

Notions of Upward Compatibility of Temporal Query Languages

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

April 2, 1997

TR-6

A TIMECENTER Technical Report

| | |
|---------------------|--|
| Title | Notions of Upward Compatibility of Temporal Query Languages |
| | Copyright © 1997 John Bair, Michael H. Böhlen, Christian S. Jensen, and Richard T. Snodgrass. All rights reserved. |
| Author(s) | John Bair, Michael H. Böhlen, Christian S. Jensen, and Richard T. Snodgrass |
| Publication History | October 1996. Manuscript. April 1997. A TIMECENTER Technical Report. |

TIMECENTER Participants

Aalborg University, Denmark

Christian S. Jensen (codirector)

Michael H. Böhlen

Renato Busatto

Heidi Gregersen

Kristian Torp

University of Arizona, USA

Richard T. Snodgrass (codirector)

Anindya Datta

Individual participants

Curtis E. Dyreson, James Cook University, Australia

Kwang W. Nam, Chungbuk National University, Korea

Keun H. Ryu, Chungbuk National University, Korea

Michael D. Soo, University of South Florida, USA

Andreas Steiner, ETH Zurich, Switzerland

Vassilis Tsotras, Polytechnic University, New York, USA

Jef Wijsen, Vrije Universiteit Brussel, Belgium

Any software made available via TIMECENTER is provided “as is” and without any express or implied warranties, including, without limitation, the implied warranty of merchantability and fitness for a particular purpose.

The TIMECENTER icon on the cover combines two “arrows.” These “arrows” are letters in the so-called *Rune* alphabet used one millennium ago by the Vikings, as well as by their predecessors and successors. The Rune alphabet (second phase) has 16 letters. They all have angular shapes and lack horizontal lines because the primary storage medium was wood. However, runes may also be found on jewelry, tools, and weapons. Runes were perceived by many as having magic, hidden powers.

The two Rune arrows in the icon denote “T” and “C,” respectively.

Abstract

Migrating applications from conventional to temporal database management technology has received scant mention in the research literature. This paper formally defines three increasingly restrictive notions of *upward compatibility* which capture properties of a temporal SQL with respect to conventional SQL that, when satisfied, provide for a smooth migration of legacy applications to a temporal system. The notions of upward compatibility dictate the semantics of conventional SQL statements and constrain the semantics of extensions to these statements. The paper evaluates the seven extant temporal extensions to SQL, all of which are shown to complicate migration through design decisions that violate one or more of these notions. We then outline how SQL-92 can be systematically extended to become a temporal query language that satisfies all three notions.

1 Introduction

A wide range of database applications manage time-varying information. These include financial applications such as portfolio management, accounting, and banking; record-keeping applications, including personnel, medical-record, and inventory; and travel applications such as airline, train, and hotel reservations and schedule management. In fact, it is difficult to identify a database application that does *not* involve time-varying data.

Currently, such applications typically use conventional relational systems. However, research in temporal data models and query languages [TCG⁺93, OS95, Sno95] clearly demonstrates that applications that manage temporal data may benefit substantially from built-in temporal support in the database management system (DBMS). The potential benefits from such support are several. Application code is substantially simplified. Due to faster development of code, which is also easier to comprehend and thus maintain, higher programmer productivity results. With built-in support, more data processing may be left to the DBMS, leading to much better performance.

There is, however, still a chasm between the approaches now used for developing such applications and the new approaches that have been proposed by the temporal database research community. While 1200 papers on temporal databases have appeared, over 300 during the last two years alone [TK96], legacy systems and the migration of such systems to new technologies has been almost totally overlooked.

This paper considers *upward compatibility*, which has been claimed to offer several potential advantages (this is related to the notion of ‘seamless’, as in “the transition from classical databases to [a temporal] model should be conceptually and literally seamless.” [GN93, p. 51]). In the context of migrating from a conventional DBMS to a temporal DBMS, upward compatibility offers an evolutionary means of introducing new technology. It provides a business enterprise with an upgrade path that preserves its investment in legacy databases. Implementers can incrementally build new features on top of existing products, by gradually learning and incorporating new language elements into their applications.

We assume that the DBMS interface is captured in a data model and thus talk about the migration of application code using an existing data model to using a new data model. We examine what it means for a temporal data model and query language to be upwardly compatible with a conventional data model, such as SQL. While the term has been used informally, we could find no formal definition (the same holds for the term ‘seamless’). This paper examines the issues behind this intuitive idea, and formalizes several increasingly restrictive notions of upward compatibility, specifically, syntactic upward compatibility, upward compatibility, and temporal upward compatibility. We evaluate the seven extant temporal extensions to SQL, including two designed by us. We show that all seven violate one or more of these useful properties. Finally, we demonstrate how SQL-92 can be extended to a temporal data model while simultaneously satisfying all three notions of upward compatibility.

2 Transitioning from Legacy Systems

We adopt the convention that a data model consists of three components, namely a set of data structures, a set of constraints on those data structures, and a language for updating and querying the data structures [TL82]. In this paper we emphasize the data structures and the data manipulation language. As we progress, it should be clear that the definitions and discussions of this section also apply to integrity constraints, although for simplicity we will not address these explicitly. 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, schemas, and associated instances, expressible by M , and QL is the set of all query and modification statements in M that may be applied to some database in DS . We use db to denote a database; a statement is denoted by s and is either a query q or a modification m (e.g., in SQL-92, any INSERT, DELETE, or UPDATE statement).

As the existing model is given, the focus is on formulating requirements to the new data model. The definitions are conceptually applicable to the transition from any data model to a new data model. However, we have found it convenient to assume that the transition is from a non-temporal to a temporal data model, specifically from the SQL-92 standard [MS93] to (some) Temporal SQL.

2.1 Upward Compatibility

Perhaps the most important concern in ensuring a smooth transition of application code from an existing data model to a new data model is to guarantee that all application code without modification will work with the new system exactly the same as with the existing system. The next two definitions are intended to capture what is needed for that to be possible.

We define a data model to be *syntactic upward compatible* with another data model if all the data structures and legal query language statements of the latter model are contained in the former model.

DEFINITION: Let $M_1 = (DS_1, QL_1)$ and $M_2 = (DS_2, QL_2)$ be two data models. Model M_1 is *syntactically upward compatible* with model M_2 if

- $\forall db_2 \in DS_2 (db_2 \in DS_1)$ and
- $\forall s_2 \in QL_2 (s_2 \in QL_1)$. □

Note that this relationship between the two data models is asymmetric, thus providing credence to the adjective ‘upward’.

Next, for a data model to be *upward compatible* with another data model, we add the requirement that all statements expressible in the existing language must evaluate to the same result in both models.

For a query language expression s and an associated database db , both legal elements of QL and DS of data model $M = (DS, QL)$, define $\langle\langle s(db) \rangle\rangle_M$ as the result of applying s to db in data model M . With this notation, we can precisely describe the requirements to a new model that guarantee uninterrupted operation of all application code.

DEFINITION: Let $M_1 = (DS_1, QL_1)$ and $M_2 = (DS_2, QL_2)$ be two data models. Model M_1 is *upward compatible* with model M_2 if

- M_1 is syntactically upward compatible with M_2 , and
- $\forall db_2 \in DS_2 (\forall s_2 \in QL_2 (\langle\langle s_2(db_2) \rangle\rangle_{M_2} = \langle\langle s_2(db_2) \rangle\rangle_{M_1}))$. □

Syntactic upward compatibility implies that all existing databases and query language statements in the old system are also legal in the new system. The second condition guarantees that all existing statements 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.

Figure 1 illustrates the relationship between Temporal SQL and SQL-92. In the figure, a conventional table¹ is denoted with a rectangle. The current state of this table is the rectangle in the upper-right corner. Whenever a modification is made to this table, the previous state is discarded; hence, at any time only the current state is available. The discarded prior states are denoted with dashed rectangles; the right-pointing arrows denote the modifications that took the table from one state to the next.

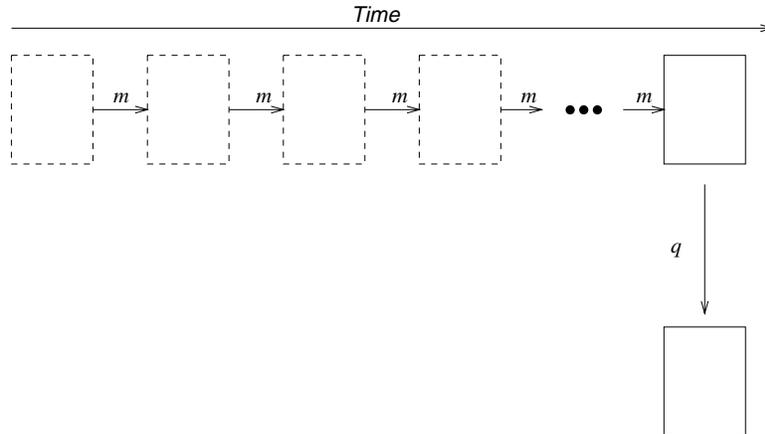


Figure 1: Upward Compatible Queries

When a query q is applied to the current state of a table, a resulting table is computed, shown as the rectangle in the bottom right corner. While this figure only concerns queries over single tables, the extension to queries over multiple tables is clear.

Upward compatibility states that (1) all instances of tables in SQL-92 are instances of tables in Temporal SQL, (2) all SQL-92 modifications to tables in SQL-92 result in the same tables when the modifications are evaluated according to Temporal SQL semantics, and (3) all SQL-92 queries result in the same tables when the queries are evaluated according to Temporal SQL.

By requiring that Temporal SQL is a strict superset (i.e., only *adding* constructs and semantics), it is relatively easy to ensure that Temporal SQL is upward compatible with SQL-92.

EXAMPLE: A company wishes to computerize its personnel records, so it creates two tables, an `Employee` table and a `Salary` table. Every employee must have a salary. These tables are populated. A high salary view identifies those employees with a monthly salary greater than \$3500. Then employee Therese is given a 10% raise. Since the salary table is a snapshot table, Therese’s previous salary is lost. These statements can be easily expressed in SQL-92. Similarly, upward compatibility ensures that these queries and modifications have precisely the same semantics in Temporal SQL. \square

2.2 Temporal Upward Compatibility

The above minimal requirements are essential to ensure a smooth transition to a new temporal data model, but they do not address all aspects of migration. Specifically, assume that an existing data model has been

¹In this paper, we use the terminology adopted in SQL: table, row, and column, rather than the terminology introduced by Codd [Cod70]: relation, tuple, and attribute.

replaced with a new temporal model. No application code has been modified, and all tables are thus snapshot tables. Upward compatibility ensures that all applications work as before, under the new temporal model.

Now, an existing or new application needs support for the temporal dimension of the data in one or more of the existing tables. This is best achieved by changing the snapshot table to become a temporal table (e.g., by using a statement of Temporal SQL).

It is undesirable to be forced to change the application code that accesses the snapshot table that is replaced by a temporal table. We formulate a requirement that states that the existing applications on snapshot tables will continue to work with no changes in functionality when the tables they access are altered to become temporal tables. Specifically, *temporal upward compatibility* requires that each query will return the same result on an associated snapshot database as on the temporal counterpart of the database. Further, this property is not affected by modifications to those temporal tables. The precise definition is given next and is explained in the following.

DEFINITION: Let $M_T = (DS_T, QL_T)$ and $M_S = (DS_S, QL_S)$ be temporal and snapshot data models, respectively. Also, let \mathcal{T} be an operator that changes the type of a snapshot table to the temporal table with the same explicit columns. Next, let m_1, m_2, \dots, m_n ($n \geq 0$) denote modification operations. With these definitions, model M_T is *temporal upward compatible* with model M_S if

- M_T is upward compatible with M_S ,
- $\forall db_S \in DS_S$ ($\mathcal{T}(db_S) \in DS_T$), and
- $\forall db_S \in DS_S$ ($\forall m_1, \dots, m_n$ ($n \geq 0$) ($\forall q_S \in QL_S$ ($\langle\langle q_S(m_n(m_{n-1}(\dots(m_1(db_S)\dots)))\rangle\rangle_{M_S} = (\langle\langle q_S(m_n(m_{n-1}(\dots(m_1(\mathcal{T}(db_S))))\rangle\rangle_{M_T})\rangle\rangle)$)). \square

Assume that, when moving to the new system, some of the existing (snapshot) tables are transformed into temporal tables without changing the existing set of (explicit) columns. This transformation is denoted by \mathcal{T} in the definition. Then the same sequence of modification statements, denoted by the m_i in the definition, is applied to the snapshot and the temporal databases. Next, consider any query in the snapshot model. Such queries are also allowed in the temporal model, due to upward compatibility being required. The definition states that any such query evaluated on the resulting temporal database, using the semantics of the temporal query language, yields the same result as when evaluated on the resulting snapshot database, now using the semantics of the snapshot query language.

Temporal upward compatibility is illustrated in Figure 2. When temporal support is added to a table, the history is preserved, and modifications over time are retained. In this figure, the rightmost dashed state was the current state when the table was made temporal. All subsequent modifications, denoted by the arrows, result in states that are retained, and thus are solid rectangles. Temporal upward compatibility ensures that the states will have identical contents to those states resulting from modifications of the snapshot table.

The query q is an SQL-92 query. Due to temporal upward compatibility the semantics of this query must not change if it is applied to a temporal table. Hence, the query only applies to the current state, and a snapshot table results.

EXAMPLE: We make both the employee and salary tables temporal. This means that all information currently in the tables is valid from today on. We add an employee. This modification to the two tables, consisting of two SQL-92 INSERT statements, respects temporal upward compatibility. Queries and views on these newly-temporal tables work exactly as before. The SQL-92 query to determine who manages the high-salaried employees returns the current information. Constraints and assertions also work exactly as before, applying to the current state and checked on database modification. \square

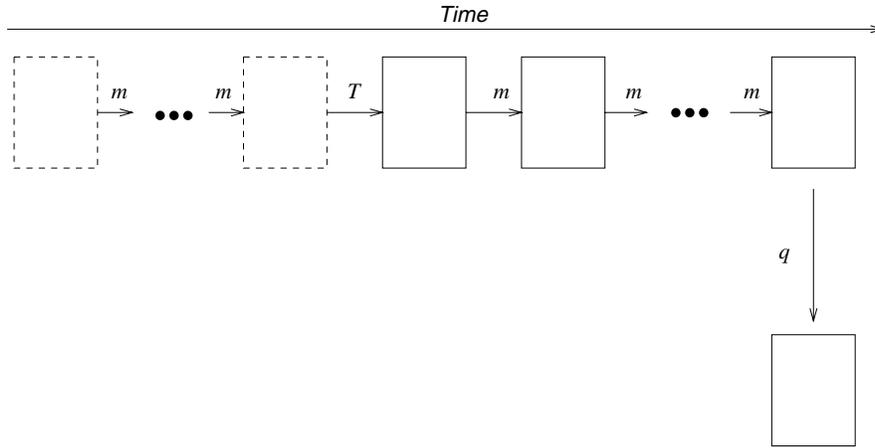


Figure 2: Temporal Upward Compatibility

There is one unfortunate ramification to the above definition. Any extension that adds constructs involving new reserved keywords will violate upward compatibility, as well as temporal upward compatibility. The reason is that the user may have previously used that keyword as an identifier. Query language statements that use the keyword as an identifier will, in the extension, be disallowed.

Reserved words are added in all temporal query languages. This phenomenon also holds for non-temporal query languages. SQL-89 defined some 115 reserved words; SQL-92 added 112 reserved words, and the draft standard SQL3 adds another 97 reserved words.

To avoid being overly restrictive, we consider upward compatibility and temporal upward compatibility to be satisfied even when reserved words are added, as long as the semantics of all statements that do not use the new reserved words is retained in the temporal model.

3 Temporal Database Management Using SQL-92

As an initial application of these notions, we first consider an approach employed frequently to implement a temporal application: emulating a time-varying table with a conventional table. As we will see, this approach does not ensure temporal upward compatibility, leading to a number of difficulties.

The underlying model will be SQL-92, that is, QL_2 is the set of SQL-92 queries and modifications. Let us return to the `Employee` table, which has three columns, `Name`, `Manager`, and `Dept`. In SQL-92, one can require that all managers be employees, by stating that `Manager` is a foreign key for `Employee.Name`. We can easily express queries such as “List those employees who are not managers,” as well as modifications, such as “Change the manager of the tools department to Bob.”

To store historical information, we wish to emulate time-varying information, and so we will use the same model, SQL-92, that is, QL_T will be the set of SQL-92 queries and modifications. We also need an operator that changes the type of a snapshot table to a ‘temporal’ table. We will define \mathcal{T} to be the following SQL-92 schema modification statements.

```
ALTER TABLE Employee ADD COLUMN Start DATE
ALTER TABLE Employee ADD COLUMN Stop DATE
```

The \mathcal{T} operator must also initialize the value of the `Start` column to be the value `CURRENT_DATE` and the value of the `Stop` column to be the value `DATE '9999-12-31'`, the largest `DATE` value. This transforms the `Employee` table into a ‘temporal’ table (in the data model M_T , which is SQL-92).

Model M_T is certainly upward compatible with M_S , as all databases in $M_T = M_S = \text{SQL-92}$. Interestingly, though, SQL-92 , along with the transformation operator just defined, is *not* temporally upward compatible with itself. As but a simple example, let q_S be the query “SELECT * FROM Employee”. It is certainly not the case that $\llbracket q_S(db_S) \rrbracket_{\text{SQL-92}} = \llbracket q_S(\mathcal{T}(db_S)) \rrbracket_{\text{SQL-92}}$. Even the schemas do not match: the schema for the result of $\llbracket q_S(db_S) \rrbracket_{\text{SQL-92}}$ has three columns, Name, Manager, and Dept, while the schema for the result of $\llbracket q_S(\mathcal{T}(db_S)) \rrbracket_{\text{SQL-92}}$ has five columns, including Start and Stop.

This violation of temporal upward compatibility has important practical ramifications. Assume that we have a 50,000-line application that manages the Employee table and other tables in a personnel database, allowing employees to be added and dropped, and the information about employees to be modified and queried in various ways. When this table is extended to store time-varying information, via the transformation \mathcal{T} discussed above, many portions of this application break.

- The constraint that all managers are employees can no longer be expressed via SQL-92 ’s foreign key constraint, which fails to take time into account. Instead, this constraint must be replaced with a complex assertion that includes in its predicate the Start and Stop columns.
- All queries must be examined, and most must be modified. Consider the query “List those employees who are not managers.” A where predicate is now required to extract the *current* managers. Also, any query that mentions ‘*’ must be modified, because the Employee table now has a different number of columns.
- Modifications must also be altered to take into account the Start and Stop columns. The modification “Change the manager of the tools department to Bob” is now quite more involved than before.

In contrast, assume that instead of attempting to emulate the time-varying aspect using conventional tables, we use a temporal data model that is provably temporally upward compatible with SQL-92 . We would be assured that *not a single line* of our 50,000-line application would have to be altered when transformation \mathcal{T} was applied to render the Employee table time-varying.

4 Temporal Query Languages

As we just saw, the fact that an emulation of temporal tables using SQL-92 is not temporally upward compatible has several unfortunate ramifications in practice. We now turn to the temporal extensions to SQL that have been defined to date. Following an overview of the evaluation, we consider each temporal SQL in turn.

4.1 Overview of Temporally Extended SQL’s

We are aware of seven temporal data models that extend SQL. We consider each of these in turn, starting with the earliest models, examining whether or not each model satisfies the requirements UC and TUC with respect to some variant of SQL, e.g., SQL-89 , SQL-92 , SQL3 , or SQL dialects of commercial DBMSs.

Ideally, we prefer to be able to independently prove that a particular temporal data model satisfies or violates a requirement. However, the available documentation of the models often is not adequately comprehensive for this to be possible. With two exceptions, only the integration of the temporal query facilities with “core” subsets of SQL are documented, and which particular SQL dialect that is being extended is also not always mentioned. This makes it hard to determine whether models are (temporal) upward compatible with “the” full SQL or some subset of “an” SQL.

Aspects related to the use of regular SQL statements—modifications, in particular—on temporal tables or a combination of temporal and non-temporal tables are typically not defined. This makes it hard to verify temporal upward compatibility.

Finally, the definition of the syntax of several of the models is quite informal and incomplete. The semantics of the models are, at best, informal and, at worst, indicated by a few examples.

For the cases where we cannot prove that a temporal data model satisfies or violates a requirement, we will report the model as satisfying (or violating) a requirement if its designers claim that the property is satisfied and we have not been able to disprove the claim with the available documentation. In addition, we will report satisfaction simply if we cannot prove dissatisfaction, again given the available documentation. Thus, we associate the following numbers with our findings, to indicate the confidence in the findings.

1. Neither satisfaction nor violation is claimed, nor can be proven.
2. Satisfaction claimed, but the claim cannot be proven nor invalidated.
3. Independently proven.

Clearly, the highest level of confidence is desired.

| Language | Reference | UC | TUC | Comments |
|-----------|----------------------------|------------------|--|---|
| TOSQL | [Ari86] | yes ² | no ³ | Extends, however, only a subset of SQL. |
| TSQL | [NA87] [NA89] [NA93] | yes ² | no ³ | Not all snapshot tables can be made temporal. Some SQL views cannot be defined on temporal tables. Automatic coalescing violates TUC. |
| HSQL | [Sar90b] [Sar93] | yes ² | no ³ | SQL queries on temporal tables return temporal tables. |
| TempSQL | [BG93] [GB93] [GN93] | yes ² | yes ³ (classical) no ³ (system) | Only a subset of SQL is considered. TUC is satisfied only for classical users. Means of specifying user types and defaults are not given. |
| IXSQL | [Lor93] [LM96] | yes ² | no ³ | Extension of SQL with a parameterized interval ADT with accompanying query-language facilities is proposed. |
| ChronoSQL | [Boe94] | yes ³ | no ³ | Temporal concepts of Datalog are carried over to SQL. Does not restrict TUC queries to the current state. |
| TSQL2 | [Sno95] | yes ³ | no ³ | Full syntax given. Semantics defined informally in SQL-standard style. |

Table 1: Summary of UC and TUC Compliance

Table 1 gives an overview of our conclusions. The first three models are documented rather sparsely for our purposes, but their designers emphasize that they satisfy upward compatibility. They do not satisfy temporal upward compatibility. The next model, TempSQL, introduces a concept of different types of users that may be used to obtain satisfaction of both compatibilities in certain circumstances. The next model is different from all the other models in that it does not provide support for implicit time; rather, it adds a parameterized abstract interval data type and associated facilities for modification and queries to SQL. ChronoSQL is one of the newer temporal models. The final model has been documented much more extensively than its predecessors, but its semantics are still given in an informal SQL-standards format. Note that no language satisfies both notions of upward compatibility.

4.2 TOSQL

TOSQL [Ari86] temporally extends a subset of an early version of SQL [AC75]. The extension is based on the TODM data model. The syntax of TOSQL is given in a BNF-like format. This syntax does not include modification statements, integrity constraints, nested queries, and queries involving aggregates using HAVING, etc. Hence, it appears that TOSQL is upward compatible with a *subset of SQL*, and is perhaps upward compatible with the full language.

It appears that the designer had a notion of temporal upward compatibility in mind when he wrote the following.

“The default options are defined such that a query that omits the temporal portion retains the standard meaning of the corresponding SQL SELECT statement.” [Ari86, p. 513]

An example two pages later states the interpretation of a conventional SQL SELECT statement “is to specify that the query relates to *current* assignments, and uses the most up-to-date data about it.” [Ari86, p. 515]. The “current assignments” refers to *now* in valid time; the “most up-to-date data” refers to *now* in transaction time.

The key phrase though is “that omits the temporal portion”. The timestamp of a table in TOSQL appears as a column named RT. A non-time-varying table would not have such a column. The conversion operator \mathcal{T} in Definition 2.2 would add this column. The problem is with queries involving ‘*’. Such queries on $\mathcal{T}(db_S)$ would return a different number of columns than queries directly on db_S . Hence, temporal upward compatibility is not satisfied.

4.3 TSQL

Navathe and Ahmed’s *temporal relational model*, TSQL, supports, in addition to conventional tables, row timestamping for valid time by attaching two mandatory timestamp columns, *Time-start* (Ts) and *Time-end* (Te) to every time-varying relational schema [MNA87, NA87, NA89, NA93]. These timestamp columns correspond to the lower and upper bounds of time intervals in which rows are continuously valid.

It is stated that TSQL is upward compatible with SQL.

“All legal SQL statements are also valid in TSQL, and such statements have identical semantics in the absence of a reference to time. [...] SQL, a subset of TSQL, remains directly applicable to non-time-varying relations in 1NF.” [NA93, p. 99].

A simplified, 1.5 page BNF-like syntax is given for TSQL [NA87]. Statements such as updates, inserts, deletes, and view definitions are not addressed in the syntax or elsewhere in the documentation. Also, the use of regular SQL queries on temporal tables is not touched upon. While this makes it hard to examine the satisfaction of TUC, there are several indications that TUC is not satisfied.

In TSQL’s data model, only tables that are in the so-called time normal form are allowed [NA87, p. 116]. Briefly, for a table to be in time normal form, it must be in Boyce-Codd normal form (disregarding the timestamp columns), and the non-key, non-timestamp columns must all be synchronous (i.e., they must change values simultaneously). As there are no such normal form requirements on snapshot tables, it follows that the \mathcal{T} operator that turns a snapshot table into a temporal table is not defined for all snapshot tables. Also, regular SQL view definitions on temporal tables are not allowed when they lead to views that are not in time normal form. This is often the case for views that are joins.

Lastly, TSQL performs automatic coalescing of *value-equivalent* rows (i.e., rows with identical non-timestamp column values) that have consecutive or overlapping timestamps. This facility leads to a violation of TUC. For example, assume that we start out with an empty snapshot table, R, and insert two identical

rows. Then `SELECT * FROM R` yields two rows. Now, we simultaneously insert the two rows into $\mathcal{T}(R)$. The most reasonable assumption is that these two rows will be given timestamps that result in them being coalesced into one row. Now, `SELECT * FROM $\mathcal{T}(R)$` yields one row.

4.4 HSQL

As the previous data model, Sarda's HDBMS also supports valid time; however, unlike the data model mentioned previously, HDBMS represent valid time in a valid-time table as a single non-atomic, implicit column [Sar90b, Sar93]. HSQL² is the query language of HDBMS.

It is emphasized that HSQL is upward compatible with respect to SQL (SQL-89, in fact).

“HSQL is a superset of the popular query language SQL.” [Sar93, p. 123]

“In fact, the standard clauses of SQL have identical meanings in HSQL.” [Sar93, p. 125]

Concerning TUC, the effects of the standard SQL insert, delete, and update statements are consistent with satisfying this requirement. However, a query `SELECT * FROM R` where R is a temporal table returns R and not the current (snapshot) state of R , as would be required in order to satisfy TUC [Sar93, pp. 126–127].

4.5 TempSQL

Gadia's TempSQL is based on a N1NF temporal data model that is value timestamped [BG93, GB93, GN93]. A column of a row may have more than one (timestamped) value. The union of the timestamps of the values of each column must be the same for all columns throughout the entire row, resulting in a homogeneous temporal table.

Conventional tables are seen as temporal tables valid at a single time instant. Thus, each column value of each row in such a temporal table is timestamped with the same instant. Integration of snapshot tables into the data model this way is proposed partly in order to obtain upward compatibility.

“By integrating it into our framework, we establish a smooth bridge for industry and its user community for migrating from classical databases to temporal databases. [...] We provide a framework for a smooth transition for industry, requiring no loss of investment in application programs developed by its user community.” [GN93, p. 32]

The particular SQL that is being extended is not identified. No BNF is given. Further, only a subset of those facilities normally associated with SQL are mentioned, with several important aspects, e.g., advanced query facilities, integrity and embedded queries, ignored. With these reservations, it is our contention that TempSQL is upward compatible with SQL. Determining whether temporal upward compatibility is satisfied is more difficult for this model than any of the other models.

TempSQL supports several types of users, e.g., system users and classical users, of a temporal DBMS. While system users have unrestricted access to the database, classical users can only access the currently valid values in the database. Thus, classical users see the current snapshots of temporal tables. Assuming that T is a temporal table, the query `SELECT * FROM T` returns T when issued by a system user and the current snapshot of T when issued by a classical user.

The absence of language syntax for specifying user types at the level of individual statements leads us to assume that, as indicated by the name, user types are fixed for individual users, and on a per-applications basis. (No information is given on how the mechanisms for different types of users interact with embedded

²In another paper, Sarda gave this extension to SQL the name TSQL [Sar90a]. We use HSQL because it was used in the most recent paper.

application programs.) Had the intention been to be able to designate individual language statements as classical or temporal, we feel that the language should have provided syntax for this. We thus think about user types as being similar to ordinary SQL privileges. This seems reasonable, as user types do restrict access to data.

The choice of the default user type matters. If all users, and thus applications, are classical by default, then it is possible to avoid modifying the legacy applications when transitioning to a TempSQL system. Having the default user type be system leads to a violation of temporal upward compatibility—legacy applications then need to be modified to indicate that they are classical.

The next issue to consider is that of the application of legacy SQL modification statements on temporal tables. As the effects of such statements persist in the current states (i.e., the states of the temporal tables valid at the (ever-increasing) current time), the statements are consistent with TempSQL satisfying temporal upward compatibility.

Our conclusion is that for classical users, temporal upward compatibility is ensured. For system users, the opposite is true. The reason is that, for a system user, a conventional SQL query over a temporal table will return a temporal table.

TempSQL is thus fine when a non-temporal application is executed on a database that has been migrated to a temporal DBMS. Where TempSQL falls short is in further migration of that application, to exploit the very useful temporal constructs of that language. This requires that the user be a system user, because a classical user is not permitted to use any of the new constructs. As soon as the user transitions from classical to system, all of the query language statements in the application must be reevaluated, and many must be substantially rewritten. Had temporal upward compatibility been ensured for all users, this jarring transition would have been much smoother.

4.6 IXSQL

IXSQL [Lor91, Lor93, LM96] differs from all the other temporal query languages in that it does not provide support for a special, built-in notion of time. Rather, IXSQL adds the ability to define columns of a parameterized interval abstract data type, and it provides special query facilities for manipulating tables with rows that have such interval values.

Actually, there exists at least two different versions of IXSQL, an early version [Lor91], and a later version [LM96]. The initial version was neither upward nor temporally upward compatible with SQL, in part because it did not permit duplicate rows in tables.

“IXSQL actually differs from the standard SQL [reference to SQL-89], in that a relation may not contain duplicate tuples.” [Lor91, p. 4]

In the remainder, we consider the later version. This version was designed to be upward compatible with SQL-92:

“IXSQL is syntactically and semantically upwards consistent with SQL2.” [LM96, p. 1]

Next, we consider temporal upward compatibility. The first step is to decide on what the meaning of \mathcal{T} should be in a model without an implicit notion of time in its tables. To be specific, let us simply assume that \mathcal{T} adds an interval-valued column to each snapshot table, with value [CURRENT_DATE, DATE '9999-12-31'] for each row. Other reasonable assumptions seem to lead to the same conclusions. The result of a legacy query such as `SELECT * FROM R` will differ from the result of `SELECT * FROM \mathcal{T} (R)`. In addition, legacy modifications to “temporal” tables will generally not be consistent with satisfying temporal upward compatibility, or they may fail altogether. In summary, legacy applications need to be rewritten when new columns are added to the tables then access.

4.7 ChronoSQL

ChronoSQL was designed and implemented as part of the ChronoLog project [Boe94]. The main purpose was to illustrate how temporal concepts developed for deductive databases can be carried over to relational databases. ChronoSQL is tightly coupled with a Datalog based language, which means that users can switch language any time.

This said it comes as no surprise that not all language features of ChronoSQL have been worked out in detail. Specifically, the temporal extension was restricted to query statements; data manipulation statements and integrity constraints were not considered. Moreover, legacy queries over temporal tables are not restricted to the current state. This clearly violates temporal upward compatibility.

Upward compatibility looks more promising. ChronoSQL adds a couple of non-mandatory syntactic constructs to SQL. No other syntactic changes are proposed. This ensures syntactic upward compatibility. Furthermore, the semantics of legacy statements over nontemporal tables remains unchanged [Boe94, p.69] meaning that upward compatibility is ensured as well.

4.8 TSQL2

TSQL2 [Sno95] is the most comprehensively documented temporal query language. Its syntax was given as an extension of the syntax of SQL-92 as presented in the official standard, and the semantics of TSQL2 was also given in the format of the SQL-92 standard. Some 500 pages of technical commentaries accompany these specifications. Upward compatibility of TSQL2 is studied in [BJS95].

In TSQL2, there are six kinds of tables: snapshot tables, valid-time event tables, valid-time state tables, transaction-time tables, bitemporal event tables, and bitemporal state tables. The first is the kind of table found in the relational model; the remaining five are temporal tables. 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, and mentions only snapshot tables 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 [SAA⁺94] 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 [Sno95] 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, TSQL2 is upwards compatible with SQL-92. However, TSQL2 is not temporally upward compatible with SQL-92, for several reasons.

The two related concepts of value-equivalent rows and duplicates will prove important. The former concept applies only to temporal tables; the latter applies to both temporal tables and to timeslices of temporal tables. To illustrate the interrelations among these concepts, consider the temporal tables depicted in Figure 3. Here, we assume rows are timestamped with periods.

Table r_1 contains no duplicates and no value-equivalent rows. Thus, no timeslices of r_1 will contain duplicates. Table r_2 contains no duplicates, but it does contain value-equivalent rows. However, as the timestamps of the value-equivalent rows are disjoint, no timeslices will contain duplicates. Table r_3 , like r_2 , contains value-equivalent rows and contains no duplicates. Unlike in r_2 , the timestamps of the value-equivalent rows are not disjoint, and thus there are timeslices of r_3 that contain duplicates. Finally, table r_4 contains duplicates and thus non-disjoint value-equivalent rows, leading again to timeslices with duplicates.

| r_1 | | | | | | |
|---|-----------|---|-------|-----------|-------|-----------|
| <table border="1" style="border-collapse: collapse; width: 60px; height: 40px;"> <tr><th style="width: 20px;">A</th><th style="width: 40px;">T</th></tr> <tr><td>a_1</td><td>[10 – 20)</td></tr> <tr><td>a_2</td><td>[15 – 50)</td></tr> </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: 60px; height: 40px;"> <tr><th style="width: 20px;">A</th><th style="width: 40px;">T</th></tr> <tr><td>a_1</td><td>[10 – 17)</td></tr> <tr><td>a_1</td><td>[17 – 20)</td></tr> <tr><td>a_2</td><td>[15 – 50)</td></tr> </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: 60px; height: 40px;"> <tr><th style="width: 20px;">A</th><th style="width: 40px;">T</th></tr> <tr><td>a_1</td><td>[10 – 20)</td></tr> <tr><td>a_1</td><td>[15 – 18)</td></tr> <tr><td>a_2</td><td>[15 – 50)</td></tr> </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: 60px; height: 40px;"> <tr><th style="width: 20px;">A</th><th style="width: 40px;">T</th></tr> <tr><td>a_1</td><td>[10 – 20)</td></tr> <tr><td>a_1</td><td>[10 – 20)</td></tr> <tr><td>a_2</td><td>[15 – 50)</td></tr> </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 Rows and Duplicates

Note that allowing value-equivalent rows does not necessarily yield duplicates in timeslices. However, if we want to have duplicates in timeslices, we must allow (non-disjoint) value-equivalent rows.

One reason that TSQL2 is not temporal upward compatible with SQL-92 is that SQL-92 tables that contain duplicates have no counterparts in TSQL2 where tables with value-equivalent rows (and thus duplicates, either in a timeslice, or in the temporal table itself) are not allowed. The definition of temporal upward compatibility (with SQL-92) requires that for every SQL-92 table r , there must exist an equivalent TSQL2 table $r^v = \mathcal{T}(r)$. 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: Let a SALARY table record (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.

| Name | Amount |
|------|--------|
| Tom | 1200 |
| Tom | 1200 |
| Tom | 800 |

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

| Name | Amount | T |
|------|--------|--------------------|
| Tom | 1200 | [1994/5 – 1995/3) |
| Tom | 1200 | [1994/8 – 1994/12) |
| Tom | 800 | [1994/11 – 1995/6) |

This table records that from May 1994 to March 1995, Tom was on one payroll and made a monthly salary of 1200; from August 1994 to December 1994 he was on another payroll where he also made 1200 per month; and from November 1994 to June 1995 he made 800 in a third job. This is not a legal TSQL2 table because it contains value-equivalent rows. □

The merit of duplicates has already been discussed heatedly (see, e.g., [Dat95, p. 109]). Doubtless, SQL-92 would be cleaner in a mathematical sense without duplicates. However, we cannot change SQL-92, so whether we like it or not, it is necessary to deal with duplicates when designing a temporally upward compatible successor to SQL-92. Specifically, for TSQL2 to be temporally upward compatible with respect to SQL-92, it must support (the equivalent of) tables containing value-equivalent rows with non-disjoint timesamps, permitting duplicates in timeslices.

As a reminder, we note that duplicates may significantly impact the results of queries. For example, the following statement computes a table that associates with every person that person's total salary.

```
SELECT Name, SUM(Amount)
FROM Salary
GROUP BY Name
```

Evaluated over the initial nontemporal `Salary` table, the query computes Tom’s salary to be \$3200. Without duplicates, the result would have been \$2000, which is unintended.

A second reason that TSQL2 is not temporally upward compatible with SQL-92 is that when the keyword `SNAPSHOT` is *not* specified in a select statement in TSQL2, a temporal table results. Hence, an SQL-92 query over a temporal table will result not in a conventional table, but rather in a temporal table.

5 Ensuring Temporal Upward Compatibility

This section explains a sequence of steps that lead to a temporal upward compatible SQL-92 extension. Implications to syntax and semantics are discussed and illustrated with examples. Temporal upward compatible extensions allow to independently migrate data structures and application code. Specifically, it permits migration of data structures without also requiring changes to application code (c.f. Definition 2.2). The examples that have been stated in prose in Section 2 are reconsidered and formulated in the temporal extension of SQL-92.

5.1 Syntax of a Temporal Upward Compatible Extension of SQL

Temporal upward compatibility does not put an upper limit on syntactic extensions to a language. It, however, defines a lower limit. First, all legacy statements must be retained. (This requirement is independently established by upward compatibility.) Second, a possibility must be provided to migrate nontemporal data structures to temporal data structures. The first requirement is met by adding (non-mandatory!) syntactic constructs to the base language. No syntactic constructs may be deleted or changed. Migrating non-temporal to temporal data structures can be achieved in different ways. We discuss two possibilities to illustrate the design space and the possible consequences to the data model.

If we want to emphasize different table types (snapshot tables, valid time tables, transaction time tables, and bitemporal tables) a reasonable syntactic choice is to extend the `<alter table action>` production of SQL-92 [MS93, p.511], by adding two options.

```

<alter table action> ::=  <add column definition>
                        |  <alter column definition>
                        |  <drop column definition>
                        |  <add table constraint definition>
                        |  <drop table constraint definition>
                        |  <add time dimension>
                        |  <drop time dimension>

```

```

<add time dimension> ::=  ADD <time dimension>

```

```

<drop time dimension> ::=  DROP <time dimension> <drop behavior>

```

```

<time dimension> ::=  VALID
                    |  TRANSACTION

```

Adding valid time turns a snapshot table into a valid time table and a transaction time table into a bitemporal table. Adding transaction time turns a snapshot table into a transaction time table and a valid time table into a bitemporal table. This is the approach chosen by TSQL2 [Sno95].

If instead we want to emphasize the conventional relational data model with tables that support time through special-purpose columns, an alternative approach would be to enhance the productions <add column definition> and <drop column definition> respectively.

```
<add column definition> ::= ADD [ COLUMN ] <column definition>
                        |   ADD [ COLUMN ] <time dimension>

<drop column definition> ::= DROP [ COLUMN ] <column name> <drop behavior>
                        |   DROP [ COLUMN ] <time dimension> <drop behavior>
```

Further syntactic alternatives can also be envisioned. It is, however, critical that all of them support the semantics discussed in the next section.

5.2 Semantics of a Temporal Upward Compatible Extension of SQL

This section discusses the semantics of various temporally upward compatible statement categories, i.e., standard SQL-92 statement categories evaluated over temporal databases. The categories include queries, views, assertions, column constraints, referential integrity constraints, insertions, deletions, and updates. This ensures a broad coverage of the functionality of a database system. Nevertheless, there are certain statement categories that are not considered explicitly, e.g., triggers. These categories do not introduce fundamentally new problems with respect to temporal upward compatibility. Instead, semantics and techniques discussed for other categories can be applied directly.

When we discuss the semantics of legacy statement categories over temporal tables we can differentiate between *non-destructive statements*, e.g., queries, views, and integrity constraints, and *modification statements*, e.g., data manipulation statements. As we will see, these two sets of categories have to be treated differently.

Below we discuss the semantics for each of the two sets of categories. Within each set all categories are analyzed and illustrated with an example.

The very first step is of course to migrate the data structures.

```
ALTER TABLE Employee ADD VALID
ALTER TABLE Salary ADD VALID
```

Both tables are turned into valid-time tables, such that all information stored in the tables can be annotated with its valid time (transaction time is discussed at the end of this section).

5.2.1 Non-destructive Statements

Non-destructive statements retrieve from or check parts of the database. They do not change the contents of the database. To get the exact same semantics that a nontemporal database would provide, we have to restrict the retrieval and checking to the current state.

Queries are implemented by adding a selection condition to the WHERE clause that selects current rows. Moreover, defaults, e.g., ‘*’ in the select clause, may not expand to include time. As an example, assume a query that determines who manages the high-salaried employees. The ‘temporal’ query is straightforward.

```
SELECT Manager
FROM   Salary AS S, Employee AS E
WHERE  S.Name = E.Name
AND    S.Amount > 3500
```

Whenever the temporal database system identifies one or more temporal in an SQL-92 statement, it must perform the actions dictated by temporal upward compatibility. In this case, it must restrict the set of rows to the current ones.

Views are similar to queries. This becomes obvious if we remember that a view is a virtual table defined by a query. The query that defines the view is enhanced along the lines outlined above. As an example, consider a view that yields high-salaried employees.

```
CREATE VIEW High_salary AS
  SELECT *
  FROM Salary
  WHERE Amount > 3500;
```

A selection condition that limits the query expression to current salaries has to be added. Moreover, the default used in the select clause has to be extended to `SELECT Name, Amount` (or an equivalent relational algebra projection) so that the valid time is not part of the result.

Integrity constraints come in different flavors. The most general form are assertions [MS93, p.211ff]. Consider the assertion that ensures that all employees get a salary, i.e., an assertion that checks that no employees without a salary exist.

```
CREATE ASSERTION CONSTRAINT Emp_has_sal CHECK
  NOT EXISTS ( SELECT *
                FROM Employee AS E
                WHERE NOT EXISTS ( SELECT *
                                   FROM Salary AS S
                                   WHERE E.Name = S.Name ) )
```

The general approach to check an assertion is to negate it and to execute it as a query, i.e.,

```
SELECT *
FROM Employee AS E
WHERE NOT EXISTS ( SELECT *
                   FROM Salary AS S
                   WHERE E.Name = S.Name )
```

If the query result is empty, i.e., if no rows are returned, the assertion is respected; otherwise it is violated. With this background, temporal upward compatible assertions can be achieved easily, because we showed above how to do so with queries.

5.2.2 Modification Statements

Modification statements change the contents of the database. An obvious (but naive) approach is to carry over the semantics from the previous section and to modify the current state. Imagine the insertion of an employee into the database.

```
INSERT INTO Employee VALUES ( 'Liliane', 'Brandt', 'Tools' )
INSERT INTO Salary VALUES ( 'Liliane', 1000 )
```

If we inserted Liliane only in the current state, subsequent queries would not return this row. When we later issue a query, time will have progressed and Liliane will no longer be in the (new) current state. Of course this is not the behavior we expect from a nontemporal database. In order to get the expected behavior, we have to make sure that Liliane remains in the changing current state. This may be achieved by using the period

from `CURRENT_DATE` to `9999-12-31` (the largest `DATE` value) as the timestamp of Liliane's tuples. But it may also be achieved using as the end point `NOBIND(CURRENT_DATE)`, where `NOBIND` has the effect of storing in the timestamp a *variable* that evaluates to `CURRENT_DATE` when accessed, rather than storing the current value of `CURRENT_DATE`. Indeed, any now-relative variable [CDI⁺94] that evaluates to a time between these two end points may be used. We will adopt the simplest choice, the date `9999-12-31`.

An equivalent observation holds for delete and update statements. Assume that we want to change the manager of the tools department to Bob.

```
UPDATE Employee
SET    Manager = 'Bob'
WHERE Dept = 'Tools'
```

If we only updated the current state, subsequent queries would not access the corrected database state. Again, we have to ensure that the update persists in the changing current state to get the exact same behavior a nontemporal database provides.

Achieving temporal upward compatibility for modification statements is slightly more complicated than achieving temporal upward compatibility for non-modification statements. The reason is that certain rows may be valid from some point in the past until some point in the future, i.e., they overlap the current time. Because temporal upward compatible statements only affect the current and future times, the modifications must not change the row during the entire time range. Let us consider each type of modification statement in turn.

Insert statements have to set the valid-time start to the current time and the valid-time end to `DATE '9999-12-31'`, as discussed above. This ensures that, until the row is deleted or modified, it will be valid.

Next we consider delete statements. Historical data, i.e., qualifying rows with a valid time end before the current time, is left untouched. Current data, i.e., qualifying rows with a valid-time start after the current time (including a valid time end equal to `DATE '9999-12-31'`), has to be deleted as of the current time. This is done by changing valid time end to the current time. For future knowledge two choices exist. If we decide not to delete it, today's future knowledge will become valid eventually. This behavior can be quite surprising for applications employing temporal upward compatibility exclusively. An alternative is to delete qualifying future knowledge. This ensures a more intuitive behavior of legacy applications, but it might not be the semantics temporal applications envision.

The most complex statements are update statements. First, rows with a valid-time start before the current time and a valid-time end after the current time (including a valid-time end equal to `DATE '9999-12-31'`) are duplicated. The valid-time end of the original row and the valid-time start of the duplicated row are set to the current time. Then the update statement is applied to all rows with a valid time start that is equal or after the current time. Again we have the choice not to update future knowledge (c.f. previous paragraph).

5.2.3 Transaction Time

With respect to temporal upward compatibility, transaction time behaves almost identically to valid time. Exactly the same semantics applies to transaction-time tables and valid-time tables.

Even bitemporal tables behave quite similarly. In non-destructive statements and insertions, both time dimensions inherit the unitemporal semantics. Deletions and updates are somewhat more complicated, due to the nature of transaction time which guarantees that at each point in time, it is possible to reconstruct previous database states. A temporal upward compatible deletion of a bitemporal row triggers the following steps.

1. Qualifying rows with a transaction-time end equal to 9999-12-31 are duplicated. The transaction-time end of the original row and the transaction-time start of the duplicated row are set to the current time.
2. The valid-time deletion is applied to qualifying rows with a transaction-time end equal to 9999-12-31.

The first step saves the current state and thus ensures reconstructability, whereas the second step performs the valid-time deletion. Update follows a similar pattern.

5.3 Implementing Temporal Upward Compatibility

Annotating information with time leaves us with several well-known possibilities at the physical level, most prominently row and column-value timestamping. It is important that, after a sequence of modifications, the current state of the database is equivalent to the state of a corresponding nontemporal database after the same sequence of modifications. We assume row timestamping and extend both tables with a `VALID` column of type `PERIOD`.

SQL-92 queries over temporal tables must restrict the set of rows to the current ones, and then must eliminate the timestamps from their results. This can be done at different levels. The highest level is SQL, which is widely accessible. The downsides are restricted optimization possibilities and limited expressive power. A lower level is the algebra level which permits more control. For example we could rewrite

$$\pi_{employee.name,city}(\text{employee}_{salary.name=employee.name}(\sigma_{amount>3500}(salary)))$$

to

$$\pi_{employee.name,city}(\sigma_{VALID \text{ overlaps } CURRENT_TIMESTAMP}(employee) \\ salary.name=employee.name(\sigma_{amount>3500}(\sigma_{VALID \text{ overlaps } CURRENT_TIMESTAMP}(salary))))$$

The selection condition `VALID overlaps CURRENT_TIMESTAMP` checks whether the valid time overlaps the current time.

We showed how integrity constraints can be mapped to queries, with the above implementation options applicable. However, in general database systems have to improve the sketched integrity checking mechanism to achieve acceptable performance. Well-known techniques are incremental consistency checking, simplification of assertions, and special-purpose checking algorithms for, e.g., column constraints and referential integrity constraints.

6 Applicability

It is instructive to consider the applicability of temporal upward compatibility in more detail. When designing larger information systems, two general approaches have been advocated. In the first approach, the system design is based on the *function* of the enterprise that the system is intended for (the “Yourdon” approach [You82]); in the second, the design is based on the *structure* of the reality that the system is about (the “Jackson” approach [Jac83]). It has been argued that the latter approach is superior because structure may remain stable when the function changes while the opposite is more often not possible. Thus, a more stable system design, needing less maintenance, is achieved when adopting the second design principle. This suggests that the data needs of an enterprise are relatively stable and only change when the actual business of the enterprise changes.

Enterprises currently use non-temporal database systems for database management, but that does not mean that enterprises manage only non-temporal data. Indeed, temporal databases are currently being managed in a wide range of applications, including, e.g., academic, accounting, budgeting, financial, insurance,

inventory, legal, medical, payroll, planning, reservation, and scientific applications. Temporal data may be accommodated by non-temporal database systems in several ways. For example, a pair of explicit time columns may encode a valid-time interval associated with a row.

Temporal database systems offer increased user-friendliness and productivity, as well as better performance, when managing temporal data. The typical situation, when replacing an underlying non-temporal DBMS with a temporal DBMS, is one where the enterprise is not changing its business, but wants the extra support offered by the temporal system for managing its temporal data. Thus, it is atypical for an enterprise to suddenly desire to record temporal information where it previously recorded only snapshot information. Such a change would be motivated by a change in the business.

The typical situation is rather more complicated. The non-temporal database system is likely to already manage some temporal data, which is encoded using snapshot tables, in an ad hoc manner. Other data will be non-temporal, recording the current state, again using snapshot tables. When adopting the new temporal DBMS, upward compatibility guarantees that it is not necessary to change the database schema or application programs. However, without changes, the benefits of the added temporal support are also limited. Only when defining new tables or modifying existing applications, can the new temporal support be exploited. The enterprise then gradually benefits from the temporal support available in the system.

Nevertheless, the concept of temporal upward compatibility is still relevant, for several reasons. First, it provides an appealing intuitive notion of a temporal table: the semantics of queries and modification are retained from snapshot tables; the only difference is that intermediate states are also retained. Second, in those cases where the snapshot table contained no historical information, temporal upward compatibility affords a natural means of migrating to temporal support. In such cases, not a single line of the application need be changed when the table is altered to be temporal. Third, snapshot tables that do contain temporal information and that have been converted to temporal tables can still be queried and modified by conventional SQL-92 statements in a consistent manner.

7 Conclusion

Upward compatibility aids in the smooth migration of applications from a conventional to a temporal data model. The definitions introduced here allow a specific temporal language to be evaluated as to the degree that it ensures upward compatibility. The extant temporal extensions to SQL are all deficient in one or more ways, rendering migration more difficult. We subsequently showed how SQL-92 can be extended to yield a temporal data model satisfying all three notions of upward compatibility. Applications can be much more easily migrated to this new data model.

The notion of temporal upward compatibility can be viewed as a form of logical data independence. In the same way that an external schema can ensure that applications are not impacted by changes to the logical schema, temporal upward compatibility ensures that applications are not impacted by a specific kind of change to the logical schema: adding or removing temporal support. Logical data independence is an important benefit provided by modern data models, in particular by the relational data model, and the specific kind discussed here provides similar advantages.

The approach we espouse here to providing temporal upward compatibility relative to SQL was adopted in the SQL/Temporal proposals [SBJ96a, SBJ96b]. These language constructs were explicitly designed to ensure upward compatibility *and* temporal upward compatibility with the entire SQL-92 standard. The constructs have been proposed to the American ANSI and international ISO SQL committees for inclusion into the next ISO SQL standard.

Several directions for further research are promising. First, there is a need for exploring different implementation alternatives for upward compatible temporal SQL extensions. Alternatives range from stand-alone implementations to implementations that maximally reuse the functionality offered by existing DBMS's with

an SQL interface. Second, it is felt that much could be learned from conducting actual case studies of the migration of legacy applications to temporal platforms. Third, the transition from explicit to implicit temporal knowledge should be investigated. Strategies must be designed to assist the user in migrating nontemporal tables with explicit time columns to temporal tables. This is essential to maximally exploit the capabilities of temporal database systems.

8 Acknowledgments

Christian S. Jensen was supported in part by the Danish Natural Science Research Council, grants 11-0089-1, and 9400911. Richard Snodgrass was supported in part by NSF grants ISI-9202244 and ISI-9632569 and by a grant from DuPont.

References

- [AC75] M.M. Astrahan and D.D. Chamberlin. Implementation of a structured English query language. *Communications of the ACM*, 18(10):580–588, October 1975.
- [Ari86] G. Ariav. A temporally oriented data model. *ACM Transactions on Database Systems*, 11(4):499–527, December 1986.
- [BG93] G. Bhargava and S. K. Gadia. Relational database systems with zero information loss. *IEEE Transactions on Knowledge and Data Engineering*, 5(1):76–87, feb 1993.
- [BJS95] M. H. Bohlen, C. S. Jensen, and R. T. Snodgrass. Evaluating the completeness of TSQL2. In S. Clifford and A. Tuzhilin, editors, *Recent Advances in Temporal Databases*, pages 153–174, Zurich, Switzerland, September 1995. Proceedings of the International Workshop on Temporal Databases, Springer Verlag.
- [Boe94] M. H. Boehlen. *The Temporal Deductive Database System Chronolog*. PhD thesis, Departement Informatik, ETH Zurich, 1994.
- [CDI⁺94] J. Clifford, C. Dyreson, T. Isakowitz, C. S. Jensen, and R. T. Snodgrass. On the semantics of “now” in temporal databases. Technical Report R-94-2047, Aalborg University, Department of Mathematics and Computer Science, Aalborg, Denmark, November 1994.
- [Cod70] E. F. Codd. A relational model of data for large shared data banks. *Communications of the ACM*, 13(6):377–387, June 1970.
- [Dat95] C. J. Date. *Relational Database Writings 1991–1994*. Addison-Wesley Publishing Company, 1995.
- [GB93] 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, TX, June 1993.
- [GN93] S. K. Gadia and S. Nair. *Temporal Databases: A Prelude to Parametric Data*, chapter 2, pages 28–66 of [TCG⁺93].
- [Jac83] Michael A. Jackson. *System Development*. Prentice-Hall International Series in Computer Science. Prentice-Hall International, Inc., 1983.

- [LM96] N.A. Lorentzos and Y.G. Mitsopoulos. SQL extension for interval data. *IEEE Transactions on Knowledge and Data Engineering*, 1996. To appear.
- [Lor91] N. A. Lorentzos. Query language for the management of interval and temporal data. TR 52, Agricultural University of Athens, Greece, 1991.
- [Lor93] N. Lorentzos. *The Interval-extended Relational Model and Its Application to Valid-time Databases*, chapter 3, pages 67–91 of [TCG⁺93].
- [MNA87] N. G. Martin, S. B. Navathe, and R. Ahmed. Dealing with temporal schema anomalies in history databases. In P. Hammersley, editor, *Proceedings of the Conference on Very Large Databases*, pages 177–184, Brighton, England, September 1987.
- [MS93] J. Melton and A. R. Simon. *Understanding the New SQL: A Complete Guide*. Morgan Kaufmann Publishers, Inc., San Mateo, CA, 1993.
- [NA87] S. B. Navathe and R. Ahmed. TSQL—a language interface for history databases. In *Proceedings of the Conference on Temporal Aspects in Information Systems*, pages 113–128, France, may 1987. AFCET.
- [NA89] S. B. Navathe and R. Ahmed. A temporal relational model and a query language. *Information Sciences*, 49:147–175, 1989.
- [NA93] S. Navathe and R. Ahmed. *Temporal Extensions to the Relational Model and SQL*, chapter 4, pages 92–109 of [TCG⁺93].
- [OS95] G. Özsoyoğlu and R. Snodgrass. Temporal and real-time databases: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 7(4), August 1995.
- [SAA⁺94] R. T. Snodgrass, I. Ahn, G. Ariav, D.S. Batory, J. Clifford, C.E. Dyreson, R. Elmasri, F. Grandi, C. S. Jensen, W. Kafer, 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. *ACM SIGMOD Record*, 23(1):65–86, March 1994.
- [Sar90a] N. Sarda. Algebra and query language for a historical data model. *The Computer Journal*, 33(1):11–18, February 1990.
- [Sar90b] N. Sarda. Extensions to SQL for historical databases. *IEEE Transactions on Knowledge and Data Engineering*, 2(2):220–230, June 1990.
- [Sar93] N. Sarda. *HSQL: A Historical Query Language*, chapter 5, pages 110–140 of [TCG⁺93].
- [SBJS96a] R. T. Snodgrass, M. H. Böhlen, C. S. Jensen, and A. Steiner. Adding valid time to SQL/Temporal. Change Proposal ANSI X3H2-96-151r1, ISO/IEC JTC1/SC21/WG3 DBL MCI-142, ISO, May 1996.
- [SBJS96b] R. T. Snodgrass, M. H. Böhlen, C. S. Jensen, and A. Steiner. Adding transaction time to SQL/Temporal. Change Proposal ANSI X3H2-96-152r1, ISO/IEC JTC1/SC21/WG3 DBL MCI-143, ISO, May 1996.
- [Sno95] R. T. Snodgrass (editor), I. Ahn, G. Ariav, D.S. Batory, J. Clifford, C.E. Dyreson, R. Elmasri, F. Grandi, C. S. Jensen, W. Kafer, N. Kline, K. Kulkarni, T.Y.C. Leung, N. Lorentzos, J.F. Roddick, A. Segev, M.D. Soo, and S.M. Sripada. *The Temporal Query Language TSQL2*. Kluwer Academic Pub., 1995.

- [TCG⁺93] A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, and R. Snodgrass (eds.). *Temporal Databases: Theory, Design, and Implementation*. Database Systems and Applications Series. Benjamin/Cummings, Redwood City, CA, 1993.
- [TK96] V. J. Tsotras and A. Kumar. Temporal database bibliography update. *SIGMOD Record*, 25(1):11, March 1996.
- [TL82] D.C. Tschritzis and F.H. Lochovsky. *Data Models*. Software Series. Prentice-Hall, 1982.
- [You82] E. Yourdon. *Managing the System Life Cycle*. Yourdon Press, 1982.