapshot reducible with respect to it. Observe that  $q^v$  being snapshot reducible with respect to  $q$  poses no syntactical restrictions on  $q^v$ . It is thus possible for  $q^v$  to be quite different from  $q$ , and  $q^v$  might be very involved. This is undesirable, as we would like the temporal model to be a

straight-forward extension of the snapshot model. Consequently, we add to the definition of temporal semi-completeness the syntactical restriction that  $q^v$  and  $q$  be syntactically similar.

**Definition 5 (temporal semi-completeness)** [5] Let  $M = (DS, QL)$  be a snapshot data model, and let  $M^v = (DS^v, QL^v)$  be a valid-time data model. Data model  $M^v$  is *temporally semi-complete with respect to* model  $M$  if and only if all three of the following conditions hold.

1. For every relation  $r$  in  $DS$ , there exists a valid-time relation  $r^v$  in  $DS^v$  and a time instant  $c$  such that  $r = \tau_c^{M^v, M}(r^v)$ .
2. For every query  $q$  in  $QL$ , there exists a query  $q^v$  in  $QL^v$  that is snapshot reducible with respect to  $q$ .
3. There exist two (possibly empty) text strings  $S_1$  and  $S_2$  such that for all pairs  $(q, q^v)$  of queries, where  $q^v$  is snapshot reducible with respect to  $q$ , query  $q^v$  is syntactically identical to  $S_1qS_2$ .  $\square$

Note that the same two strings  $S_1$  and  $S_2$  must apply to all  $(q, q^v)$  pairs. The strings represent particular syntactical constructs in the language  $QL^v$ .

If the valid-time data model treats valid-time relations as a new type of relation, as does TSQL2, it is possible to use the same syntactical constructs (i.e.,  $q^v$  and  $q$  are identical) for querying snapshot and valid-time relations. In this case, the type of a relation determines the meaning of the syntactical construct.

Temporal semi-completeness of a valid-time data model with respect to a snapshot data model guarantees that the temporal model is a straightforward extension of the snapshot model. Temporal semi-completeness is limited in the sense that it covers only those queries in the temporal data model that are snapshot reducible to a query in the snapshot data model. Most often, a temporal data model allows for the formulation of other queries as well.

## 4.2 TSQL2 and Temporal Semi-Completeness with Respect to SQL-92

This section identifies where TSQL2 falls short in fulfilling the requirements of temporal semi-completeness. The two related concepts of value-equivalent tuples and duplicates will prove important in this section. The former concept applies only to temporal relations; the latter applies to both valid-time relations and timeslices of valid-time relations. To illustrate the interrelations among these concepts, consider the valid-time relations depicted in Figure 3.

Relation  $r_1$  contains no duplicates and no value-equivalent tuples. Thus, no timeslices of  $r_1$  will contain duplicates. Relation  $r_2$  contains no duplicates, but it does contain value-equivalent tuples. However, as the timestamps of the value-equivalent tuples are disjoint, no timeslices will contain duplicates. Relation  $r_3$ , like

$r_1$						
<table border="1" style="border-collapse: collapse; width: 100px; height: 60px;"> <thead> <tr><th style="padding: 2px 10px;">A</th><th style="padding: 2px 10px;">T</th></tr> </thead> <tbody> <tr><td style="padding: 2px 10px;"><math>a_1</math></td><td style="padding: 2px 10px;">10–20</td></tr> <tr><td style="padding: 2px 10px;"><math>a_2</math></td><td style="padding: 2px 10px;">15–50</td></tr> </tbody> </table>	A	T	$a_1$	10–20	$a_2$	15–50
A	T					
$a_1$	10–20					
$a_2$	15–50					

$r_2$								
<table border="1" style="border-collapse: collapse; width: 100px; height: 60px;"> <thead> <tr><th style="padding: 2px 10px;">A</th><th style="padding: 2px 10px;">T</th></tr> </thead> <tbody> <tr><td style="padding: 2px 10px;"><math>a_1</math></td><td style="padding: 2px 10px;">10–17</td></tr> <tr><td style="padding: 2px 10px;"><math>a_1</math></td><td style="padding: 2px 10px;">17–20</td></tr> <tr><td style="padding: 2px 10px;"><math>a_2</math></td><td style="padding: 2px 10px;">15–50</td></tr> </tbody> </table>	A	T	$a_1$	10–17	$a_1$	17–20	$a_2$	15–50
A	T							
$a_1$	10–17							
$a_1$	17–20							
$a_2$	15–50							

$r_3$								
<table border="1" style="border-collapse: collapse; width: 100px; height: 60px;"> <thead> <tr><th style="padding: 2px 10px;">A</th><th style="padding: 2px 10px;">T</th></tr> </thead> <tbody> <tr><td style="padding: 2px 10px;"><math>a_1</math></td><td style="padding: 2px 10px;">10–20</td></tr> <tr><td style="padding: 2px 10px;"><math>a_1</math></td><td style="padding: 2px 10px;">15–18</td></tr> <tr><td style="padding: 2px 10px;"><math>a_2</math></td><td style="padding: 2px 10px;">15–50</td></tr> </tbody> </table>	A	T	$a_1$	10–20	$a_1$	15–18	$a_2$	15–50
A	T							
$a_1$	10–20							
$a_1$	15–18							
$a_2$	15–50							

$r_4$								
<table border="1" style="border-collapse: collapse; width: 100px; height: 60px;"> <thead> <tr><th style="padding: 2px 10px;">A</th><th style="padding: 2px 10px;">T</th></tr> </thead> <tbody> <tr><td style="padding: 2px 10px;"><math>a_1</math></td><td style="padding: 2px 10px;">10–20</td></tr> <tr><td style="padding: 2px 10px;"><math>a_1</math></td><td style="padding: 2px 10px;">10–20</td></tr> <tr><td style="padding: 2px 10px;"><math>a_2</math></td><td style="padding: 2px 10px;">15–50</td></tr> </tbody> </table>	A	T	$a_1$	10–20	$a_1$	10–20	$a_2$	15–50
A	T							
$a_1$	10–20							
$a_1$	10–20							
$a_2$	15–50							

Figure 3: Illustration of Value-equivalent Tuples and Duplicates

$r_2$ , contains no duplicates, but contains value-equivalent tuples. Unlike in  $r_2$ , the timestamps of the value-equivalent tuples are not disjoint and thus there are timeslices of  $r_3$  that contain duplicates. Finally, relation  $r_4$  contains duplicates and thus non-disjoint value-equivalent tuples, leading again to timeslices with duplicates. Note that allowing value-equivalent tuples does not necessarily yield duplicates in timeslices. However, if we want to have duplicates in timeslices, we must allow (non-disjoint) value-equivalent tuples.

### Lack of Duplicates in TSQL2

One reason why TSQL2 is not temporally semi-complete with respect to SQL-92 is that SQL-92 relations that contain duplicates have no counterparts in TSQL2 where relations with value-equivalent tuples (and thus duplicates, either in a timeslice, or in the temporal relation itself) are not allowed. Definition 5 requires that for every SQL-92 relation  $r$ , there must exist a TSQL2 relation  $r^v$  and a time instant  $c$  such that  $\tau_c^{TSQL2, SQL-92}(r^v) = r$ . However, it is not possible to find an  $r^v$  in TSQL2 for  $r$ 's in SQL-92 that contain duplicates. An example illustrates this.

**Example 4** Let salary relation, *salary\_entry*, be given that records (current) monthly incomes of persons. Assume that the person Tom has three incomes because he has three jobs. In two jobs, he makes 1200, and in one he makes 800. This can be represented in SQL-92 as follows.

*salary\_entry*

Name	Amount
Tom	1200
Tom	1200
Tom	800

No timeslice of a TSQL2 relation can yield this relation. The following is a reasonable attempt at adding valid time to the SQL-92 relation to obtain a TSQL2 relation.



If TSQL2 is to be semi-complete with respect to SQL-92, there must be valid-time queries in TSQL2 that are snapshot reducible with respect to the two queries above and are similar to them. Indeed, the default valid clause of TSQL2 was designed to make those two valid-time queries be identical to the two queries above. The valid-time queries are given below, with the implicit default valid clauses shown.

<pre> SELECT r5.a VALID INTER- SECT (VALID(r5), VALID(r6)) FROM r5, r6 WHERE r5.a=r6.a </pre>	<pre> SELECT r5.a VALID VALID(r5) FROM r5 WHERE EXISTS (SELECT * VALID VALID(r6) FROM r6 WHERE r5.a=r6.a) </pre>
---	--

The query to the left behaves as expected. The valid clause states that the valid time of a result tuple is the intersection of the valid times of the argument tuples from  $r_5$  and  $r_6$ . This means that the left-hand-side valid-time query is snapshot reducible with respect to the left-hand-side snapshot query. The result ( $result_1$ ) of the query for two sample instances of  $r_5$  and  $r_6$  is shown in Figure 4. The situation gets

$r_5$	$r_6$	$result_1$	$result_2$																
<table border="1" style="border-collapse: collapse; text-align: center;"> <tr><td style="padding: 2px 10px;">A</td><td style="padding: 2px 10px;">T</td></tr> <tr><td style="padding: 2px 10px;"><math>a_1</math></td><td style="padding: 2px 10px;">5-9</td></tr> </table>	A	T	$a_1$	5-9	<table border="1" style="border-collapse: collapse; text-align: center;"> <tr><td style="padding: 2px 10px;">A</td><td style="padding: 2px 10px;">T</td></tr> <tr><td style="padding: 2px 10px;"><math>a_1</math></td><td style="padding: 2px 10px;">7-10</td></tr> </table>	A	T	$a_1$	7-10	<table border="1" style="border-collapse: collapse; text-align: center;"> <tr><td style="padding: 2px 10px;">A</td><td style="padding: 2px 10px;">T</td></tr> <tr><td style="padding: 2px 10px;"><math>a_1</math></td><td style="padding: 2px 10px;">7-9</td></tr> </table>	A	T	$a_1$	7-9	<table border="1" style="border-collapse: collapse; text-align: center;"> <tr><td style="padding: 2px 10px;">A</td><td style="padding: 2px 10px;">T</td></tr> <tr><td style="padding: 2px 10px;"><math>a_1</math></td><td style="padding: 2px 10px;">5-9</td></tr> </table>	A	T	$a_1$	5-9
A	T																		
$a_1$	5-9																		
A	T																		
$a_1$	7-10																		
A	T																		
$a_1$	7-9																		
A	T																		
$a_1$	5-9																		

Figure 4: Computing a Valid-time Join Without or With a Subquery

more complicated when we consider the query to the right. The outermost valid clause implies that the valid time of a result tuples is equivalent to the valid time of the argument tuple from  $r_5$  (see  $result_2$  in Figure 4 for an example). This means that the right-hand-side valid-time query is not snapshot reducible with respect to the right-hand-side snapshot query. TSQL2 thus lacks a valid-time query that is snapshot reducible with respect to and is a simple syntactic extension of the right-hand-side snapshot query. Consequently, TSQL2 is not temporally semi-complete with respect to SQL-92.

## Summary

We have identified two reasons why TSQL2 is not temporally semi-complete with respect to SQL-92. The first is that, while duplicates are allowed in SQL-92, value-equivalent tuples are not allowed in TSQL2. The second reason is that the valid clause in TSQL2 is not sufficiently powerful to ensure that all SQL-92 queries



have similar, snapshot reducible counterparts in TSQL2. We showed this for nested queries. We conjecture that there are also some problems with aggregation, grouping, and ordering.

## 5 Temporal Completeness

Temporal semi-completeness poses useful restrictions on temporal data models. However, temporal semi-completeness poses restrictions on only a subset of the queries that are generally expressible in temporal data models. For example, it does not cover queries that retrieve information concerning relationships between perceived states of the world at *different* points in time. Furthermore, temporal semi-completeness does not say anything about the format of valid time. Both aspects are accounted for by the notion of a temporally complete data model.

### 5.1 Definition

**Definition 6 (temporal completeness)** [5] A valid-time data model  $M^v = (DS^v, QL^v)$  is *temporally complete with respect to* a snapshot data model  $M = (DS, QL)$  if and only if all five of the following conditions hold.

1.  $M^v$  is temporally semi-complete with respect to  $M$ .
2. For every snapshot reducible query  $q^v$  in  $QL^v$ , it is possible to override snapshot reducibility, either by dropping the syntactic extensions that enforce snapshot reducibility (c.f., Definition 5) or by modifying  $q^v$  syntactically to  $S_1qS_2$ , where  $S_1$  and  $S_2$  are (possibly empty) text strings that depend on  $QL^v$  but not on  $q^v$ . Overriding snapshot reducibility means to evaluate a query without interpreting valid times.
3. The name of a valid-time relation within a statement can be syntactically substituted (perhaps with other syntactic modifications and additions, such as parentheses) with a query  $q^v$  in  $QL^v$  that defines the respective valid-time relation without changing the semantics of the statement. The syntactic modifications must depend on  $QL^v$  only, not on  $q^v$ .
4. Allen's temporal relationships [2] can be used between (a) temporal attributes of stored valid-time relations (i.e., valid time attributes and explicit temporal attributes), (b) implicitly computed valid times associated with temporally semi-complete (sub)queries, and (c) temporal constants.
5. It is possible to retrieve and constrain (a) maximal continuous valid-time periods and (b) valid times as specified by the user. □

First, we require that temporally complete languages are temporally semi-complete. This accounts for queries that can be answered by examining (sequences of) snapshots. Overriding snapshot reducibility accounts for a fundamental principle in

databases, namely that a query should treat the elements of a database as uninterpreted objects [8, p.158]. Section 5.2 provides an example that illustrates this. The third condition ensures that the syntactic construct that is used to enforce snapshot reducibility can be applied not only to whole queries, but also to subqueries. In other words, a temporally complete query may consist of several temporally semi-complete queries. Allen's operators are necessary to state arbitrary temporal relationships. (They were proven to exhaustively describe the relationships between periods [2]. However, other, equally expressive operators are possible as well.) Note that there are different timestamps that are of interest in a temporal database: temporal attributes of base relations, implicitly computed valid times, and temporal constants. We require that all of them can be used together as operands to Allen's operators. Finally the database system has to support maximal continuous periods and valid times as specified by the user. Both kinds of timestamps have been shown necessary in answering temporal queries [29]. It must be possible to retrieve and constrain (i.e., use as operands of functions and predicates) either kind of timestamp.

We emphasize that the notions of temporal semi-completeness and temporal completeness go beyond approaches that define the completeness in terms of an algebra (i.e., by requiring a temporal language to have the same expressive power as an algebra). For example, temporal semi-completeness (and thus temporal completeness) may, depending on the language it is with respect to, cover aggregates, grouping, null values, ordering, and duplicates.

## 5.2 TSQL2 and Temporal Completeness with Respect to SQL-92

In order to qualify for temporal completeness, a temporal query language must fulfill the five requirements listed in Definition 6. We first must modify TSQL2 to make it temporally semi-complete. To ensure temporal completeness, it must in addition be possible to override snapshot reducibility. The valid clause in TSQL2 is intended for this purpose, but as its scope does not extend to set operations such as EXCEPT and UNION, the clause cannot override snapshot reducibility for them, either.

The third condition is that a temporal language must allow a valid-time query to appear in a larger query everywhere a valid-time relation name may appear, so that if the valid-time query computes the named relation, the two forms of the larger queries compute the same result. This feature is provided by table expressions, which were introduced in SQL-92 [22, p.178] and carried over to TSQL2.

The fourth requirement is satisfied by the where clause which is enhanced with temporal predicates that have the same expressive power as Allen's predicates. Temporal attributes of base relations, implicitly computed valid times (e.g., valid times computed by table expressions), and temporal constants can be used as

operands to these predicates.

Finally, a temporal language must support maximal continuous valid-time periods and valid times as specified by the user. In the second subsection we will see that TSQL2 ignores the user-specified valid time format.

### Overriding Snapshot Reducibility

In TSQL2, the `valid` clause can be used to override snapshot reducibility—if no `valid` clause is specified, the semantics defaults to valid-time intersection. However, the scope of the `valid` clause does not include set operations and, therefore, it is not possible to override valid-time semantics associated with these operations. As an example, suppose the valid-time relations  $r_5$  and  $r_6$  of Figure 4. In TSQL2, it is not possible to use `EXCEPT` to retrieve all tuples in  $r_5$  that are not in  $r_6$ . Snapshot reducibility is hard-wired into `EXCEPT`, which means that TSQL2 always yields  $result_3$  rather than  $result_4$ , see next.

A	T
$a_1$	5–7

A	T
$a_1$	5–9

### Beyond Coalescing

The last point of Definition 6 requires that a temporal query language be able to retrieve and constrain (a) maximal continuous valid-time periods and (b) valid times as specified by the user. First, TSQL2 falls short in doing this at the outermost level of queries. The results of queries are always coalesced relations, i.e., relations where value-equivalent tuples are eliminated by combining their valid timestamps. This also holds for an individual select statement which may be part of a larger query.

Second, TSQL2 relations are constrained to contain coalesced tuples. To exemplify why this may be a problem, consider relations  $r_1$  and  $r_2$  of Figure 3. We may envision that it is significant to a user whether the explicit attribute value  $a$  is associated with one single timestamp, 10–20, or is associated with two separate timestamps, 10–17 and 17–20. These two relations may mean different things to a user. However,  $r_2$  is not a legal TSQL2 relation, and if the user inserts tuples  $\langle a_1, 10-17 \rangle$  and  $\langle a_1, 17-20 \rangle$  into a TSQL2 relation, the tuples will be coalesced, and relation  $r_1$  will be the result. Put differently, TSQL2 does not consider the difference between  $r_1$  and  $r_2$  (and  $r_3$  and  $r_4$ ) important and thus only admits coalesced relations.

Temporal completeness requires that TSQL2 respects the valid times as provided by the user. If the user provides two intervals for attribute value  $a$ , TSQL2 must maintain those two intervals and cannot simply coalesce them. Clearly, this

matters for queries. For example, the query “Does there exist an entry with a valid time identical to 10–17” should return “yes” if applied to  $r_2$  (because the user has inserted a tuple with this valid time into  $r_2$ ) and “no” if applied to  $r_1$  (because the user has not inserted a tuple with this valid time into  $r_1$ ).

Currently, TSQL2 is a point-based [9], or a snapshot-equivalence preserving [20], temporal query language that uses time intervals at the representational level to achieve a reasonable performance. Changing TSQL2 to respect the valid times as specified by the users represents a substantial conceptual change to TSQL2. It may be argued that admitting uncoalesced relations represents a complication, but it also adds to its expressiveness. With implicit coalescing, users do not have to be concerned with the valid times, but they also cannot associate special semantics with valid times (c.f., Section 4.2).

## Summary

Apart from not being temporally semi-complete, two aspects prevent TSQL2 from being temporally complete. First, it is not possible to override the temporal semantics of set operations. Second, implicit coalescing prevents TSQL2 from respecting valid times as provided by the users.

## 6 Summary and Future Research

This paper has evaluated the consensus temporal query language TSQL2 using existing notions of completeness, some of which were further formalized in the paper.

In consistency with its design goals, TSQL2 was shown to be upwards compatible with SQL-92 and thus to be relationally complete. TSQL2 was also characterized as temporally ungrouped and not temporally grouped. The evaluation of the temporal semi-completeness of TSQL2 with respect to SQL-92 pointed to two important deficiencies: not all SQL-92 relations can be produced taking timeslices of TSQL2 temporal relations, and not all SQL-92 queries have a similar temporal counterpart in TSQL2. Without these deficiencies, TSQL2 would be a “cleaner” extension of SQL-92. The evaluation of temporal completeness of TSQL2 with respect to SQL-92 pointed to two additional problems: with set operations in TSQL2 queries, it is not possible to freely control the valid timestamps of result tuples, and TSQL2 does not respect the valid timestamps of tuples as entered by the users (because value-equivalent tuples are coalesced).

As future research, it would be interesting to use additional completeness notions in the evaluation of TSQL2. Also, a comparative study of completeness notions for temporal databases that sheds light on their interrelations and practical implications, and perhaps leads to new completeness notions, would be worthwhile.

## 7 Acknowledgments

This work was conducted while the first author visited the University of Arizona. He was supported in part by the Swiss NSF. The second author was supported in part by the Danish Natural Science Research Council through grants 11-1089-1, 11-0061-1, and 9400911. The third author was supported in part by NSF grant IRI-9302244.

We thank John Baer, Jan Chomicki, and Charles Kline for interesting discussions and for sharing their knowledge of temporal databases.

## References

- [1] P. Atzeni and P. P. Chen. Completeness of Query Languages for the Entity-Relationship Model. In P. P. Chen, editor, *Proceedings of the Second International Conference on Entity-Relationship Approach*, pages 111-123, October 1981.
- [2] J. F. Allen. Maintaining Knowledge about Temporal Intervals. *Communications of the ACM*, 16(11):832-843, 1983.
- [3] A. V. Aho and J. D. Ullman. Universality of data retrieval languages. In *Proceedings of the 6th ACM Symposium on Principles of Programming Languages*, pages 110-117, January 1979.
- [4] F. Bancilhon. On the Completeness of Query Languages for Relational Databases. In *Proceedings of the 7th Symposium on Mathematical Foundations of Computer Science*. Lecture Notes in Computer Science, Springer Verlag, September 1978.
- [5] M. Böhlen and R. Marti. On the Completeness of Temporal Database Query Languages. *Proceedings of the First International Conference on Temporal Logic*, pages 283-300, July 1994.
- [6] M. Böhlen. *The Temporal Deductive Database System ChronoLog*. Ph.D. thesis, Departement Informatik, ETH Zürich, 1994.
- [7] J. Clifford, A. Croker, and A. Tuzhilin. On Completeness of Historical Relational Query Languages. *ACM Transactions on Database Systems*, 19(1):64-116, March 1994.
- [8] A. K. Chandra and D. Harel. Computable Queries for Relational Data Bases. *Journal of Computer and System Sciences*, 21(2):156-178, October 1980.
- [9] J. Chomicki. Temporal Query Languages: a Survey. *Proceedings of the First International Conference on Temporal Logic*, pages 506-534, July 1994.
- [10] E. F. Codd. A Relational Model of Data for Large Shared Data Banks. *Communications of the ACM*, 13(6):377-387, June 1970.

- [11] E. F. Codd. Relational Completeness of Data Base Sublanguages. *Courant Computer Symposia Series*, 6:65–98, 1972.
- [12] C. J. Date. *Relational Database Writings 1991–1994*. Addison-Wesley Publishing Company, 1995.
- [13] S. K. Gadia. A Homogeneous Relational Model and Query Languages for Temporal Databases. *ACM Transactions on Database Systems*, 13(4):418–448, December 1988.
- [14] S. K. Gadia and G. Bhargava. SQL-like Seamless Query of Temporal Data. In R. T. Snodgrass, editor, *Proceedings of the International Workshop on an Infrastructure for Temporal Databases*, Arlington, Texas, June 1993.
- [15] C. S. Jensen, J. Clifford, R. Elmasri, S. K. Gadia, P. Hayes, and S. Jajodia (eds.). A Glossary of Temporal Database Concepts. *SIGMOD RECORD*, 23(1):52–64, March 1994.
- [16] C. S. Jensen. A Consensus Test Suite of Temporal Database Queries. Technical Report R 93-2034, Department of Mathematics and Computer, Institute for Electronic Systems, Fredrik Bajers Vej 7, DK 9220 Aalborg, Denmark, November 1993.
- [17] G. Jaeschke and H.-J. Schek. Remarks on the Algebra of Non First Normal Form Relations. In *Proceedings of the ACM Symposium on Principles of Database Systems*, pages 124–138, March 1982.
- [18] C. S. Jensen and R. T. Snodgrass. Proposal of a Data Model for the Temporal Structured Query Language. Technical Report 37, University of Arizona, TempIS Technical Report, July 1992.
- [19] C. S. Jensen, M. D. Soo, and R. T. Snodgrass. Extending Normal Forms to Temporal Relations. Technical Report TR 92-17, University of Arizona, July 1992.
- [20] C. S. Jensen, M. D. Soo, and R. T. Snodgrass. Unifying Temporal Models via a Conceptual Model. *Information Systems*, 19(7):513–547, 1994.
- [21] L. E. McKenzie and R. T. Snodgrass. Evaluation of Relational Algebras Incorporating the Time Dimension in Databases. *ACM Computing Surveys*, 23(4):501–543, 1991.
- [22] J. Melton and A. R. Simon. *Understanding the new SQL: A Complete Guide*. Morgan Kaufmann Publishers, 1993.
- [23] P. O’Neil. *Database: Principles, Programming, and Performance*. Morgan Kaufmann, San Francisco, 1994.
- [24] G. Özsoyoğlu and R. T. Snodgrass. Temporal and Real-Time Databases: A Survey. *IEEE Transactions on Knowledge and Data Engineering*, 7(4), August 1995.

- [25] R. T. Snodgrass, I. Ahn, G. Ariav, D. Batory, J. Clifford, C. E. Dyreson, R. Elmasri, F. Grandi, C. S. Jensen, W. Käfer, N. Kline, K. Kulkarni, T. Y. C. Leung, N. Lorentzos, J. F. Roddick, A. Segev, M. D. Soo, and S. M. Sripada. TSQL2 Language Specification. *SIGMOD RECORD*, 23(1):65–86, March 1994.
- [26] R. T. Snodgrass, I. Ahn, G. Ariav, P. Bayer, J. Clifford, C. Dyreson, F. Grandi, L. Hermosilla, C. S. Jensen, W. Käfer, N. Kline, T. Y. C. Leung, N. Lorentzos, Y. Mitsopoulos, J. F. Roddick, M. Soo, and S. M. Sripada. An Evaluation of TSQL2. In *TSQL2 Commentary*. The TSQL2 Language Design Committee, September 1994.
- [27] B. Schueler. Update reconsidered. In G. M. Nijssen (ed.), *Architecture and Models in Data Base Management Systems*. North Holland Publishing Co., 1977.
- [28] R. T. Snodgrass. The Temporal Query Language TQuel. *ACM Transactions on Database Systems*, 12(2):247–298, 1987.
- [29] S. M. Sripada. *Temporal Reasoning in Deductive Databases*. Ph.D. thesis, Imperial College of Science and Technology, University of London, 1991.
- [30] A. Tuzhilin and J. Clifford. A Temporal Relational Algebra as a Basis for Temporal Relational Completeness. In D. McLeod, R. Sacks-Davis, and H.-J. Schek, editors, *Proceedings of the International Conference on Very Large Databases*, pages 13–23, August 1990.